HP OpenView Storage Data Protector Administrator's Guide

Manual Edition: October 2005

Manufacturing Part Number: B6960-90106

Release A.05.50

Legal Notices

Contents

Contents

Contents

Contents

7. Copying Data

8. Restore

Contents

11. Managing the Data Protector Internal Database

12. Disaster Recovery

Contents

Contents

Contents

Contents

Printing History

The manual printing date and part number indicate its current edition. The printing date will change when a new edition is printed. Minor changes may be made at reprint without changing the printing date. The manual part number will change when extensive changes are made.

Manual updates may be issued between editions to correct errors or document product changes. To ensure that you receive the updated or new editions, you should subscribe to the appropriate product support service. See your HP sales representative for details.

Table 1 **Edition History**

Part Number	Manual Edition	Product
B6960-90078	May 2003	Data Protector Release A.05.10
B6960-90106	October 2004	Data Protector Release A.05.50
PDF format only	January 2005	Data Protector Release A.05.50 with patches DPWIN_00113, PHSS_32330, PHSS_32344, and DPSOL_00131.
PDF format only	May 2005	Data Protector Release A.05.50 with patch DPWIN_00137.
PDF format only	August 2005	Data Protector Release A.05.50 with patch DPWIN_00153.
PDF format only	September 2005	Data Protector Release A.05.50 with patches PHSS_33192, and DPSOL_00165.

Table 1 Edition History

Part Number	Manual Edition	Product
PDF format only	November 2005	Data Protector Release A.05.50 with patch DPWIN_00169.

Conventions

The following typographical conventions are used in this manual.

Table 2

Convention	Meaning	Example
Italic	Book or manual titles, and manual page names	Refer to the *HP OpenView Storage Data Protector Integration Guide* for more information.
	Provides emphasis	You *must* follow these steps.
	Specifies a variable that you must supply when entering a command	At the prompt type: `rlogin` *`your_name`* where you supply your login name.
Bold	New terms	The Data Protector **Cell Manager** is the main ...
`Computer`	Text and items on the computer screen	The system replies: `Press Enter`
	Command names	Use the `grep` command ...
	File and directory names	`/usr/bin/X11`
	Process names	Check to see if `Data Protector Inet` is running.
	Window/dialog box names	In the `Backup Options` dialog box...
	Text that you must enter	At the prompt, type: `ls -l`
Keycap	Keyboard keys	Press **Return**.

Data Protector provides a cross-platform (Windows and UNIX) graphical user interface.

Figure 1 Data Protector Graphical User Interface

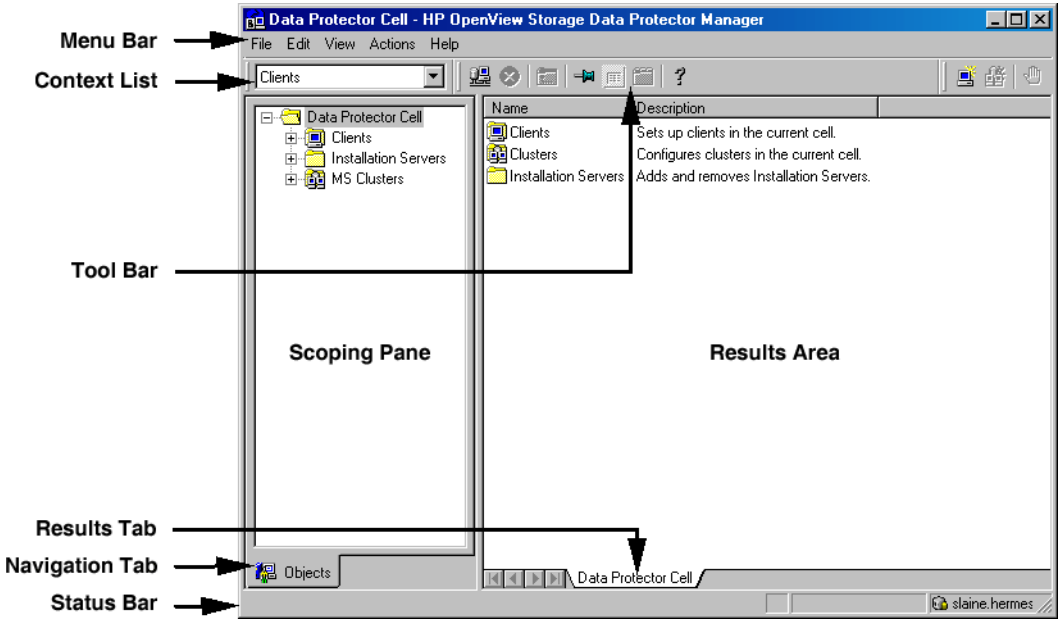

Contact Information

General Information

General information about Data Protector can be found at

http://www.hp.com/go/dataprotector

Technical Support

Technical support information can be found at the HP Electronic Support Centers at

http://support.openview.hp.com/support.jsp

http://www.hp.com/support

Information about the latest Data Protector patches can be found at

http://support.openview.hp.com/patches/patch_index.jsp

For information on the Data Protector required patches, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

HP does not support third-party hardware and software. Contact the respective vendor for support.

Documentation Feedback

Your comments on the documentation help us to understand and meet your needs. You can provide feedback at

http://ovweb.external.hp.com/lpe/doc_serv/

Training Information

For information on currently available HP OpenView training, see the HP OpenView World Wide Web site at

http://www.openview.hp.com/training/

Follow the links to obtain information about scheduled classes, training at customer sites, and class registration.

Data Protector Documentation

Data Protector documentation comes in the form of manuals and online Help.

Manuals

Data Protector manuals are available in printed format and in PDF format. Install the PDF files during the Data Protector setup procedure by selecting the `User Interface` component on Windows or the `OB2-DOCS` component on UNIX. Once installed, the manuals reside in the *<Data_Protector_home>*`\docs` directory on Windows and in the `/opt/omni/doc/C/` directory on UNIX. You can also find the manuals in PDF format at http://ovweb.external.hp.com/lpe/doc_serv/

HP OpenView Storage Data Protector Concepts Guide

This manual describes Data Protector concepts and provides background information on how Data Protector works. It is intended to be used with the task-oriented *HP OpenView Storage Data Protector Administrator's Guide*.

HP OpenView Storage Data Protector Administrator's Guide

This manual describes typical configuration and administration tasks performed by a backup administrator, such as device configuration, media management, configuring a backup, and restoring data.

HP OpenView Storage Data Protector Installation and Licensing Guide

This manual describes how to install the Data Protector software, taking into account the operating system and architecture of your environment. This manual also gives details on how to upgrade Data Protector, as well as how to obtain the proper licenses for your environment.

HP OpenView Storage Data Protector Integration Guide

This manual describes how to configure and use Data Protector to back up and restore various databases and applications. It is intended for backup administrators or operators. There are four versions of this manual:

- *HP OpenView Storage Data Protector Integration Guide for Microsoft Applications: SQL Server 7/2000, Exchange Server 5.x, Exchange Server 2000/2003, and Volume Shadow Copy Service*

This manual describes the integrations of Data Protector with the following Microsoft applications: Microsoft Exchange Server 2000/2003, Microsoft Exchange Server 5.x, Microsoft SQL Server 7/2000, and Volume Shadow Copy Service.

- *HP OpenView Storage Data Protector Integration Guide for Oracle and SAP*

  This manual describes the integrations of Data Protector with Oracle, SAP R3, and SAP DB.

- *HP OpenView Storage Data Protector Integration Guide for IBM Applications: Informix, DB2, and Lotus Notes/Domino*

  This manual describes the integrations of Data Protector with the following IBM applications: Informix, IBM DB2, and Lotus Notes/Domino.

- *HP OpenView Storage Data Protector Integration Guide for Sybase, Network Node Manager, and Network Data Management Protocol*

  This manual describes the integrations of Data Protector with Sybase, Network Node Manager, and Network Data Management Protocol.

HP OpenView Storage Data Protector Integration Guide for HP OpenView

This manual describes how to install, configure, and use the integration of Data Protector with HP OpenView Service Information Portal, HP OpenView Service Desk, and HP OpenView Reporter. It is intended for backup administrators. It discusses how to use the OpenView applications for Data Protector service management.

HP OpenView Storage Data Protector Integration Guide for HP OpenView Operations for UNIX

This manual describes how to monitor and manage the health and performance of the Data Protector environment with HP OpenView Operations (OVO), HP OpenView Service Navigator, and HP OpenView Performance (OVP) on UNIX.

HP OpenView Storage Data Protector Integration Guide for HP OpenView Operations for Windows

This manual describes how to monitor and manage the health and performance of the Data Protector environment with HP OpenView Operations (OVO), HP OpenView Service Navigator, and HP OpenView Performance (OVP) on Windows.

HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide

This manual describes Data Protector zero downtime backup and instant recovery concepts and provides background information on how Data Protector works in a zero downtime backup environment. It is intended to be used with the task-oriented *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* and the *HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide*.

HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide

This manual describes how to configure and use the integration of Data Protector with HP StorageWorks Virtual Array, HP StorageWorks Enterprise Virtual Array, EMC Symmetrix Remote Data Facility and TimeFinder, and HP StorageWorks Disk Array XP. It is intended for backup administrators or operators. It covers the zero downtime backup, instant recovery, and the restore of filesystems and disk images.

HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide

This manual describes how to configure and use Data Protector to perform zero downtime backup, instant recovery, and standard restore of Oracle, SAP R/3, Microsoft Exchange Server 2000/2003, and Microsoft SQL Server 2000 databases. The manual also describes how to configure and use Data Protector to perform backup and restore using the Microsoft Volume Shadow Copy Service.

HP OpenView Storage Data Protector MPE/iX System User Guide

This manual describes how to configure MPE/iX clients and how to back up and restore MPE/iX data.

HP OpenView Storage Data Protector Media Operations User's Guide

This manual provides tracking and management of offline storage media. It is intended for network administrators responsible for maintaining and backing up systems. It describes the tasks of installing and configuring the application, performing daily media operations and producing reports.

HP OpenView Storage Data Protector Software Release Notes

This manual gives a description of new features of HP OpenView Storage Data Protector A.05.50. It also provides information on supported configurations (devices, platforms and online database integrations, SAN, and ZDB), required patches, and limitations, as well as known problems and workarounds. An updated version of the supported configurations is available at http://www.openview.hp.com/products/datapro/spec_0001.html.

Online Help

Data Protector provides context-sensitive (F1) Help and Help Topics for Windows and UNIX platforms.

In This Book

The *HP OpenView Storage Data Protector Administrator's Guide* describes how to configure and use the Data Protector network backup product. You must properly install Data Protector before you can configure it.

NOTE

This manual describes Data Protector functionality without specific information on particular licensing requirements. Some Data Protector functionality is subject to specific licenses. The related information is covered in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Audience

This manual is intended for network administrators responsible for maintaining and backing up systems on the network.

Conceptual information can be found in the *HP OpenView Storage Data Protector Concepts Guide,* which is recommended to fully understand the fundamentals and the model of Data Protector.

Organization

The manual is organized as follows:

1 Introducing Data Protector

In This Chapter

This chapter contains some general principles on how Data Protector works, covered in these sections:

The Data Protector Cell Environment

The Data Protector **cell** is a network environment containing a **Cell Manager**, **clients**, and **backup devices**. The Cell Manager has the main Data Protector control software installed and is the central point from which the cell is administered and backup and restore operations are controlled. Systems that are to be backed up can be added to the cell and set up as Data Protector clients. When Data Protector performs a backup of data from these clients, it saves the data to media (such as magnetic tapes, or hard disks) contained within backup devices.

The **Data Protector Internal Database (IDB)** keeps track of the files backed up, making it is easy to browse and restore them, either singly or collectively.

The **Cell Manager** is the main control center for the cell and contains the IDB. It runs the core Data Protector software and the Session Manager, which starts and stops backup and restore sessions and writes session information to the IDB.

Any system within a chosen cell environment can be set up as a Data Protector **client**. Essentially, a client is a system that can be backed up, a system connected to a backup device with which the backup data can be saved, or both. The role of the client depends on whether it has a Disk Agent or a Media Agent installed.

A client that will be backed up using Data Protector must have a **Disk Agent** installed. Data Protector controls the access to the disk. The Disk Agent lets you back up information from, or restore information to, the client system.

A client system with connected backup devices must have a **Media Agent** installed. This software controls the access to the backup device. A Media Agent controls reading from and writing to a backup device's media.

A **backup device** performs the actual recording of backup data to a recording medium, and the retrieval of restore data from a medium.

The physical object upon which the data is recorded, such as a DAT tape or a hard disk, is called the backup **medium**.

NOTE For further information on these terms, or on the principles of Data Protector operation, see the *HP OpenView Storage Data Protector Concepts Guide*.

How a Backup Session Works

A backup session starts either when a backup is requested through the user interface, or when a scheduled backup is initiated. During this session, Data Protector backs up the requested filesystems and disks to the specified media.

1. The Cell Manager determines the type of session that has been requested (backup) and starts the appropriate Session Manager.
2. The Session Manager reads the backup specification and determines what needs to be backed up and which devices to use.
3. The Session Manager then starts a Media Agent for each media drive that will be used and a Disk Agent for each disk that will be read.
4. The Monitor window appears. This window lets you respond to mount requests and view the progress of a backup session.
5. The Disk Agents start sending data to a Media Agent.
6. If more than one Disk Agent is used, the Disk Agents send data to a Media Agent concurrently and a Media Agent places the data on the medium.
7. As each block of data is written to the medium, a Media Agent sends information to the Session Manager about what has been backed up. The Session Manager uses this information to update the catalog of backed-up files in the IDB.

How a Restore Session Works

A restore session starts when a restore is requested. During this session, Data Protector performs a restore of requested files and disks from the media.

1. You specify which filesystems to restore and how to restore them, using the Data Protector user interface.

2. The Cell Manager determines the type of session that has been requested (restore), and starts the appropriate Session Manager.
3. The Session Manager then determines which filesystems or directories to restore, which devices to use, and what restore options have been specified.
4. The Session Manager starts the appropriate Disk Agent and Media Agent. For example, a Media Agent is started for the media (tape) drive that will be used and a Disk Agent is started for the disk to which the data will be restored.
5. The `Monitor` window appears. This window lets you respond to mount requests and view the progress of a restore session.
6. A Media Agent starts sending data to the Disk Agent.
7. The Session Manager then updates the IDB and the Disk Agent writes the data to the disk.

Using the Data Protector User Interface

There is one Data Protector user interface, available on Windows and UNIX platforms. It consists of the Data Protector graphical user interface (GUI) and the command-line interface.

Using the Data Protector user interface, you can perform all Data Protector tasks.

Figure 1-1 HP OpenView Storage Data Protector Graphical User Interface

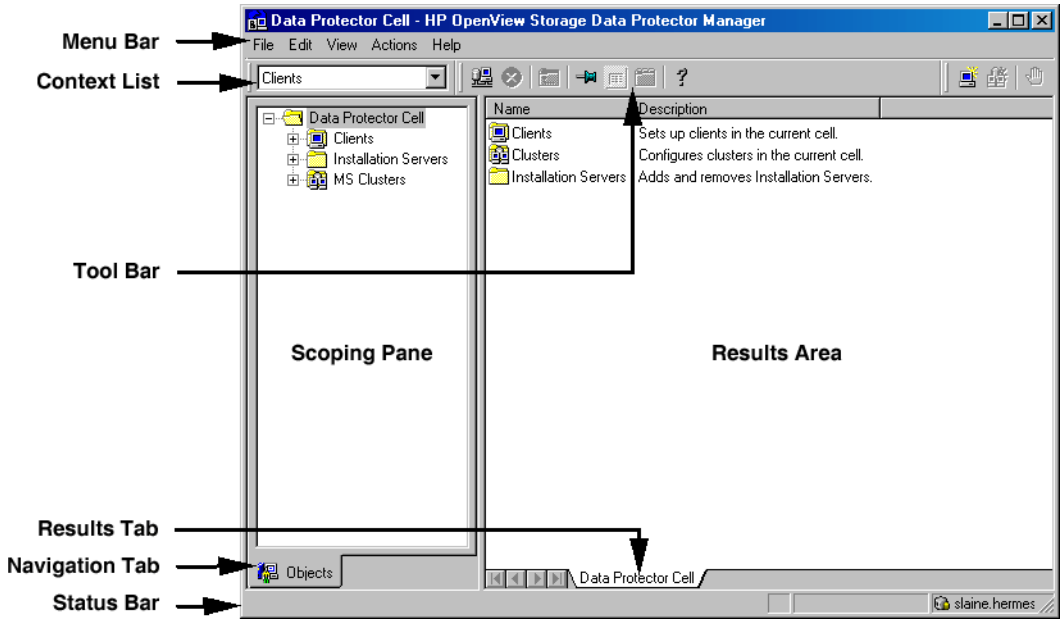

Graphical User Interface

The Data Protector graphical user interface (GUI) uses features such as buttons and text boxes to enhance usability. Whenever possible, drop-down lists are provided to allow you to select from a list instead of typing in your selection. In addition, a comprehensive online Help system provides information about each window and each task.

Depending on the user rights, you can either use the GUI to access the complete Data Protector functionality or to access only specific contexts.

For more information on user rights, refer to “Data Protector User Rights” on page 129.

For more information on Data Protector contexts, refer to “Context List” on page 9.

For more information on internationalization and localization settings, refer to “Settings for the File Name Encoding in GUI” on page 622.

Starting GUI on Windows Platforms

To start the Data Protector GUI on Windows platforms, do one of the following:

- Click `Start` on the Windows desktop and click `Data Protector Manager` from the `HP OpenView Storage Data Protector` program group to start the GUI for the complete Data Protector functionality.
- Use the `manager` command to start the GUI for the complete Data Protector functionality.

  Context-specific options for this command enable you to start one or more Data Protector contexts. For example, the command
  `manager -backup -restore`
  starts the Data Protector `Backup` and `Restore` contexts.

  To specify the Cell Manager you want to connect to, use the following command: `manager -server <Cell Manager_name>`.

  For more information on these commands, refer to the *`omnigui`* man page.

Starting GUI on UNIX Platforms

To use the Data Protector GUI functionality on those UNIX Cell Manager platforms where the Data Protector GUI is not supported, use the `omniusers` command to create a remote user account on the Cell Manager. You can then use the created user account on any other system with the Data Protector GUI installed to start the GUI and connect to the Cell Manager. Refer to the `omniusers` man page for details, and to

the *HP OpenView Storage Data Protector Software Release Notes* for details on supported operating system versions/releases for the user interface. For more information on local language support and the usage of non-ASCII characters in file names, refer to *HP OpenView Storage Data Protector Administrator's Guide*.

For all other UNIX systems where the Data Protector GUI is supported, enter:

xomni	to start the GUI with the complete Data Protector functionality
xomniadmin	to start the administration (configuration) of clients, users, reports, and the IDB GUI
xomnibackup	to start the backup GUI
xomnicellmon	to start the MoM cell monitoring GUI
xomnicopy	to start the copy GUI
xomnimm	to start the media and devices management GUI
xomnimonitor	to start the monitoring a single cell GUI
xomnirestore	to start the restore GUI
xomniinstrec	To start the instant recovery GUI. A special license is needed to start this GUI. Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide* and *HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide* for more information on the instant recovery functionality and to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information on Data Protector licenses.
xomnimom	to start the Manager-of-Managers GUI

For more information on these commands, refer to the `omnigui` man page.

Printing from the Data Protector Graphical User Interface

Data Protector lets you print from the GUI. You can print session messages, reports, event logs, and various lists (for example, lists of configured clients and devices). Generally, you can print anything displayed as a list in the Results Area, and the online Help topics. However, you are not able to print any of the `Properties`. Instead, you

can use the Data Protector reporting functionality to configure various reports about your backup environment. For more information on reporting, refer to “Data Protector Reporting” on page 388.

Prerequisite

You must have a printer already configured on your system.

When you click `Print` on HP-UX, you can choose among predefined printers. Note that if you do not have a proper printer driver installed, you are not able to print. In that case, choose a PS printer and select the `Print to file` option. You can then send the generated file to the PS printer using the UNIX `lp` command from the UNIX terminal.

On Windows, however, a displayed printer in the `Select Printer` window means that the printer is already configured on your system and you are able to print.

For detailed steps on printing, refer to online Help, index keyword “printing from GUI”.

Elements of the Data Protector Graphical User Interface

For the visual representation of the GUI elements, refer to Figure 1-1 on page 6.

Context List

The **Context List** is a drop-down list, from which you can select the management contexts described below:

`Clients`	Controls all of the client systems in the current Data Protector cell. You can add, delete, and monitor any client within the cell.
`Users`	Adds and removes users, user groups and their rights.
`Devices & Media`	Controls device and media maintenance and access to media which store the data.
`Backup`	Controls which data is to be backed up, where, and how.
`Copy`	Controls which data is to be copied, where, and how.
`Monitor`	Allows you to monitor sessions that are in progress.

`Restore`	Controls which data is to be restored, where, and how.
`Instant Recovery`	Controls the split mirror instant recovery processes. A special license is needed to display this context. Refer to the *HP OpenView Storage Data Protector HP StorageWorks Disk Array XP Integration Guide* for more information on the instant recovery functionality and to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information on Data Protector licenses.
`Reporting`	Allows you to get information on your cell configuration, backup specifications, media and media pools, as well as on specific sessions and objects.
`Internal Database`	Allows you to get information on the IDB storage capacity, database objects, and sessions.

Scoping Pane

The **Scoping Pane** provides a tree of items that can be selected to open a view. Selecting an item in the Scoping Pane displays information in the **Results Area**.

Results Area

Selecting an item in the Scoping Pane displays corresponding information in the **Results Area**. If you click `Clients` in the Scoping Pane, the Results Area displays a list of all the clients within your cell.

Navigation Tabs

Navigation Tabs appear at the bottom of the Scoping Pane. These tabs allow you to switch between the two possible item list views in the Scoping Pane: **Objects** and **Tasks**. Not every Scoping Pane has these views.

Tab Name *What the Tab Displays in the Scoping Pane*

`Objects`	A hierarchical presentation of data, similar to the directory tree in Windows Explorer. For example, in the Devices & Media context, the Scoping Pane will display the list of devices and media configured with Data Protector.
`Tasks`	A list of tasks that you can perform. Clicking a task displays a wizard that will walk you through an entire task, such as backing up a file.

Results Tab

The name on the **Results Tab** corresponds to the name of the item currently selected in the Scoping Pane. You can click the Pin icon on the toolbar to make this view "stick" and keep it available for the future. For example, if you need to use the GUI to look up some other information, but you want to continue with the previous view later, you can access this view by selecting the "pinned" tab.

You can remove one or more tabs by right-clicking the area and selecting `Remove Tab` or `Remove Other Tabs`.

The Command-Line Interface

The command-line interface (CLI) follows the standard UNIX format for commands and options and provides complete Data Protector functionality. You can use these commands in scripts to speed up your commonly performed tasks.

The `omniintro` man page lists all supported Data Protector commands, as well as differences between commands on the UNIX and Windows platforms.

See also “Data Protector Commands” on page A-7.

Data Protector Online Resources

Information about Data Protector is available in this manual and in the online Help system. This manual contains the information you need to plan and administer your Data Protector network, and information on some more commonly performed tasks. The online Help system contains the information you need to perform all available tasks.

The following Data Protector online resources are available:

Help Topics	Online Help with task instructions and reference information. You can select topics by using the contents list, index, or search facility.
Help Navigator	Context-sensitive Help that provides detailed help on the current task.
Online Documentation	Online manuals in PDF format that can be read with the Adobe Acrobat Reader.
Data Protector on the Web	Opens your Web browser to the Data Protector homepage, where more information about Data Protector can be found.
Online Support	Opens your Web browser to the HP OpenView interactive Online Support service page.
About	Displays version and copyright information for Data Protector, as well as licensing information.

You can access the online resources by either using the `Help` drop-down menu or the `Help` buttons provided on the Data Protector windows.

Hyperlinks (cross-references) to additional information and definitions help you navigate through online Help. You click the hyperlinked word or phrase to jump to the new topic. Hyperlinked words and phrases are marked with either solid underlining or different color.

Starting and Using the Help Navigator

The Help Navigator provides context-sensitive online Help, which can be used to find information about the current GUI panel or task.

If the GUI concerned is running on Windows, the Help Navigator is dynamic: Once it is started, its contents automatically change as you go to the next page of the wizard or to another view in the Data Protector user interface.

To start the Help Navigator, either:

- Press **F1**
- Click `Help Navigator` from the Help menu, or
- Click the Help Navigator icon (the question mark) on the button bar

Using the Online Manuals

Data Protector provides online manuals in PDF format that can be read using the Adobe Acrobat Reader. Once installed, the online manuals reside in the *<Data_Protector_home>*\docs directory (Windows) or the /opt/omni/doc/C directory (HP-UX or Solaris) on the Cell Manager system.

Using Microsoft Management Console (MMC)

On Windows systems, it is possible to integrate the Data Protector GUI with the Microsoft Management Console.

The `Microsoft Management Console (MMC)` is a Graphical User Interface (GUI) that lets you manage and run your administrative tools within a common interface environment. You can add already installed software, hardware, or network management applications to the console, where the primary type of tool that can be added to the console is called a **snap-in**.

The Data Protector snap-in, known as **OB2_Snap,** provides a basic integration of Data Protector and the MMC. Using OB2_Snap, you can go to the Data Protector home page or to Data Protector Web/Java Reporting. You can also start the Data Protector GUI on Windows from the MMC.

Proceed as follows to add **OB2_Snap** to the MMC.

1. Download the MMC software from http://www.microsoft.com/downloads/.
2. From the Windows desktop, click `Start`, and then select `Run`.

3. In the Open text box, enter mmc to open the Microsoft Management Console window.
4. From the Console menu, select Add/Remove Snap-in. In the Standalone property page of the Add/Remove Snap-in window, click Add.
5. In the Add Standalone Snap-in window, select HP OpenView Storage Data Protector. Click Close to exit the window, then click OK to get back to the Microsoft Management Console window.

   The HP OpenView Storage Data Protector item will be displayed under Console Root. Once you have added the applications to MMC, save the file as <*Console_Name*>.msc .

Overview of Tasks to Set Up Data Protector

Although configuring Data Protector is easy, some advanced planning will help you configure the environment and optimize your backups. This section provides an overview of the global tasks to set up a backup environment.

Depending on the size and complexity of your environment, you may not need to go through all these steps.

1. Analyze your network and organizational structure. Decide which systems need to be backed up. For more information refer to the *HP OpenView Storage Data Protector Concepts Guide*.
2. Check whether there are any special applications and databases which you want to back up, such as Microsoft Exchange, Microsoft SQL, Oracle, SAP R/3, or others. Data Protector provides specific integrations with these products.

   Refer to the *HP OpenView Storage Data Protector Integration Guide* for instructions on how to configure the integrations.
3. Decide on the configuration of your Data Protector cell, such as:
   - the system to be your Cell Manager
   - systems on which you want to install the user interface
   - local backup versus network backup
   - systems to control backup devices and libraries
   - type of connection, LAN and/or SAN
4. Purchase the required Data Protector licenses for your setup. This way you obtain the passwords you will need to install.

   Alternatively, you can operate Data Protector using an instant-on password. However, this is valid only for 60 days from the date of installation. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.
5. Consider security aspects:
   - Analyze security considerations. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

- Consider which user groups you need to configure.

6. Decide how you want to structure your backups:
   - Which media pools would you like to have, and how will they be used?
   - Which devices will be used, and how?
   - How many copies of each backup do you want?
   - How many backup specifications do you want to have, and how should they be grouped?
7. Install the Data Protector Cell Manager and Installation Server(s). Then use the Data Protector GUI to distribute Data Protector agents to other systems. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for detailed instructions.
8. Configure backup devices. See Chapter 2, “Configuring and Using Backup Devices,” on page 17.
9. Configure media pools and prepare the media. See Chapter 5, “Managing Media,” on page 143.
10. Configure backup specifications, including backup of the IDB. See Chapter 6, “Backup,” on page 195.
11. Configure reports, if required. See Chapter 9, “Monitoring, Reporting, Notifications, and the Event Log,” on page 379.
12. Prepare for disaster recovery. See Chapter 12, “Disaster Recovery,” on page 513.
13. Become familiar with tasks such as:
    - Handling failed backups
    - Performing restores
    - Duplicating backed up data and vaulting media
    - Testing disaster recovery
    - Maintaining the IDB

2 Configuring and Using Backup Devices

In This Chapter

This chapter includes information on the following topics:

NOTE Backup devices (like tape drives) are subject to specific Data Protector licenses. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.

Configuring Backup Devices

Preparation of a backup device consists of connecting the device to the system and knowing which of the (working) associated device files (SCSI address) is to be used. To configure a device:

1. Connect the device to a computer. Refer to the documentation that comes with the device.
2. Make sure that you have done the following:

UNIX Systems

Create or find the device filename for a device connected to a UNIX system. For detailed steps, refer to the online Help index keyword “creating device filenames” or “finding device filenames”. For further information, refer to Appendix B of the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Windows Systems

Provide a SCSI address and load the driver that will be used with a device connected to a Windows system.

For tape drives, the Windows native tape driver can be:

- unloaded (preferred) or
- loaded

  The device filename depends on whether a Windows native tape driver is used with a particular tape drive.

On how to obtain the SCSI address, refer to the online Help index keyword “creating SCSI addresses”. For further information, refer to Appendix B of the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Windows Robotics Drivers	On Windows, disable the Removable Storage Service or Windows medium changer (robotics) driver before you configure the robotics device with Data Protector. For detailed steps, refer to the online Help index keyword “robotics drivers”. For further information, refer to Appendix B of the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

3. Boot the system to have the device recognized by the system.
4. Configure the device, as described in the following sections, so that you can use it with Data Protector.
5. Prepare the media that you want to use with your backups. On how to format media, refer to “Formatting Media” on page 155.

   A default media pool is used with each device so that you do not have to create one. If you want to create your own media pool, refer to “Creating a Media Pool” on page 149.

   Figure 2-1 on page 22 shows the relationship between the backup specification, devices, and media pools. The devices are referred to in the backup specification, while each device is linked to a default media pool. This media pool can be changed in the backup specification.

Figure 2-1 How Backup Specifications, Devices, and Media Pools Relate

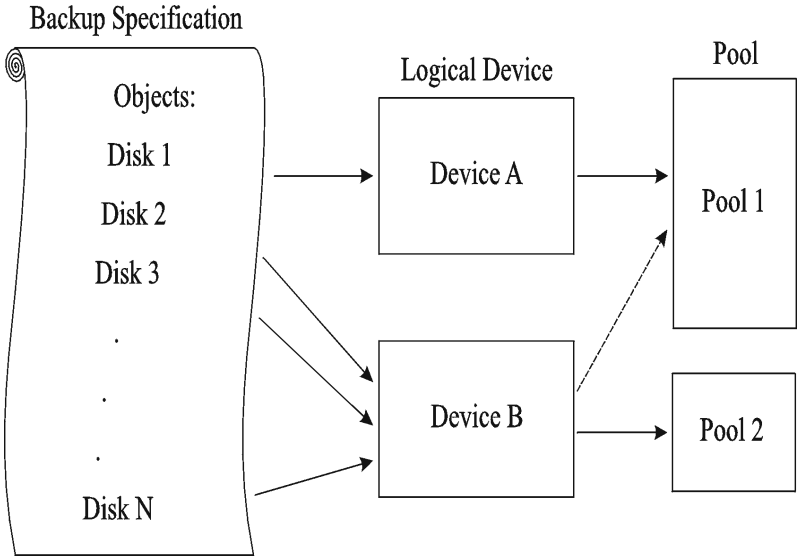

Configuring Standalone Devices

What Are Standalone Devices?

Standalone devices are simple backup devices with one drive that reads from or writes to one medium at a time. They are used for small-scale backups. When a medium is full, the operator has to manually replace it with a new medium so that the backup can proceed. Standalone devices are not appropriate for large, unattended backups.

Data Protector provides simple configuration and management of media used in standalone backup devices.

How to Configure a Standalone Device

Once you have prepared the device for configuration as described in "Configuring Backup Devices" on page 20, configure a standalone device so that you can use it with Data Protector. In the `Devices & Media` context, right-click `Devices` and click `Add Device`. In the `Add Device` wizard, specify the `Standalone` device type. Refer to Figure 2-2.

For detailed steps and examples, refer to the online Help index keyword "configuring standalone devices".

Data Protector supports a specific set of backup devices. Refer to the *HP OpenView Storage Data Protector Software Release Notes* for a detailed list of supported devices and their corresponding media types.

In case you want to use a device that is not in the list of supported devices, refer to "Support of New Devices" on page 48.

TIP

You can also let Data Protector automatically configure most common devices. You still need to prepare the media for a backup session, but Data Protector determines the name, policy, media type, media policy, and the device's SCSI address or device file. For detailed steps, refer to the online Help index keyword "autoconfiguring backup devices".

Figure 2-2 **Specifying Device Type and Name**

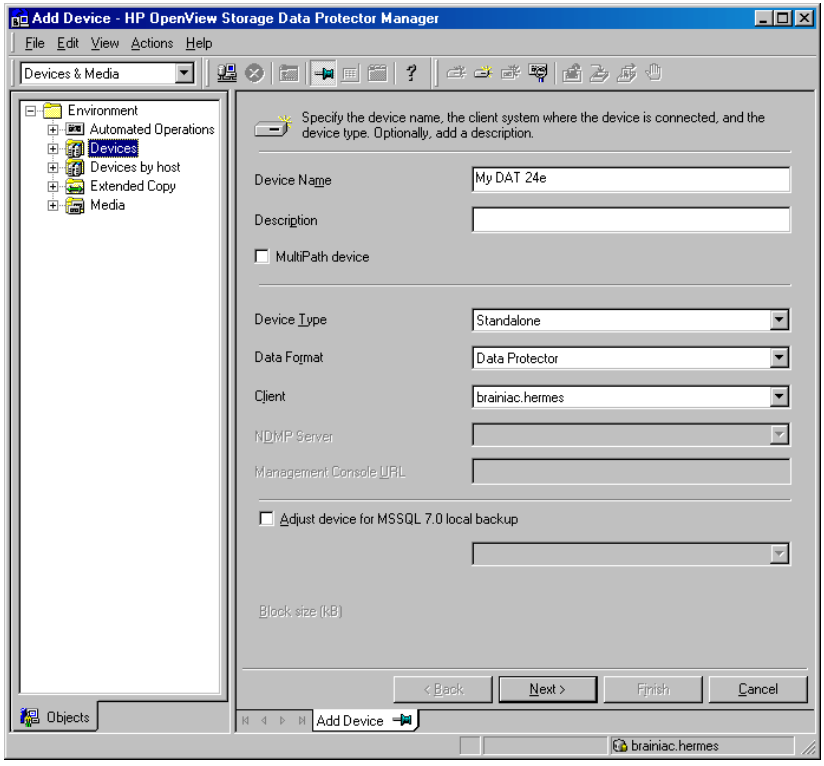

Configuring Device Chains

Data Protector allows you to configure standalone devices of the same type into **device chains**. When a medium in one device becomes full, the backup automatically continues on the medium in the next device in the device chain. Device chains are possible for only one Media Agent, that is, you can connect a device chain to only one system.

The configuration is the same as for a standalone device, except that you enter multiple SCSI addresses (on Windows) or device filenames (on UNIX).

NOTE

The order in which the devices are added determines the order in which Data Protector uses them.

When all of the media in a device chain are full, Data Protector issues a mount request. The operator must replace the medium in the *first device* with a new medium, format it, and then confirm the mount request. Data Protector can immediately use media that are recognized and unprotected. Data Protector can also use blank media, so that you do not have to format them.

Configuring Library Devices

What Are Library Devices?

SCSI library devices, also called autoloaders, are large backup devices. They contain a number of media cartridges in a device's repository and can have multiple drives handling multiple media at the same time. Most library devices also allow you to configure automatic drive cleaning, which is performed by Data Protector when the drive gets dirty. Refer to "Drive Cleaning" on page 82.

A library device has a SCSI ID for each drive in the device, and one for the library robotic mechanism. This mechanism moves media from slots to drives and back again. For example, a library with four drives has five SCSI IDs, four for the drives and one for the robotic mechanism.

Slot Number

Each slot in the device's repository holds one medium. Data Protector assigns a number to each slot, starting from 1. When managing a library, you refer to the slots using their numbers. For example, a library with 48 repository slots has slot numbers 1, 2, 3, 4, 5, 6...47, 48.

Drive Index

The drive index identifies the mechanical position of the drive in the library. Refer to Figure 2-3.

The index number is relevant for the robotics control. The robot knows only index numbers and has no information about the SCSI address of the drive. The drive index is a sequential integer (starting from 1) which has to be coupled with the SCSI address of this drive. For example, for a four-drive library, the drive indexes are 1,2,3,4.

If you have only one drive in the library, the drive index is 1.

Drive SCSI Address

The drive index must match the corresponding SCSI address. This means that you need to configure the pairs as follows:

SCSI address_A for index 1,
SCSI address_B for index 2, and so on.

NOTE It is not necessary to configure all drives for use with Data Protector. You can configure one media pool for all drives, or have an independent media pool for each drive. It is recommended that you use the default media pool when configuring a device.

Figure 2-3 Drive Index to SCSI Address Mapping

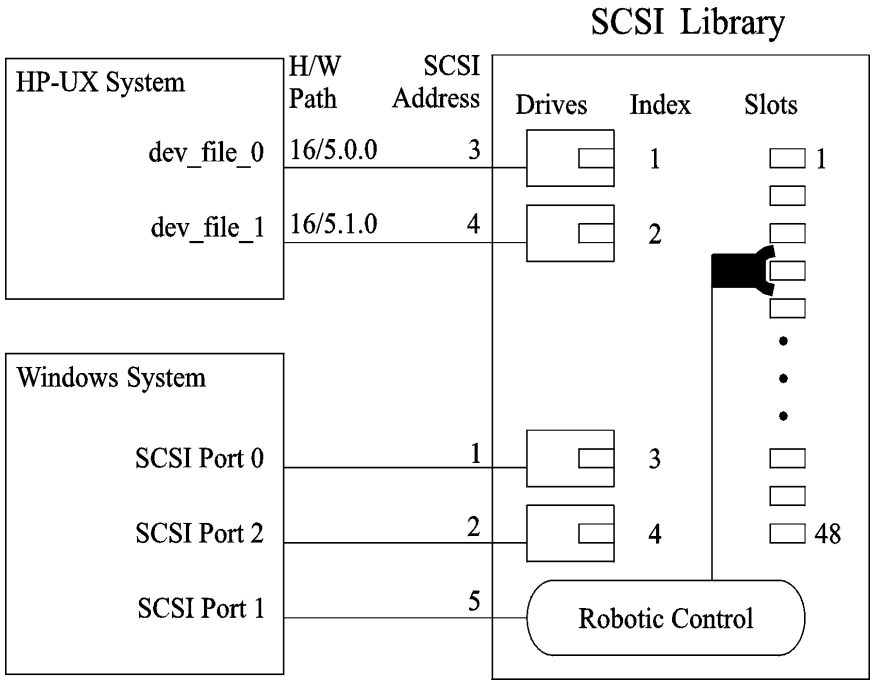

How to Configure a Library Device Once you have prepared the device for configuration as described in "Configuring Backup Devices" on page 20, configure a library device, including its drive(s). The `Add Device` wizard guides you through both configurations. For detailed steps and examples, refer to the online Help index keyword "configuring SCSI libraries".

TIP

You can also have Data Protector automatically configure the library devices for you. You still need to prepare the media for a backup session, but Data Protector determines the name, policy, media type, media policy, and the device file or SCSI address of the device, and also configures the drive and slots. For detailed steps, refer to the online Help index keyword “autoconfiguring backup devices”.

To verify the device configuration, right-click the created drive, and choose `Scan Medium`. If the device is configured correctly, Data Protector will be able to load, read, and unload media in the slots.

What’s Next?

If you have configured all the backup devices you want to use with Data Protector, do the following:

- Add media to the media pools that you will use with the newly configured device. Refer to “Adding Media to a Media Pool” on page 154.
- If you want to configure a cleaning tape, refer to “Drive Cleaning” on page 82.
- If your device uses barcodes, refer to “Activating Barcode Support” on page 87.
- Configure a backup for your data. Refer to Chapter 6, “Backup,” on page 195.

Configuring Libraries with Multiple Systems

You can configure a library so that each drive receives data from a different system running a Data Protector Media Agent (the General Media Agent or the NDMP Media Agent). The library robotics control is still performed by one system that have either the General Media Agent or the NDMP Media Agent installed. This improves performance in high-end environments by allowing local backup, instead of having to move the data over the network.

Prerequisites

- Each client system that you want to use with the drives in the library must have a Data Protector Media Agent (the General Media Agent or the NDMP Media Agent) component installed.
- You need to have connected the backup device to the system, and a working device file (SCSI address) must exist before you can configure the device for use with Data Protector.

For more information on multi-drive support, see the *HP OpenView Storage Data Protector Concepts Guide*.

How to Configure Libraries with Multiple Systems

Configure a library as described in “Configuring Library Devices” on page 26. When you are prompted to configure drives in the library, specify the client system that you want to use with each drive. For detailed steps, refer to the online Help index keyword “configuring libraries for multiple systems”.

TIP

To verify the device configuration, select a range of slots from the library and then click `Scan` from the `Actions` menu. If the device is configured correctly, Data Protector will be able to load, read, and unload media back into the slots.

What’s Next?

If you have configured all the backup devices you want to use with Data Protector, do the following:

- Add media to the media pools that you will use with the newly configured device. See “Adding Media to a Media Pool” on page 154.
- If you want to configure a cleaning tape, see “Drive Cleaning” on page 82.

- If your device uses barcodes, see “Activating Barcode Support” on page 87.
- Configure a backup. See Chapter 6, “Backup,” on page 195.

Configuring Magazine Devices

What Are Magazine Devices?

Magazine devices group a number of media into a single unit called a **magazine**. A magazine allows you to handle large amounts of data more easily than when using many individual media.

Data Protector allows you to perform media management tasks on magazines as sets, or on a single medium.

Prerequisite

Create at least one media pool with the `Magazine Support` option set. See “Adding Media to a Media Pool” on page 154.

How to Configure a Magazine Device

Magazines must be configured as libraries. Select the `SCSI Library` device type in the `Add Device` wizard. The media pool to which magazines belong needs to have the `Magazine Support` option selected. For detailed steps, refer to the online Help index keyword “configuring SCSI libraries”.

TIP

You can also let Data Protector automatically configure your device for you. You still need to prepare the media for a backup session, but Data Protector determines the name, policy, media type, media policy, and the device file or SCSI address of the device, and also configures the drive and slots. For detailed steps, refer to the online Help index keyword “autoconfiguring backup devices”.

To verify the device configuration, right-click the created drive, and then choose `Scan`. If the device is configured correctly, Data Protector will be able to load, read, and unload media in the slots.

What’s Next?

If you have configured all the backup devices you want to use with Data Protector, do the following:

- If you want to configure a cleaning tape, refer to “Drive Cleaning” on page 82.
- If your device uses barcodes, refer to “Activating Barcode Support” on page 87.
- Configure a backup for your data. Refer to Chapter 6, “Backup,” on page 195.

Configuring Stacker Devices

What Are Stacker Devices?

A stacker is a single device with one drive and sequentially accessed media. Using stacker devices requires more human media management than using a small library. Data Protector provides simple configuration and management of media used in stacker backup devices.

How to Configure a Stacker Device?

To create a stacker device, specify the `Stacker` device type in the `Add Device` wizard. For detailed steps, refer to the online Help index keyword “configuring stacker devices”.

Stacker Device Media Management

The operations scan, verify, or format have to be run separately on each medium in a stacker device. When performing these operations, use the `Eject medium after operation` option, in order to have each medium loaded automatically (only the first medium should be loaded manually). When all the tapes in the stacker magazine are used, the magazine must be unmounted manually and the next one inserted.

Stackers load media in sequential order, therefore a `Loose` media allocation policy is recommended. A `Strict` policy would require media to be loaded in the same order as they are to be used.

Example

1. Manually load the first medium.
2. Run format/verify/scan (with `Eject after operation` enabled) -- (next tape will be loaded automatically).
3. Repeat step 2 until all tapes are finished.
4. When all the tapes in the stacker magazine are used, unmount the magazine manually and insert the next one.

NOTE

If a medium is not properly loaded, Data Protector will abort the session.

Backup and Restore with Stacker Devices

Only the first medium has to be manually loaded. When a tape is full, it is ejected and the next tape is loaded automatically. When all the tapes are used in a stacker magazine, the magazine has to be unmounted manually and the next one has to be inserted. Again the first tape has to be loaded manually into the drive.

NOTE A backup or restore session will not be aborted if media are not present, but a mount request will be issued instead. The whole session will not be aborted if a user does not change stacker magazines within a time out period.

Configuring ADIC/GRAU DAS and STK ACS Libraries

This section assumes that you have already physically configured the ADIC/GRAU or STK ACS library. If you have not done so, refer to the documentation that comes with the ADIC/GRAU or STK ACS library for instructions on configuring the library. For a list of supported software versions, refer to *HP OpenView Storage Data Protector Software Release Notes*.

NOTE

The ADIC/GRAU and STK functionality is subject to specific Data Protector licenses. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.

Typically, the Data Protector and ADIC/GRAU DAS or STK ACS libraries are used in complex environments where the amount of backed up data is exceptionally large and, therefore, so is the amount of media needed to store the data. The ADIC/GRAU and STK ACS libraries are not only capable of managing large amounts of media, they are also capable of managing media used by different applications, not just Data Protector.

Data Protector provides full support for the ADIC/GRAU DAS and the STK ACS Library Systems. Since these libraries manage media used by different applications, you have to configure which media you want to use with Data Protector, which media you want to track, and which drives you want to use with Data Protector. The Figure 2-4 on page 35 and Figure 2-5 on page 35 illustrate the two library integrations.

Figure 2-4 **Data Protector and ADIC/GRAU DAS Library**

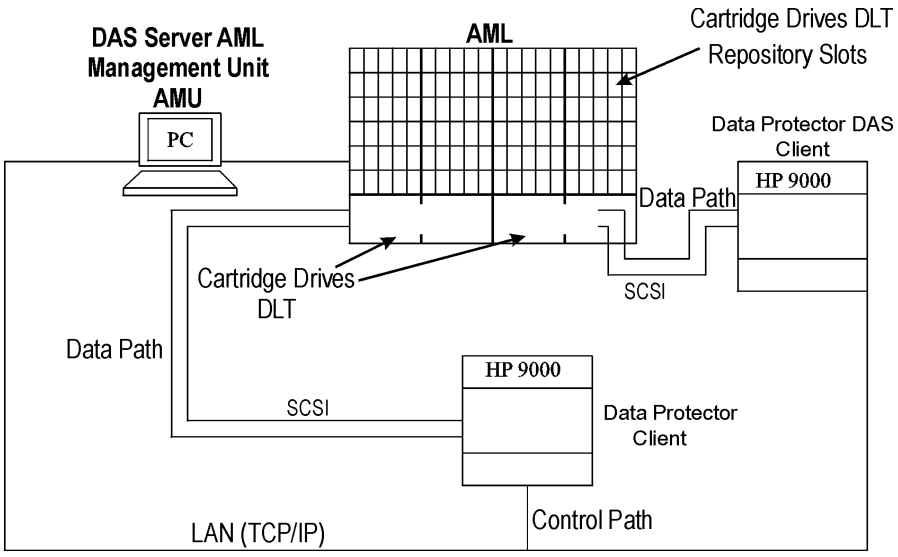

Figure 2-5 **Data Protector and StorageTek ACS Library**

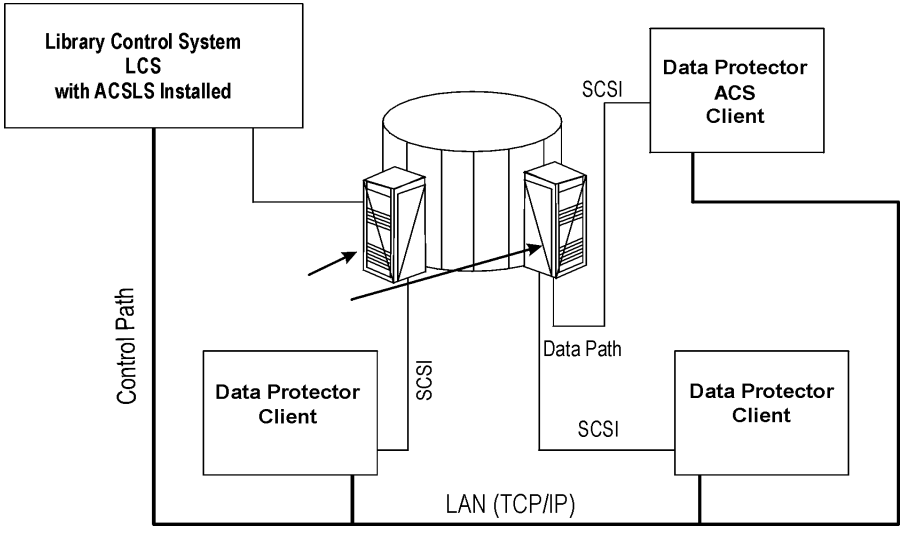

Configuration Basics

A Data Protector Media Agent software (the General Media Agent or the NDMP Media Agent) has to be installed on the client that accesses the library robotics through the GRAU/ADIC DAS or STK ACS Library Server and on the clients that access the drives.

Media Management Basics

In the ADIC/GRAU DAS and STK ACS library devices, a medium is identified by its volume serial number, or volser. The volser, similar to a barcode, uniquely identifies each medium during its life.

Media in the library can be used by many applications, not just by Data Protector, so that you have to know which applications use which media to prevent them from being overwritten.

Ideally, you will use the ADIC/GRAU or ACS library with Data Protector exclusively and let Data Protector manage the complete library, but if you have other applications using the library, you should take care to assign non-intersecting subsets of media to Data Protector and other applications. Also, note that Data Protector does not make use of scratch pools but maintains its own independent media allocation policy. This implies that if a specific medium has been allocated to Data Protector (added to an Data Protector media pool), it remains under Data Protector's control during its lifetime or until it is removed from the Data Protector media pool.

IMPORTANT

For each type of media you have to configure a library in Data Protector. While the ADIC/GRAU or STK ACS system can store many physically different types of media, Data Protector can only recognize a library with a single type of media in it. Therefore you have to create a *logical* Data Protector library for every media type. Refer also to "The Data Protector Query Operation Used with ADIC/GRAU DAS or STK ACS Libraries" on page 38.

The actual physical location of a medium is maintained by the DAS Server (in the ADIC/GRAU library) or the ACS Server (in the STK ACS library), not Data Protector. The DAS or ACS Server tracks the location

using its volser. When a medium is moved around the repository, it is not assigned to the same physical slot each time. Therefore, you cannot rely on the slot number when handling the media, but on the barcode (volser).

For media in the device's repository, Data Protector displays the location as **resident**. For media stored outside the device's repository, Data Protector displays the location as **non-resident**.

NOTE

Data Protector will not overwrite media containing data in a recognizable format. However, Data Protector can not guarantee that Data Protector data on tapes will not be overwritten by some other application using the same media. We recommend that you make sure that media used by Data Protector are not used by any other application, and vice versa. Refer also to "The Data Protector Query Operation Used with ADIC/GRAU DAS or STK ACS Libraries" on page 38.

Tracking Media

Data Protector tracks both Data Protector and non-Data Protector media. For media in a recognizable format, Data Protector displays the format as the media type, such as **tar**. For media in a non-recognizable format, Data Protector displays **foreign** as the media type.

Labeling Media

Data Protector labels each medium used by Data Protector with the unique medium label and medium ID. Both are stored in the IDB and allow Data Protector to manage the medium. The medium ID is assigned by Data Protector, while the medium label is combined from your description and the volser of this medium.

Although you can change the label and exclude the barcode number, this is not recommended. In this case you should manually keep track of the actual barcode and the medium label you assigned to the medium.

Initializing Other Formats

If Data Protector recognizes some other media data format or media that have been used by another application, it will not initialize these media unless the `Force Operation` option is selected. Data Protector recognizes the following data formats and media used by other applications: tar, cpio, Fbackup, FileSys, Ansi and OmniStorage.

Drive Cleaning Support

The ADIC/GRAU DAS and STK ACS libraries can automatically clean their drives after the drive has been used a set number of times. This is not recommended, as library built-in drive cleaning interrupts the

session, causing it to fail. If you want to use the library's cleaning functionality, you have to ensure that drive cleaning is performed when no Data Protector sessions are running.

For more information on drive cleaning methods, refer to "Drive Cleaning" on page 82.

The Data Protector Query Operation Used with ADIC/GRAU DAS or STK ACS Libraries

For more information on the Data Protector query operation, refer to "Scanning Media in a Device" on page 175.

When the Data Protector query operation is started, *all* the media configured on the DAS or ACS Library Server is queried, even in cases when these media are configured in Data Protector as belonging to several *logical* ADIC/GRAU DAS or STK ACS libraries (for the same physical library).

Additionally, the Data Protector query operation queries also the media configured on the DAS or ACS Library Server that are configured to be used with applications other than Data Protector. The consequence is that after the query operation is started from Data Protector, the media belonging to other *logical* ADIC/GRAU DAS or STK ACS libraries than the one for which the query operation was started, are moved to the *logical* ADIC/GRAU DAS or STK ACS library for which the query operation was started.

Therefore, with ADIC/GRAU DAS or STK ACS libraries, it is not recommended to use the Data Protector query operation. It is recommended to add volsers manually using the Data Protector add volsers operation instead of synchronizing the IDB using the Data Protector query operation.

NOTE

The information in this section does not apply in case of ADIC/GRAU DAS libraries, when logical libraries are not configured using Data Protector, but using the ADIC/GRAU DAS utilities. If several logical libraries are configured using the ADIC/GRAU DAS utilities, the Data Protector query operation can safely be used on such libraries.

For information on how to add the volsers manually, refer to "Adding Volsers Manually" on page 189.

Additional Media Management Tips

Remember the following list of tips when you begin to use Data Protector with the GRAU DAS or STK ACS device.

- Create at least one media pool for each media type, for example, one for 4mm and one for 3480 media type. Depending on your environment, you may want to create more media pools, for example, one for each department. See *HP OpenView Storage Data Protector Concepts Guide* for more information on how to plan your media pools.
- Use Data Protector commands to handle media. If you handle media manually using ADIC/GRAU DAS or STK ACS commands, Data Protector will not be able to track the changes in location or information on the media.
- Manage the whole library with Data Protector. This provides single-point administration where you can track Data Protector and non-Data Protector media in the library.
- Make sure that Data Protector and other applications do not use the same set of media.

Installation

Refer to “Installing a Media Agent to Use the ADIC/GRAU Library or the StorageTek Library” section of the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to prepare an ADIC/GRAU or STK ACS library to work with Data Protector. Then continue with the instruction in the following section.

Configuration

Follow the procedure below to configure an ADIC/GRAU or STK ACS library in Data Protector:

1. In the `HP OpenView Storage Data Protector Manager` switch to the `Devices & Media` context. In the Scoping Pane, right-click `Devices` and then click `Add Device`.
2. Optionally, select `MultiPath device`.
3. Enter the `Device Name` and `Description` for the device.

4. In the Device type text box, choose the GRAU DAS Library or the StorageTek ACS Library.

   If the MultiPath device option is not selected, select the name of the Media Agent client that will access ADIC/GRAU or STK ACS robotics.

5. Optionally, enter a valid URL of the library management console in the Management Console URL text box. This will enable you to invoke a web browser and load the management console interface directly from the Data Protector GUI.

   Click Next.

6. If you are configuring an ADIC/GRAU library, enter the hostname of the DAS Server in the DAS Server text box.

   If you are configuring an STK ACS library, enter the hostname of the ACS Library Server in the ACSLM Hostname text box.

   For multipath devices, select also the client name and click Add to add the path to the list of configured paths.

7. In the Busy drive handling drop-down list, select one of the possible actions.

   Click Next.

8. If you are configuring an ADIC/GRAU library, specify the Import and Export Areas for the library.

   If you are configuring an STK ACS library, specify the CAPs for the library.

   Click Next.

9. In the Media Type drop-down list, choose media type for the library.

   Click Finish and then Yes to configure the drives for the library.

10. In the Device Name text box, enter a name for the drive.

    In the Description text box, enter a description for the drive.

    Optionally, select MultiPath Device.

    If the MultiPath device option is not selected, select the name of the Media Agent client that will access ADIC/GRAU or STK ACS robotics.

    Click Next.

11. In `Data Drive` drop-down list, enter the SCSI address of the device.

    For multipath devices, select also the name of the Media Agent client that will access ADIC/GRAU or STK ACS robotics and click `Add` to add the path to the list of configured paths.

    Select `Automatically discover changed SCSI address` to enable automatic discovery of changed SCSI addresses.

    In the `Drive Name` text box, enter the ADIC/GRAU Drive name or STK ACS drive ID you obtained during the installation of a Media Agent as described in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

    Click `Next`.

12. From the `Default Media Pool` drop-down list, select the media pool, which will be the default media pool for the device that you are configuring.

    Select `Advanced Options` to change `Concurrency` and other settings as necessary.

    If needed, set the `Use lock name` option under the `Other` tab. Refer to “Device Locking” on page 93 for more information on this option.

    Click `Finish`.

Create a library for each type of media that you will use with Data Protector.

Configuring a Library for Mixed Media

A mixed media library contains media of several types, such as DLT and magneto-optical. It uses identical robotics to move all the media (regardless of media type) between slots and drives.

In order to use this library functionality, configure several (sub)libraries: one library definition per media type.

To take full advantage of this feature, perform the following steps:

- Configure at least one media pool (or use the default pool) per media type.
- Configure the library robotics once per media type, including the slot range for that media type. Make sure the robotics control (SCSI path on Windows systems or device file on UNIX systems) for each of the library robotic definitions resides on the same host and that they are identical.
- Configure all the drives for a media type and link them to the related library robotic and media pool. Make sure the drive index is unique for each physical device, regardless of media type.

Configuring Multiple Paths to Physical Devices in the SAN Environment

You can configure more than one path, that is the client name and SCSI address (device file on UNIX) to a physical device as a single logical device - multipath device.

For a more detailed description of the multipath concept, refer to the *HP OpenView Storage Data Protector Concepts Guide*.

Manual Configuration

You can enable multiple paths for single physical devices by manually configuring the device using the device configuration wizard.

For detailed steps on how to configure multipath devices, refer to the online Help index keyword “configuring, multiple paths to devices”.

Figure 2-6 Enabling Multipath Devices

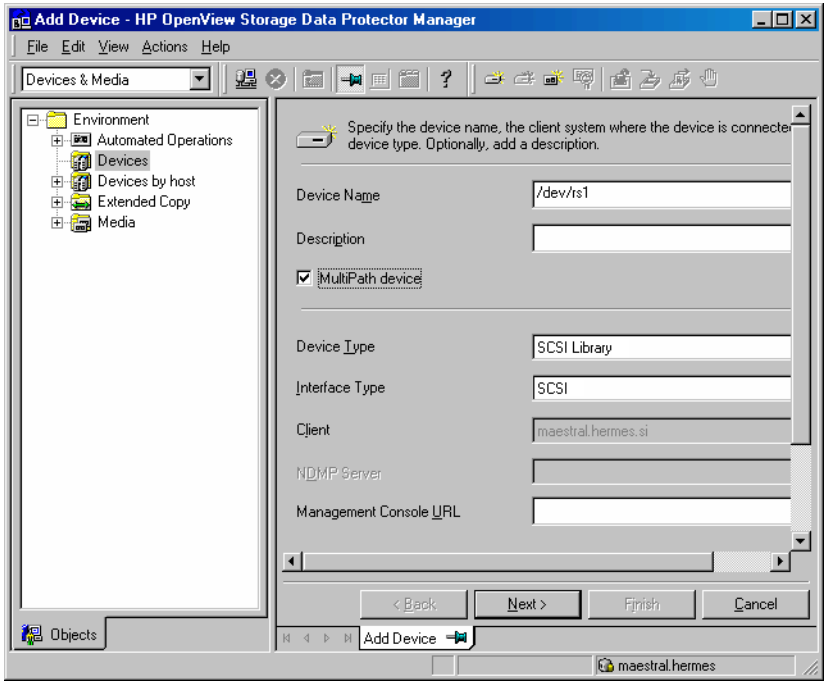

Automatic Configuration

Alternatively, you can configure multipath devices automatically during device autoconfiguration or by using the `sanconf` command:

- By default, the device autoconfiguration wizard will configure all physical devices that are connected to more than one client as multipath devices. For details refer to the online Help index keyword "configuring, multiple paths to devices"
- By default or if the `-no_multipath` option is given, `sanconf` does *not* configure multipath devices – a separate logical device will be configured for each path. To configure all paths pointing to a single physical device as a single multipath device, run `sanconf` with the `-multipath` option.

  For details, refer to "Automatic Configuration of Libraries in a SAN Environment Using the sanconf Command" on page 68 and to the `sanconf` man page.

Limitations

The following limitations apply:

- Multiple paths are not supported for NDMP devices and Jukebox libraries.
- Device chains are not supported for multipath devices.
- When multipath devices are used, the `libtab` file functionality is disabled.

Configuring Devices for Direct Backup

This section provides the configuration steps for backup devices used in a direct backup environment. Refer to the *HP OpenView Storage Data Protector Concepts Guide* for a more detailed information on direct backup concepts.

Direct backup is a Data Protector backup solution in a SAN environment. Please read the section "Shared Devices in the SAN Environment" on page 51 for general information on SAN environments. Note that the direct backup device configuration steps differ from the configuration steps described in the mentioned section, and are given in this section.

A direct backup environment consists of the following:

- a SAN network
- internal or external Fibre Channel bridge(s) (FC bridge)
- backup device(s) connected to FC bridge(s) (standalone or SCSI library)
- physical XCopy engine(s) (present in an FC Bridge)
- disk array(s) assuring point-in-time stability of data (HP StorageWorks Disk Array XP)
- application system(s) connected to the disk array original disk(s)
- backup system(s) connected to the disk array mirror disk(s) and controlling the SCSI library robotics and SCSI library/standalone device drives

An internal FC bridge is embedded in the backup device, whereas an external FC bridge resides at any point in the SAN.

A backup device used in the direct backup environment is identified by the World Wide Name (WWN) of the Fibre Channel bridge that it is connected to or embedded in the backup device, and by the device (standalone device) or drive (SCSI library) Logical Unit Number (LUN) as seen on the SAN. If a SCSI library is used, its robotics does not have to be connected to a FC bridge.

Backup Device Auto-detection

The XCopy engine must reside in the FC bridge to which the backup device or drive is connected (external FC bridge), or in the internal FC bridge. A backup device that is used with direct backup functionality is auto-detected whenever a direct backup session is started. Even if auto-detection is used, the WWN and the LUN parameters must be entered manually when configuring a device; the LUN must be reconfigured every time the LUN changes.

XCopy Engine

There can be more than one physical XCopy engine in a direct backup environment. Each of these physical XCopy engines can have more logical XCopy engines configured and assigned. Which of these logical XCopy engines will be used in a direct backup session is specified in the direct backup specification by specifying the backup device(s) to be used and assigning them a logical XCopy engine. The physical XCopy engine behind the logical XCopy engine specified in the backup specification must be configured for the backup system specified in the backup specification.

The following types of backup devices are supported for a Data Protector direct backup:

- standalone devices
- SCSI libraries

Configuration Procedure

Refer to the following online Help index keywords and perform any necessary steps before configuring backup devices as described later in this section:

- online Help index keyword “preparing backup devices”
- online Help index keyword “configuring direct backup environment”

Configuring a backup device for direct backup consists of the following:

1. Configuring a standalone device or SCSI library.
2. Configuring XCopy engines.
3. If direct library access will be used, configuring the `libtab` file.

Configuring Standalone Devices

Refer to the online Help index keyword “configuring standalone devices for direct backup” for detailed information on how to configure a standalone device for a direct backup.

Configuring SCSI Libraries

Refer to the online Help index keyword “configuring SCSI libraries for direct backup” for detailed information on how to configure a SCSI library for a direct backup.

Configuring XCopy Engines

Refer to the online Help index keyword “configuring XCopy engine” for detailed information on how to configure an XCopy engine.

Configuring the libtab File

Configuration of the `libtab` file is necessary only if direct library access is to be used.

Refer to the “Manually Configuring the libtab Files” on page 63 for detailed information on how to configure the `libtab` file.

Support of New Devices

To use a device that is not listed as supported in the *HP OpenView Storage Data Protector Software Release Notes*, download the latest software package for the `scsitab` file from the HP OpenView World Wide Web site at http://www.hp.com/go/dataprotector.

IMPORTANT Modifying the `scsitab` file is not supported.

After you have downloaded the `scsitab` software package, follow the installation procedure provided with it.

The `scsitab` file must be located on the system to which the device is connected, in the following location:

- `<Data_Protector_home>\scsitab` on Windows platforms
- `/opt/omni/scsitab` on HP-UX and Solaris platforms
- `/usr/omni/scsitab` on other UNIX platforms

If you receive an error message while configuring your device, please contact HP Support to get information about when the device will be supported.

Using Several Drive Types in a Library

Using several drive types of a similar technology like DLT 4000/7000/8000 (the same is true within the DDS family) in the same library can lead to problems when you use the media in any drive, but do not ensure a common format on all media.

For example, at restore time, a DLT 4000 cannot read a tape that has been written with a DLT 8000 (highest density). Compressed and non-compressed media cannot be used interchangeably.

To avoid these kind of problems, you can either use a common density setting for all your media, or you can separate your media pools. Both of these solutions are described in the following sections.

Same Density Setting

This method uses a common format on all media, which allows you to use all media interchangeably in any drive.

For devices used on Windows systems, consult the drive documentation for information about using a specific write density.

On UNIX systems, you can set the density for drives by selecting the related device filenames and using them in the device definitions. The density must be set at the same value. For example, in case of DLT 4000 and DLT 7000 drives, the DLT 4000 drive density should be set.

Make sure the block size setting of the devices used is the same. This setting in the device definition must be used at the time the media are formatted.

The free pool concept can be used as desired.

During a restore, any drive can be used with any media.

On HP-UX, you can set the density of a drive when creating the device filename. See Appendix B, “Creating the Device Files on HP-UX”, in the *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information.

Different Media Pools (on UNIX and Windows)

This method separates the media used by one group of drives from the media used by another group of drives, allowing you to better optimize drive and media usage.

On Windows and UNIX systems, you can configure separate media pools for different groups of drives. This allows you to use different density settings for different drive types. For example, you could create a DLT 4000 pool and a DLT 8000 pool.

The related setting in the device definition must be used at the time the media are formatted. For example, the media in the pool for the DLT 8000 highest density must be formatted by a DLT 8000 in highest density setting.

The free pool concept cannot be used across such pools. This would not identify media from the other pool to the devices correctly; they would be seen as “foreign” media. The free pool concept can at most be used only in a single pool (like the DLT 8000 pool), in case the same media type (DLT) is written in an incompatible way.

Care must be taken during restore, since media from a given pool can only be used with related devices.

To configure new media pools, refer to the online Help index keyword “configuring media pools”.

To modify media pool settings for a drive, modify the drive properties. For detailed steps, refer to the online Help index keyword “modifying, media pools”.

Shared Devices in the SAN Environment

This section describes some of the basic concepts of Storage Area Networks (SANs). For further conceptual information, see the *HP OpenView Storage Data Protector Concepts Guide*.

The concepts and instructions provided here are the following:

- Device locking when the library is accessed exclusively by Data Protector
- Using the Data Protector user interface to configure the library robotics and drives
- Locking library robotics and drives
- Direct versus indirect library access

What Is a SAN?

A Storage Area Network (SAN) is a network dedicated to data storage, based on high-speed Fibre Channel technology. A SAN lets you offload storage operations from application servers to a separate network. Data Protector supports this technology by enabling multiple hosts to share storage devices connected over a SAN, which allows multiple systems to be connected to multiple devices. This is done by defining the same physical device multiple times, for example, once on every system that needs access to the device.

Key Concepts

There are some key concepts to consider when using Data Protector in a SAN environment:

- Each system can have its own (pseudo-)local device, although the devices are typically shared among several systems. This applies to individual drives, as well as to the robotics in libraries.
- Take care to prevent several systems from writing to the same device at the same time. Access to devices needs to be synchronized between all systems. This is done using locking mechanisms.
- SAN technology provides an excellent way of managing library robotics from multiple systems. It creates the ability to manage the robotics directly, as long as the requests sent to the robotics are synchronized among all the systems involved.

Figure 2-7 Multiple System to Multiple Device Connectivity in SAN

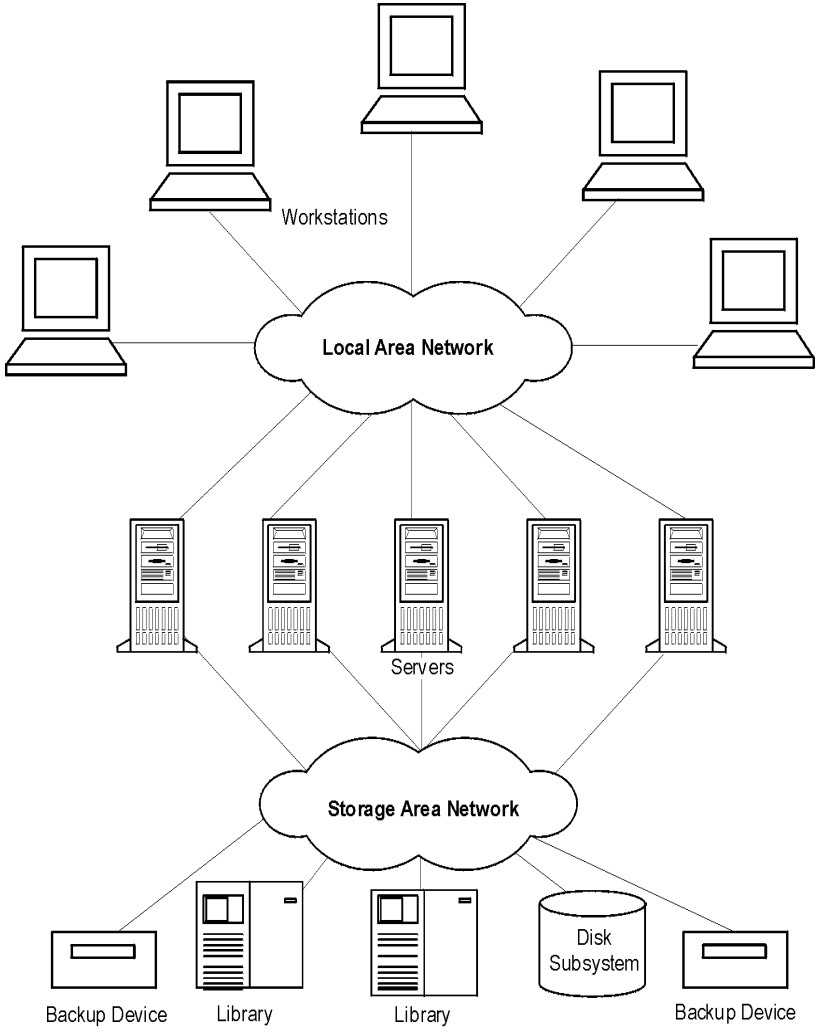

Using FC-AL SANs with LIP

Using tape devices in Fibre Channel Arbitrated Loops (FC-ALs) may cause certain anomalies that could abort a backup session. This problem arises because the FC-AL performs a Loop Initialization Protocol (LIP) whenever a new FC link is connected or disconnected, or whenever a system connected to the FC-AL is rebooted. This re-initialization of the FC-AL causes running backups to be aborted. Such terminated jobs should be restarted.

When a LIP occurs on the FC-AL Loop, any utility with an active I/O process shows an I/O error. For backup utilities attempting to use a shared tape, an I/O error causes failure of the current backup session, causing active tapes to be rewound and unloaded, and the backup session to abort.

To avoid these problems, take the following precautions:

- Do not add new devices or remove devices from the Arbitrated Loop while backup sessions are running.
- Do not touch FC components while backup sessions are running. The static charge can cause a LIP.
- Do not use `discovery` on Windows or `ioscan` on HP-UX system since these also cause a LIP.

Locking Devices Used Exclusively by Data Protector

If Data Protector is the only application that uses a drive, but that same drive needs to be used by several systems, Device Locking has to be used.

If Data Protector is the only application that uses a robotics control from several systems, Data Protector handles this internally, provided that the library control is in the same cell as all the systems that need to control it. In such a case, all synchronization of access to the device is managed by Data Protector internal control.

Locking Devices Used by Multiple Applications

If Data Protector and at least one other application want to use the same device from several systems, the same (generic) device locking mechanism has to be used by each application. This mechanism needs to work across several applications. This mode is not currently supported by Data Protector. Should this be required, operational rules must ensure exclusive access to all devices from only one application at a time.

Direct Library Access Concept

With direct library access, every system sends control commands directly to the library robotics. Therefore, a system does not depend on any other system in order to function.

With direct library access, when multiple systems send commands to the same library, the sequence of such communication has to be coordinated. Therefore, every library definition is associated by default with a host controlling the library robotics. If another host requests that a medium be moved, Data Protector will first access the system specified in the library definition for performing the move. If the system is not available, direct access from the local host to the library robotics is used if the `libtab` file is set. All of this is done in a transparent manner within Data Protector.

If direct library access is enabled for multipath devices, local paths (paths on the destination client) are used for library control first, regardless of the configured path order. The `libtab` file is ignored with multipath devices.

Indirect Library Access Concept

With indirect library access, only one system (the default robotics control system) sends robotic control commands that are initiated from Data Protector. Any other system that requests a robotics function forwards the request to the robotics control system, which then sends the actual command to the robotics. This is the default setting, and is done in a transparent manner within Data Protector for all requests from Data Protector.

Configuration Overview

This section provides an overview of the steps involved in configuring your system. It includes the following topics:

- Configuration goals

  This section specifies the mixed SAN environment to be configured.

- Configuration methods

  This section outlines the configuration methods that need to be performed for UNIX, Windows, and mixed SAN environments.

- Autoconfiguring the devices

  This section outlines the device autoconfiguration specifics in a SAN environment.

- Manually configuring the robotics

  This section describes how you can manually configure the library robotics so that they can be used in a SAN environment.

- Manually configuring the devices

  This section describes the steps that need to be performed to configure the drives. It also explains when `Lock Names` and direct access should be used.

- Manually configuring the `libtab` file

  This section describes the purpose and usage of the `libtab` file. Examples of `libtab` files are also provided.

- Simplified configuration using the `sanconf` command

  This section describes the `sanconf` command, which simplifies configuration of libraries in a SAN environment.

Configuration Goals

The SAN environment can range from one host using a library to several hosts using several libraries. The hosts can run on several operating system platforms. In the example below, the SAN environment is made up of the following systems:

- two Windows XP systems (the Windows XP system `dupin` is used as the default host to control the library robotics)
- one Windows 2000 system
- one HP-UX 11.11 system
- one bridge
- one switch
- one library with 4 HP LTO Ultrium drives and 40 slots

Figure 2-8 SAN Environment Configuration

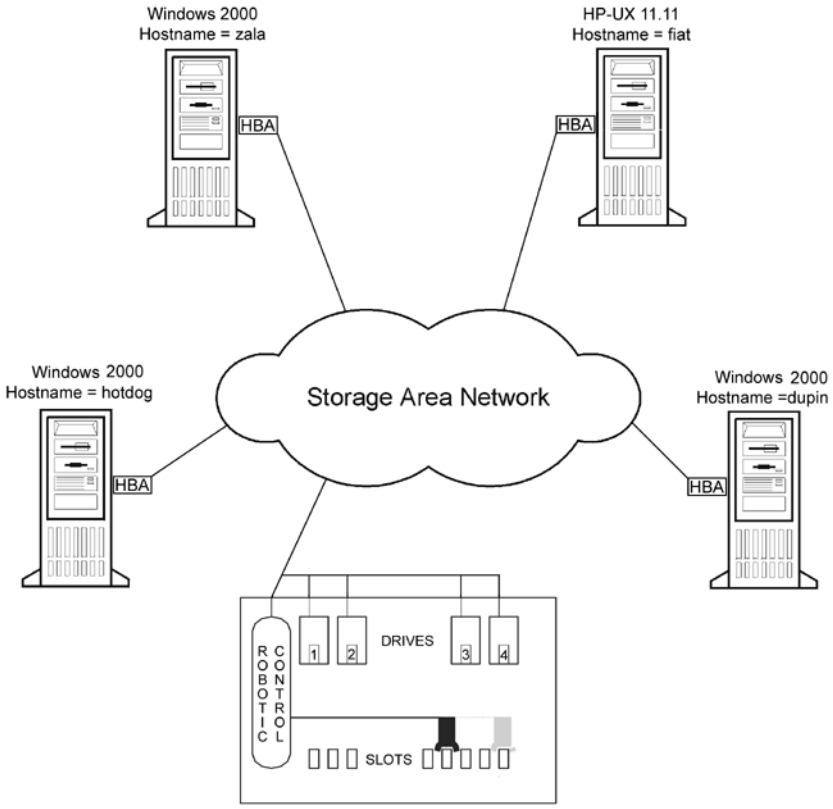

Because the library is attached to several systems that can access its drives directly, you need to configure as many drives on each host as you want to use from that host. In this case, all four physical drives are to be used from each host.

From a Data Protector perspective, the goal is as follows.

- On each host that is to share the library robotics, create a library robotics definition for each host. If there is only one host that is controlling the robotics, the library definition is created only for the default robotics control host.
- On each host that is to participate in sharing the same (tape) drives in the library:

  — Create a device definition for each device to be used.

— Use a lock name if the (physical) device will be used by another host as well (shared device).

— Optionally, select direct access if you want to use this functionality. If you use it, ensure that the `libtab` file is set up on that host.

Configuration Methods

There are three configuration methods that depend on the platforms that participate in the SAN configuration:

- You can use the Data Protector device autoconfiguration functionality to configure devices in a SAN environment using the GUI. Device autoconfiguration is supported on the following operating systems: Windows, HP-UX, Solaris, Linux, Novell NetWare, Tru64, and AIX. Refer to “Device Autoconfiguration” on page 57.
- You can use the `sanconf` command to automatically configure devices in a SAN environment using the command line. For more information, refer to “Configuration Using the sanconf Command” on page 58.
- If your environment contains systems that do not support device autoconfiguration, configure your devices manually. For more information, refer to “Manually Configuring the Library” on page 59.

Device Autoconfiguration

The Data Protector autoconfiguration functionality provides automated device and library configuration on multiple hosts in a SAN environment.

Limitations

Autoconfiguration cannot be used to configure the following devices in a SAN environment:

- mixed media libraries
- DAS or ACSLS libraries
- NDMP devices

Data Protector discovers the backup devices connected to your environment. For library devices, Data Protector determines the number of slots, the media type, and which drives belong to the library. Data

Protector then configures the device by setting up a logical name, a Lock Name, the media type, and the device file or SCSI address of the device, as well as the drive and slots.

During the autoconfiguration procedure, you can choose the libraries and devices you want to be configured and on which hosts. In case different hosts use tape drives in one library, this library will be visible from each host, multiple hosts can share tape devices, and one host (`Control Host`) will control the robotics.

NOTE

When you introduce a new host into a SAN environment, the configured libraries and devices will not be updated automatically.

To use an existing library on a new host, delete this library and autoconfigure a new library with the same name on the new host.

To add devices to an existing library, you can delete the library, and autoconfigure a library with the same name and new drives on a new host, or you can manually add the drives to the library. For the manual configuration procedure, refer to the online Help index keyword “configuring, devices in SAN environment”.

When autoconfiguring libraries while the Removable Storage service is running, drives and robotics (exchangers) will not be combined correctly.

For detailed steps on the autoconfiguration procedure, refer to the online Help index keyword “autoconfiguring backup devices”.

Configuration Using the sanconf Command

Configuration of devices in a SAN environment is automated using the `sanconf` command.

The `sanconf` command performs the following actions:

- Configuration of the default robotic control host.
- Configuration of the devices (tape drives) on all hosts by simply providing a list of hosts. This includes configuration of lock names.

For details about how to use the `sanconf` command, refer to “Automatic Configuration of Libraries in a SAN Environment Using the sanconf Command” on page 68, and to the `sanconf` man page.

Manually Configuring the Library

You first need to configure the library robotics control on a host, which acts as the default robotics control system. This host will be used to manage media movements, regardless of which other host requests a media move.

This is done in order to prevent conflicts in the robotics if several hosts request a media move at the same time. Only if the hosts fail, and direct access is enabled, is the robotics control performed by the local host requesting the media move.

Prerequisite

Before configuring Data Protector devices in a SAN environment, a Media Agent must be installed on each host that needs to communicate with the shared library. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information on installing a Media Agent.

Configuring the Library Robotics

To create the library itself, refer to “Configuring Library Devices” on page 26 or to the online Help index keyword “configuring, devices in SAN environment”.

For robotics control, you can use any host within the SAN; here the system `dupin.company.com` is used. The library robotics will be controlled by that host, unless the host is unavailable and direct access is enabled as explained in detail in “Enabling Direct Access” on page 62.

Configuring the Library Robotics in a Cluster

If you want the robotic control to be managed by a cluster, you need to make sure that:

- The robotics control exists on each cluster node.
- The virtual cluster name is used in the library robotics configuration.
- The common robotics and device filenames are installed either using the `mksf` command or using the `libtab` file. For information on how to configure the `libtab` file, refer to “Manually Configuring the libtab Files” on page 63.

After you have configured the library robotics, create the drives.

Manually Configuring the Devices (Drives)

You need to configure each device (tape drive) on each host from which you want to use the device.

Lock Names must be used to prevent the same device from being used by several hosts at the same time. Optionally, the “direct access” mode can be selected.

Configuring Drives

As will be seen shortly, it helps to follow a drive naming convention similar to the following:

LibraryLogicalName_DriveIndex_Hostname, for example SAN_LIB_2_computer_1.

The drive naming convention shows its benefits during backup specification creation. Whenever you configure a backup on any host, all you have to do now is to use the drive that is configured on that host, since the drive includes the host name in its name.

Table 2-1 Device Locking for Drives

Environment Conditions	Required Action
The drive is used by only one system and Data Protector only	No locking is necessary. Leave the fields blank, for example, Lock Name = blank
The drive is used by several systems (SAN), Data Protector is the only application accessing the drive	Use device locking (define a Lock Name) as described in the section “Device Locking” on page 93
The drive is used by several systems and several applications (not only by Data Protector)	Use device locking (define a Lock Name) and ensure that operational rules provide exclusive access to all devices from only one application at a time

Defining Lock Names

Using Lock Names is necessary in a SAN environment. This prevents collisions on the device caused by several systems talking to it at the same time. It is recommended to use the following convention for Lock Names:

LibraryLogicalName_DdriveIndex, for example SAN_LIB_D1.

Figure 2-9 Setting Advanced Options

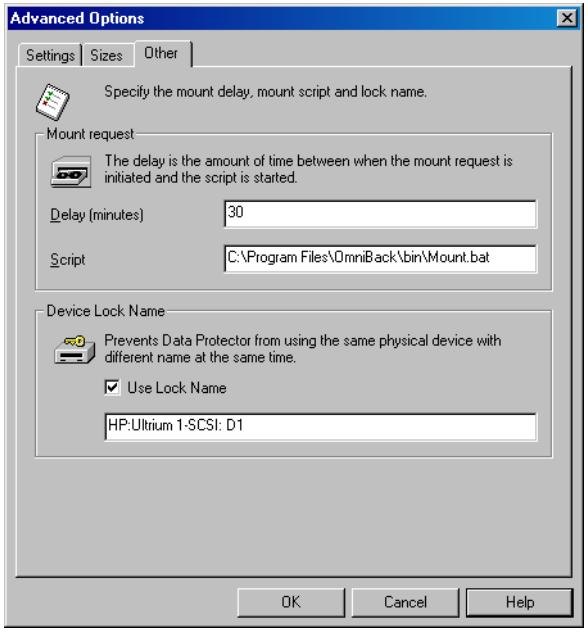

When you are setting the locking name of a drive, use the same lock name for the same physical drive when using it in the device definition on another host.

Figure 2-10 Summary of Device Definitions Using Lock Names

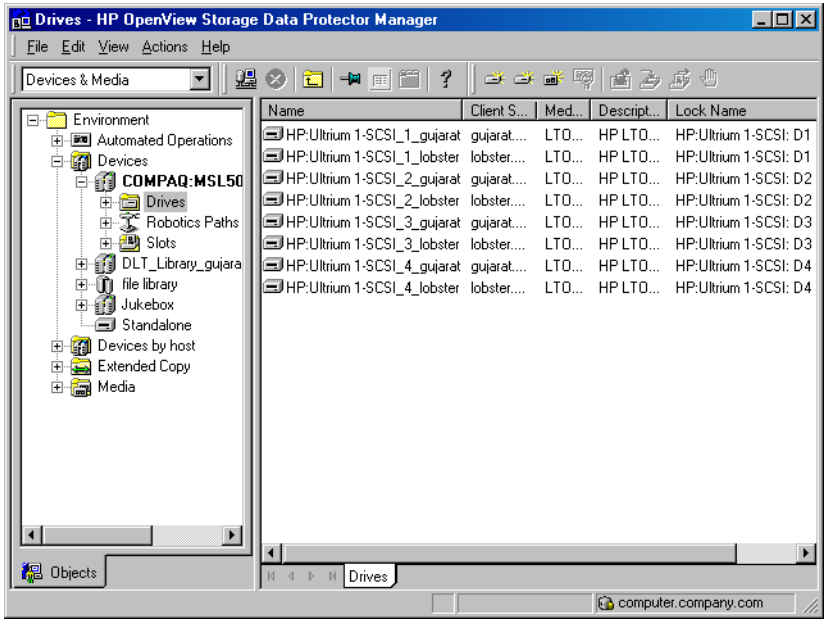

Enabling Direct Access

The Direct Access mechanism always uses the default robotics control host first for media movements, but if this fails, Data Protector uses direct access, if enabled.

To enable direct access, select the `Use direct library access` option (see Figure 2-11 on page 63) and configure the `libtab` file on every host on which you want to use direct access.

NOTE

If direct access is enabled for multipath devices, local paths (paths on the destination client) are used for library control first, regardless of the configured order.

Figure 2-11 Selecting Direct Access

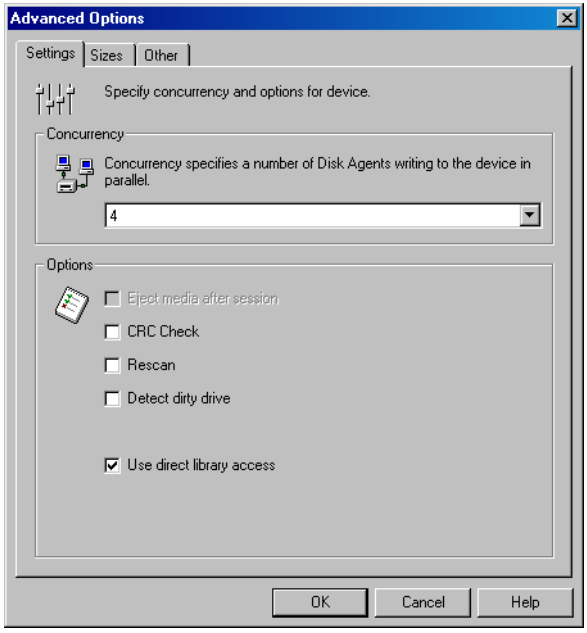

Manually Configuring the libtab Files

The purpose of the `libtab` files is to map the library robotic control access so that it also works on the “direct access requesting system”, since here the local control path is likely to be different from the one used on the default library robotic control system.

NOTE The `libtab` file is ignored for multipath devices.

You need to create a `libtab` file for every Windows and UNIX system client host that needs “direct access” to the library robotics, and is not identical to the system configured as the default library robotics control system.

On each system requesting direct access, a plain text file with the following format must be provided:

<FullyQualifiedHostname> <DeviceFile | SCSIPath> <DeviceName>

- *<FullyQualifiedHostname>* is the name of the client host demanding direct access control for the library robotics. If the host is part of a cluster, the node name should be used.
- *<DeviceFile | SCSIPath>* is the control path to the library robotic driver on this client host.
- *<DeviceName>* is the name of the device definition used on this client host.

You need one line per device for which you request direct access.

The `libtab` file is located on:

- *<Data_Protector_home>*`\libtab` on Windows systems
- `/opt/omni/.libtab` on HP-UX and Solaris systems
- `/usr/omni/.libtab` on other UNIX systems

Example files follow for all systems involved. Definitions are separated by blank lines, which are ignored. Since the default library robotics are defined on the host `dupin.company.com`, no `libtab` file is needed on this system.

TIP

It is possible to have only one `libtab` file that includes definitions for all systems involved and is distributed to all such systems. In this case, when a specific system needs “direct access” to the library robotic, the definitions for other systems are ignored and only the definitions for the system are used.

Example `libtab` file on zala

Example of the `libtab` file on host `zala.company.com` (Windows):

```
zala.company.com  scsi:2:0:2:0  SAN_LIB_1_zala
zala.company.com  scsi:2:0:2:0  SAN_LIB_2_zala
zala.company.com  scsi:2:0:2:0  SAN_LIB_3_zala
zala.company.com  scsi:2:0:2:0  SAN_LIB_4_zala
```

Example `libtab` file on oda

Example of the `libtab` file on host `oda.company.com` (HP-UX):

```
oda.company.com  /dev/spt/lib  SAN_LIB_1_computer_2
oda.company.com  /dev/spt/lib  SAN_LIB_2_computer_2
oda.company.com  /dev/spt/lib  SAN_LIB_3_computer_2
```

```
oda.company.com  /dev/spt/lib  SAN_LIB_4_computer_2
```

Example `libtab` file on donat

Example of the `libtab` file on host `donat.company.com` (Solaris):

```
donat.company.com /dev/rsst6 SAN_LIB_1_sample
donat.company.com /dev/rsst6 SAN_LIB_2_sample
donat.company.com /dev/rsst6 SAN_LIB_3_sample
donat.company.com /dev/rsst6 SAN_LIB_4_sample
```

NOTE

If the host is part of a cluster, `<FullyQualifiedHostname>` must be the virtual host name, and `<DeviceFile | SCSIPath>` must refer to the local node.

Shared Devices and MC/ServiceGuard

If you are using Data Protector with MC/ServiceGuard for clustering, you can implement the integration in a SAN environment. Since clustering is based on sharing resources such as network names, disks, and tapes among nodes, Fibre Channel and SAN are well suited as enabling technologies for storage device sharing.

This section explains how to create the necessary device files, how to configure the virtual host, how to configure static and floating drives, and how to use the Data Protector GUI to configure the integration for use in a SAN environment.

Configuration Basics

Nodes in a cluster can share SAN-connected devices in order to perform a "LAN-free" backup of an application running in a cluster. Cluster-aware applications can, at any time, run on any node in a cluster since they run on the virtual host. In order to perform a LAN-free local backup of such an application, you need to configure the logical device with a virtual hostname instead of a real node name.

You can configure as many logical devices for a single physical device as you need, but you have to use the same `Lock Name` for all devices.

In order to share a device among multiple systems, configure one logical device for each system on which you want to use the device locally.

Refer to the following documents for detailed information:

- B3935-90015 *MC/ServiceGuard Version A.11.05 Release Notes*
- B3936-90026 *Managing MC/ServiceGuard*, Sixth Edition

Configuring Drives

Floating Drives

Drives that should be accessible from both hosts, depending on which host the package is running, have to be configured based on the virtual host.

Table 2-2 How to Configure a Floating Drive

Hostname	`node_Appl`
Device Control Path	`/dev/rmt/st3m`
Lock Name	`Libl_Drive_1`

Static Drives

The drives can still be used in the standard way using the static hostname and the local device file (you can use the local HP-UX device file). The local drives should be configured on the node. For example:

Table 2-3 How to Configure a Static Drive

Hostname	`Host_A`
Device Control Path	`/dev/rmt/0m`
Lock Name	`Lib1_Drive_1`

The previous examples for floating and static drives show the device identified by `/dev/rmt/0m` and `/dev/rmt/st3m`. Both device files refer to the same physical devices, and therefore the lock name (`Lib1_Drive_1`) is identical.

Automatic Discovery of Changed SCSI Addresses

In a SAN environment, SCSI addresses (device files on UNIX) can change dynamically and therefore cause backup sessions to fail. Data Protector can internally store serial numbers of SAN devices and check them every time the device is used. If the device address does not match the address saved in the IDB, a process of device path discovery is started. If the new path is found, the IDB is updated.

Limitations

- Automatic discovery of changed SCSI addresses cannot be enabled for devices configured as a chain of standalone devices.
- Automatic discovery of changed SCSI addresses must not be used if the device does not have serial numbers.

To enable automatic discovery of changed SCSI addresses, select `Automatically discover changed SCSI address` in the second step of the Add device wizard or in the `Control` property page of the device. Refer to Figure 2-12.

NOTE

To reconfigure the device, delete the device serial number. If you do not delete the serial number, Data Protector will search for the old serial number.

Figure 2-12 Automatic Discovery of Changed SCSI Devices

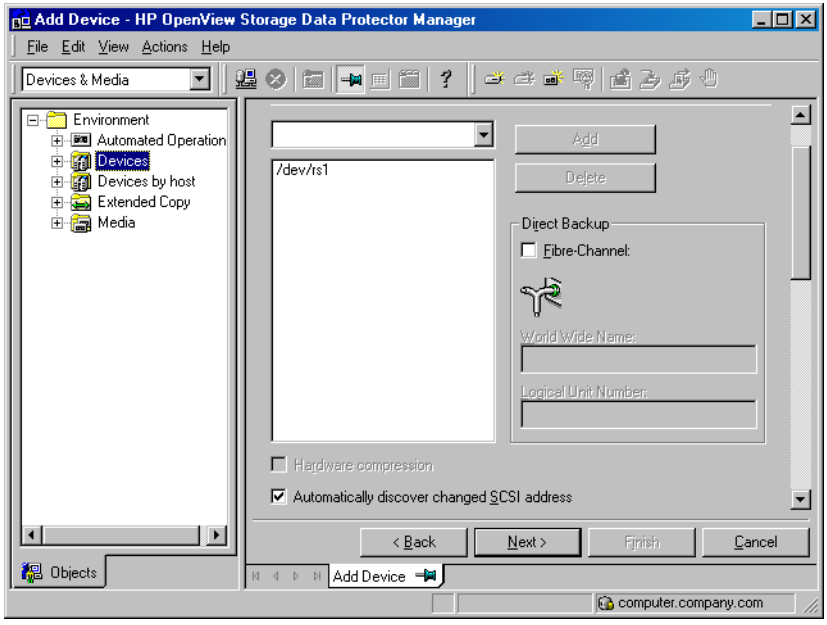

Automatic Configuration of Libraries in a SAN Environment Using the sanconf Command

The `sanconf` command is a utility that provides easier configuration of libraries in SAN environments. It can automatically configure a library within a SAN environment by gathering information on drives from multiple clients and configuring them into a single library.

The `sanconf` command can be run on the Data Protector Cell Manager or on Data Protector clients. It resides in the *<Data_Protector_home>*`\bin` directory on Windows and in the `/opt/omni/lbin` directory on HP-UX and Solaris clients.

You can perform the following tasks using the `sanconf` command:

- Scan the specified Data Protector clients and gather information on SCSI addresses of drives and robotic controllers connected to the clients in the SAN environment. For more information, refer to “Gathering Device Information on Clients” on page 70.
- Create or modify an existing library or drive settings for given clients using the information gathered during the scan of Data Protector clients. For more information, refer to “Library Device Configuration” on page 71.
- Remove drives on the specified clients from a library. For more information, refer to “Removing the Configuration” on page 78.

All `sanconf` sessions are logged to *<Data_Protector_home>*`\log\sanconf.log` on Windows or to `/var/opt/omni/log/sanconf.log` on HP-UX and Solaris systems. The log file can be used for troubleshooting.

Limitations and Recommendations

This section describes limitations and recommendations specific to the `sanconf` command.

Limitations

The limitations of the `sanconf` command are the following:

- `sanconf` is available on the following platforms:
  - Windows
  - HP-UX
  - Solaris
- `sanconf` can detect and configure devices connected to clients running on the following platforms:
  - Windows
  - HP-UX
  - Solaris
  - Linux
  - Novell NetWare
  - Tru64
  - AIX
- For a full list of libraries that are supported with `sanconf`, refer to the Support Matrices in the *HP OpenView Storage Data Protector Software Release Notes*, also available at http://www.openview.hp.com/products/datapro/spec_0001.html.
- `sanconf` does *not* support the following features:
  - Placing spare drives in drive slots.
  - Mixing drive types; for example, combinations of DLT, 9840, and LTO drives.
  - Configuring clients that are not in the Data Protector cell where `sanconf` is run.
  - Configuring clients that are currently unavailable. Configuration of such clients is possible only when the configuration of the library is performed using a configuration file that includes information gathered by scanning the clients.

Recommendation

It is recommended that only one driver is configured on a system for a specific device.

Gathering Device Information on Clients

You can obtain information on SCSI addresses of drives and robotic controllers connected to clients in the SAN environment by running the `sanconf -list[_devices]` command. `sanconf` scans the specified Data Protector clients, gathers the information and uploads it to the Media Management Database on the Cell Manager. The information about the configuration and scan data can also be stored in a **configuration file**. You can use this configuration file when you are configuring the library.

The syntax for gathering device information on clients is as follows:

Syntax

```
sanconf -list[_devices][<ListFileName>]
        [-hosts <host_1> [<host_2>... ] |
        -hostsfile <HostsFileName>]
```

When the command is executed, messages appear on the screen showing the progress of client scanning. A summary is listed when the scan is done.

If the *<ListFileName>* is specified, then the configuration and scan information are stored in the configuration file. The configuration file can be used with the `-configure` option. For more information on this option, refer to “Library Device Configuration” on page 71

The `sanconf` command requires a list of clients to scan. This list can be specified using the `-hosts` or `-hostsfile` option. If neither option is specified, all clients within the current cell are scanned.

If the `-hosts` option is specified, then all clients listed are scanned. If the `-hostsfile` option is specified, then all clients listed in the *<HostsFileName>* file are scanned. *<HostsFileName>* must be an ASCII file containing the list of clients that you want to scan. Every client must be specified in a separate line. It is recommended that all clients are specified in the clients list before you save the scan information to the configuration file

NOTE

You have to scan all clients that you want to configure, those that can see the robotics and those that can see the drives.

Example

The following is an example of scanning the devices on hosts "host01", "host02" and "host03":

```
sanconf -list_devices -hosts host01 host02 host03

================== SUMMARY REPORT ==================
LIBRARY="HP:C7200-8000" serial="ABC0000123"
        on hosts: host01 host02 host03
        DRIVE: index=1; name="QUANTUM:DLT8000";
serial="A0B1C3D4E5"
        DRIVE: index=2; name="QUANTUM:DLT8000";
serial="F6G7H8I9J0"
=====================================================
```

Library Device Configuration

A library device is configured using the `sanconf -configure` command.

The syntax for configuring a tape library is as follows:

Syntax

```
sanconf -configure [<ListFileName>]
       -library <LibrarySerialNumber> <LibraryName>
       [<RoboticControlHostName>]
       [<DeviceTypeNumber> | ".<DeviceTypeExtension>"]
       [-hosts <host_1> [<host_2> ] | -hostsfile
       <HostsFileName>]
       [-drive_template <DriveTemplateFileName>]
       [-library_template <LibraryTemplateFileName>]
       [-[no_]multipath]]
```

If the *<ListFileName>* parameter is not specified, the `sanconf` command scans all clients within the current cell. If the parameter is specified, `sanconf` assumes that all information on clients is already present in the configuration file. In such a case the scan is not performed. If a client is not scanned, the library will be configured without drives connected to this host.

By default or if the `-no_multipath` option is given, `sanconf` configures a separate logical device for each path. If the `-multipath` option is used, `sanconf` configures all paths pointing to a single physical device as a single multipath device.

Initial Configuration of Library Device

For the initial configuration you have to run the `sanconf` command using the `-configure` option. `sanconf` starts `devbra` on all specified clients. `devbra` queries devices for configuration information on physical libraries and drives. It is recommended that the information on clients is first stored using the `-list[_devices]` *<ListFileName>* option and then the library and drives are configured using `-configure` *<ListFileName>*.

When the *<ListFileName>* parameter is used, the `sanconf` command only reads the configuration file. In this case no scan of the clients is performed.

When you use the `-library` parameter, the `sanconf` command creates a logical library in the system and all devices on all specified clients. The library serial number can be obtained using the `-list_devices` option.

Example

In the following example, the `sanconf` command scans the specified clients and then creates a logical library named "SAN_STORE" with robotics configured on client "host33" and drives for that library configured on clients "host01", "host02" and "host03".

```
sanconf -configure -library MPC0100013 SAN_STORE host33
       -hosts host01 host02 host03
```

Example

In the following example, the `sanconf` command first scans the SAN environment for the configuration information on the specified clients "host01", "host02", "host03", and "host33". This information is then saved in the `mySAN.cfg` file.

```
sanconf -list_devices mySAN.cfg
        -hosts host01 host02 host03 host33

sanconf -configure mySAN.cfg -library MPC0100013 SAN_STORE
        host33 -hosts host01 host02 host03
```

The second command uses the information stored in the `mySAN.cfg` file and creates a logical library named "SAN_STORE" with robotics configured on client host33 and drives for the library configured on clients "host01", "host02", and "host03".

If the *<RoboticControlHostName>* parameter is specified, then this library will be configured with robotics controlled on the specified client. If *<RoboticControlHostName>* is not specified, the Cell Manager is used as the robotics control host.

If the *<DeviceTypeNumber>* parameter or the ".*<DeviceTypeExtension>*" parameter is used, the drives of that type are configured in the library. For the supported drive types and their parameters, see Table 2-4 on page 73. If neither of these parameters is specified, the DLT drive type is used as default. When drives in the library are not of the same type as specified, an error is reported. You can run the *<DeviceTypeNumber>* parameter without specifying the *<RoboticControlHostName>* parameter.

Table 2-4 Device Type Numbers and the Corresponding Extensions

Device Type Numbers	Device Type Extensions
1	.DDS
2	.QIC
3	.EXA
4	.AIT
5	.3480
6	.RDSK
7	.REGFILE
8	.9840
9	.TAPE
10	.DLT
11	.D3
12	.3590
13	.LTO
14	.SDLT
15	.VXA

Table 2-4 Device Type Numbers and the Corresponding Extensions

Device Type Numbers	Device Type Extensions
16	.DTF
17	.9940
18	.SAIT
19	.3592

To configure drives only on the specified clients, use the `-configure` option with either the `-hosts` or `-hostsfile` option. If neither option is specified, the drives will be configured on all clients in the cell. If the `-hosts` option is specified, then the drives are configured on all listed clients. If the `-hostsfile` option is specified, then the drives are configured on all clients that are listed in the *<HostsFileName>* file. *<HostsFileName>* must be an ASCII file containing all clients on which you want to configure the drives of the library. Every client must be specified in a separate line.

To configure multipath devices, add the following at the beginning of the *<HostsFileName>* file:

```
<OPTIONS>
-multipath
</OPTIONS>
```

For multipath devices, the path order is determined by the order in the given list or file, for example if you use `-hosts Client2 Client1 Client3`, then the first path will contain `Client2`.

Example

To specify a specific drive type when configuring a library using the `sanconf` command, run the following command:

```
sanconf -configure -library MPC0100013 SAN_STORE host33
        ".9840" -hosts host01 host02
```

This command creates a library named SAN_STORE with robotics configured on client host33 and STK drives configured on clients "host01" and "host02". The drives are named as follows:

```
SAN_STORE_1_host01

SAN_STORE_1_host02
```

```
SAN_STORE_2_host01

SAN_STORE_2_host02

...
```

You can alter the default configuration of all configured drives in a library using the `-drive_template <DriveTemplateFileName>` option. The parameters must be specified in an ASCII file with each parameter in a separate line. The drive template supports the following parameters:

`VERIFY`	This parameter corresponds to the `CRC Check` option in the Data Protector GUI.
`CLEANME`	This parameter corresponds to the `Detect dirty drive` option in the Data Protector GUI.
`RESCAN`	This parameter corresponds to the `Rescan` option in the Data Protector GUI.

You can also alter the default configuration of the library using the `-library_template <LibraryTemplateFileName>` option. The parameters must be specified in an ASCII file with each parameter in a separate line. The library template supports the following parameters:

`BARCODEREADER`	This parameter corresponds to the `Barcode reader support` option in the Data Protector GUI.
`BUSYDRIVETOSLOT`	This parameter corresponds the `Busy drive handling: Eject medium` option in the Data Protector GUI.
`BUSYDRIVETOMAILSLOT`	This parameter enables the `Busy drive handling: Eject medium to mail slot` option specified in the Data Protector GUI.

IMPORTANT

You can alter the default configuration of the library only at the initial configuration. After the library is configured, you can no longer change the configuration of the library using the `sanconf` command.

Example

In this example, the `sanconf` command scans all clients in the cell and then creates a logical library named SAN_STORE with robotics configured on client host33. The library and drives are also configured with the parameters specified in the files `DriveTemplate.txt` and `LibraryTemplate.txt`.

```
sanconf -configure -library MPC0100013 SAN_STORE host33
       -drive_template DriveTemplate.txt
       -library_template LibraryTemplate.txt
```

Reconfiguration of Library Device

To reconfigure an existing library, use the `-configure` option. The `sanconf` command can be used the same way as with the initial configuration except for the clients that already have drives configured on. It is recommended that configuration information is first stored in the configuration file in case of configuration failure. It is also recommended that a different filename is used so that the initial configuration can be restored without any complications. `sanconf` reuses the custom settings when reconfiguring a library.

If you are reconfiguring a non-multipath library as a multipath library, the library is changed to a multipath library and the library control host is used as the first path. Non-multipath drives are not changed or removed. Instead, new multipath drives are created. Only multipath drives are modified.

If you are reconfiguring a multipath library as a non-multipath library, the library is changed to a non-multipath library and only one path is created. Multipath drives contained inside a multipath library however, are not changed. Instead, new drives are created. Only non-multipath drives are modified.

Reconfiguring When Switch Changes SCSI Addresses

You can reconfigure or modify an existing logical library when, for example, SCSI addresses change after a switch was rebooted. You can run the `sanconf` command with the same parameters that were used in the initial configuration. See the examples below.

Example 1

To reconfigure a library device named "SAN_STORE" with serial number "MPC0100013" with devices on hosts "host01", "host02" and "host03", and robotic control on host "host33", after the SCSI addresses had changed, run the following command:

```
sanconf -configure -library MPC0100013 SAN_STORE host33
       -hosts host01 host02 host03
```

Example 2

To reconfigure the same library device in similar situation using the configuration file "mySAN.cfg", run the following commands:

```
sanconf -list_devices mySAN.cfg
       -hosts host01 host02 host03 host33
```

```
sanconf -configure mySAN.cfg -library MPC0100013 SAN_STORE
       host33 -hosts host01 host02 host03
```

Reconfiguring When New Clients Are Added

The `sanconf` command can add drives to an existing library on specified clients. It is recommended that all new clients are scanned first for configuration data and then configured in the existing database.

Example 1

To reconfigure a library device named "SAN_STORE" with serial number "MPC0100013" using the `-hosts` option, when new clients "host04" and "host05" are added to the cell, run the following command:

```
sanconf -configure -library MPC0100013 SAN_STORE host33
       -hosts host04 host05
```

Example 2

To reconfigure the same library device using the configuration file "myNewHostsSAN.cfg", run the following commands:

```
sanconf -list_devices myNewHostsSAN.cfg -hosts host04 host05
```

```
sanconf -configure myNewHostsSAN.cfg -library MPC0100013
       SAN_STORE host33 -hosts host04 host05
```

Example 3 - Multipath

To reconfigure the library from the previous examples as a multipath library, run the following command:

```
sanconf -configure -library MPC0100013 SAN_STORE host33
       -hosts host04 host05 -multipath
```

Removing the Configuration

Removing Drives from a Library Device

Run the `sanconf` command using the `-remove_drives` option to remove the configured drives that reside in a library on the specified clients. Drives that are configured as multipath drives are not removed.

The following syntax for the removal of the configuration is provided:

Syntax

```
sanconf -remove_drives <LibraryName>
        [-hosts <host_1> [<host_2>... ] |
        -hostsfile <HostsFileName>]
```

NOTE No rescanning is required for this operation.

Example To remove all drives in the library named "SAN_STORE" that are configured on clients "host04" and "host05", run the following command:

```
sanconf -remove_drives SAN_STORE -hosts host04 host05
```

Removing Paths or Drives for a Given Client

To remove paths or drives from library devices for a given client, run the `sanconf` command with the `-remove hosts` option.

The following syntax for the removal of the configuration is provided:

Syntax

```
sanconf -remove_hosts <LibraryName>
        [-hosts <host_1> [<host_2>... ] |
        -hostsfile <HostsFileName>]
         [-[no_]multipath]
```

NOTE No rescanning is required for this operation.

All paths containing the specified hosts are removed. However, if the specified clients cover all paths of the library, no paths are removed from this library, instead a warning is displayed.

Use the -[no]_multipath option to determine the behavior of this option for multipath and non-multipath devices:

- To remove paths from *multipath* devices (libraries and drives), add the -multipath option. Non-multipath devices are not affected.
- To remove *non-multipath* devices configured on the given client, add the -no_multipath option. Multipath devices are not affected.
- To remove paths from multipath devices *and* to remove non-multipath devices configured on the given client, run the command *without* the -no_multipath and -multipath options.

Example - Multipath

To remove all paths that are configured on clients host04 and host05 from the multipath library named SAN_STORE and from the drives it contains, run the following command:

```
sanconf -remove_hosts SAN_STORE -hosts host04 host05
-multipath
```

Example

To remove all non-multipath drives from a library named SAN_STORE1 that are configured on clients host02 and host03, run the following command:

```
sanconf -remove_hosts SAN_STORE1 -hosts host02 host03
-no_multipath
```

Removing the Entire Library Device

The sanconf command as a configuration utility does not provide an option to delete an entire library from the IDB. You can perform this action by running the omniupload command with the following option:

```
omniupload -remove_library <LibraryName>
```

For additional information on the omniupload command, refer to the omniupload man page.

The sanconf Command-Related omnirc File Variable

The sanconf command uses the following omnirc file variable:

OB2SANCONFSCSITIMEOUT=s

Default value: 20 seconds

This variable sets the time-out value for `sanconf` related operations and is used on Windows systems. It must be set on all clients affected by `sanconf`, before running the command.

For more information on the location and usage of the `omnirc` file, refer to "Using Omnirc Options" on page 615.

Drive Naming Convention

All drives created with the `sanconf` command are named automatically by `sanconf`.

IMPORTANT Drive names must not be changed manually because the reconfiguration will not work.

Depending on whether the device is configured in multipath mode or not, the following drive naming convention must be followed:

- For *non-multipath* devices:

  `libname_index_host`

  `libname_index_busindex_host`

  The `busindex` number is used only if there is more than one path for the drive.

- For *multipath* devices:

  `libname_index`

Examples A library named "SAN_STORE" with a drive on a client "host01" that is in position 2 in the physical library uses the following drive name:

`SAN_STORE_2_host01`

If you have a dual HBA environment, the bus number is also appended. If there is another drive located on the client "host01" that has an index 2 in the library and is seen on SCSI adapter bus 4, it will be named:

`SAN_STORE_2_4_host01`

Example - Multipath

A library named "SAN_STORE" with a drive on a client "host01" that is in position 2 in the physical library uses the following drive name:

```
SAN_STORE_2
```

If you have a dual HBA environment, the bus number is not appended. If there is another drive located on the client "host01" that has an index 2 in the library and is seen on SCSI adapter bus 4, only a new path will be added.

Drive Locking Mechanism

After a successful configuration in the SAN environment there are multiple logical drives created that represent one physical drive. Data Protector features a locking mechanism which prevents users from starting simultaneous backups on several logical drives that map to the same physical drive. The key part of this mechanism are **lock names**.

A lock name can be used by all logical drives that represent one physical drive. The `sanconf` command automatically creates lock names for drives that it is configuring. A lock name consists of the drive vendor ID string, the product ID string and the product serial number.

Example

An HP DLT 8000 drive with vendor ID "HP", product ID "DLT8000", and serial number "A1B2C3D4E5" will be automatically configured with the following lock name:

```
HP:DLT8000:A1B2C3D4E5
```

You must not change the lock names that were created by the `sanconf` command. All other logical drives that are created manually and represent physical drives that have been configured by `sanconf` must also use lock names created by `sanconf`.

Drive Cleaning

There are several methods for cleaning dirty drives:

- Library built-in cleaning mechanism

  Some tape libraries have a functionality for cleaning drives automatically when a drive requests head cleaning. When the library detects a dirty drive, it automatically loads a cleaning tape. However, Data Protector is not notified of this action. This interrupts any active session, causing it to fail.

  This hardware-managed cleaning procedure is not recommended, since it is not compatible with Data Protector. Use automatic drive cleaning managed by Data Protector instead.

- Automatic drive cleaning managed by Data Protector

  Data Protector provides automatic cleaning for most devices using cleaning tapes. For SCSI libraries and magazine devices, you can define which slots contains cleaning tapes. A dirty drive sends the cleaning request, and Data Protector uses the cleaning tape to clean the drive.

  This method prevents failed sessions due to dirty drives, provided that suitable media are available for backup. Refer to “Configuring Automatic Drive Cleaning” on page 84.

- Manual cleaning

  If automatic drive cleaning is not configured, you need to clean the dirty drive manually. If Data Protector detects a dirty drive, a cleaning request appears in the session monitor window. You then have to manually insert a cleaning tape into the drive.

A special tape-cleaning cartridge with slightly abrasive tape is used to clean the head. Once loaded, the drive recognizes this special tape cartridge and starts cleaning the head.

Limitations

- Data Protector does not support the diagnostic vendor-unique SCSI command for performing drive cleaning with cleaning tapes stored in one of the special cleaning tape storage slots. These special cleaning tape storage slots are not accessible using the normal SCSI commands, and therefore cannot be used with the automatic drive cleaning managed by Data Protector. Configure the standard slot(s)

to store cleaning tape(s).

- Detection and use of cleaning tapes depends on the system platform where a Media Agent is running. See the *HP OpenView Storage Data Protector Software Release Notes* for further information.
- You should not use another kind of device management application if you configure automatic drive cleaning managed by Data Protector, as this may cause unexpected results. This is due to the “cleanme” request being cleared as it is read, depending on the specific device type and vendor.
- Automatic drive cleaning for logical libraries with a shared cleaning tape is not supported. Each logical library needs to have its specific cleaning tape configured.

Conditions for Automatic Cleaning

Automatic drive cleaning is supported for libraries with barcode support, as well as for those without barcode support.

The following conditions must be met for automatic cleaning:

- In a library without barcode support, a cleaning-tape slot has been configured in the Data Protector device definition and contains a cleaning-tape cartridge. The cleaning-tape slot must be configured together with the other library slots.
- In a library with barcode support, the cleaning tape has a barcode label with “CLN” as its prefix. Further, barcode support must be enabled. Refer to “Activating Barcode Support” on page 87.
- The configured drive has the `Detect Dirty Drive` option enabled.

When Data Protector receives notification that the drive needs cleaning, it automatically loads the cleaning tape, cleans the drive, and then resumes the session.

All cleaning activities are logged in the following file:

- on Windows: `<Data_Protector_home>\log\server\cleaning.log`
- on UNIX: `/var/opt/omni/server/log/cleaning.log`

Configuring Automatic Drive Cleaning

The configuration of automatic drive cleaning is performed in two steps:

1. Enable dirty drive detection. This needs to be done for all device types (standalone and libraries). This enables Data Protector to recognize the event issued by the drive.
2. Configure a slot for the cleaning tape in the library or magazine device.

Enabling Dirty Drive Detection

To enable dirty drive detection, select the `Detect dirty drive` advanced option in the `Settings` property page for the drive. For detailed steps, refer to the online Help index keyword “configuring drive cleaning”.

Configuring a Slot for a Cleaning Tape

To configure a slot for a cleaning tape in a SCSI library, click the `Cleaning Slot` option and select an existing slot in the drop-down list in the `Repository` property page for the device. Note that you cannot configure a cleaning slot for libraries with the barcode reader support enabled. For detailed steps, refer to the online Help index keyword “configuring drive cleaning”.

Testing the Drive Cleaning Configuration

To test if drive cleaning has been successfully configured, do the following:

Preparation

1. Log on to the system where a Media Agent for the drive is installed.
2. Change to the Data Protector `tmp` directory:
   - on HP-UX and Solaris systems: `/var/opt/omni/tmp/`
   - on other UNIX systems: `/usr/omni/tmp/`
   - on Windows systems: `<Data_Protector_home>\tmp\`
   - on Novell NetWare systems: `\usr\omni\tmp\`
3. Create an ASCII file named `simtab` on Windows systems or `.simtab` on UNIX systems. Consider the following when creating this file:
   - The field separators should be a single ASCII character (tab or space)

- The logical device name cannot be quoted and cannot contain spaces (e.g. “test drive”)

The content of the `simtab`/`.simtab` file should be the following:

```
CLEANME  <file_name> <drive_name>
```

Where *<file_name>* is the name of the file you will use to simulate a dirty drive, and *<drive_name>* is the name of the drive you want to test.

You can add multiple entries for various drives. Do not add any directories in front of the name of the file.

Testing the Configuration

In order to test your configuration, do the following:

1. In the Data Protector `tmp` directory, create an empty file that will be used to simulate a dirty drive. Use the same name as in the `simtab` or `.simtab` file.
2. Start a backup using the drive you are testing.

Data Protector behaves as though the selected drive were dirty and performs the cleaning action.

To stop simulating dirty drive behavior for the specific drive, delete the file used for simulation.

Busy Drive Handling

Data Protector expects drives to be empty, i.e., there should not be a medium in the drive unless a restore or backup is currently active. Several factors can cause a medium to still be in a drive, for instance, if the medium was used with a different application and not removed, or if the system writing the data to the tape (a Media Agent) failed during the backup. The next backup using this drive has to deal with this situation. Data Protector can respond automatically in several ways. The response is configurable via the library option `Busy Drive Handling`.

The following options are available:

Abort The backup will be aborted (default).

Eject Data Protector will eject the medium from the drive and put it in any empty slot.

Eject to mail slot Data Protector will eject the medium from the drive and put it in the library mail slot (CAP).

If the backup continues automatically, select **Eject**. Because the tape is moved to an unknown slot, the library should be scanned before the next backup.

Activating Barcode Support

If a SCSI library device uses media with barcodes, Data Protector can use barcodes by providing the following support:

- Recognition of cleaning tapes with a CLN prefix.
- Reference to media by their barcodes. Data Protector adds a barcode to the Data Protector media label in the IDB. You can optionally use the barcode as medium label and write it to the medium header on the tape during the initialization of the medium.
- Quickly scanning the media in the slots of the library repository using media barcodes. This is considerably faster than scanning a repository without the barcode functionality. In the `Action` menu, click `Barcode Scan` to scan the library repository for media.

Activate barcode support by selecting the `Barcode reader support` option from the `Control` property page of the device. If you select the `Use barcode as medium label on initialization option`, then the `Use barcode` option is going to be enabled by default when initializing media using this library. Refer to Figure 2-13 on page 88. For detailed steps, refer to the online Help index keyword “activating barcode reader support”.

NOTE All barcodes in a cell must be unique, regardless of the type of media or the fact that there are multiple libraries.

Figure 2-13 Activating Barcode Reader Support

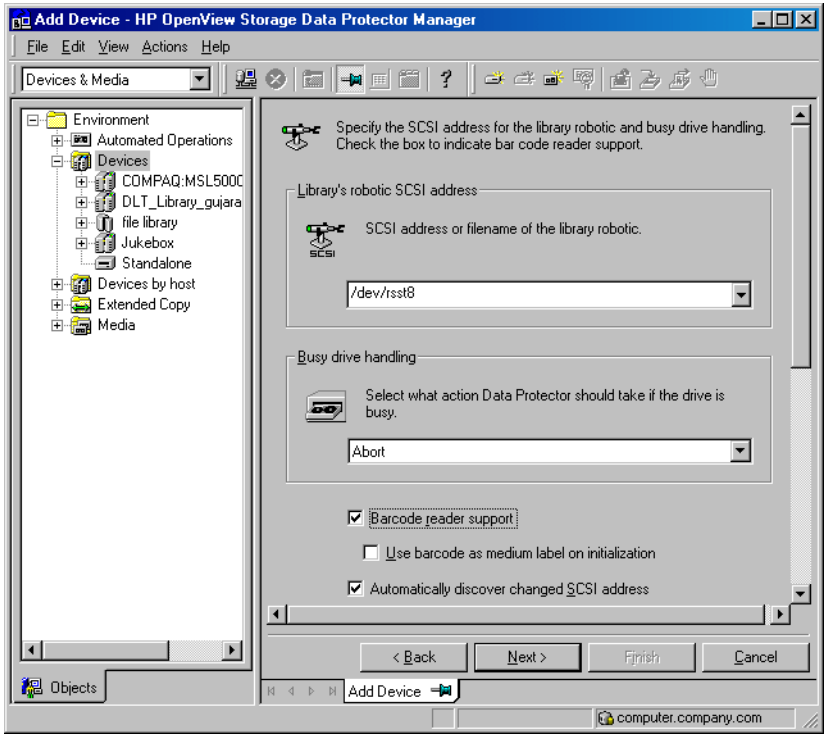

Disabling a Backup Device

Disabling a backup device is useful when the device is damaged or in maintenance mode.

If you disable a backup device, all subsequent backups skip this device. The next available device defined in the list of devices for the backup specification is used, provided that load balancing has been selected. All devices using the same lock name as the disabled device are also disabled.

This lets you avoid backups that fail due to a device needing service, while keeping other devices available (and configured) for backup.

How to Disable a Device

Disable a backup device by selecting the `Disable device` option in the `Settings` property page of the device or drive. Refer to Figure 2-14. For detailed steps, refer to the online Help index keyword “disabling backup devices”.

How to Restart the Device

To resume using the device for backups, deselect the `Disable device` option.

Figure 2-14 Disable Device

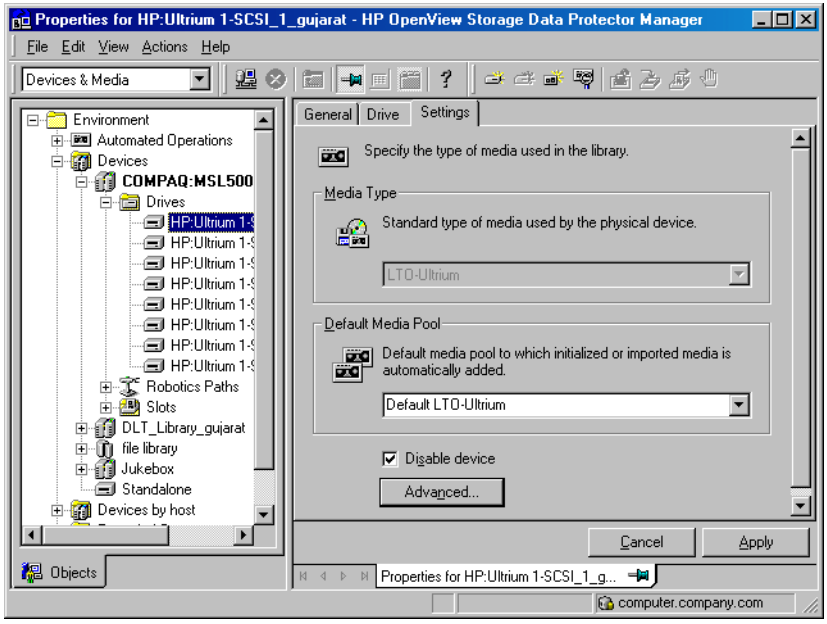

Removing a Backup Device

By removing a backup device from the Data Protector configuration, you stop using this device for backup or restore. Make sure that you remove the device from all backup specifications that use the device. Otherwise the backup or restore will fail.

TIP

Also, if you are not using a certain backup device with Data Protector anymore, you may want to remove the Media Agent software from the system. This can be done using the Client context.

How to Remove a Backup Device

To remove a backup device, delete it from the `Devices & Media` context. For detailed steps, refer to the online Help index keyword “deleting backup devices”.

Renaming a Backup Device

When you rename a backup device, the device is no longer used under its old name for backup or restore.

IMPORTANT

Make sure that you remove the device's old name from all backup specifications that use the device. Otherwise, Data Protector tries to back up to or restore from a device that does not exist, and the session fails.

How to Rename a Backup Device

Rename a backup device in the `General` property page of the device. For detailed steps, refer to the online Help index keyword “renaming backup devices”.

Device Locking

Internal Locking

The internal locking of backup devices prevents two Data Protector sessions from accessing the same physical device at the same time. For example, if one backup session is using a particular device, all other backup/restore sessions must wait for this device to become free before starting to use it. When a backup or restore session starts, the Data Protector locks the device, the drive, and the slot used for that session.

Media sessions performing media operations such as initialization, scanning, verifying, copying, or importing also lock devices. During that time, no other operations can lock and use the device. If a media session cannot obtain a lock, the operation fails, and you have to retry the operation at a later time.

Locking When a Mount Request Is Issued

During a mount request of a backup or restore session, Data Protector allows the device to be used for media management operations, such as formatting a new medium.

When the mount request is confirmed, the backup or restore session locks the device again and continues with the session.

Locking with Data Protector

You can configure the same physical device many times with different characteristics, simply by configuring devices with different device names.

Since the internal locking operates on logical devices rather than on physical devices, a collision can occur if you specify one device name in one backup specification and another device name for the same physical device in another backup specification. Depending on the backup schedule, this may result in Data Protector trying to use the same physical device in several backup sessions at the same time. This can also happen when two device names are used in other operations, such as backup and restore, backup and scan, and so on.

To prevent this collision, you can specify a virtual lock name in both device configurations. Data Protector then uses this lock name to check if the device is available, thus preventing collisions.

If you configure two Data Protector backup devices that actually point to the same physical device, you are advised to specify the `Lock Name` in the advanced options for the two logical devices. `Lock Name` is the name that

Data Protector recognizes in order to lock the device before starting backup and restore sessions. Both logical devices need to have the same lock name. Refer to “Shared Devices in the SAN Environment” on page 51 for example on how to use `Lock Name`.

How to Lock a Device

Lock a backup device by selecting the `Use Lock Name` advanced option from the Settings property page for the device, and then entering the lock name of your choice. For detailed steps, refer to the online Help index keyword “locking backup devices”.

Device Concurrency, Segment Size, and Block Size

Streaming

To maximize a device's performance, it has to be kept streaming. A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes writing to the tape, and so on. In other words, if the data rate written to the tape is less than or equal to the data rate which can be delivered to the device by the computer system, the device is streaming. Device streaming is also dependent on other factors such as network load and the block size of the data written to the backup device in one operation.

For additional information on device concurrency, segment size, and block size, see the Media Management chapter in the *HP OpenView Storage Data Protector Concepts Guide*.

Changing Concurrency

Data Protector provides a default number of Disk Agents that are started for each device. Increasing the number of Disk Agents sending data to a Media Agent at the same time improves device streaming.

In the `Advanced Options` dialog box of a specific device, set the `Concurrency` to the maximum number of Disk Agents allowed to feed data to each Media Agent. See Figure 2-15 on page 96. For detailed steps, refer to the online Help index keyword "concurrency".

Concurrency can also be set in the backup specification. The concurrency set in the backup specification will take precedence over the concurrency set in the device definition. See Figure 2-16 on page 97. For detailed steps, refer to the online Help index keyword "concurrency".

Figure 2-15 Advanced Options Dialog Box: Concurrency

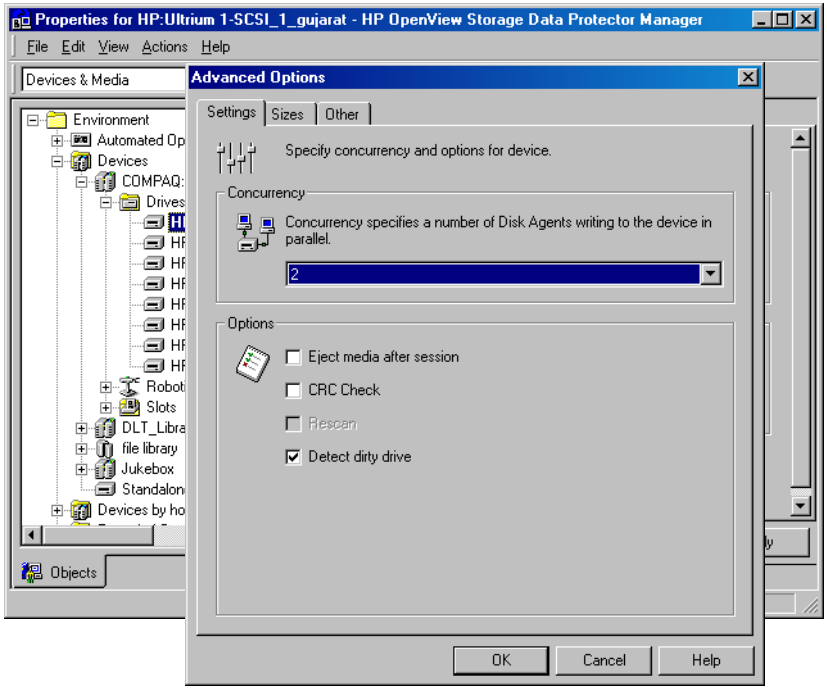

Figure 2-16 Device Properties Dialog Box: Concurrency

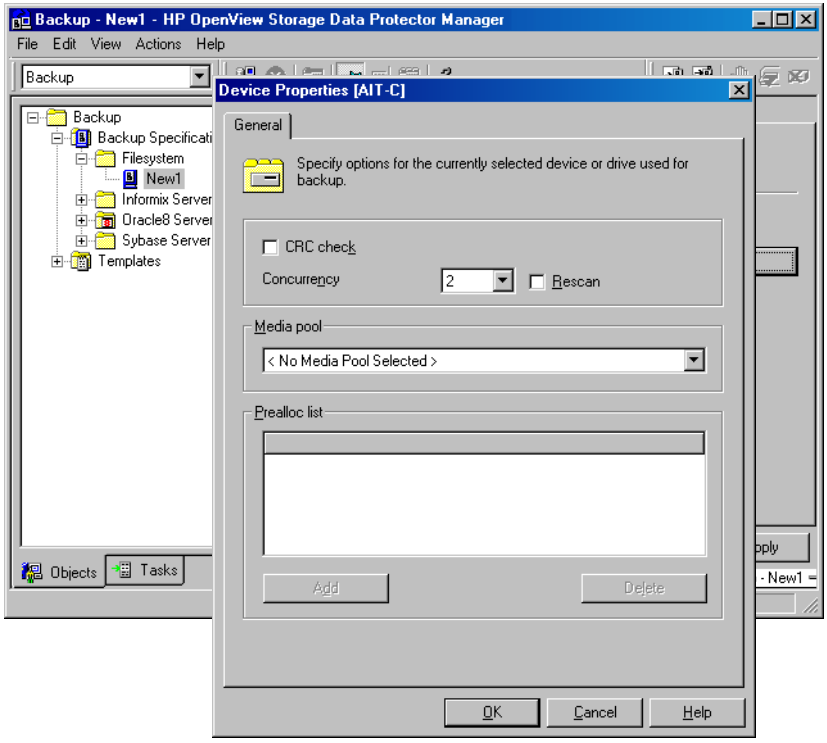

Changing Segment Size

Segment size is related to the size of data areas which Data Protector uses in writing data to the media. It is user-configurable for each device. Note that a smaller segment size consumes media space because each segment has a file mark which takes up space on a medium. A larger number of file marks results in faster restores, because a Media Agent can quickly locate the segment containing the data to be restored.

Optimal segment size depends on the media type used in the device and the kind of data to be backed up. The average number of segments per tape is 50. The default segment size can be calculated by dividing the native capacity of a tape by 50. The maximum catalog size is limited to a fixed number (12 MB) for all media types.

Data Protector finishes a segment when the first limit is reached. When backing up a large number of small files, the media catalog limit is reached faster, which can result in smaller segment sizes.

You can change the segment size in the `Advanced Options` dialog box of a specific device. For detailed steps, refer to the online Help index keyword “segment size”.

Changing the Number of Buffers

Data Protector Media Agents and Disk Agents use memory buffers during data transfer. Memory is divided into a number of buffer areas. Values from 1 - 32 may be specified.

Each buffer area consists of 8 Disk Agent buffers, which are of the same size as the block size configured for the device. The default device block size is 64 KB.

You can change the number of buffers by changing the `Advanced Option` properties of the selected drive. For detailed steps, refer to the online Help index keyword “number of Disk Agent buffers”.

Block Size

When a device receives data, it processes it using a device-type-specific (DDS, DLT) block size.

NOTE

Each backup device (drive) has a block size. A restore adjusts to block size.

Before Changing Block Size in Data Protector

Data Protector uses a default device block size for each device type. The block size applies to all devices created by Data Protector and to Media Agents running on the different platforms.

The device block size is written on a media header so that Data Protector knows the size to be used. If the device block size differs from the medium’s block size, an error occurs.

You can change the device block size in the Data Protector GUI. However, before changing the block size you need to check the supported block size of the host adapter.

The minimum block size for old SCSI cards, such as the Adaptec 2940, was 56 KB. Currently, the minimum block size that is mainly used with newer SCSI cards is 64 KB.

You can increase the maximum block size on a Windows Media Agent client by modifying its Registry. For information on how to modify the block size, see the example in “Changing Block Size on Windows Media Agent” on page A-51.

Before changing the block size for a particular SCSI card, refer to the SCSI vendor documentation or contact the vendor support.

Changing the Block Size in Data Protector

You can set the block size in the `Advanced Options` dialog box of a specific device. For detailed steps, refer to the online Help index keyword “block size”.

Device Performance Tuning

Block Size

Every logical device can be configured to process data in units of a specific size (**block size**). Different devices have different default block sizes, which are safe (all sessions are completed successfully) but they may not be optimal. By adjusting the block size, you can enhance the performance of Data Protector sessions.

The optimal block size value depends on your environment:

- Hardware (devices, bridges, switches,...)
- Firmware
- Software (operating system, drivers, firewall,...)

To achieve the best results, first optimize your environment by installing the latest drivers and firmware, optimize your network, and so on.

How to Find the Optimal Block Size

To find the optimal block size, perform different tests by running usual Data Protector tasks (backup, restore, copy, and so on) with different block size values and measure the performance.

Note that once you have changed the device block size, you cannot restore old backups (with the old block size) with this device anymore. Therefore, keep your old logical devices and media pools intact to be able to restore the data from the old media and create new logical devices and media pools with different block size values for testing purposes.

Limitations

Before changing the default block size, consider the following limitations:

- Disaster recovery: To be able to perform an offline recovery for EADR/OBDR (see “Enhanced Automated Disaster Recovery of a Windows System” on page 539 and “One Button Disaster Recovery of a Windows System” on page 550), back up your data using the default 64KB block size.
- Library: If you are using several drive types of similar technology in the same library, the drives must have the same block size.
- SCSI adapters: Check if the selected block size is supported by the host SCSI adapter the device is connected to.

- Object copy functionality: The destination devices must have the same or larger block size than the source devices.
- Mirroring: Block size of devices must not decrease within a mirror chain. The devices used for writing mirror 1 must have the same or a larger block size than the devices used for backup; the devices used for writing mirror 2 must have the same or a larger block size than the devices used for writing mirror 1, and so on.
- For other limitations, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

Changing the Block Size in Data Protector

You can set the block size in the `Advanced Options` dialog box of a specific device. For detailed steps, refer to the online Help index keyword “block size”.

3 Configuring and Using Disk-Based Devices

In This Chapter

This chapter includes information on the following topics:

NOTE Backup devices are subject to specific Data Protector licenses. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.

Overview

Data Protector has a selection of devices designed to perform backup and restore to and from disks. These devices are referred to as disk-based devices because they are designed to back up data to disk as opposed to tape. The devices vary in their functional sophistication and expected uses. The devices are:

Standalone File Device

The standalone file device is the simplest disk-based device. It is configured manually. It is not possible to change the properties of the device once it has been created. The recommended maximum capacity of data that can be stored in the standalone file device is up to 2 TB. The main problem with using this device is that once it has been created it is not possible to reconfigure the device during use. For detailed information about the standalone file device see “Standalone File and File Jukebox Devices” on page B-1.

File Jukebox

The file jukebox device is a logical equivalent of a tape stack. It contains slots whose size is defined by the user during the initial device configuration. This device is configured manually. The file jukebox properties can be altered while it is being used. The recommended maximum data storage capacity of this device is limited only by the amount of data that can be stored in a filesystem by the operating system on which the file jukebox is running. For detailed information about the file jukebox device see “Standalone File and File Jukebox Devices” on page B-1.

File Library Device

The file library device is the most sophisticated disk-based device. It is designed to carry out unattended back up and restore of large amounts of data. It can be automatically configured using a wizard in the Data Protector GUI. As with the file jukebox, the recommended maximum storage capacity of this device is limited only by the amount of data that can be stored in a filesystem by the operating system on which the file library device is running.

Recommended Use

Of the three disk-based devices, the file library device is recommended for use as an unattended back up device. The standalone file device was originally developed to hold small amounts of data. It has very simple functionality. The standalone file device was subsequently extended to be the file jukebox device, this is slightly more sophisticated, but is still less flexible than the file library device. For more information about the configuration and use of the standalone file and file jukebox devices, see “Standalone File and File Jukebox Devices” on page B-1. The file library device functionality is documented in the following sections of this chapter.

Introduction to the File Library Device Functionality

File Library Device Directory Structure

The file library is structured like a file system with directories containing files. The directory is created by you when initially configuring the file library device. The files are called file depots. The diagram below illustrates the file library directory structure.

Figure 3-1 **File Library Device Directory Structure**

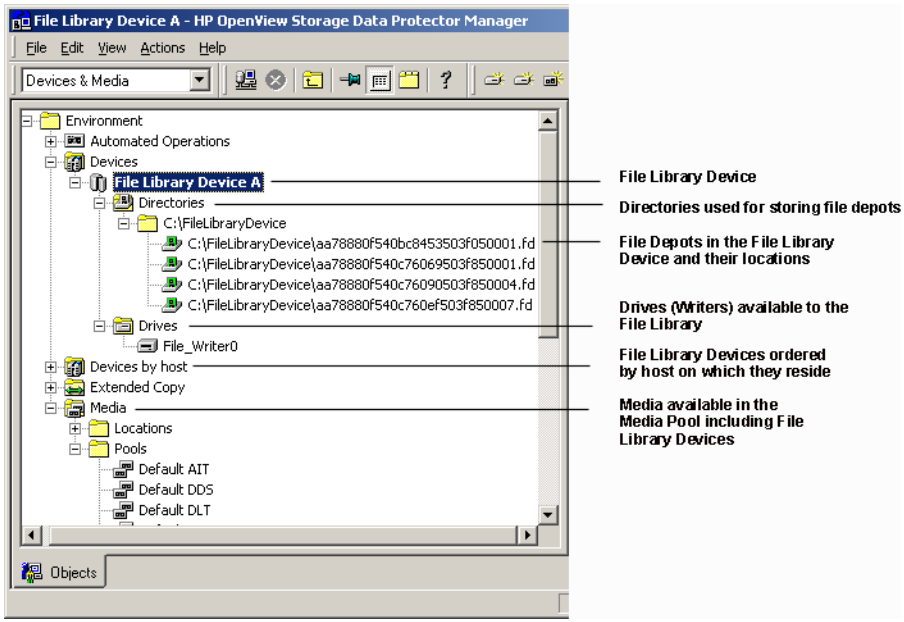

What is a File Depot?

A file depot is a part of the file library where data is stored. A file depot is automatically created each time a backup or copy is made to the file library device.

NOTE In this chapter the term back up is used to cover both a copy session and a backup session. Each of these operations saves data to the device.

File Depot Creation When the first backup is started using the file library device file depots are automatically created in the device by Data Protector. Data Protector creates one file depot for each data backup session made using the device. If the amount of data being backed up is larger than the default maximum file depot size Data Protector creates more than a single file depot for a backup session.

File Depot Name The name of each file depot is a unique identifier which is automatically generated by the system.

Data Protector also adds a media identifier to the file depot. This identifies the file depot as a media in the media pool. The identifier added to media helps to identify a particular backup session when performing a restore. The identifier can be seen in the `Description` field of the `Results Area` and when the file depot properties are viewed. For more information about accessing file depot properties refer to the online Help index keyword “setting file library device properties”.

Figure 3-2 Multiple File Depots

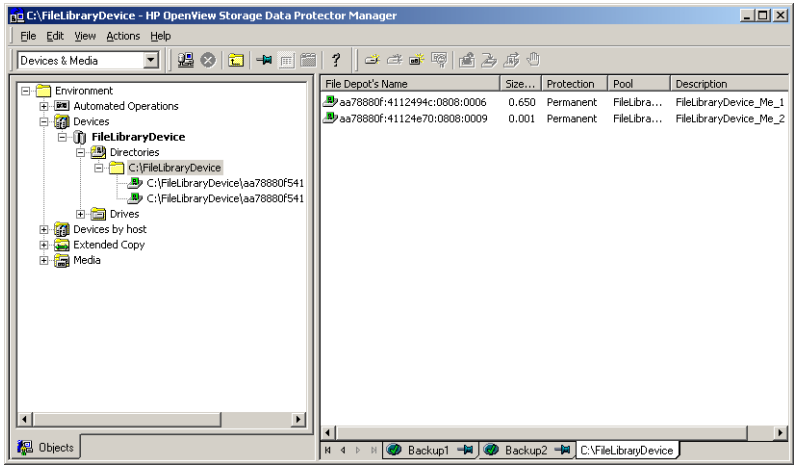

NOTE

Note that if the file depot has been recycled the file depot name may disappear from the GUI although the file depot icon is still visible in the GUI. For more information about this see “Recycling and Deletion” on page 124 of this section.

File Depot Size

The size of file depots is defined when you initially create the file library device. During this process you specify all sizing properties for the device, including the maximum size of the file depots. The sizing properties of the file depots, although only entered once, are globally applied to each file depot. If the size of data to be backed up within one session is larger than the originally specified file depot size, Data Protector automatically creates more file depots until the allocated disk space for the file library device has been consumed.

On Windows the maximum recommended file depot/slot size is 50 GB, although the file library device has been tested on Windows with file depots of up to 600 GB. On Unix the maximum recommended file depot size is 2 TB.

File Depot Space Consumption

Data Protector automatically creates file depots until there is no more disk space available for the device. The amount of space which must stay free for the file library device is defined in the device properties when the

device is initially being set up. See “Setting the File Library Device Properties” on page 114 for information about setting file depot properties.

Disk Full Handling

If the total disk space available to the file library device goes below a user specified level, a notification appears in the event log.

If there is insufficient disk space for a backup, a space information message is displayed. See “File Library Device Disk Full” on page 702 for information about how to deal with the disk being full.

Viewing the Contents of the File Library Device

It is possible to see the contents of the file library device from two viewpoints:

- Devices
- Media

Devices View

Using the devices view you can see a listing of all the devices registered in the internal database. This view enables you to perform all possible operations on the file library device. It is this view which is used to create and configure file library devices. For more information about the types of operations you can perform in the Devices view see “File Library Devices View” on page 117.

Media View

The media view shows you a listing of all the media assigned to a particular media pool. The media pool is a collection of all the media of the same type which can be used to create backups. For more information about media pools see Chapter 5, “Managing Media,” on page 143 of this manual. For information about the type of operations you can perform on file library devices in the Media view see “File Library Media View” on page 120.

Creation and Configuration of the File Library Device

File library device configuration defines the essential characteristics of the device. Device creation and configuration are carried out using a wizard in the Data Protector GUI.

Configuring the File Library Device

Configuring a file library device consists of the following steps:

1. Creating one or more directories on the disk(s) for use by the file library device.
2. Defining the device properties using the Data Protector GUI wizard.

After the device has been created and a backup has been made to the device file depots will be created in the device. These contain the backed up data.

File Library Device Properties

The properties of the file library device can be changed at any time after the device is created. See “Setting the File Library Device Properties” on page 114 for more information about setting and modifying the file library device properties.

Media Pool

A new media pool is created for each new file library device unless otherwise specified in the device specification. For more information about media pools see Chapter 5, “Managing Media,” on page 143 in this manual.

Maximum Size of File Library Devices

A file library device consists of a set of file depots therefore the maximum size of the file library device is equal to the maximum size of the file system. Check the operating system documentation for information about the maximum file sizes the particular operating system can manage.

Number of Devices Per Disk

The file library device can include one or several directories. Only one directory can be located on a file system.

In situations where the file depots are located on a variety of disks, it is not recommended to put file depots from two different file library devices on a single disk. This is owing to the fact that if the properties are different, it can cause a conflict in Data Protector. An example of a conflict situation would be if one file library device has properties which specify that the remaining disk space for the file depot should be 20 MB and the other file library device specifies that the disk where the file depot resides should have 10 MB of remaining space.

Figure 3-3 File Depots Properties Conflict

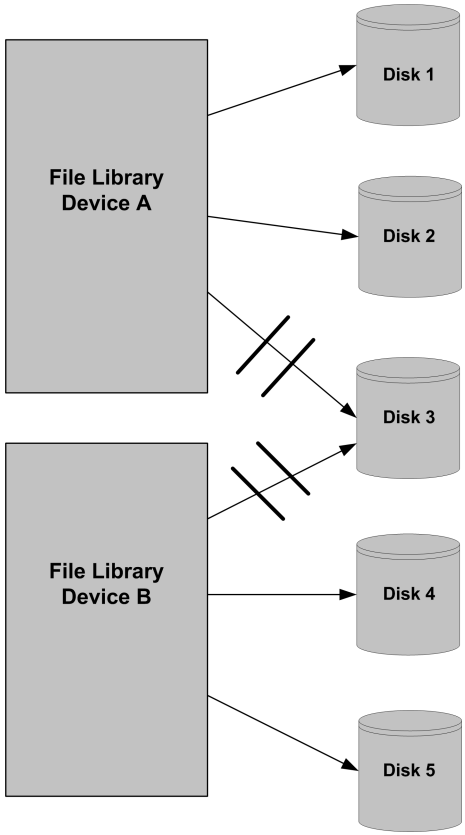

Creating a File Library Device Using the File Library Device Wizard

You can configure the file library device using the file library device wizard in the Data Protector GUI.

Prerequisites

Before the file library device can be created the following prerequisites need to be satisfied:

- The disk on which the file library device will reside must be visible in the file system in which the file library device resides.
- The directory in which the contents of the file library device are to be created must exist on the disk where the file library device will reside.
- If you want to create the file library device on a Windows filesystem, disable the compression option before creating the device.

How To Create a Directory for the File Library

To create a file library device, create a directory on either a local or shared disk where you would like the device to reside, for example `C:\FileLibrary`.

You can locate the file library on a network drive or NFS mounted filesystem on UNIX. The network shared disk needs to be mapped to a drive:

```
S:\datastore\My_FileLibrary
```

not a server name:

```
\\hostname\share_name
```

Shared network drives are not shown in the `Browse Drives` dialog where you enter the file library path. You need to enter the path to shared network drives yourself.

HP recommends that the disk where the file library device resides is local to the Media Agent. If not, file library device performance could be slow.

CAUTION

It is critical that the directory created for the file library is not deleted from the disk. If it is deleted, any data within the file library device will be lost.

IMPORTANT

For a backup to the device to work, change the Data Protector Inet account on the Media Agent client to establish the required permissions to access the shared disk you want to back up.

How To Define the Device Characteristics

Start Data Protector and select the `Devices & Media` context. Right-click `Devices` and click `Add Device`. For a detailed description of how to add a new file library device refer to the online Help index keyword “configuring file library devices”.

How To Define Device Properties

During the course of configuring the file library device it is possible to set the device sizing properties. For more information about this see “Setting the File Library Device Properties” on page 114.

What’s Next?

At the time the file library is created it does not contain any file depots. File depots are created in the file library as needed when a backup is made to the device.

Setting the File Library Device Properties

File Library Device properties are set during device creation. Properties can be changed at any time during device use.

What Properties Can Be Defined?

There are two types of properties that can be set for the file library device:

- Device properties
- Media pool properties

Device properties

The device properties determine all of the devices’ characteristics. The device properties include the device name, the host the device resides on and the media pool to which the device is assigned. The media type has the value “File” and cannot be changed. For detailed information about accessing the device properties refer to the online Help index keyword “setting file library device properties”.

Device properties also determine the amount of disk space which must be available for the device to operate. Device properties are specified using the Properties dialog. The Properties dialog is accessible through the device properties pane.

The Properties Dialog

The `Properties` dialog of the `File Library Device Wizard` enables you to specify the sizing properties of the file library device.

Data Protector applies the properties specified in the `Properties` dialog to each file depot created in the file library device after the device properties were changed. The properties of any file depots created before the subsequent device properties changes will not be affected.

Figure 3-4 Properties Dialog

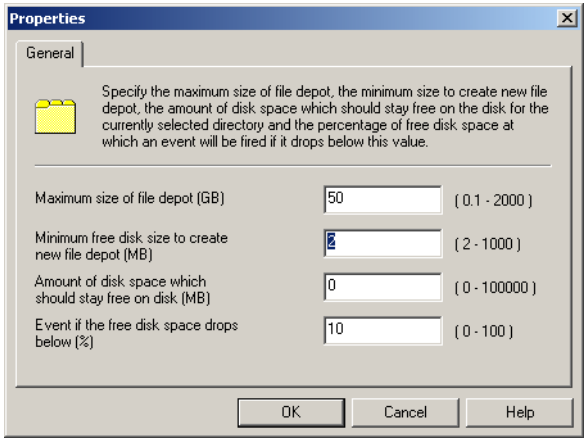

Table 3-1 Properties Dialog Possible Values

Field Name	Default Value	Minimum Value	Maximum Value
Maximum size of file depot (GB)	50	0.1	2 TB minus 1 MB
Minimum free disk space to create new file depot (MB)	2	2	1000

Table 3-1 Properties Dialog Possible Values

Field Name	Default Value	Minimum Value	Maximum Value
Amount of disk space which should stay free on disk (MB)	0	0	10^5 (100000)
Event if the free disk space drops below (%)	10	0	100

Calculating the Properties

Data Protector calculates the file library device properties as shown in the following example:

Disk Properties

Total directory size = 100 GB

Used directory space = 20 GB

File Library Device Properties

Free disk space setting = 30 GB

Available Writable Space Calculation

Available space minus Free disk space 80 GB - 30 GB = 50 GB

Available writable space = 50 GB

Data Protector performs the calculation above before using a file depot. If the amount of space is insufficient, Data Protector sends a message to notify you. This check is made at the start of each back up session. For information on how to deal with a lack of disk space, see “File Library Device Disk Full” on page 702.

Changing the File Library Device

Once the file library device has been created, Data Protector provides a range of facilities for both checking and changing the device configuration.

File Library Devices View

The view of the file library device enables you to modify the behavior and properties of the device and the file depots in the device.

Accessing the Devices View

To access the devices view do as follows:

1. Open Data Protector and choose `Devices and Media` in the `Context List`.
2. Expand `Devices` and click the desired device.
3. To display directories or drives of the device, click the appropriate item.
4. Select the name of the file library device in which you are interested from the list of devices.

Accessible Items

The following are the GUI items it is possible to access in the `Devices` view.

- Directories
- Drives

Directories

The Directories item shows you the directories contained in the file library device. It is from here that it is possible to access the file library depots contained within the directory or directories.

Drives

The Drive item of the file library shows which drives write to the file library device. A drive is a ‘writer’ which passes data to the backup device.

Performing Operations on Drives

The Drives part of the Devices view enables you to perform operations on the drive where the File Library resides. Each menu item is described in this section.

Scan

Scanning a drive updates the information in the IDB about the File depots in the file library. Use this option if you have moved one of the file depots to another location.

NOTE

Drive scanning is only available if the device contains file depots, which means that a backup has been performed to this device.

Delete

The Delete option removes the drive from the IDB. The file library device cannot be used if it has no drives. If you have deleted all drives, you will need to create new ones to be able to use the device for backup and restore.

Properties

The properties option enables you to see the general properties of the drive. This includes the device name using the drive, the device type and the host where the drive resides. It is also possible to check and modify the properties of the writers (drives) to the device such as cyclic redundancy checking, block and segment sizes of data, the delay between when a mount request comes to the drive and when writing begins and a locking name for the drive.

Performing Operations on the File Library

Using the right-click menu on a file library device name in the Devices view enables you to delete a device or view its properties. Each menu item is described in this section.

Delete

The Delete option enables you to delete the file library from the IDB. It is only possible to do this when the device does not contain protected data. For more information see “Recycling and Deletion” on page 124 or refer to the online Help index keyword “deleting file library devices”.

Properties

The Properties option enables you to see and change the device configuration. You can use this option to rename the device, change the sizing properties of each directory in the device and remove the directory

from the IDB. For further information about changing the sizing properties of file depots (directories) see “Setting the File Library Device Properties” on page 114.

Performing Operations on File Depots in the Devices View

Using the right-click menu on a file depot in a file library, you can scan, format, import, export, or recycle a file depot, import the catalog, or view the properties of a file depot.

Scan

Data Protector provides the scan facility for running a diagnostic on each file depot. If you select a file depot and then run a scan on it, Data Protector provides you with information about the full path to the depot, the type of medium stored in the depot and the label the depot has in the IDB.

NOTE

It is not possible to scan a file depot in the file library device until the first backup has been made to the device.

Format

Formatting a file depot prepares it for use with Data Protector by saving information about it in the Data Protector IDB. The type of information stored is its media ID, description and location. When you format a file depot you also specify to which media pool it will belong.

Under normal circumstances, it will not be necessary to format file depots because they are formatted by default when they are created. If a medium from another device is imported into the file library and you wish to change its characteristics for example, assign a different block size, you will have to format the medium.

Import

You can import file depots which have been exported from the file library. It is only possible to import file depots which previously belonged to the file library device and which have previously been exported.

Export

File depot export is used to remove the contents of the depot from the IDB. This utility is used as part of depot deletion. For more information about deleting file depots see open the Help key index and enter the keywords “recycling a file depot”.

Recycle

Recycling file depots sets the data protection on the file depot to 'None'. The next time a backup is made to the file library device this space will be overwritten with the backed up data. For more information about recycling file depots open the Help key index and enter the keywords "recycling a file depot".

Import Catalog

If a file depot's catalog protection has expired, you cannot browse the items available for restore in the Data Protector GUI. Importing the catalog from the IDB enables you to browse those items again.

Import Catalog can also be used in where the log level on the file depot was set to a low level and you want to increase it to "log all". Import Catalog allows you to re-set the log level to a higher one.

Properties

The properties option enables you to see all the available information about each file depot. It is possible for you to modify the file depots' properties. For information about modifying properties see "Setting the File Library Device Properties" on page 114 or open the help key index and enter the keywords "setting file library device properties".

File Library Media View

The Media view enables you to modify the contents of the device from the media pool point of view. A media pool represents a set of media of the same type that you use for backups. A new media pool is created for each file library device with the naming convention: `<LibraryName>_MediaPool`.

Media Pool Setup

It is possible to change the setup of a media pool. For complete information about managing media in Data Protector see Chapter 5, "Managing Media," on page 143.

The media view enables you to perform a variety of operations on media belonging to a certain media pool and the individual file depots within the devices.

Accessing the Media View

To access the media view do as follows:

1. Open Data Protector and choose `Devices and Media` in the `Context List`.
2. Click on `Media`.
3. Click on `Pools`.

4. Double-click on the name of the file library media pool in which you are interested.

Performing Operations on a Media Pool Used By a File Library Device

This section describes each menu item that appears allowing you to perform operations on the media pool to which the File Library belongs.

Format

It is possible to select a medium from a media pool and format the medium. Formatting a medium saves information about it in the Data Protector IDB.

NOTE

You cannot format a file library device until after the first backup has been made to it. This is because before this point, the device does not contain any file depots, and you cannot create them manually. File depots created during backup are the equivalent of a medium. According to the file library device's media pool media allocation policy, newly formatted media are automatically deleted.

When formatting a medium do not format it to a size greater than that allowed by the file system on which the device resides. For example, if you use Windows FAT as the file system, the file size is limited to 4 GB.

Import

Importing a medium to the media pool imports a medium into the present media pool.

This option is normally only used when you have exported a file depot. It can be used to re-import file depots which were exported in preparation for device deletion. For more information about deleting file library devices see "Recycling and Deletion" on page 124 or refer to the online Help index keywords "deleting file library devices".

NOTE

If you want to import media from a file library residing on a host other than the target host, it can only be done to a jukebox device. Refer to the online Help index keywords "importing file library media" to see the complete procedure.

Delete

The Delete option enables you to remove the entire media pool. This is only possible when the file library device connected to the media pool does not contain any protected data. For information about changing the data protection level on a piece of data within the file library device refer to the online Help index keywords “deleting file library devices”.

Select Media

The Select Media option is used for copying the contents of media to another location. This option can only be used when both the source and target media are of the same type. In the case of the file library device, the Select Media option can only be used to copy data from the disk on which it resides to another disk.

Properties

The Properties option enables you to see information about the media pool. This includes information such as which file library the media pool is used by, what media allocation policies are being used in the pool, the amount of time the media will be considered valid for backup maximum number of times the media may be overwritten during its life. It is possible to change the properties of the media pool using the ‘Properties’ option. For more information open the Help key index and type in the keywords “properties of”.

Performing Operations in the Media View

This section describes each menu item that enables you to perform operations on media. in this case file depots, in the media view.

Export

Exporting a file depot removes information about the file depot from the IDB. Data Protector no longer recognizes that the file depot exists. The depot information is however still retained, and can later be imported if it is necessary to recover the file depot. Export is a precursor to deleting a file library device. For more information about this see “Performing Operations on File Depots in the Devices View” on page 119 or open the Help key index and type in the keywords “deleting file library devices”.

Change Location

The Change Location option enables you to insert information about the current location of the medium. This information will not be evaluated by Data Protector, it is for user information purposes only.

Recycle

Recycling a medium changes the level of protection that the data held in the medium has to ‘None’. This means that the space used by the medium on the disk can be re-used during the next backup. Recycling a medium is part of deleting a file library device. For further information

about this see “Recycling and Deletion” on page 124 of this chapter or open the Help key index and enter the keywords “deleting file library devices”.

Move To Pool This option enables you to move the media from one media pool to another media pool containing media of the same type. You need this option if you want to reorganize the backups and re-designate the purpose of each media pool.

Copy It is possible to copy the contents of a file depot to another media. This can be useful when you want to copy the data elsewhere for archiving purposes, or keep a copy at another location for restore. This option only works when the target media is of the same type as the source media. In the case of the file library device the target media must be of type disk as the source media type is a disk.

Verify Verifying a file depot is useful when you want to check whether the media can still be restored. Verify checks a variety of aspects of the file depot’s quality, such as all identification information, it reads all the blocks of information and checks that cyclic redundancy checking was used when the file depot was originally written.

Import Catalog This option is used when the catalog data generated when the file depot was created has expired. Import catalog re-imports catalog data from the media to the IDB. Catalog data is required to generate a list of media for browsing when you are performing a restore.

Select Media This function is used to search and select particular media, in this case file depots, without browsing the entire list of media in the media pool.

Properties The properties option shows you the properties of the media within the media pool. This option provides you with information about the media including the name of the media, the media ID within the IDB, the name of the media pool it is allocated to, a list of the objects that are stored in it and details of how often it is used. You can use the Properties option to change the media description string and to allocate the order in which media in a media pool are used for backup.

NOTE By default, a media pool for a file library device has a media allocation usage policy of non-appendable. This means that pre-allocation of media for backup is not possible because by default, before making a backup,

the file library device first checks to see whether there is any unprotected data in the device. If there is unprotected data in the device, it is deleted and the new data is then written to the device. However, the media pre-allocation policy still works if you set it to appendable. For information about changing the media allocation policy refer to the online Help index keyword “preallocating media”.

Recycling and Deletion

Disk space can be freed by recycling and deleting file depots or entire file library devices.

Recycling

If you want to re-use the space allocated to the file library device you can recycle either individual file depots or all of the file depots in a file library. This means that the disk space occupied by the recycled item can be recovered and used in the next backup. This is done by deleting the unprotected file depot(s) and creating new ones.

Exporting

Export is the option used to remove an item from the Data Protector IDB. This option can only be used when there is no data protection on the item, therefore, to use the Export option it is first necessary to recycle the item.

Once a file depot has been exported its name disappears, and only the depot icon is visible in the Data Protector GUI as shown below.

Figure 3-5 Exported File Depot Label

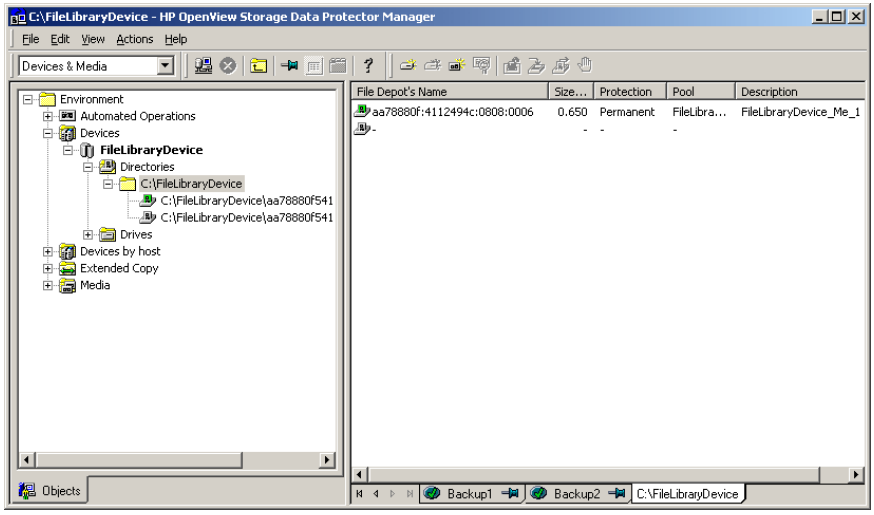

It is possible to delete the depot icon after exporting the file depot. This deletes the icon from the list on the GUI, but does not physically delete the file depot from the IDB. For detailed steps describing how to delete file library device icons refer to the online Help index keywords “deleting file library devices”.

Deletion

Deletion is a three phase process. Firstly data protection must be removed from the item to be deleted. This is the recycling process. After this, the item must be exported from the Data Protector IDB using the Export option. Finally the item can be deleted. It is possible to use these options both on a single item and on an entire file library.

Deletion of the file library is only possible when the device does not contain any protected data. The file library device contents are held in file depots. Before you can delete an entire file library device you need to first modify the level of data protection on each file depot and then export each file depot. At this point you can delete the device.

For detailed steps describing how to delete file library devices, refer to the online Help index keywords “deleting file library devices”.

File Library Device Command-Line Interface Options

There are two utilities that you can use to manipulate the file library device from the command-line interface (CLI), these are:

- `omniupload`
- `omnidownload`

`omniupload` is used to create the file library device, and `omnidownload` is used during device deletion. For more information about the way these commands are used see the `omniupload` and `omnidownload` man pages.

4 Configuring Users and User Groups

In This Chapter

This chapter explains how to configure both user groups and individual users. It contains information about the following subjects:

Data Protector User Rights

Data Protector users have the user rights of the user group they belong to. For example, all members of the Admin user group have the rights of the Data Protector Admin user group.

When configuring a Windows user in a Data Protector cell running the Cell Manager on the HP-UX or Solaris platform, the user has to be configured with the Domain Name or the wildcard group `"*"`.

The Data Protector user rights are described below:

Clients configuration	Allows the user to install and update Data Protector software on client systems.
User configuration	Allows the user to add, delete, and modify users and user groups. Note that this is a powerful right.
Device configuration	Allows the user to create, configure, delete, modify, and rename devices. This includes the ability to add a mount request script to a logical device.
Media configuration	Allows the user to manage media pools and the media in the pools and to work with media in libraries, including ejecting and entering media.
Reporting and notifications	Allows the user to create Data Protector reports. To use Web Reporting you also need a Java user under the Applet domain in the Admin user group.
Start backup	Allows users to back up their own data as well as monitor and abort their own sessions.

Start backup specification	Allows the user to perform a backup using a backup specification, so that the user can back up objects listed in any backup specification and can also modify existing backups.
Save backup specification	Allows the user to create, schedule, modify, and save any backup specification.
Back up as root	Allows the user to back up any object with the rights of the root login on UNIX clients. This user right is effective only for UNIX clients. It is required to run any backup on Novell NetWare clients. This user right has no effect on Windows systems. Backups on Windows always run in the context of the Data Protector `Inet` service.
Switch session ownership	Allows the user to specify the owner of the backup specification under which the backup is started. By default, the owner is the user who started the backup. Scheduled backups are started as root on a UNIX Cell Manager and under the Cell Manager account on Windows systems. This user right is appropriate if the Start backup specification user right is enabled. See “Ownership: Who Will Be Able to Restore?” on page 279 for more details.
Monitor	Allows the user to view information about any active session in the cell and to access the IDB to view past sessions.
Abort	Allows the user to abort any active session in the cell.

Mount request	Allows the user to respond to mount requests for any active session in the cell.
Start restore	Allows users to restore their own data as well as monitor and abort their own restore sessions. Users that have this user right are able to view their own objects and public objects on the Cell Manager.
Restore to other clients	Allows the user to restore an object to a system other than the one from where the object was backed up.
Restore from other users	Allows the user to restore objects belonging to another user. It is effective only for UNIX clients.
Restore as root	Allows the user to restore objects with the rights of the root UNIX user. Note that this is a powerful right that can affect the security of your system. This user right is required to run any restore on Novell NetWare clients. This user right has no effect on Windows systems. Restores on Windows always run in the context of the Data Protector `Inet` service.
See private objects	Allows the user to view and restore objects that were backed up as private.

Predefined Data Protector Users and User Groups

The following default groups are provided: Admin, Operator, and User.

User Rights	Admin	Operator	User
Clients configuration	Y		
User configuration	Y		
Device configuration	Y		
Media configuration	Y	Y	
Reporting and notifications	Y		
Start backup	Y	Y	
Start backup specification	Y	Y	
Save backup specification	Y		
Back up as root	Y		
Switch session ownership	Y	Y	
Monitor	Y	Y	
Abort	Y	Y	
Mount request	Y	Y	
Start restore	Y	Y	Y
Restore to other clients	Y		
Restore from other users	Y	Y	
Restore as root	Y		
See private objects	Y	Y	

TIP To see the exact user rights for each user group, select the group, right-click it, and select `Properties` from the menu.

The user rights you have set on the Cell Manager determine the availability of the Data Protector Cell Manager GUI or GUI contexts to the computer from which you connect to the Cell Manager. For example, if you have only the `Start Restore` user right set, then only the `Restore` context is available when you install the User Interface component.

After the initial installation, all default user groups are empty except for the Admin group. Data Protector adds the following users to the Admin group:

- Cell Manager on UNIX:
  - The `root` user on the Cell Manager (`root`, `<any group>`, `<Cell Manager host>`).

    This user account should not be changed. It is required for proper operation of the `CRS` daemon and other processes on the Cell Manager. Only the `root` user on the Cell Manager is initially allowed to administer the cell. To administer the cell from any other client, add a new user first.
  - The `java` user (`java`, `applet`, `webreporting`).

    This user account enables Web reporting. It needs to be modified when certain security settings are changed. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information.
- Cell Manager on Windows:
  - The CRS service account, as specified during the Data Protector installation (limited to the Cell Manager).

    The CRS service account should remain unchanged unless you modify the logon parameters of the CRS service. It is required for proper operation of the CRS service and other processes on the Cell Manager.
  - The local system account on the Cell Manager (`SYSTEM`, `NT AUTHORITY`, `<Cell Manager host>`).

    This account is provided in case the `CRS` service is configured to log on as the local system account.

— The user who installed the Cell Manager (the initial cell administrator; allowed from any client).

This user is configured as the initial cell administrator and can administer the cell from any client. It is recommended to modify this user account after the Data Protector installation is complete. Specify the client from which you will administer the cell instead of allowing access from any client. If you will be using another account, add this account and then remove the "initial cell administrator" or allow it only from the Cell Manager host.

— The `java` user (`java`, `applet`, `webreporting`)

This user account enables Web reporting. It needs to be modified when certain security settings are changed. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information.

IMPORTANT Admin capabilities are very powerful. A member of the Data Protector Admin user group has system administrator capabilities for the whole cell.

It is recommended to define specific groups for each type of users in an environment to minimize the set of rights assigned to them.

IMPORTANT For further information on security refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Adding or Deleting a User Group

The default Data Protector user groups are sufficient for most needs. However, due to security reasons it is recommended to define specific groups for each type of users in an environment to minimize the set of rights assigned to them. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for security considerations.

Adding a User Group

1. In the `Data Protector Manager`, switch to the `Users` context.
2. In the Scoping Pane, right-click `Users`, and then click `Add User Group`. The `Add User Group` wizard appears.
3. Follow the wizard. For further information, press **F1**.

Figure 4-1 Adding New User Groups

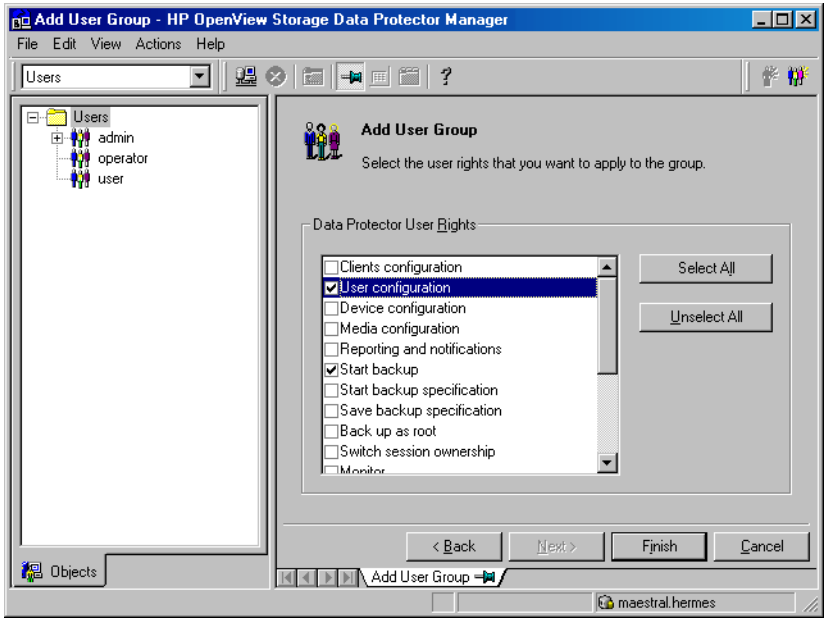

Deleting a User Group

1. In the `Data Protector Manager`, switch to the `Users` context.
2. In the Scoping Pane, expand `Users` to display the user groups.
3. Right-click the user group to be deleted and click `Delete`.
4. Confirm the action.

Adding or Deleting a User

After the product installation, the following users are configured in the Admin user group:

- UNIX root user on UNIX systems
- Windows administrator on Windows systems
- The user performing the installation

By adding a new user to one of the Data Protector user groups you assign this user the rights of that particular group. See “Data Protector User Rights” on page 129 for a description of the user rights.

NOTE Before you can start using the Data Protector GUI on the client system, add a user from that system to an appropriate Data Protector user group on the Cell Manager.

You can configure users from both UNIX and Windows environments.

UNIX users are defined by their login name, UNIX user group, and the system from which they log on. A wildcard character (*) can be used.

Windows users are defined by their logon name, Windows user group (domain), and the system from which they log on. A wildcard character (*) can be used.

To add a user, do the following:

1. In the `Data Protector Manager`, switch to the `Users` context.
2. In the Scoping Pane, expand `Users`.
3. Right-click the group to which you want to add a user, or from which you want to delete a user, and then click `Add/Delete Users` to open the wizard.

Figure 4-2 **Adding New Users**

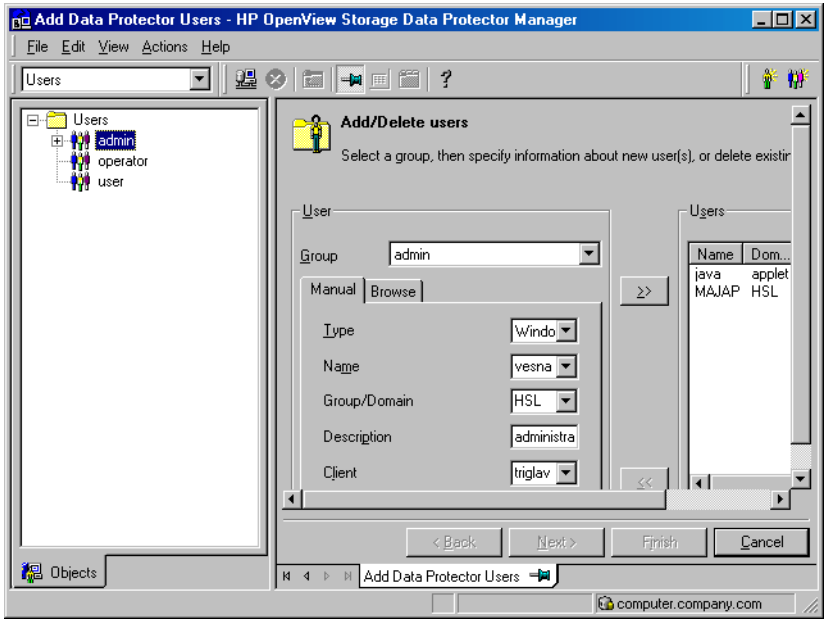

For further information, press **F1**.

Modifying a User

You can change the properties of an existing user, or move the user from one user group to another.

NOTE You cannot change user rights for individual users, but only for the entire user group.

Changing User Properties

To modify a user's properties, follow these steps:

1. In the `Data Protector Manager`, switch to the `Users` context.
2. In the Scoping Pane, expand `Users`, and click the user group to which the user belongs.
3. Right-click the user and click `Properties` to open the user's property page.

   For further information, press **F1**.

Moving a User to Another User Group

To change the user rights of an individual user, move the user to another user group.

1. In the `Data Protector Manager`, switch to the `Users` context.
2. In the Scoping Pane, expand `Users`, and click the user group to which the user belongs.
3. In the Results Area, right-click the user and click `Move`.

   For further information, click `Help`.

Changing User Group Rights

Users have the rights of the groups to which they belong. So, changing the user rights of the user group changes the user rights for all users in that group. You can change the rights of user groups and, in doing so, change the rights of each user within that group. You cannot change the rights of the Admin user group, however.

NOTE You can also modify the properties of each user within a group, for example the domain to which the user belongs, the user's real name, and the user's user group.

The following steps explain how to change user group rights, and consequently, the rights of each user in the group:

1. In the `Data Protector Manager`, switch to the `Users` context.
2. Browse for and select the user group whose rights you want to change.

NOTE If you select a group that does not have any users in it, the Results Area will display the properties for the group. If you select a group that has users in it, the Results Area will list the users in the group. You can also modify properties of each user in a user group by clicking on the user whose properties you want to modify.

3. Right-click the user group you selected, and then click `Properties`. The properties for the user group appear in the Results Area.
4. Click the `User Rights` tab to display the list of rights available to this group.

   For further information, press **F1**.

Example User Configurations

This section gives some examples of typical user configurations.

Allowing Users to Restore Their Own Files

This restore policy allows all or just selected users to restore their own data. It provides sufficient security and may relieve the backup operator from doing a number of restore operations.

When to Use This Policy

- When the users have sufficient knowledge to handle restores. You need to provide some way of training the users on basic backup concepts and restore operations.
- You use library backup devices with media of all most recent backups. The Data Protector User group by default does not allow users to handle mount requests for needed media. The users will still need an intervention from the backup operator in case of a mount request.

What Needs to Be Done?

1. Add the users who will be allowed to restore their own data to the Data Protector `users` user group. For additional security, you may limit the access to Data Protector for these users to a specific system only.
2. Install the Data Protector User Interface on the systems the users are using. Data Protector automatically checks the user rights and allows restore functionality only.
3. When you configure backup of the user systems, make backups visible to the users by setting it to public.

Enabling Users to Back Up Their Systems

Data Protector differentiates between the user's right to configure a backup and the user's right to run an already configured backup.

To create rights for a user to run their own backup, follow these steps:

1. Create a new user group or modify the existing group so that it has the `Start backup` user right.
2. Add the users who will be able to configure their own backups to this user group.

3. Change the owner of the backup configuration so that the users will be able to start these backups. See Figure 4-1 on page 135.

5 Managing Media

In This Chapter

This chapter gives detailed information on how to manage your media, including:

Overview of Data Protector Media Management

Data Protector provides a powerful media managing functionality that allows simple and efficient management of a large number of media.

NOTE

Data Protector recognizes and uses different format types to write data to media. For limitations incurred, refer to “Using Different Media Format Types” on page 192.

- Grouping media into logical groups called media pools, which enable you to manage large sets of media without having to worry about each individual medium.
- Data Protector keeps track of all media and the status of each medium, including data protection expiration time, availability of media for backup, and a catalog of what has been backed up to each medium.
- Fully automated operation. If Data Protector has control of enough media in the library devices, the media management functionality enables backups to run without the need for an operator to handle the media.
- Automated media rotation policies, so that you do not have to enforce policies manually.
- The ability to explicitly define which media and which devices you want to use for backup.
- Optimized media management for specific device types, such as standalone, magazine, library devices, and large silo devices.
- Automatic recognition of Data Protector media and other popular tape formats.
- Recognition and support of barcodes on large library and silo devices with barcode support.
- Recognition, tracking, viewing, and handling of media used by Data Protector in large library and silo devices.

- The ability to store information about media in a central place and share this information among several Data Protector cells.
- Support for **media vaulting**, also known as **archiving** or **off-site storage**.
- Interactive or automated creation of additional copies of the data on the media.

Information about the media used is stored in the IDB.

For more information on media management, see the *HP OpenView Storage Data Protector Concepts Guide*.

Figure 5-1 indicates the relationship among the components, backup specification, devices, and media pools. The pool is used during a backup session. A default pool is part of the device definition. However, a different pool can be specified in the backup specification.

Figure 5-1 **How the Media Pool Relates to Other Components**

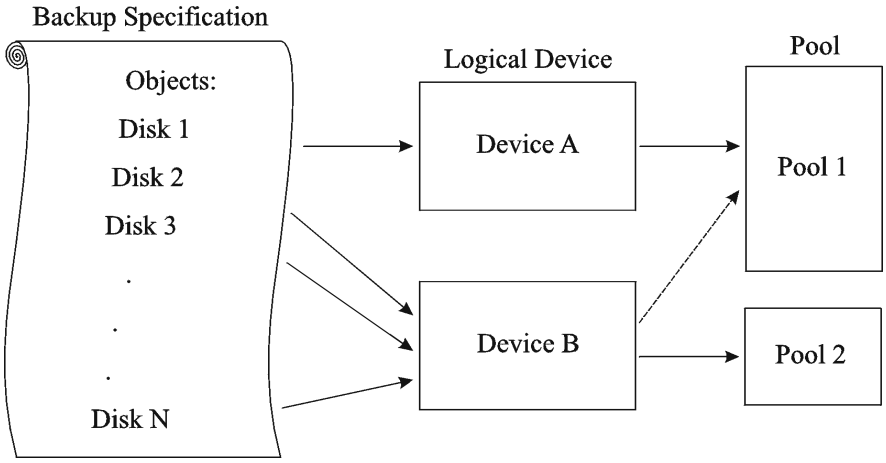

Media Life Cycle

A typical media life cycle consists of the following steps:

1. Preparing media for backup. This includes formatting media for use with Data Protector and assigning media to a media pool. The media pool is used to track these media. See the following topics for detailed information:

   “Creating a Media Pool” on page 149.

   “Adding Media to a Media Pool” on page 154.

2. Using media for backups. This includes how the media are selected for a backup, what media condition factors are checked (for example, the number of overwrites), how new backups are appended to the media, and when data on the media can be overwritten.
3. Vaulting media to a safe place (vault).

   You can use one of Data Protector’s data duplication methods to make copies of the backed up data for vaulting purposes.

4. Recycling media once data on the media is not needed anymore. These media can then be reused.

5. Retiring Media. Once the medium has expired (according to its maximum usage criteria), it is marked as Poor and no longer used by Data Protector. See “Factors Influencing the Condition of Media” on page 179 for more information.

Details are explained in the following sections.

Creating a Media Pool

What Is a Media Pool?

A media pool represents a set of media of the same type (for example DLT) used for backup, with the same usage policy and properties. For example, you may have one media pool for regular backup, one for archive backup, and one for each department.

What Is a Free Pool?

A free pool is an auxiliary source of media of the same type (for example, DLT) for use when all free media in a regular pool run out. This helps to avoid failed backups due to unavailable media.

Media are moved between regular and free pools in two events:

- Allocation. Media are moved from a free pool to a regular pool.
- Deallocation. Media are moved from a regular pool to a free pool. You can specify in the GUI whether deallocation is performed automatically.

Protected (allocated, used) media belong to a specific regular pool (such as a SAP pool), while free Data Protector media can be automatically moved to a free pool. This free pool is later used for allocation of free media to a specific regular pool during backup, when needed.

See the *HP OpenView Storage Data Protector Concepts Guide* for more information on media pools.

Default Media Pools

Data Protector provides default media pools for each media type that you can use in your initial configuration, for example: Default_DDS.

If you do not want to create a media pool at this time and to use the default media pools instead, go to “Adding Media to a Media Pool” on page 154 for instructions.

How to Create a Media Pool

Create a new media pool in the `Devices & Media` context using the `Add Media Pool` wizard. For detailed steps, refer to the online Help index keyword “adding media pools”.

What’s Next?

The next step is to add media that you want to use for backup to the media pool. See “Adding Media to a Media Pool” on page 154 for instructions.

Properties of a Media Pool

This section describes the properties of a media pool. You specify them when you are configuring the media pool. Some of the properties can be modified later.

Pool Name

A media pool name identifies a media pool. It can be up to 32 characters long, including spaces. You should assign a meaningful name that will help you identify the media pool later, for example, your department name.

Description

A description is optional and helps you to identify the media pool. It can contain any characters and can be up to 80 characters long.

Media Type

Data Protector shows you a list of available media types for your configuration.

For an up-to-date list of the supported media types, refer to the *Device Support Matrix* in the *HP OpenView Storage Data Protector Software Release Notes*.

Once you select the media type, Data Protector calculates the available space on the media for that media pool. This calculation is based on the selected media type.

Media Allocation Policy

The media allocation policy defines the order in which media are accessed within a media pool, so that media wear out evenly.

For more information on how Data Protector selects media for backup, see “Selecting Media for Backup” on page 166.

Strict Directs Data Protector to require a specific medium. The medium has to be already formatted for use with Data Protector. If this policy is used, Data Protector does not format media. This allocation policy should be used with library devices to prevent accidental overwrite of non-Data Protector media in the library and where even usage of media has priority.

Loose Directs Data Protector to accept any suitable medium in the pool except a medium in *poor* condition or a protected medium. This option is combined with the `Allocate unformatted media first` option.

If `InitOnLoosePolicy` is set to 1 (by default, it is set to 0) media that are unrecognized by Data Protector (new media) are automatically formatted. This policy is preferred if you want unattended backup to succeed, as it maximizes the number of media Data Protector can choose from.

Unformatted media first This is a modification of the `Loose` policy. If selected, this policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library. This is recommended if Data Protector is the only application using the library and you want to have even usage of all media.

Use free pool Directs Data Protector to search in the free pool for suitable media in addition to the regular pool. By default, this option is OFF.

See "Selecting Media for Backup" on page 166 for detailed information.

See the *HP OpenView Storage Data Protector Concepts Guide* for more information on free pools.

Media Usage Policy

Media usage policy controls how new backups are added to already used media.

Appendable A backup session starts writing data to the space remaining on the last medium used in the previous backup session. Subsequent media needed in this session are written from the beginning of the tape, hence only unprotected or new tapes can be used. Data may be appended from any backup specification to any other backup specification. Appending media conserves media space but can add complexity to a restore operation, because one medium can contain data from several backup sessions.

Non-Appendable A backup session writes data beginning at the first position on the first available medium for backup.

Appendable on incrementals only The first medium used in a backup session is appended to only if an incremental backup is performed. If several appendable media are available in the pool, the least recently written to

medium is used first. If additional media are needed during the same backup session, they must be free and not contain any protected backups. This media usage policy will create media which will contain a full backup, followed by any number of incremental backups.

NOTE If you use the append functionality and the backup requires more than one medium, only the first medium used can contain backed up data from a previous session. Subsequently, Data Protector will use empty or unprotected media only.

See “Appending Backups to Media” on page 163 and “Selecting Media for Backup” on page 166 for more information.

Magazine Support

Magazine support allows you to use a set of media configured as magazines. A backup device used with these media must have support for magazines, such as the HP 12000e.

You can set this option when you configure a new media pool.

See the following sections for more information:

- “Configuring Magazine Devices” on page 31 for instructions on how to configure a magazine device.
- “Formatting Media” on page 155 for instructions on how to format a full magazine or a single medium in the magazine.
- “Importing Media” on page 159 for instructions on how to import a full magazine or a single medium.

Media Condition Factors

Media condition factors define the status of the media, thus determining how long media can be reliably used for backup. If a pool uses the free pool option, the media condition factors are inherited from the free pool. Data Protector calculates the status of media in use via media condition factors. The two media condition factors you can select are:

Medium valid for The age of a medium is calculated as the number of months that have elapsed since it was formatted. Once a medium is older than the threshold number of months, it is marked as `poor`. The default threshold is 36 months.

Maximum number of overwrites The usage of a medium is defined as the number of overwrites from the beginning of the medium. Once the medium has more than the threshold number of overwrites, it is marked as `poor`. The default threshold is 250 overwrites, except for DDS tapes, for which it is 100 overwrites.

For more information on how media condition factors are calculated, see “Changing How Media Condition Is Calculated” on page 180.

Adding Media to a Media Pool

Once you have created a media pool, you have to add the media that you want to use for backup to this media pool.

How to Add the Unused Media

To add unused media to the media pool, see “Formatting Media” on page 155. If your media allocation policy for the media pool is set to `loose`, formatting media as a separate step is not required. If `InitOnLoosePolicy` is set to 1 (by default, it is set to 0), the media are formatted before the backup session in which they are used. See “Media Usage Policy” on page 151 for more information.

How to Add Used Media

To import previously used Data Protector media without overwriting them, see “Importing Media” on page 159.

To add used non-Data Protector media to the media pool, you have to reformat them. See “Formatting Media” on page 155.

For more information on how Data Protector handles media used by other applications, see “Recognizing Other Data Formats” on page 157.

Labeling Media

Data Protector labels each medium with a unique media label and medium ID. Both are stored in the IDB and allow Data Protector to manage the medium. The medium ID is assigned by Data Protector. The media label is a combination of the user-defined description and the barcode of the medium (if the medium has a barcode and the `Barcode Reader Support` option is enabled). The barcode is displayed as a prefix of the medium description. For example, `[CW8279]Default DLT_1` is a media label with the `Default DLT_1` description and the `CW8279` barcode. You can optionally write the barcode as medium label to the medium header on the tape during the initialization of the medium.

In the Data Protector GUI, you can sort media by media label. You do this by clicking the `Media label` field in the Results Area.

What’s Next?

Once you have added media to the media pool, you can select data that you want to back up. Refer to Chapter 6, “Backup,” on page 195 for instructions.

Formatting Media

What Is Formatting Media?

Formatting media prepares them for use with Data Protector by saving the information about the media (media IDs, description and location) in the IDB, and also writes this information on the medium itself (medium header). When you format media, you also specify to which media pool the media belong.

When to Format Media

You need to format media before the media can be used for backup. If the media are not formatted before backup and the `Loose` media allocation policy is defined for the media pool, and the global variable `InitOnLoosePolicy` is set to 1 (default is `0`) Data Protector automatically formats new media when they are selected for backup. In this case, the media are labeled with default values or with the barcode if the `Use barcode as medium label on initialization` option is selected in the properties of your library.

Non-Data Protector media must be formatted before backup.

Recognition of Other Formats

Data Protector recognizes common media formats, if the medium was already in use. See “Recognizing Other Data Formats” on page 157 for detailed information.

Formatting with Padding Blocks

You can extend the size of the medium header and fill it up with incompressible data, padding blocks. This becomes useful when creating media copies. The padding blocks are not copied to the target medium. This way you make sure that the target medium does not reach the end of the tape before the source medium.

Tape padding is not required if you copy backed up data using the object copy functionality.

Tape padding is disabled by default. To enable it, set the `OB2BLKPADDING_n` variable in the `omnirc` file on the system with the backup device connected. For more information, see “Using Omnirc Options” on page 615.

How to Format Media

To format media, browse for the specific device, media pool, or library slot in the `Devices & Media` context, right-click it and click `Format`. For detailed steps, refer to the online Help index keyword “formatting media”.

If you use library devices, you can select multiple slots using the `Ctrl` key and format several media in a single step. You can optionally use the barcode as medium label and write it to the medium header during the initialization of the medium. For detailed steps, refer to the online Help index keyword “formatting media in library devices”.

TIP

To format media used by other applications, use the `Force Operation` option. Data Protector protected media cannot be re-formatted using this option. You have to first remove the protection. See “Recycling Media” on page 169 for more information.

NOTE

When selecting the `Medium Size` option, choose between `Default` and `Specify MB`. If you have chosen the `Default` medium size, the estimated and not the real size of the media is shown. Be aware that the total media size is set for non-compressed media. Hardware compression of the device may double the space on the media. The correct media size is shown when the media are full.

What’s Next?

Once you have formatted your media, you may use the media for backup. See Chapter 6, “Backup,” on page 195 for more information on how to configure backups.

Formatting Media in a Magazine

If you are using a device with magazine support, Data Protector allows you to format all media or a single medium in the magazine.

How to Format a Full Magazine

To format a full magazine, browse for the media pool used for the device, right-click it and click `Format Magazine`. For detailed steps, refer to the online Help index keyword “formatting media in magazines”.

How to Format a Single Medium in a Magazine

To format a single medium in a magazine, browse for the media pool used for the device, right-click it and click `Format`. For detailed steps, refer to the online Help index keyword “formatting a single medium in magazines”.

TIP

To format media used by other applications, use the `Force Operation` option. Data Protector protected media cannot be re-formatted using this option. You have to first remove the protection. See "Recycling Media" on page 169 for more information.

What's Next?

Once you have formatted your media, you may use these media for backup. See Chapter 6, "Backup," on page 195 for more information on how to configure backups.

Recognizing Other Data Formats

Recognized Formats

To prevent accidental overwrite of data already written to the media, Data Protector recognizes a number of different tape formats:

Table 5-1 Data Protector Media Format Categories

Media Format	Data Protector Behavior
unknown or new	`Loose` Policy: formatted and used for backup only if the global variable `InitOnLoosePolicy` is set to 1 `Strict` Policy: not used for backup
media written with compression, now used without compression	
media written without compression, now used with compression	
foreign Data Protector (from another cell)	not used for backup unless imported or formatted with the `Force Operation` option
tar, cpio, OmniStorage, OmniBack I, ANSI label, filesystem	not used for backup unless formatted with the `Force Operation` option
Data Protector unprotected media	used for backup
Data Protector protected media	used for appending backups

NOTE Do not rely on Data Protector to recognize other media types, as recognition depends on the platforms you use.

NOTE If you try to read from a medium that was written using hardware compression with a device that does not support hardware compression, Data Protector cannot recognize the medium and read the data. Therefore, the medium will be treated as unknown or new.

Importing Media

Importing media adds media already used by Data Protector to a media pool, without losing the data on the media. Media used by Data Protector are media that were formatted by Data Protector, but exported from the Data Protector cell.

Importing a medium writes detailed information about backed up data on the medium to the IDB, so that you can later browse it for a restore.

Use media import when moving your media between Data Protector cells.

This operation is not available for media in free pools.

NOTE

Attribute information such as object or media size will not be reconstructed during import. Thus the size of the imported objects will be shown as 0 KB.

Importing can take a considerable amount of time, depending on the device and media used.

IMPORTANT

Import all media used in one backup session at once. If you add only some media from the backup session, you will not be able to restore data spanning to other media.

How to Import Media

To import media, browse for the specific device, media pool or library slot in the `Devices & Media` context, right-click it and click `Import`. For detailed steps, refer to the online Help index keyword “importing media”.

If you use library devices, you can select multiple slots using the `Ctrl` key and import several media in a single step. Refer to Figure 5-2. For detailed steps, refer to the online Help index keyword “importing media in library devices”.

Figure 5-2 **Import Multiple Media**

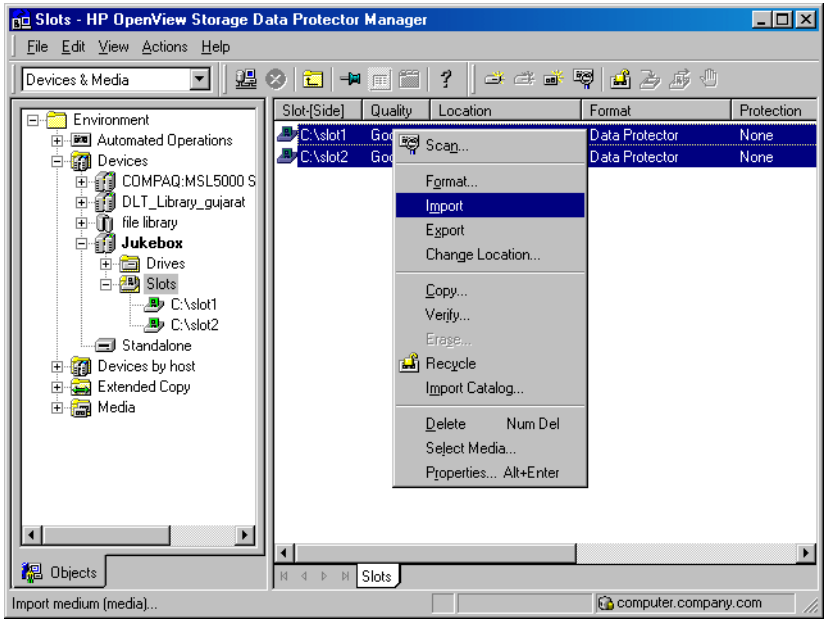

Importing the Catalog from Media

Importing the catalog from a medium writes the information about file versions into the IDB, enabling you to browse files and directories for restore.

Use `Import Catalog` if the catalog protection for a particular object has expired and you can no longer browse its files and directories.

This operation is not available for media in free pools.

How to Import the Catalog from Media

To import the catalog from a medium, browse for the specific medium, device or library slot in the `Devices & Media` context, right-click it and click `Import Catalog`. Refer to Figure 5-3. For detailed steps, refer to the online Help index keyword “importing catalogs from media”.

Figure 5-3 Import Catalog

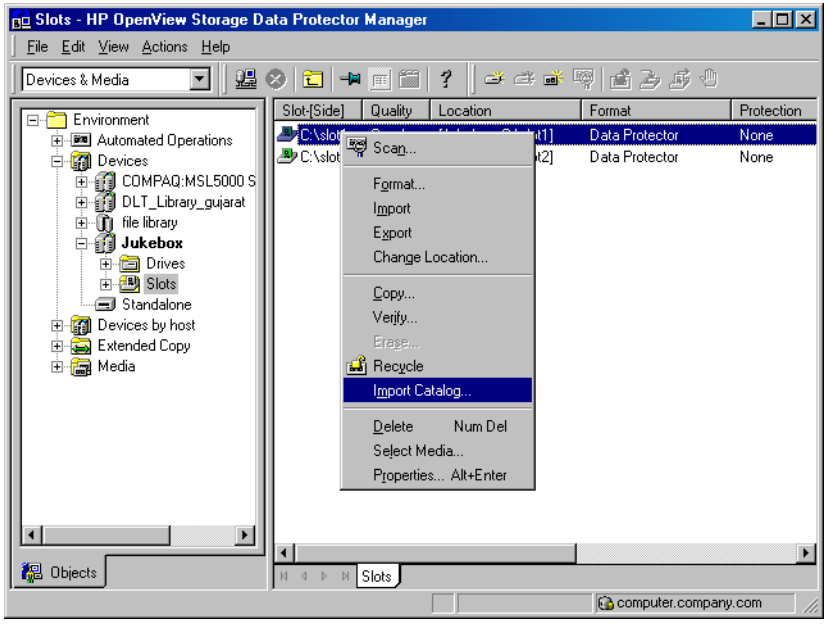

Importing Media in a Magazine Device

If you use a device with magazine support, Data Protector allows you to import all media or a single medium into the magazine.

Prerequisite

The media pool for the magazine device must be configured with the `Magazine Support` option enabled.

How to Import All Media

To import all media in a magazine device, expand the media pool used for that device in the `Devices & Media` context, right-click the `Magazines` item and then click `Import Magazine`. Refer to Figure 5-4. For detailed steps, refer to the online Help index keyword “importing media in magazines”.

Figure 5-4 Import Magazine

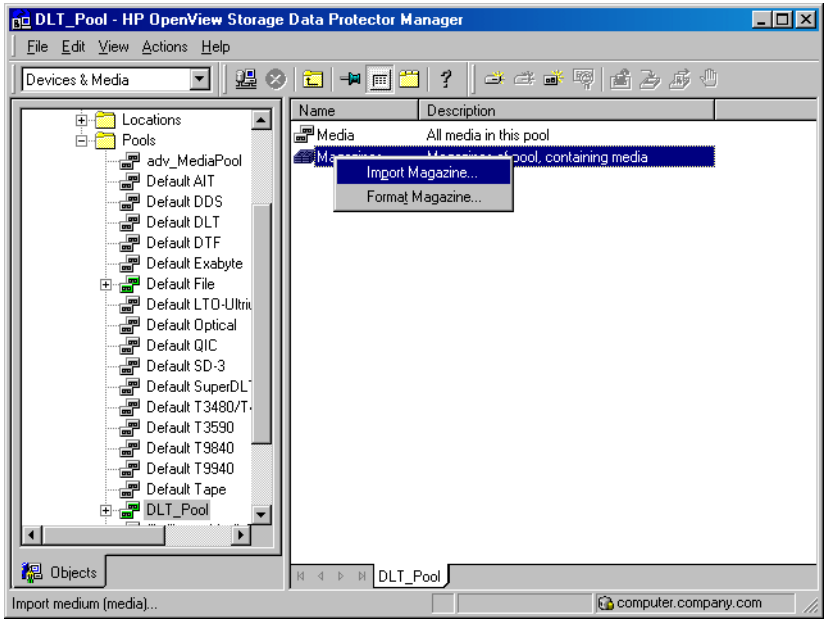

How to Import a Single Medium into a Magazine

To import a single medium into a magazine device, expand the media pool used for that device in the `Devices & Media` context, select the specific magazine, right-click the `Media` item and then click `Import`. For detailed steps, refer to the online Help index keyword “importing a single medium in a magazine”.

What’s Next?

Once you have imported the media, you may use these media for backup. See Chapter 6, “Backup,” on page 195 for more information on how to configure backups.

Appending Backups to Media

Data Protector allows you to add new backups to media which already contain backups. This method conserves media space.

Limitation

Backups cannot be appended on media used in Travan devices.

The appendable media usage policy can be selected when configuring a media pool. Appendable media contain some currently protected objects; the media must be in good condition and must not be full.

If several devices are used with load balancing, the appendable concept applies on a per device basis, that is, each device uses an appendable medium (if available) as the first medium in a backup session. The backup sessions appending data on the same medium do not have to use the same backup specification.

Two alternatives of appendable media usage policies are available:

- **Appendable:** The first medium used in a backup session uses the space remaining on the medium from the previous backup session. If several appendable media are available in the pool, the least recently used medium is used first. If additional media are needed during the same backup session, they must be free and not contain any protected backups. For this media usage policy, the type of backup (full or incremental backup) can be mixed in any order on the media.
- **Appendable on incrementals only:** The first medium used in a backup session is appended to only if an incremental backup is performed. If several appendable media are available in the pool, the least recently used medium is used first. If additional media are needed during the same backup session, they must be free and not contain any protected backups. This media usage policy will create media which will contain a full backup, followed by any number of incremental backups.

TIP

If you want to create tapes which contain only one full backup and the incremental backups related to the same client, configure Data Protector as follows:

- Configure one pool per client with the media usage policy `Appendable on Incrementals only`.
- Link a different pool to each client in the backup specification, or create a separate backup specification per client.

This is a method to create media containing restore chains. Be aware that occasionally media will be created which contain incremental backups only.

See “Media Usage Policy” on page 151 for a description of media usage policy options like `Appendable`.

See “Selecting Media for Backup” on page 166 for more information on how the media usage policy influences how media are selected for backup.

To modify the settings later, open the properties for the media pool.

Using a Pre-Allocation List of Media for Backup

You can specify the order in which media from a media pool will be used for backup. This order is called a **pre-allocation list**. You specify the pre-allocation list when configuring a backup. The purpose of a pre-allocation list is to control which media will be used for a backup session. You have to match the pre-allocation list with the available media before each backup.

You can also preallocate media when using the object copy functionality.

Depending on the allocation policy of the media pool, Data Protector behaves in two different ways:

- If the pre-allocation list is used in combination with the `Strict` media allocation policy, Data Protector expects the media in a backup device to be available in that order. If the media are not available, Data Protector issues a mount request. If the media mentioned in the pre-allocation list are loaded in a SCSI exchanger, Data Protector handles the media sequence automatically.
- If the pre-allocation list is used in combination with the Loose media allocation policy, media in the pre-allocation list are used first. If the media are not available, any suitable media in the library are used.

Preallocating Media for Backup

On how to preallocate media for backup, refer to the online Help index keyword “preallocating media”.

Selecting Media for Backup

Data Protector media management automatically selects the most appropriate media for backup. This section explains various factors that influence how media are selected for backup.

Media Allocation Policy

You can influence how media are selected for backup using the **media allocation** policy. You can specify a `Loose` policy where any suitable media are used for backup, or a `Strict` policy where specific media have to be available in a predefined order.

See “Media Allocation Policy” on page 150 for more information.

Pre-Allocating Media

You can specify the order in which media from a media pool will be used for backup. This order is called a **pre-allocation list**. For more information, see “Using a Pre-Allocation List of Media for Backup” on page 165.

Media Condition

The condition of the media also influences which media are selected for backup. For example, media in good condition are used for backup before media in fair condition. Media in poor condition are not used for backup.

CAUTION

Media that are marked as fair will only be used if there are no protected objects on the media. Otherwise, a mount request is issued, and data might be lost before backup completes.

See “Factors Influencing the Condition of Media” on page 179 for more information.

Media Usage

The media usage policy also influences which media are selected for backup. See “Media Usage Policy” on page 151 and “Appending Backups to Media” on page 163 for a detailed description.

Media Selection

This section describes the criteria Data Protector uses to select media for backup.

Media in poor condition are not used for backup. Media in fair condition are used only if no media in good condition are available. If available, media in good condition are used.

Media are always selected first from the specified pool and (optionally) from the free pool.

Table 5-2 How Media Are Selected for Backup

Allocation Policy	Allocate Unformatted Media First	Data Protector Selection Order
Loose	OFF	1. Pre-allocation list (if specified) 2. Appendable (as set in usage policy) 3. Unprotected Data Protector media 4. Unformatted media 5. Fair media
Loose	ON	1. Pre-allocation list (if specified) 2. Appendable (as set in usage policy) 3. Unformatted media 4. Unprotected Data Protector media 5. Fair media
Strict	(Not applicable)	1. Pre-allocation list (if specified) 2. Appendable (as set in usage policy) 3. Unprotected Data Protector media 4. Fair media

Setting Data Protection for Media

Data Protector keeps track of data on every medium used. When configuring a backup, you can protect your data from being overwritten by newer backups for a specified time. This protection is on a session basis: if data from several sessions is on the same media, the longest protection defines protection of the media. See “Data Protection: Specifying How Long Data Is Kept on the Media” on page 272 for detailed information.

You can also re-use the media by removing their protection. See “Recycling Media” on page 169 for more information.

Recycling Media

Data Protector keeps track of data on every medium used. When configuring a backup, you protect your data from being overwritten by newer backups for a specified time. See Chapter 6, “Backup,” on page 195 for detailed information.

Keep in mind that on all media there may be data from several backup sessions. Each session can contain data from several backup objects (file systems).

Recycling removes the data protection from all backed up data on the medium, thus allowing Data Protector to overwrite it during one of the next backups. Recycling does not actually change the data on the medium, it only tells Data Protector that this data is not protected anymore. This option is not available for media in free pools.

For instructions on how to change the protection of a specific session or an object, see Chapter 11, “Managing the Data Protector Internal Database,” on page 457.

How to Recycle Media

In the `Devices & Media` context, browse for a medium, right-click it and click `Recycle`. For detailed steps, refer to the online Help index keyword “recycling media”.

Moving Media to Another Pool

Data Protector lets you move a medium from one media pool to another media pool of the same media type.

You need this feature if you want to reorganize the backups and rearrange the purpose of each pool. It is also useful when you want to use the medium in a device which is the default device of another media pool.

How to Move Media to Another Pool

In the `Devices & Media` context, browse for a medium, right-click it and click `Move to Pool`. For detailed steps, refer to the online Help index keyword “moving media”.

Moving Media Using a Free Pool

When using a free pool, media are moved in two instances:

- When media are selected (allocated) for backup, they are moved from a free pool to a regular pool.
- When the media protection has expired, media are moved from a regular pool to a free pool.

This behavior depends on the free pool options selected.

For further information see “Creating a Media Pool” on page 149.

Exporting Media from Data Protector

What Is Exporting Media?

Exporting (removing) a medium removes the information about the medium and its contents from the IDB. Data Protector no longer recognizes that this medium exists. The medium and the data it contains remain unchanged. You can import the medium later, thus re-reading the information about data on the medium back to the IDB. See “Importing Media” on page 159 for instructions.

When to Export Media

If you want to move media to another cell, you have to export the media from one cell and import them to another.

Media that contain protected data cannot be removed. You have to recycle the media first. See “Recycling Media” on page 169 for instructions.

TIP

Export all the media from a backup session. If a backup session spans several media and you do not remove all of them, you will not be able to restore data; Data Protector still recognizes that data exists on the media, but the media will not be available anymore.

How to Export Media

In the `Devices & Media` context, browse for a medium, right-click it and click `Export`. For detailed steps, refer to the online Help index keyword “exporting media”.

What’s Next?

See “Adding Media to a Media Pool” on page 154 if you want to add media to another pool or move them to another cell.

See “Importing Media” on page 159 if you want to import media into another cell.

Modifying Media Locations

What Is a Location?

The media location helps you to physically locate the media. If an object version exists on more than one media set, the media location also enables you to influence the selection of the media set used for a restore.

You enter the location when you format the media. The initial location information is written on the media and to the IDB.

You should modify the location whenever you move media to a different place, such as to off-site storage, for example, “Shelf 4-Box 3”. The revised location information is only written to the IDB.

Data Protector allows you to create a list of pre-defined locations to simplify vaulting and archiving (also known as off-site storage). See “Vaulting Media” on page 186 for more information.

NOTE

When you modify a location, Data Protector modifies the location in the IDB and not on the medium itself.

If you export and import media again, the location information in the IDB is replaced with the location stored on the media.

TIP

You can modify the location of multiple media at the same time. This is useful for vaulting (archiving) purposes. See “Vaulting Media” on page 186.

How to Modify Media Location

Modify media location in the `General` property page for the medium. For detailed steps, refer to the online Help index keyword “modifying media location”.

How to Set Media Location Priority

In the `Devices & Media` context, expand `Media`, then `Locations`, and click a location. For detailed steps, refer to the online Help index keyword “setting media location priority”.

Modifying Media Descriptions

What Is a Description?

The media description helps you identify media. You can define a media description when you format new media. The initial description is written on the media and to the IDB.

If media were auto-formatted during backup, you may want to change the automatically-created description to something better suited to your needs. The revised description information will only be written to the IDB.

NOTE

When you modify a media description, Data Protector modifies the description in the IDB and *not* on the medium itself.

Therefore, if you export and import media that have not been updated, the description in the IDB is replaced with the description from the media.

Media Label

The media label is composed of the user-defined description and the barcode of the medium (if the medium has a barcode and the Barcode Reader Support option is enabled). For example, [CW8279]Default DLT_1 is the media label with the Default DLT_1 description and the CW8279 barcode. If the media description is changed, the descriptive part of the media label is changed too, but the barcode part remains the same.

How to Modify a Media Description

Modify a media description in the General property page for the medium. For detailed steps, refer to the online Help index keyword “modifying, media descriptions”.

Verifying Data on a Medium

What Is Verifying?

Verifying a medium shows whether the data on the medium is valid. It also updates the information about the medium in the IDB, such as medium condition.

Data Protector performs the following:

- Checks the Data Protector headers with information about the medium (medium ID, description, and location.)
- Reads all blocks on the medium.
- If the CRC (Cyclic Redundancy Check) option was used while writing to the medium, Data Protector recalculates the CRC and compares it to the one stored on the medium.

  If the CRC option was not used, and the verify operation passed, this means that all the data on the medium has been read. The medium did not cause a read error, so the hardware status of the tape is at the very least acceptable. This level of check can be viewed as partial.

  Additionally, if the CRC option was used, the backup data itself is consistent within each block. This level of check has a high level of reliability.

NOTE

Depending on the backup devices and media you use, this task can take a considerable amount of time to complete.

When to Verify Media

If errors were reported during backup, you can verify the medium to check whether the backup is usable.

How to Verify Data on a Medium

In the `Devices & Media` context, browse for a medium, right-click it, and click `Verify`. For detailed steps, refer to the online Help index keyword “verifying media”.

Scanning Media in a Device

What Is Scanning?

You scan a device to update Data Protector information about the media in the device or library.

- In a standalone device, you scan a medium in a drive.
- In a library device, you scan media in the selected slots.
- With ADIC/GRAU DAS or STK ACS libraries, Data Protector queries an ADIC/GRAU DAS or an STK ACSLM Server and then synchronizes the information in the IDB with information returned from the Server.

IMPORTANT

With ADIC/GRAU DAS and STK ACS libraries, when several *logical* libraries are configured for the same physical library, it is not recommended to query the DAS or STK ACSLM Server. Add volsers manually. For more information refer to “The Data Protector Query Operation Used with ADIC/GRAU DAS or STK ACS Libraries” on page 38. For information on how to add volsers manually, refer to “Adding Volsers Manually” on page 189.

With ADIC/GRAU DAS libraries, however, when logical libraries are not configured using Data Protector, but using the ADIC/GRAU DAS utilities, the Data Protector query operation can safely be used on such libraries.

Limitation

Volsers scan may not complete successfully if the ADIC/GRAU library is configured with more than 3970 volsers in a repository. A workaround for this problem is to configure multiple logical ADIC/GRAU libraries in order to separate the slots from the large repository into several smaller repositories.

When to Scan the Device

You have to scan the device when you change the location of media (enter, eject) manually without using the Data Protector commands. This creates inconsistencies with the information in the IDB, because Data Protector cannot track the actual location of the media.

Scanning loads media from all the selected slots into a drive, checks the format of media, displays the media header information, and updates the information about the repository in the IDB.

NOTE

Depending on the number of selected slots, scanning may take a considerable amount of time. Data Protector has to load a medium from each slot into a drive and read the medium header with information about the medium.

How to Scan Media in a Device

Scan media in a device by selecting the device and clicking `Scan` from the Actions menu. For detailed steps, refer to the online Help index keyword “scanning backup devices”.

If you are using a library device, you can scan several media in a single action. However, you can only use one drive. For detailed steps, refer to the online Help index keyword “scanning drives in library devices”.

Barcode Scan

To scan a library with barcode support, use the `Barcode Scan` option. Data Protector only checks the barcode on the medium and updates the information in the IDB.

Querying ADIC/GRAU DAS and STK ACSLM Servers

If you want to get information about a repository in the GRAU DAS or STK ACS library from the Server, you can query the DAS or ACSLM Server. A query responds with the contents of the media database on the DAS or ACSLM Server, and then synchronizes the information in the IDB with what is actually in the repository.

This is especially useful if you were using GRAU DAS or STK ACS commands to manage media, as this results in inconsistencies with the IDB - Data Protector does not know the latest status of media in the library repository. For detailed steps, refer to the online Help index keyword “querying ADIC/GRAU & StorageTek Hosts”.

IMPORTANT

With ADIC/GRAU DAS and STK ACS libraries, when several *logical* libraries are configured for the same physical library, it is not recommended to query the DAS or STK ACSLM Server. Add volsers manually. For more information refer to “The Data Protector Query Operation Used with ADIC/GRAU DAS or STK ACS Libraries” on page 38. For information on how to add volsers manually, refer to “Adding Volsers Manually” on page 189.

With ADIC/GRAU DAS libraries, however, when logical libraries are not configured using Data Protector, but using the ADIC/GRAU DAS utilities, the Data Protector query operation can safely be used on such libraries.

Checking the Condition of a Medium

Data Protector allows you to view information about the usage and condition of a medium. The condition of the medium affects the ability to write to the medium and read the data contained on it. This helps you determine when the medium has to be replaced. See “Factors Influencing the Condition of Media” on page 179 for a description of when to change your media.

Use the `Info` property page of a medium to view information about the medium quality (condition). Refer to Figure 5-5.

Figure 5-5 Information on Media

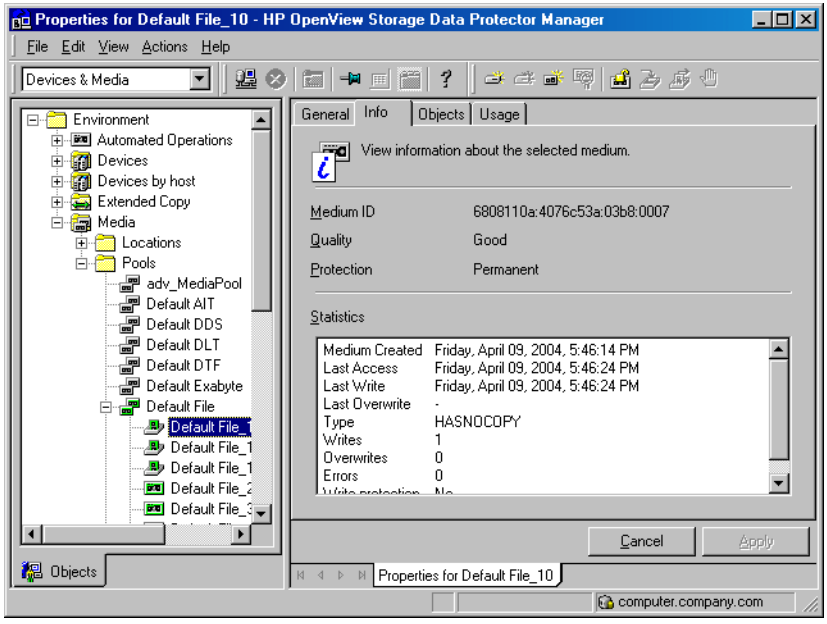

Selection of Backup Media

Media condition influences how media are selected for backup. Media in good condition are selected before media in fair condition. Media in poor condition are never selected. See “Selecting Media for Backup” on page 166 for details.

Factors Influencing the Condition of Media

Data Protector uses **media condition factors** to calculate the condition of the media. The condition of the media in a media pool determines the condition of the media pool. For example, as soon as one medium in a pool is `poor`, the whole media pool is `poor`. When media that are in poor condition are removed from the pool, the pool status reverts to either `fair` or `good` status.

The condition of a media pool indicates the reliability of that media pool for backups. For example, a backup to old or worn media is more likely to have read/write errors.

Media Condition Factors

The two media condition factors you can select are:

Medium valid for. The age of a medium is calculated as the number of months that have elapsed since the medium was formatted. Once a medium is older than the threshold number of months, it is marked as `poor`. The default threshold is 36 months.

Maximum number of overwrites. The usage of a medium is defined as the number of overwrites at the beginning of the medium. Once the medium has more than the threshold number of overwrites, it is marked as `poor`. The default threshold is 250 overwrites, except for DDS, which is set up with a default of 100 overwrites.

Device Error and Media Condition

If a device fails during backup, the media used for backup in this device are marked as `poor`. This prevents future errors if the problem was caused by the bad media.

If this error was due to a dirty drive, clean the drive and verify the medium to reset its condition.

It is recommended that you investigate if media marked `poor` appear in a pool. You can use `Verify` to get more information on each medium's condition. It is not recommended to simply recycle the medium.

Statuses of Media and Media Pools

Media or media pools can have three statuses, based on the media condition factors:

Good. Less than 80% of the threshold for age or usage.

Fair. 81 to 100% of the threshold for age or usage.

Poor. Exceeds 100% of the threshold for age or usage, or read/write errors have occurred on this medium.

See below for information on how to change the media condition factors.

Changing How Media Condition Is Calculated

When you add a medium to a media pool, you can define the media condition factors that are used to calculate the condition of the medium.

IMPORTANT For Data Protector to accurately calculate the condition of the media, use new media when adding media to the media pool.

Change the media condition factors using the `Condition` property page for the media pool. These condition factors are set for the entire media pool.

Figure 5-6 The Media Condition Property Page

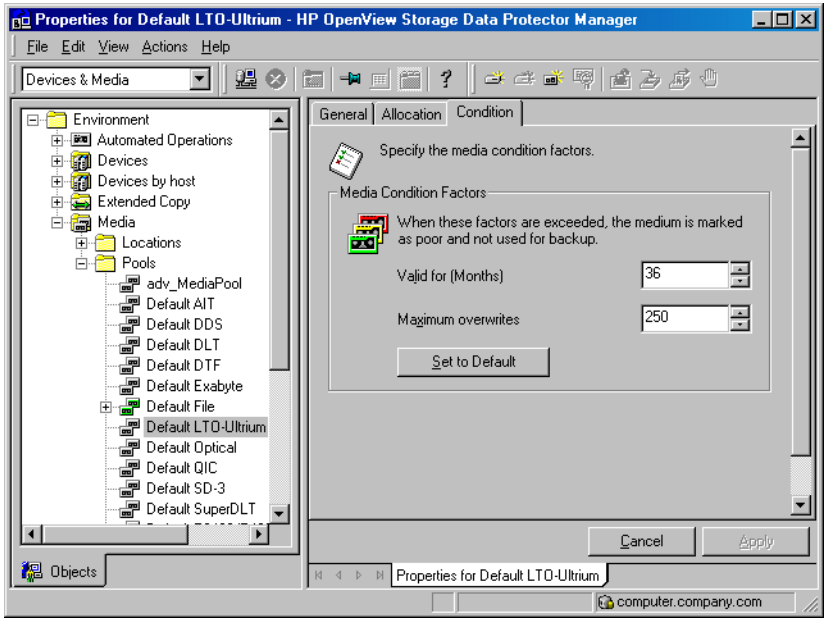

Searching for and Selecting a Medium

Use this function to locate and select specific media without having to browse through the entire list of media.

Media selection is especially useful for vaulting purposes, for example, selecting all media older than 14 days and moving them to a vault. See “Vaulting Media” on page 186 for more information

How to Search for and Select Media

In the `Devices & Media` context, browse for a media pool or a library device, right-click it, and click `Select Media`. For detailed steps, refer to the online Help index keyword “searching for media”.

Entering a Medium into a Device

Data Protector allows you to physically enter media into a library device. You can select the slot that you want to use. Entering and ejecting media does not affect the media pool to which they belong.

IMPORTANT

It is recommended that you use Data Protector to handle the media in the device. This keeps the information about the media in the IDB up to date. If you enter media into the device manually using the device's controls, the information in the IDB is not consistent, and you have to scan the device to update this information. See "Scanning Media in a Device" on page 175 for instructions.

TIP

You can enter multiple media into a device in a single action. See the instructions below.

How to Enter Media into a Device

1. In the `Data Protector Manager`, switch to the `Devices & Media` context.
2. In the `Scoping Pane`, click `Devices`. The list of configured devices is displayed in the `Results Area`.
3. In the list of configured devices, click the name of the library, then expand it to display the `Drives` and `Slots` items.
4. Click `Slots` to display the list of slots.
5. Right-click the slot (or multiple slots) where you want to enter the media, and then click `Enter Medium`.

   A session starts that will prompt you to insert additional media into the device as needed.

What's Next?

If you want to add media to a media pool, see "Adding Media to a Media Pool" on page 154 for more information.

Ejecting a Medium from a Device

Data Protector allows you to physically eject media from the device. When used with library devices, media are moved to the specified slot. You can select the slot that you want to use.

IMPORTANT

It is recommended that you use Data Protector to handle the media in the device. This keeps the information about the media in the IDB up to date. If you eject media from the device manually using the device's controls, the information in the IDB is not consistent, and you have to scan the device to update this information. See "Scanning Media in a Device" on page 175 for instructions.

Bulk Eject of Media

You can eject multiple media from a library in a single action. Data Protector instructs you to remove media from a mail slot when the mail slot becomes full, to free up space for other media selected for ejection.

Predefined Eject of Media

Some operations include the possibility of ejecting the media automatically when the session finishes. For example, when you copy media, you can specify whether the media will be ejected after the session.

When media cannot be ejected because the mail slot is full, Data Protector retries the operation until the mail slot becomes free or until the predefined time limit expires. During this retry, the robotics are accessible to other sessions.

During the eject execution, none of the specified media can be used by other sessions.

Limitation

On Novell NetWare, Bulk Eject functionality is not supported.

How to Eject Media

In the `Devices & Media` context, eject media by right-clicking a medium/slot (or multiple media/slots) and then clicking `Eject`. For detailed steps, refer to the online Help index keyword "ejecting media".

TIP Ejecting of media can be scheduled. Refer to “Scheduled Eject of Media” on page 184 for details.

What’s Next? If you want to put media in a vault, see “Vaulting Media” on page 186 for more information.

Scheduled Eject of Media

Data Protector allows you to schedule the ejection of specific media through the reporting mechanism. The scheduled ejection of media is linked to a specific report made using the external send method. This method enables you to send the report to a user-definable external script, which can then parse the report and execute the ejection of media (using `omnimm -eject` command).

Prerequisite A program or script must be created on the Cell Manager to perform the ejection, and any applicable interpreters must also be installed on the Cell Manager. A Perl script is used in this example.

Overview You can set up and schedule a report group so that it creates a report and sends it as an input to a script. Such a report group should be set up so that it lists the media you want to eject (for example, the List of Media Report) by specifying the report parameters, so that the report contains only the media you want to eject. When the Report Group is started (as the result of a schedule or as triggered by a notification, for example the End of Session notification), Data Protector starts the script with the report result as an input for the script. The script then parses the report and performs the ejection of the specified media by using the Data Protector `omnimm` CLI command.

Notification on Mail Slots Full By default, the Event Log Viewer will notify you if you need to remove media from mail slots in order to continue the eject operation. This situation will arise when there are more media to be ejected than there are empty mail slots in a library. Refer to Chapter 9, “Monitoring, Reporting, Notifications, and the Event Log,” on page 379 for more information on Data Protector notifications.

If media are not removed from the mail slots after a default time span, and there are still media to be ejected, the `omnimm` command aborts the operation. You can change the default time span in the `omnirc` file. Refer to “Using Omnirc Options” on page 615.

For an example of configuring scheduled ejection of media, refer to “Example of Scheduled Eject of Media” on page A-15.

Vaulting Media

What Is Vaulting?

Vaulting is a process of moving media with important information to a safe place where they are kept for a specified period of time. The safe place for media is often called a **vault**. This is also known as off-site storage.

Vaulting and Data Protector

Data Protector supports vaulting on various levels:

- Enables setting up of data protection and catalog protection policies.
- Enables easy selection and ejection of media from the library.
- The media location function tells you the physical location where the media are stored.
- A report shows media used for backup within a specified time frame.
- A report shows which backup specifications have used specified media during the backup.
- A report shows media stored at a specific location with data protection expiring at a specific time.
- Displays a list of media needed for a restore and the physical locations where the media are stored.
- Enables filtering of media from the media view based on specific criteria, such as time written to the media or media with expired protection.

Implementing Vaulting

How you implement vaulting depends on your company's backup strategy and policies for handling data and media. Generally, it consists of the following steps:

1. Specify the desired data protection and catalog protection policies when configuring the backup of data.

   In a backup specification, you can choose to create multiple copies of the same data during the backup. See "Configuring a Backup" on page 198.

2. Configure a vault in Data Protector. Essentially, this means specifying a name for the vault that you will use for the media, such as Vault_1. See "Configuring Vaults" on page 187.

3. After a backup is done, you can create additional copies of the backed up data for vaulting purposes.

   To copy specific backed up objects, use the object copy functionality. See “Configuring Object Copy” on page 323. To create identical copies of media, use the media copy functionality. See “Copying Media” on page 331.

4. Select the media that you want to store in the vault, change the location of the media, eject the media, and store them in the vault.
5. Select the media that you want to remove from the vault, such as media with expired data protection. You can get a list of such media using the List of Media report. For how to generate this report, see “Running Individual Reports” on page 410.
6. Enter the media into the library, scan them, and then change the location field.
7. Establish the appropriate media maintenance policy for media in the vault.

Configuring Vaults

Data Protector allows you to create a list of pre-defined vault locations that you often use. This simplifies entering locations when you move media to the vault.

In the `Devices & Media` context, click `Locations` from the `Edit` menu. For detailed steps, refer to the online Help index keyword “configuring lists of vaults”.

Moving Media to a Vault

Depending on your company’s policies, you can move the original media to a vault directly, or you can create copies of the backed up data and move the copies.

Moving media to a vault consists of two steps:

1. Select media that you want to move and change the location for the media. See “Modifying Media Locations” on page 172.
2. Eject the media from the device and move them to the vault. See “Ejecting a Medium from a Device” on page 183.

Restoring from Media in a Vault

Restoring media from a vault is no different from restoring from any other media. Depending on how your data and catalog protection policies are defined, you may need to take some additional steps:

1. Identify the media needed for restore.
2. Take the media from a vault, enter the media in the library, and scan them.
3. If the catalog protection for the media is still valid, restore data by selecting what you want to restore, using the Data Protector user interface.

   If the catalog protection for the media has expired, Data Protector may not have detailed information about the backed up data. You can restore by manually specifying the files or directories that you want to restore, or use the `List from media` functionality.

TIP

To re-read the detailed information about files and directories from the media once the catalog protection has expired, export the media and import them back, specifying that you want to read the detail catalog data. Now you will be able to browse files and directories in the Data Protector user interface again.

Adding Volsers Manually

With ADIC/GRAU DAS or STK ACS libraries, you can manually add volsers to a library configured in Data Protector instead of querying the library. For more information on querying ADIC/GRAU DAS or STK ACS libraries refer to “Scanning Media in a Device” on page 175.

With ADIC/GRAU DAS or STK ACS libraries, when several *logical* libraries are configured for the same physical library, this is the recommended way of adding volsers to a library configured in Data Protector. With ADIC/GRAU DAS libraries, however, when logical libraries are not configured using Data Protector, but using the ADIC/GRAU DAS utilities, the Data Protector query operation can safely be used on such libraries instead of adding volsers manually. Refer to “The Data Protector Query Operation Used with ADIC/GRAU DAS or STK ACS Libraries” on page 38 for more information.

For detailed steps, refer to the online Help index keyword “adding volsers manually”.

Removing Slots or Volsers

Data Protector provides full support of handling slots and media in media pools used by libraries. Deleting a slot prevents Data Protector from using and accessing the slot in the repository. Information about the slot is removed from the IDB.

This action does not affect volsers in the GRAU DAS library but only removes specific media from the IDB. Therefore, Data Protector does not know that these media exist and does not use them.

For detailed steps on removing slots or volsers, refer to the online Help index keyword “slots”.

Detection of Write-Protected Media

Data Protector can detect and handle media that has been mechanically protected by setting the write protection switch on.

NOTE It is recommended not to use write-protected media with Data Protector.

The following operations can detect and handle write-protected media:

- Read-only operations, such as: list, scan, and verify.

  Read-only operations detect the write-protected media and proceed without any warnings.

- Write operations, such as: initialize, erase, and backup.

  Write operations detect the write-protected media and either abort the session or skip the write-protected media. Backup sessions treat write-protected media as unusable media and behave according to the media allocation policy. If the allocation policy is `strict`, a mount request is issued. If the allocation policy is `loose`, the medium is skipped.

The detection of a write-protected medium and all changes to the write-protection state of the medium are logged to the `media.log` file.

Using Different Media Format Types

Data Protector recognizes and uses two different format types to write data to media:

- Data Protector (for backup devices that are under direct Data Protector control)
- NDMP (for backup devices that are connected to NDMP servers)

Both format types use different Data Protector Media Agent components to communicate with backup devices.

Limitations

Take into account the following limitations, when using different media format types:

- Media that are written by one format type will be recognized as blank or as foreign in a backup device that uses a different format type.
- You cannot back up objects using different format types on the same medium.
- You cannot have two different Data Protector Media Agent components installed on the same system.
- It is strongly recommended that you use different media pools for different media format types.

Modifying Views in the Media Management Window

You can customize the information you see about the media in the Media Management window. This enables you to always see the information you need.

To customize your view, do the following:

1. Open the global options file.

   On the UNIX Cell Manager:

   `/etc/opt/omni/server/options/global`

   On the Windows Cell Manager:

   `<Data_Protector_home>\Config\server\Options\global`

2. Customize the attributes that are to be displayed in the library or media management view by specifying the corresponding token strings.

6 Backup

In This Chapter

This chapter explains how to back up your data. It also describes some advanced Data Protector features.

For information on how to back up database applications such as Oracle, SAP R/3, MS Exchange, MS SQL, Informix, IBM DB2 UDB or Sybase, refer to the *HP OpenView Storage Data Protector Integration Guide*.

For information on how to back up the Data Protector internal database (IDB), see "Configuring the Database Backup" on page 474.

For information on how to install and configure Data Protector management applications, see Chapter 15, "Integrations with Other Applications," on page 731.

NOTE Backup devices (such as tape drives) are subject to specific Data Protector licenses. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.

Configuring a Backup

A backup is a process that creates a copy of system data on backup media. This copy is stored and kept for future use in case the original is destroyed or corrupted.

Prerequisites

- You need to have a Disk Agent installed on every system that is to be backed up, unless you use NFS (on UNIX) or Network Share Backup (on Windows) for backing up these systems.
- You need to have at least one backup device configured in the Data Protector cell.
- You need to have media prepared for your backup.
- You need to have appropriate user rights for performing a backup.

Backup Configuration

Configuring a backup consists of the following steps:

1. Selecting what to back up - the data sources on the Disk Agent clients.
2. Selecting where to back up to - the backup devices connected to Media Agent clients.
3. Selecting how many additional backup copies to create and which backup devices to use - the object mirror functionality.
4. Selecting how to back up - backup options.
5. Optionally, you can schedule an unattended backup.

You specify these options when creating a **backup specification**. Refer to “Creating a Backup Specification” on page 199.

At a specified time, Data Protector starts the backup session based on the backup specification. A **backup object** is a backup unit that contains all items selected for backup from one disk volume (logical disk or mount point). The selected items can be any number of files, directories, or the entire disk or mount point. A backup object is uniquely defined by the client and mountpoint where it resides, by a description, and by the object type (for example filesystem or Oracle).

During the backup session, Data Protector reads the objects, transfers data through the network, and writes them to the media residing in the devices.

The backup specification defines the devices to be used and, optionally, the media pools. If no media pool is specified, the default media pool, which is a part of the device specification, is used.

A backup specification can be as simple as backing up one disk to a standalone DDS drive, or as complex as specifying a backup for 40 large servers to a tape library with 8 drives.

A **backup session** is based on the backup specification, and can be started interactively. During the backup session, Data Protector reads the backup objects, transfers their data through the network, and writes them to the media residing in the devices.

Figure 6-1 Backup Session

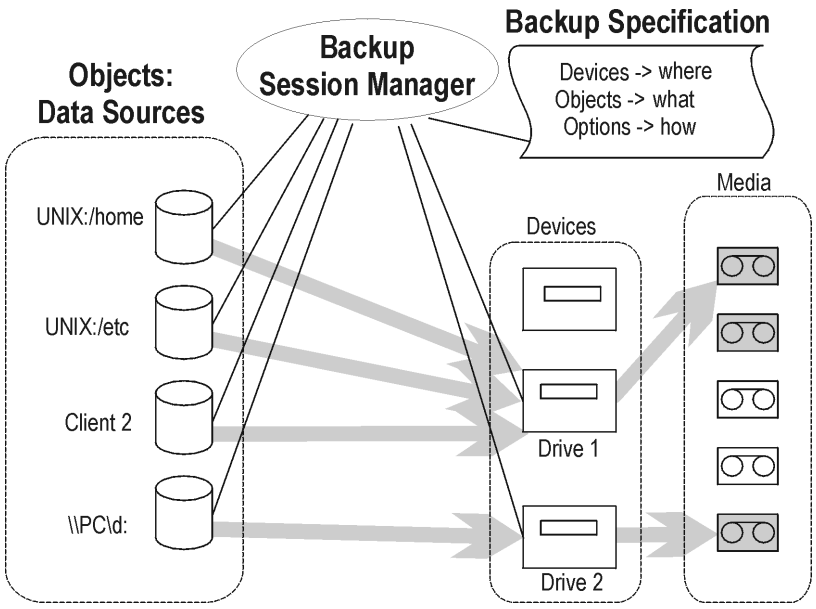

Creating a Backup Specification

You can configure a backup specification using the Data Protector user interface. A backup specification defines the client systems, drives, directories, and files to be backed up, the devices or drives to be used, the number of additional backup copies (mirrors), the backup options for all objects in the specification, and the days and times when you want backups to be performed.

You can create multiple backup specifications by copying an existing specification and then modifying one of the copies.

Data Protector provides default options that are suitable for most cases. To customize the behavior, use Data Protector backup options.

Keep the following key points in mind when you run a backup session:

Key Points

- The backup type (full or incremental) is the same for the whole backup session. All data in a group is backed up using the same backup type.
- A backup object can be added to multiple backup specifications. For example, you may have one backup specification for full backups, one for incremental backups, one for a departmental backup, and one for the archive backup. You can give a description for each object. It is important that you choose the description carefully, because this lets you differentiate among various backups from the same filesystem.
- Objects or clients can be grouped into one backup specification if the media and the backups are managed in the same way, or if media are put into a vault.
- If many backup specifications exist or are planned, you should structure them in groups of backup specifications. If the groups are structured along common option settings (how to back up), then you can apply the backup templates efficiently.
- The Data Protector GUI can display a limited number of backup specifications. The number of backup specifications depends on the size of their parameters (name, group, ownership information and information about whether the backup specification is load balanced or not). This size should not exceed 80 Kb.

Example of Creating a Backup Specification

The following example shows how to create a backup specification for a filesystem and how to start the backup interactively.

1. In the `HP OpenView Storage Data Protector Manager` window, switch to the `Backup` context.
2. In the Scoping Pane, expand `Backup`, and then double-click `Backup Specifications`.
3. In the Results Area, right-click `Filesystem`, and then click `Add Backup`. The `Create New Backup` dialog box appears.

4. In the Create New Backup dialog box, select the `Blank Filesystem Backup` template, and then click `OK` to start the Backup wizard. See Figure 6-2 on page 201.

Figure 6-2 Create New Backup Dialog Box

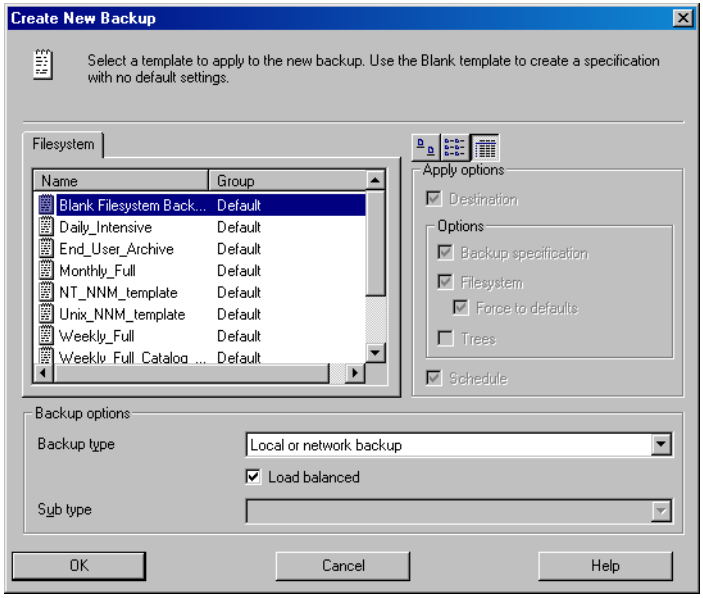

5. Select what you want to back up. Figure 6-3 on page 202 shows data sources selected for backup. Click `Next` to proceed.

Figure 6-3 Source Page of the Backup Wizard

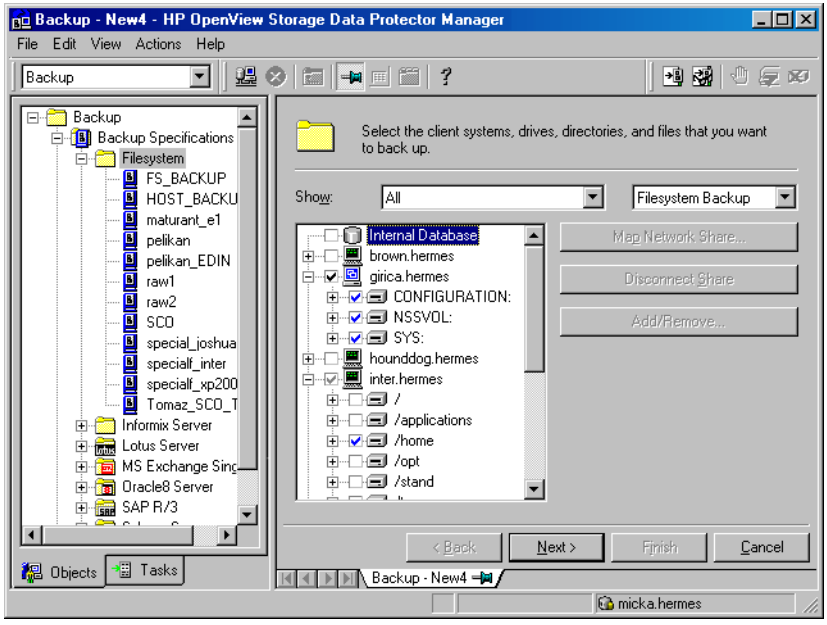

Browsing of Windows systems is not supported in the UNIX GUI, therefore only `Filesystem backup` option is supported for viewing backup objects using the UNIX GUI. `Network share backup` is a Windows specific option for backing up Windows shared disks and is available only in the Windows GUI.

6. Select the device(s) that will be used to back up your data. See Figure 6-4 on page 203.

   You can also specify whether you want to create additional copies (mirrors) of the backup during the backup session. Specify the desired number of mirrors by clicking the `Add mirror` and `Remove mirror` buttons. Select separate devices for the backup and for each mirror.

   For detailed information on the object mirror functionality, see “Object Mirroring” on page 329.

Click `Next` to proceed.

Figure 6-4 Device Page of the Backup Wizard

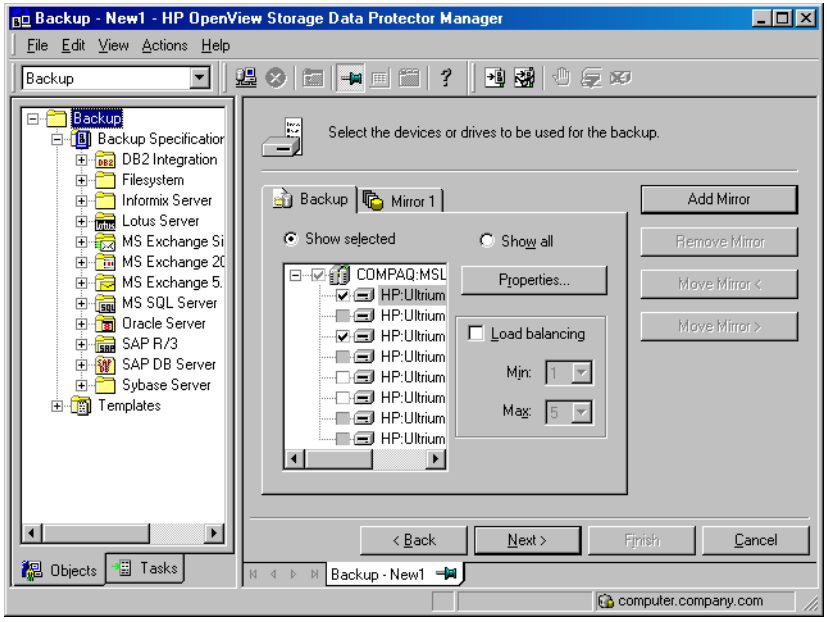

7. Select backup options. See “Using Backup Options” on page 269 for details. Click `Next`.
8. In the `Schedule` page, you can schedule the backup. See “Scheduling Unattended Backups” on page 250 for more information. Click `Next`.

9. In the `Backup Object Summary` page, you can review the backup options. See Figure 6-5 on page 204. Click `Next`.

Figure 6-5 Backup Object Summary Page

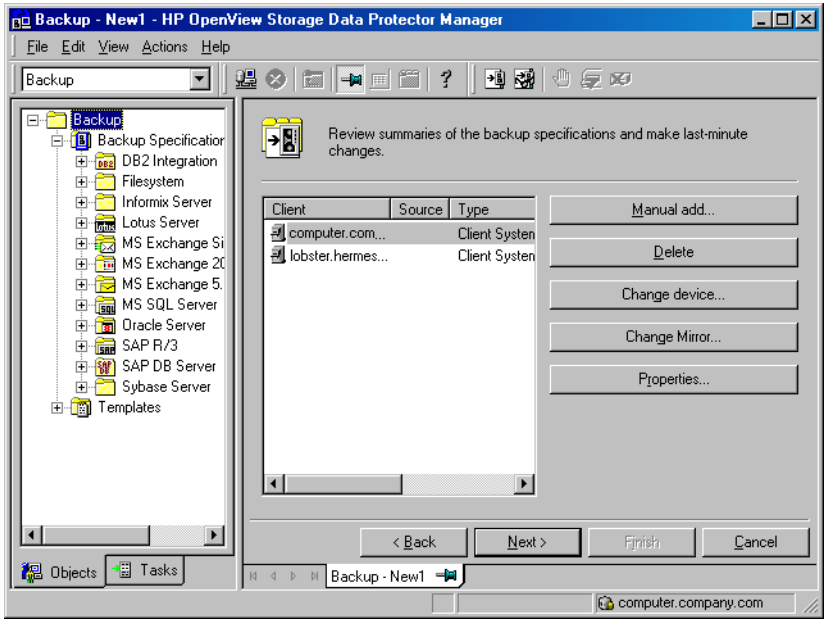

10. In the final page of the Backup wizard, you can save the backup specification, start the interactive backup, or preview the backup. See Figure 6-6 on page 205.

    It is recommended to save the backup specification so that you can schedule or modify it later.

Figure 6-6 Final Page of the Backup Wizard

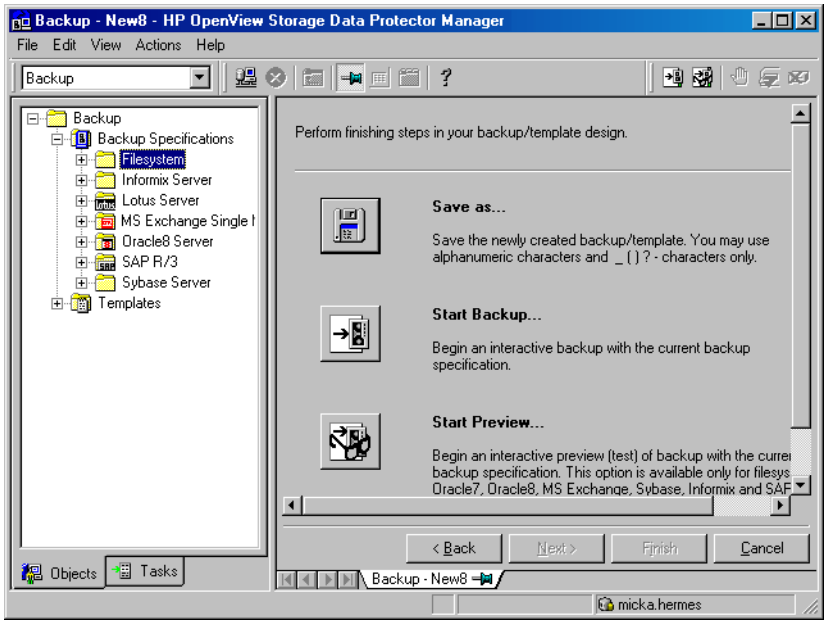

11. Click `Start Backup` to run the backup interactively. The `Start Backup` dialog box appears.

NOTE

During a backup, you may be prompted to add more media to continue your backup. This is called a mount request. See “Responding to Mount Requests” on page 384 for more detailed information.

Backing Up UNIX Systems

You can install a Disk Agent on every UNIX system in order to back it up. Alternatively, you may use the Network Filesystem (NFS) to back up data from systems that do not have a Disk Agent.

See “Backing Up Disks Using NFS” on page 209 for details.

See the *HP OpenView Storage Data Protector Installation and Licensing Guide* or online Help for instructions on how to install a Disk Agent.

See the *HP OpenView Storage Data Protector Software Release Notes* for a complete list of supported platforms and known limitations.

Backing Up UNIX Filesystems

Limitations

The maximum size of the files you can back up depends on operating system and filesystem limitations. For file size limitations, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

Data Protector backs up the directory structure, regular files, and special files. Special files are character device files, block device files, UNIX domain sockets, FIFO files, HP-UX network special files, and XENIX specially-named files.

Softlinks and mountpoints are not followed, and are backed up as softlinks and ordinary empty directories, respectively.

If there are multiple hardlinks referencing the same file, the file is backed up only once. You can change this by setting the `Backup hardlinks as files` option, as explained in “List of Data Protector Backup Options” on page 280.

Basic ACLs (file permission attributes) and time attributes are backed up together with the files on all platforms. However, the support for extended ACLs is limited on some platforms. For details, refer to the *HP OpenView Storage Data Protector Software Release Notes*. The time of the last access to each file is saved before reading the file and then returned to the original value after the file is backed up. This behavior can be changed by setting the `Do not preserve access time attributes` option, as explained in “Using Backup Options” on page 269.

Data Protector provides a sophisticated mechanism for incremental backups. To determine which files have changed, the Data Protector Disk Agent checks when each was last modified. This method keeps Data Protector from detecting moved files, as moving the file does not change the modification time.

NOTE

During a backup session, each file being backed up is opened and read. Therefore, the access time of the file is changed after the backup. Unless the `Do not preserve access time attributes` backup option is set, the access time attribute is set to its original value. OFF is the default value. If this option is set, moved files on UNIX clients are included in the incremental backup, because detection is based on the `inode` modification time.

Selecting Specific Files or Directories

For each filesystem, you can restrict the backup to specific directory trees. For each directory tree you can:

- Exclude any sub-tree or file
- Back up files that match a specific wildcard pattern
- Skip files that match a specific wildcard pattern

Some files are permanently in use, for example, by database applications. These files should be excluded from ordinary filesystem backup and should be backed up in a special way. This is also true for the IDB itself.

Therefore, exclude the IDB directories `/var/opt/omni/server/db40` and `/etc/opt/omni` on UNIX Cell Managers from standard filesystem backups to ensure the consistency of data.

For detailed information on how to back up the IDB, see “Configuring the Database Backup” on page 474.

You should also exclude temporary directories.

How to Back Up UNIX Files

Back up UNIX files using the procedure described in “Example of Creating a Backup Specification” on page 200.

See also “Using Backup Options” on page 269 for information on using and structuring your backup options.

Backing Up Clients Using Disk Discovery

How Are Disks Discovered?

If you specify a client backup with **disk discovery**, Data Protector contacts the client at backup time and finds all filesystems on the disks that are attached to that system. Only mounted disks are identified using the `mount` command. Then Data Protector backs up each filesystem identified as a regular filesystem, except for NFS, CD mounted filesystems, and removable volumes. The description for each filesystem object is generated and the filesystem mountpoint is appended to the description of the client backup.

When to Use Disk Discovery

This backup type is recommended under the following conditions:

- If you back up workstations with relatively small disks that are frequently mounted or unmounted.
- If you would like to back up the data following a mountpoint into one directory, regardless of how many filesystems are mounted. For example, `/home/data`, where `/home/data/disk1 and /home/data/newdisk/disk2` can be mounted or unmounted frequently and independently of each other.

You can use disk discovery by specifying the client as a data source. If another disk is mounted later, it will be included in the backup.

In contrast to a filesystem backup, where you have to specify any newly added disk or mounted filesystem that is not yet specified in the backup specification, this is unnecessary if you use disk discovery.

To create a backup specification that will define a disk discovery backup, follow the procedure described in “Example of Creating a Backup Specification” on page 200.

Once you get to the `Source` property page of the Backup wizard, click the check box next to the client. This selects the entire client to be backed up, as shown in Figure 6-7.

Figure 6-7 Selecting an Entire Client to Be Backed Up

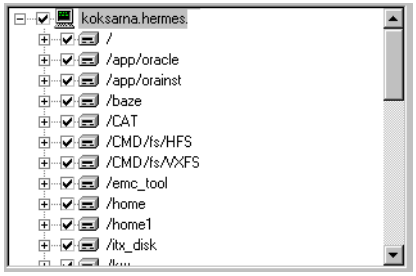

NOTE

Selecting all of the client's drives is not the same as selecting the check box next to the client name, which is the procedure for a Disk Discovery backup.

When you perform a client backup, all the files and directories that belong to the root (/) mountpoint are automatically backed up. Therefore, you cannot exclude the root in the backup specification. If you want to exclude the root, perform a filesystem backup.

To check the configured backup type, see the `Backup Object Summary` property page. Under the `Type` label, you will see `Client System` if you have configured a Disk Discovery backup and `Filesystem` if only the drives have been selected.

Also see "Using Backup Options" on page 269 for information on structuring your backup specifications.

Backing Up Disks Using NFS

What Is NFS?

NFS (Network Filesystem) is a distributed file system protocol on UNIX that allows a computer to access files over a network as though they were on its local disks.

When to Use NFS Backup

Use NFS backup in either of the following situations:

- A system to be backed up is not a part of the Data Protector cell or does not have a Disk Agent installed.

- You want to back up system platforms that are not supported by Data Protector.

Example

Figure 6-8 shows a typical configuration. You want to back up the filesystem `/home/special` from system `wiz`, which is not part of the Data Protector cell and has no Data Protector software installed. But the filesystem is mounted as `/nfs/wiz/home/special` on a Data Protector client `bserv`.

To back up this filesystem using NFS, follow the same procedure as if you were backing up any other filesystem on `bserv`, except that you have to manually type `/nfs/wiz/home/special` as a mountpoint. Only local filesystems can be browsed.

Figure 6-8 NFS Environment

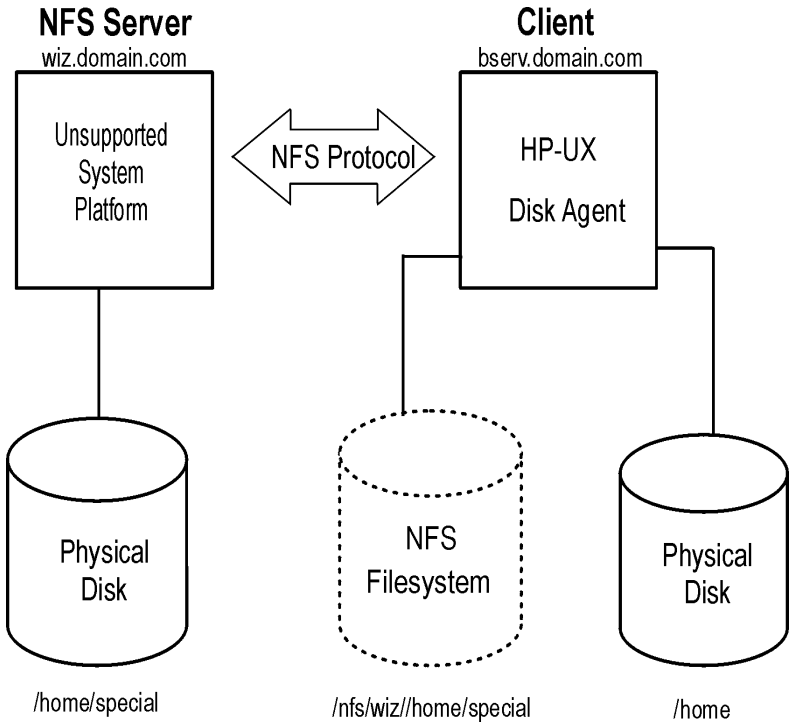

Limitations

- You can back up NFS mounted volumes only on HP-UX and Solaris clients. You can not back up soft links, or character and device files.

- ACL (access control list) attributes are not preserved. NFS does not support ACLs on remote files. Individual manual entries specify the behavior of various system calls, library calls, and commands. When transferring a file with optional entries over the network or manipulating a remote file, the optional entries may be unexpectedly deleted.

NOTE It is recommended to have root permission on mounted NFS filesystems.

To back up a filesystem using NFS, follow the procedure described in "Example of Creating a Backup Specification" on page 200 until you get to the `Backup Object Summary` page of the wizard. Proceed as follows:

1. In the `Backup Object Summary` page, click `Manual Add`.
2. Click the `UNIX Filesystem` button, and then click `Next`.
3. In the `General Selection` page, select a client and manually add the mount point in the `Mountpoint` text box. See online Help for details.

Backing Up UNIX Disks as Disk Image Objects

What Is a Disk Image Backup?

A **disk image backup** is a high-speed backup of disks, disk partitions, or logical volumes without tracking the file and directory structure stored on these data sources. Data Protector stores the disk image structure at the character level.

When to Use a Disk Image Backup

Use a disk image backup in any of the following situations:

- You have lots of small files and a high backup speed is required.
- A full disk backup is needed, for example, for disaster recovery or before a major software update.
- A direct disk-to-disk connection is not possible and you want to duplicate a filesystem to another disk. The latter must be identical to the original disk.

Where to Find Rawdisk Sections

On the HP-UX and Solaris systems, the rawdisk sections are usually listed in the `/dev/rdsk` directory. On HP-UX, raw logical volumes can be found in `/dev/vg<number>`. The first letter of the new logical volume must be `r`, for instance `/dev/vg01/rlvol1`.

IMPORTANT

Unmount a disk before a disk image backup and mount it later. You can use `pre-` and `post- exec` commands for this purpose. See “Examples of Pre-Exec and Post-Exec Commands for UNIX” on page A-21.

To back up a disk image object, follow the procedure described in “Example of Creating a Backup Specification” on page 200 until you get to the `Backup Object Summary` page of the wizard. Proceed as follows:

1. In the `Backup Object Summary` page, click `Manual Add`.
2. Click the `Disk image object` button, and then click `Next`.
3. In the `General Selection` page, select a client and manually add the mount point in the `Mountpoint` text box. See online Help for details.

Backing Up Windows Systems

Prerequisites

You have to install a Disk Agent on at least one Windows computer in the Data Protector cell. This computer then becomes a Disk Agent client.

Files that do not reside on Disk Agent clients can be backed up if they share their disks with Disk Agent clients. It is better to install a Disk Agent on every Windows system that you want to back up.

See “Backing Up Windows Shared Disks” on page 230 for details.

See the *HP OpenView Storage Data Protector Installation and Licensing Guide* or online Help for instructions on how to install a Disk Agent.

See the *HP OpenView Storage Data Protector Software Release Notes* for a complete list of supported system platforms.

Limitation

- To run a VSS filesystem backup, your system must have at least one NTFS filesystem.

Backing Up Filesystems (Logical Disk Drives)

Selecting Backup Objects

Select a file, a directory, or a logical disk drive for backup in the Backup wizard.

See “Example of Creating a Backup Specification” on page 200 and “Using Backup Options” on page 269 for details.

What Is Backed Up?

A filesystem backup of a disk drive involves reading the directory structure and the contents of the files on the selected disk drive. The following data is also backed up along with the data in the file:

- Full Unicode filenames
- FAT16, FAT32, VFAT, and NTFS attributes

  Once a file is backed up, its archive attribute is cleared. You can change this behavior by setting the `Do not use archive attribute` option among the `Advanced` filesystem backup options in the backup specification. See online Help for details.
- NTFS alternate data streams
- NTFS security data

- Directory share information

  If a directory is shared over network, share information is by default backed up. During the restore, share information is restored by default and directory is shared on the network after the restore. You can change this behavior by unselecting the `Backup share information for directories` option in the Filesystem options window.

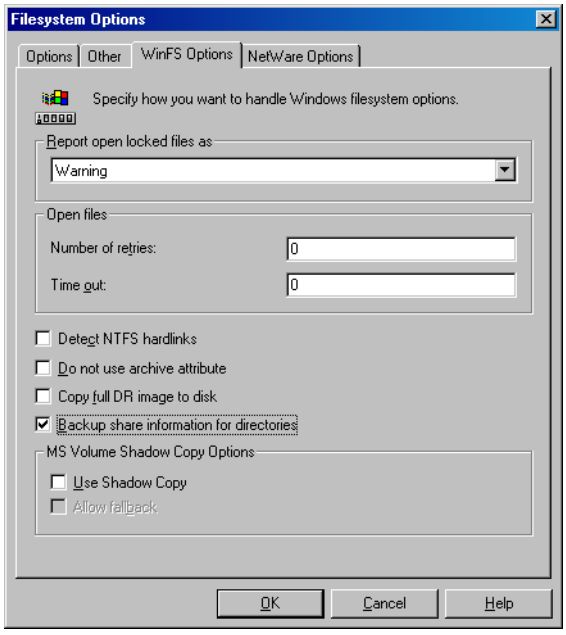

What Is Not Backed Up?

In the backup specification, you can specify the files to be excluded from or skipped by the backup. The list of these files is also known as a **private exclusion list**.

See “Object Options” on page 284 and online Help for more information on how to exclude or skip files and directories.

In addition to the private exclusion list, Data Protector by default excludes the following:

- The *<Data_Protector_home>*`\log` and *<Data_Protector_home>*`\tmp` directories from a Windows client or Cell Manager backup.

- The *<Data_Protector_home>*\db40 directory from a Windows Cell Manager backup.

  For example, the *<Data_Protector_home>*\db40 directory is excluded from the Cell Manager backup even though it was selected in the backup specification. This is because the *<Data_Protector_home>*\db40 directory contains the IDB, which must be backed up in a special way to ensure data consistency. See "Configuring the Database Backup" on page 474 for details.

The skipped file is the Pagefile.sys system file. Before starting a backup, Data Protector reads the list of excluded and skipped files from the following Registry keys:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView
\\OmniBack II\Agents\FileSystem\Exclude
```

```
HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView
\\OmniBack II\Agents\FileSystem\Skip
```

NTFS 3.x Filesystem Features

The NTFS 3.x filesystem has introduced new file attributes and concepts, which can be summarized as follows:

- The NTFS 3.x filesystem supports **reparse points.** The **volume mount points**, **Single Instance Storage (SIS)**, and **directory junctions** are based on the reparse point concept.
- The NTFS 3.x filesystem supports **sparse files** as an efficient way of reducing the amount of allocated disk space.
- The NTFS 3.x filesystem supports the **Object IDs** that are backed up by Data Protector along with other alternate data streams.
- Some of the NTFS 3.x filesystem-specific features are controlled by system services that maintain their own data records. These data structures are backed up as a part of CONFIGURATION.

  See "Backing Up CONFIGURATION" on page 218 and "Backing Up the Windows Services" on page 223 for details.
- The Microsoft-encrypted NTFS 3.x files are backed up and restored encrypted, but their contents can only be properly viewed when they are decrypted. Refer to the *HP OpenView Storage Data Protector Software Release Notes* for details about related limitations.

VSS Filesystem Backup

Volume Shadow Copy service (VSS) is implemented on the Windows Server 2003 operating system. This service provides an additional Windows filesystem backup, where the level of data integrity is slightly increased compared to traditional backup of active volume.

To prepare for creation of the shadow copy, all I/O activity is stopped by the VSS mechanism. When the shadow copy is created, Data Protector starts its normal backup procedure, except that the source volume is replaced by the newly created shadow copy. If the shadow copy creation fails, Data Protector can proceed with the normal filesystem backup, if the `Allow Fallback` option was specified in the backup specification.

During the VSS filesystem backup the consistency of data is improved in comparison with the non-VSS filesystem backup. VSS allows you to create shadow copy backups of volumes and exact point-in-time copies of files, including all open files. In this way, the files changed during the backup are copied correctly.

The advantages of the VSS filesystem backup are the following:

- A computer can be backed up while applications and services are running. Therefore, the applications can continue writing data to the volume during a backup.
- Open files are no longer skipped during the backup process, because they appear closed on the shadow copy volume at the time of the shadow copy creation.
- Backups can be performed at any time without locking out users.
- There is little or no impact on performance of the application system during the backup process.

For VSS filesystem backup related options, refer to “Using Backup Options” on page 269. Also refer to the *HP OpenView Storage Data Protector Concepts Guide* for details on the VSS concepts.

Reparse Points

Basically, reparse points are plain filesystem objects with a unique tag attached, known as a reparse point ID. The NTFS 3.x directories or files can contain a reparse point, which typically imitates the contents by directing to data from another location.

When Data Protector encounters reparse points, the reparse point IDs are not followed by default, what is also known as backing up raw reparse points. This affects the way you configure your backups:

- ✓ If you configure a backup using Disk Delivery, all data will be backed up once.
- ✓ If you back up filesystems or drives containing reparse points, ensure that the data pointed to by a reparse point gets backed up. For example, the Windows **directory junctions** reparse points are not followed, so the junctions have to be backed up separately. SIS reparse points are exceptions.

  The **Single Instance Storage (SIS)** service regularly checks the files on a disk. If the service detects several identical files, it replaces them with the reparse points and stores the data into a common repository. In this way, the disk space usage is reduced.

Reparse points let you mount logical volumes as disk drives. Data Protector treats the mounted volumes as though they were ordinary drives, so that they are visible as selectable objects for backup.

Sparse Files

Sparse files contain many zero data sets as opposed to, for example, compressed files. At backup time, Data Protector automatically skips zero-parts, so that the media space on the backup device is allocated for non-zero parts only.

UNIX and Windows sparse files are not compatible.

Manual Definition of Multiple Disk Agents

If you want to back up one mount point through multiple Disk Agents (DA), you have to specify each object separately using the `Manual add` functionality. Give a new description to each object and use the `Trees/Exclude` option in the `Manual add` wizard to specify the path for an object. Refer to Figure 6-9.

In addition, consider the following:

- You have to manually define the data area split, taking care to avoid overlapping the same data.
- If more than one DA is concurrently accessing the same mount point, which is defined as one disk, the data transfer speed will drop. This can be different when using disk arrays.

Figure 6-9 **Specifying Objects Using Manual Add**

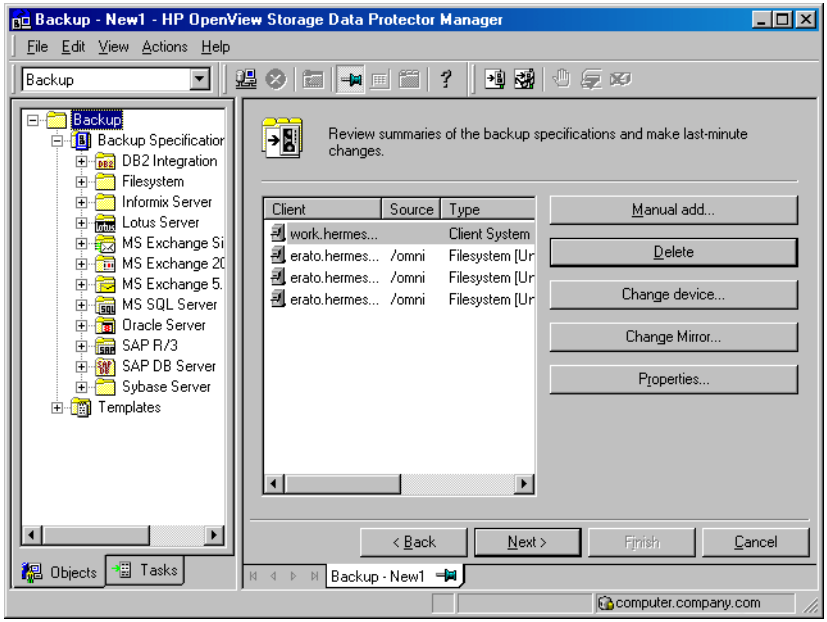

For detailed steps, refer to the online Help index keyword “concurrency”.

Backing Up CONFIGURATION

The Data Protector `CONFIGURATION` object is a set of data structures maintained by the Windows operating system that are not treated as a part of a filesystem backup when you, for example, select logical drives such as `C:` or `D:` for the backup.

Windows CONFIGURATION

`CONFIGURATION` consists of the following objects:

- `EventLog`
- `Profiles`
- `Registry`
- `SystemRecoveryData`
- `EISA Utility Partition`
- `WINS`, `DHCP` (on the Windows TCP/IP protocol servers)

- `QuotaInformation`, `RemovableStorageManagementDatabase`, and `FileReplicationService`.
- The **System State** services

  See “Backing Up the Windows System State” on page 220.
- `DNSServerDatabase`

  See “Backing Up WINS, DHCP, and DNS” on page 223.
- `SysVol`

  `SysVol` is a shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.
- `IIS`

  Microsoft Internet Information Server is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).

Figure 6-10 Windows CONFIGURATION

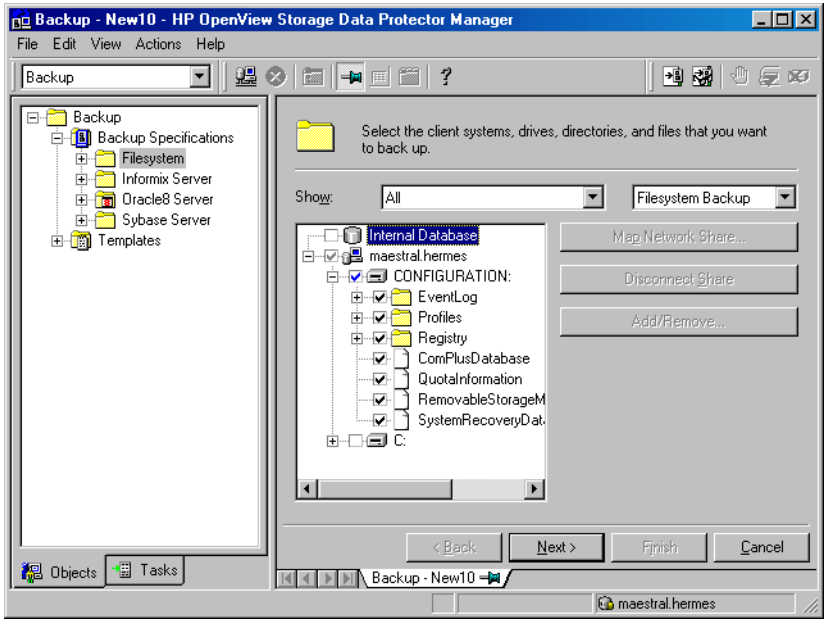

CONFIGURATION varies among different versions of Microsoft Windows.

Backing Up CONFIGURATION

Only one CONFIGURATION backup can run on a system at the time. You have to expand a client and select its CONFIGURATION in the Backup wizard.

See “Example of Creating a Backup Specification” on page 200, and Figure 6-10.

Backing Up the Windows System State

The Windows System State consists of several elements related to various aspects of Windows. They are structured under their respective Windows backup object. The Windows System State includes the following:

- Registry and ComPlusDatabase
- The following boot files: Ntldr.exe, Ntdetect.com and boot.ini

- The `System Volume Information` directory, which keeps data accessed by the `System File Protection (SFP)` service.

Provided that the services were installed and configured, the `System State` data of a Windows Server system also includes:

- `ActiveDirectoryService`
- `CertificateServer`
- `TerminalServiceDatabase`

See “Example of Creating a Backup Specification” on page 200 for a detailed backup procedure. Figure 6-11 shows how to select `System State` in the Backup wizard.

Figure 6-11 System State on Windows

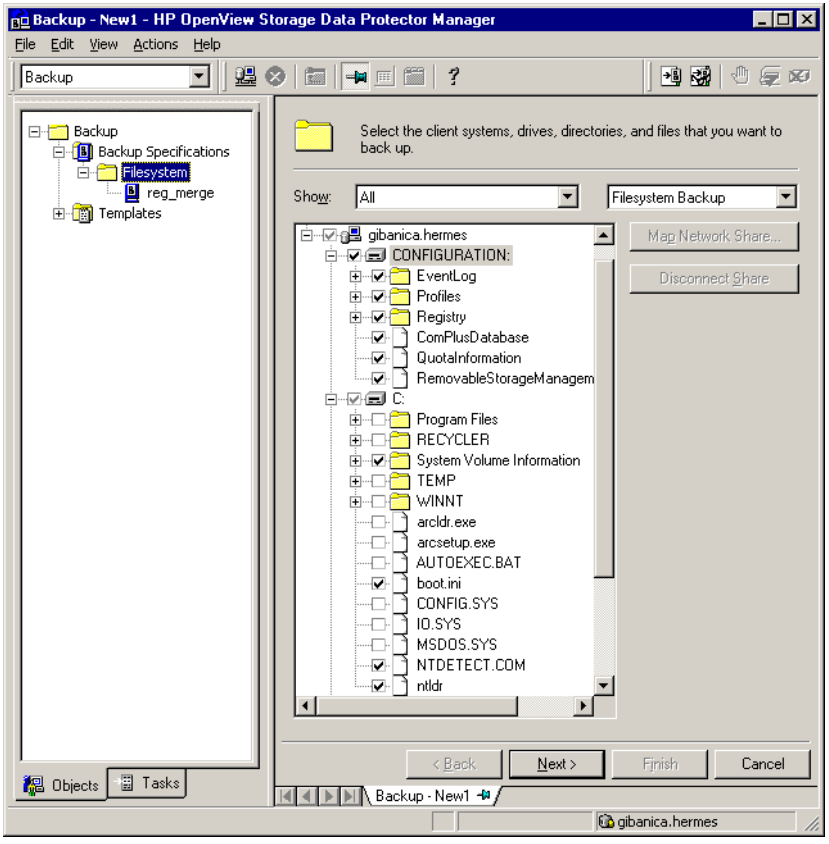

Backing Up the Windows Registry

The database repository of information containing the Windows system configuration is known as the `Registry`. The Windows `Registry` is important for the system operation, and must be backed up regularly.

The `Registry` can be backed up as a part of `CONFIGURATION`, or separately by selecting the `Registry` folder as shown in Figure 6-12.

Figure 6-12 **Backing Up the Windows Registry**

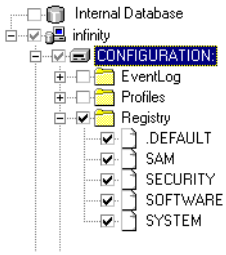

Backing Up WINS, DHCP, and DNS

WINS, DHCP, DNS Servers

In TCP/IP networks, the following services can be configured and run on Windows servers:

- **WINS Server**

  This service, also known as Windows Internet Name Service, is a dynamic replicated database service that can register and resolve NetBIOS names to IP addresses used on a TCP/IP network.

  To back up this database, select `WINS` in the Backup wizard.

- **DHCP Server**

  This service provides dynamic IP address assignment and network configuration for Dynamic Host Configuration Protocol (DHCP) clients.

  To back up this database, select `DHCP` in the Backup wizard.

- **DNS Server**

  This service runs on a Domain Name System server and maintains its own database. A DNS Server answers queries and updates requests for DNS names.

  To back up this database, select `DNSServerDatabase` in the Backup wizard.

Backing Up the Windows Services

Backing up the Windows services means backing up the data structures used by these services. A particular database is exported (dumped) into a file, which is then backed up. The services are always backed up, if `CONFIGURATION` was selected in the Backup wizard.

NOTE

A service has to be up and running so that Data Protector can detect it and show it as a selectable item in the Backup wizard. If a service is not running at backup time, the corresponding backup object will fail. See “Managing Failed Backups” on page 311 for more information.

To back up a specific service, you can select the corresponding folder under the `CONFIGURATION` backup object.

Figure 6-13 Backing Up Windows Services

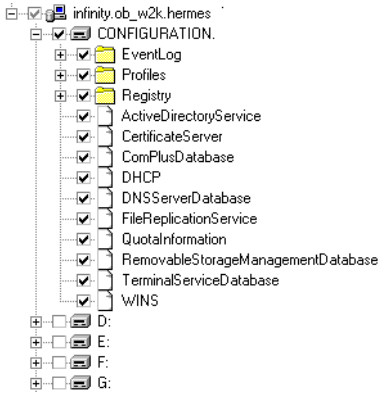

See also “Example of Creating a Backup Specification” on page 200 for a step-by-step procedure.

Data Protector can detect and back up the following Windows services:

- COM+ Event System

  This service provides automatic distribution of events to subscribing COM+ components. To back up this database, select the `ComPlusDatabase` in the Backup wizard.

- Removable Storage

  This service manages removable media, drives, and libraries. To back up this database, select `RemovableStorageManagementDatabase` in the Backup wizard.

IMPORTANT

You can back up the Removable Storage database, but this service is not used for Data Protector media management. The native robotics driver used with robotics media changers has to be disabled before a device is configured by Data Protector.

Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.

- Active Directory Service

  Active Directory Service is the Windows server directory service that enables you to manage data structures distributed over a network. For example, Active Directory Service stores information about user accounts, passwords, phone numbers, profiles, and installed services. It provides the methods for storing directory data and making this data available to network users and administrators.

  To back up the Active Directory data structures that are stored on the local system, select `ActiveDirectoryService` in the Backup wizard.

- Terminal Services

  These services provide a multi-session environment that allows client systems to access a virtual Windows desktop session and Windows-based programs running on the server.

  To back up this database, select `TerminalServiceDatabase` in the Backup wizard.

- Certificate Services

  These services issue, revoke, and manage certificates employed in public key-based cryptography technologies. To back up this database, select `CertificateServer` in the Backup wizard.

  For example, if you use Active Directory to publish Certificate Revocation Lists (CLRs), back up the Active Directory services along with the Certificate Services.

- Remote Storage Service

  Remote Storage Service (RSS) is used to automatically move infrequently accessed files from local to remote storage. Remote files are recalled automatically when the file is opened. Although RSS

databases are part of System State data, you must back them up manually. Refer to “Backing Up a Remote Storage Service Database” on page 226.

- System File Protection Service

  System File Protection (SFP) service scans and verifies the versions of all protected system files after you restart your computer. If the SFP service discovers that a protected file has been overwritten, it retrieves the correct version of the file and then replaces the incorrect file. Data Protector enables you to back up and then restore protected files without overwriting them. The protected files can be backed up using the `Move Busy Files` option in a standard filesystem backup procedure.

- DNS, DHCP, and WINS

  See “Backing Up WINS, DHCP, and DNS” on page 223.

Backing Up the DFS

Data Protector backs up the Windows Distributed File System (DFS) as part of one of the following:

- Windows Registry, if the DFS is configured in a standalone mode.
- Windows Active Directory, if the DFS is configured in a domain mode.

Backing Up a Remote Storage Service Database

Data Protector allows you to back up the Remote Storage Server (RSS) database by following the standard filesystem backup procedure. The RSS databases must be backed up offline. You can stop and restart the Remote Storage Service using `pre-` and `post-exec` scripts, or you can perform this manually before and after the backup. Use the following commands:

```
net stop/start “Remote Storage Engine”
net stop/start “Remote Storage File”
```

The RSS databases are located in the following directories:

```
<%SystemRoot%>\System32\RemoteStorage
<%SystemRoot%>\System32\NtmsData
```

Backing Up Windows User Profiles, Event Logs, and User Disk Quotas

User Profiles

A `User Profile` contains information about a user configuration. This includes the profile components, such as desktop settings, screen colors, and network connections. When a user logs on, the user profile is loaded and the Windows environment is set accordingly.

The user profile data resides in the `\Documents and Settings` directory.

These directories contain all user profiles that are configured on the system and backed up by Data Protector. If a system is configured for multiple users, a separate user profile belongs to each defined user. For example, the `All Users` and `Default User` profiles contain the profile components common to all defined users and those assigned to a newly created user.

Data Protector reads the location of the profiles from the following Registry keys:

```
HKEY_USERS\.DEFAULT\Software\Microsoft\Windows\\
CurrentVersion\Explorer\Shell Folders
```

where information about common profile components resides.

```
HKEY_USERS\.DEFAULT\Software\Microsoft\Windows\\
CurrentVersion\Explorer\User Shell Folders
```

NOTE

If you back up `CONFIGURATION` and the whole Windows system partition as a filesystem, the `Profiles` are backed up twice; as part of a filesystem backup and as part of `CONFIGURATION`. To avoid this, exclude the profile data (see above for location) from the filesystem backup.

See also “Warnings When Backing Up System Disks” on page 311.

Event Logs

Event logs are files where the Windows operating system saves information about events, such as starting and stopping services or the logging on and logging off of a user.

User Disk Quotas

User Disk Quotas enable enhanced tracking and control over disk space usage on Windows.

Data Protector backs up user disk quotas on the whole system and for all configured users at a time.

Event Logs, User Profiles, and User Disk Quotas are always backed up if CONFIGURATION was selected in the Backup wizard.

See Figure 6-10, “Windows CONFIGURATION.”, and refer to “Example of Creating a Backup Specification” on page 200 for a step-by-step procedure.

Backing Up Windows Clients Using Disk Discovery

You can use disk discovery by specifying the client as a data source. If another disk is added later, it will be included in the backup.

How Are Disks Discovered?

If you specify a client backup with disk discovery, Data Protector contacts the client and discovers all logical disk drives that belong to physical disks on the client, except for CDs and removable drives. Then it backs up the CONFIGURATION folder and each discovered logical drive as a regular filesystem. The description text of each filesystem object will be generated by appending the drive letter in square brackets to the description of the Client Backup.

When to Use Disk Discovery

This backup type is recommended under the following circumstances:

- When backing up systems with relatively small disks
- When performing a whole system backup to prepare for disaster recovery
- When the number of disks connected to the system varies.

For a client backup with disk discovery, it is not possible to select only specific directory trees, because this implies a single logical drive backup. It is, however, possible to exclude any directory from the backup.

How to Perform a Backup

To perform a Windows client backup, you have to create a backup specification as described in “Example of Creating a Backup Specification” on page 200.

In the Source property page of the Backup wizard, select the check box next to the client name to obtain the disk discovery functionality. Then, follow the wizard.

NOTE

Selecting all of the client's drives is not the same as selecting the check box next to the client name, which is the procedure for a Disk Discovery backup.

To check the configured backup type, see the `Backup Object Summary` property page. Under the `Type` label, you will see `Client System` if you have configured a Disk Discovery backup or `Filesystem` if only the disks have been selected.

Figure 6-14 Selecting the Client

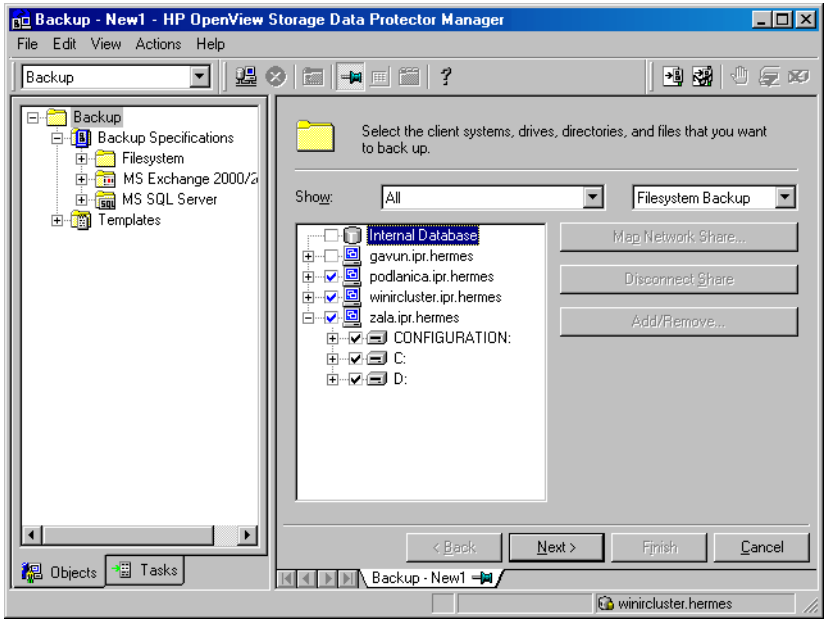

See "Using Backup Options" on page 269 for information on using and structuring your backup specifications.

Backing Up Windows Shared Disks

Data Protector allows you to back up data on Windows shared disks. You have to use a regular Disk Agent client, which can then be used to back up other remote systems via shared disks. Then you can configure a backup specification.

NOTE

Backup using the shared disk method is a workaround for backing up systems which cannot be backed up otherwise. It is not recommended to use it as the main backup approach.

When to Use Shared Disks Backup

Use shared disks backup in either of the following situations:

- The remote system does not belong to the Data Protector cell and does not have the Data Protector Disk Agent installed.
- The platform to be backed up is not directly supported by Data Protector, for example, Windows 3.11.

TIP

To reduce the network load, a Disk Agent client should be a Media Agent client as well. Otherwise, data is transferred over the network twice.

You can use one Windows client to manage backups and restores involving shared disks or other remote systems. Backup performance may be reduced if you start too many backups at a time, since one Disk Agent is started for each backed up disk. In this case, you should configure additional Disk Agent clients to increase the backup speed.

Limitation

Backing up writers that store their data on network shared volumes using the VSS functionality is not supported.

Requirements

- Use the Windows GUI, because browsing of Windows systems is not supported in the UNIX GUI.
- You have to map the shared drives using the Backup wizard.

IMPORTANT

The Disk Agent client must have the `Inet` service configured using an account with access to the shared disks. This must be a specific user account, not the system account. See “Setting the User Account for the Data Protector Inet Service” on page 232 for more information on how to use the appropriate logon account.

Once you have set the user account for the `Inet` service, you can back up the shared disks as though they were residing on the local system.

How to Perform a Windows Shared Disks Backup

1. In the `Data Protector Manager`, switch to the `Backup` context.
2. Expand the `Backups` item, and then double-click `Backup Specifications`.
3. Right-click `Filesystem`, and then `Add Backup`.
4. In the `Create New Backup` dialog box, select one of the available templates, and then click `OK` to open the wizard.
5. On the first page of the wizard, in the drop-down list, select `Network Share Backup`.
6. Click `Map Network Share`. The `Browse Network Shares` dialog box opens.
7. In the `Client System` drop-down list, select the client with the Disk Agent that will be used to back up the remote system.
8. Select the shared disk. It appears in the `Share Name` text box.
9. Enter the required information. See online Help for details.

NOTE

If a directory is shared over network, the share information is by default backed up. During the restore, share information is restored by default and directory is shared on the network after the restore. You can change this behavior by unselecting the `Backup share information for directories` option in the Filesystem options window.

Setting the User Account for the Data Protector Inet Service

The following procedure describes how to change the user account used by the `Data Protector Inet` service to access disks that belong to remote computers. This account must have permission to access both the local client and the remote shared disks. It must be a specific user account, not the system account.

Proceed as follows to change the user account on a Windows Disk Agent client:

1. In the `Control Panel`, click `Administrative Tools`, and then double-click `Services`.
2. Scroll down the list of services and select `Data Protector Inet`.
3. Under the `General` property page, click `Stop`. Then select the `Log On` tab.

Figure 6-15 Inet General Property Page on Windows

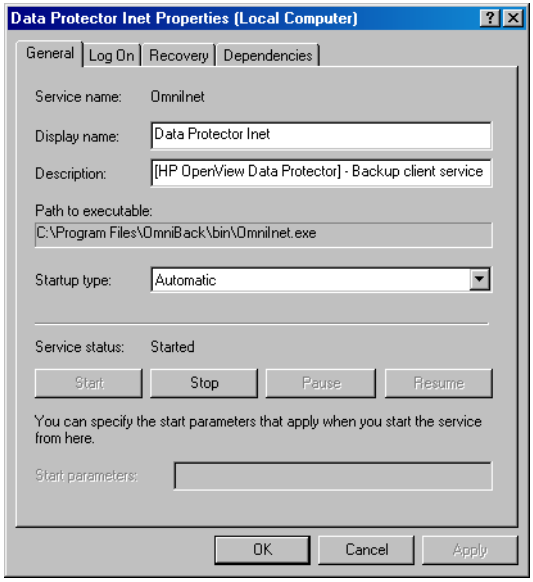

4. In the `Log On As` area, select the `This Account` button.
5. Enter or browse for the account that has the correct permission to access the shared disks you want to back up.

6. Enter the password, then confirm it.

Figure 6-16 Inet Logon option on Windows

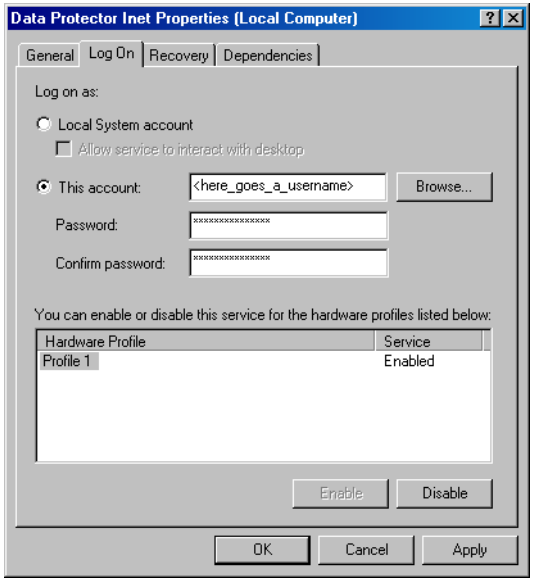

7. Click `Apply` to apply the changes and then restart the service by clicking the `Start` button in the `General` property page.

Backing Up Windows Disks as Disk Image Objects

What Is a Disk Image Backup?

A disk image backup is a high-speed backup of disks, disk partitions, or logical volumes without tracking the file and directory structure stored on these data sources.

When to Use a Disk Image Backup

Use a disk image backup in the following situations:

- You have lots of small files and a high backup speed is required.
- A full disk backup is needed, for example, for disaster recovery or before a major software update.
- A direct disk-to-disk connection is not possible and you want to duplicate a filesystem to another disk. The latter must be identical to the original disk.

How to Specify a Disk Image Section

You can specify a disk image section in two ways. In case of a zero downtime backup (snapshot or split mirror), you must use the second way.

- `\\.\<drive_letter>`, for example: `\\.\E:`
- `\\.\PHYSICALDRIVE#`,

  where # is the current number of the disk you want to back up.

  For example: `\\.\PHYSICALDRIVE3`

Where to Find a Disk Number (Physical Drive Number)

On Windows systems, you can find the current numbers of your disks (as well as the drive letters) by clicking `Control Panel`, `Administrative Tools`, `Computer Management`, `Storage`, `Disk Management`.

Figure 6-17 **The Numbers Representing Disks (Physical Drive Number) on a Windows System**

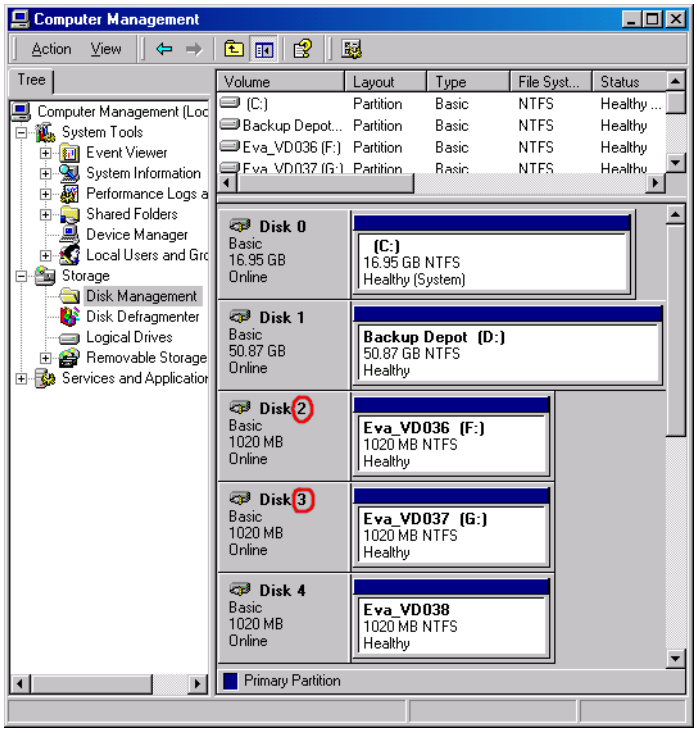

NOTE The numbers representing disks can change if the system is rebooted.

Limitation A disk image backup fails if a file on the target system is open, since Data Protector cannot lock the file.

How to Perform a Disk Image Backup To perform a disk image backup, use the `Manual add` function from the `Backup Object Summary` page. For detailed steps, refer to the online Help index keyword “backing up, disk images”.

Figure 6-18 The Manual Add Functionality

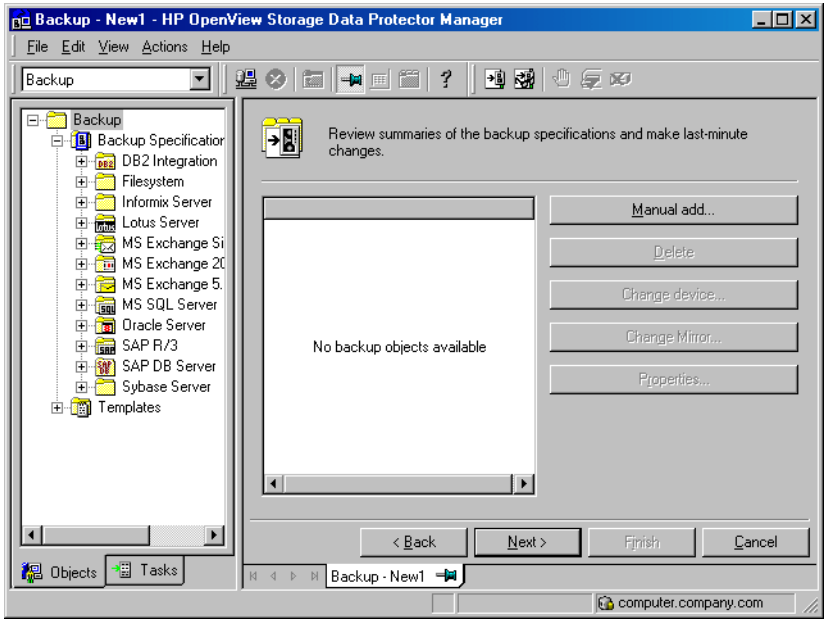

Backing Up Novell NetWare Systems

This section describes how to back up Novell NetWare filesystems and NDS/eDirectory.

Backing Up Novell NetWare Filesystems (Volumes)

Prerequisites

To back up data on a Novell NetWare system, install the Novell NetWare Disk Agent on the Novell NetWare system. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.

To use backup devices connected to a Novell NetWare system with Data Protector, install the General Media Agent on the Novell NetWare system. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.

How to Back Up a Novell NetWare System

To back up Novell Netware filesystems, proceed as follows:

1. In the `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context.
2. Expand `Backup`, right-click `Backup Specifications`, and then click `Add Backup`.
3. In the `Create New Backup` dialog box, select one of the available templates, and then click `OK` to open the wizard.
4. In the drop-down list, select `Filesystem Backup`.
5. Expand the client whose data you want to back up.
6. Select the backup objects. Follow the wizard to select a backup device.
7. In the next wizard page, click the `Advanced Filesystem Options` tab to open the `Filesystem Options` dialog box. Click the `NetWare Options` tab to set backup options. Refer to “Object Options” on page 284 for a description of the options.
8. Follow the wizard and then save and start your backup.

What Is Backed Up?

The directory structure and the files are backed up as well as the following filesystem information:

- Four Name Space information: DOS, Mac, NFS, Long

- Trustee information
- Inherited right mask
- File and directory attributes
- Time attributes (creation date/time, modification date/time, last accessed date/time, last modified date/time, last archived date)
- Owner
- Owning name space
- Search mode
- Volume or directory space restrictions. To back up volume restrictions, select the whole volume object for backup.

Server Specific Information is backed up separately as a part of the CONFIGURATION object.

After backing up each file, the file's archive flag is cleared and the archive time is set.

What Is Not Backed Up?

- Files that are opened for shared access with the `Denied read` option enabled cannot be backed up by Data Protector. You can set the `Number of retries` option to increase the probability of the file being backed up. This option is only useful if the applications operate in such a mode that they use a certain file and then release it after a certain time.
- System files that are in Queue directories are not backed up.
- All files that belong to NDS/eDirectory are skipped. You can back up NDS/eDirectory separately.
- Extended attributes (which can be installed as a NetWare addition) are not backed up.

Limitations

- The following backup options are unavailable for NetWare:
  - `Software cmpression`
  - `Do not preserve access time attributes`
  - `Lock files during backup`
- Only files not exceeding 4 GB can be backed up on NetWare 5.1 due to the Novell Storage Management Services (SMS) limitation. Other Novell NetWare systems do not have file size limitations.

- Data Protector cannot back up **moved files** during incremental backup sessions.

  To determine the files that have been modified, the Data Protector Disk Agent checks the last modification time of each file. This method prevents Data Protector from detecting moved files, as moving the file does not change the modification time.

  See the *HP OpenView Storage Data Protector Concepts Guide* for details about incremental backups.

To allow users to run backups on the Novell NetWare system, grant them the `Backup as Root` user right. See Chapter 4, "Configuring Users and User Groups," on page 127 for details on how to change user rights.

Selecting Specific Files or Directories

For each filesystem, you can restrict the backup to specific directory trees. For each directory tree you can do the following:

- Exclude any sub-tree or file
- Back up files that match a specific wildcard pattern
- Skip files that match a specific wildcard pattern

Backing Up the CONFIGURATION

Data Protector enables you to back up a special data structure known as `CONFIGURATION`, which consists of the following components, as shown in Figure 6-19 (NetWare 5.x and 6.x).

CONFIGURATION Components

- Server Specific Info
- Schema
- Root

To back up the `CONFIGURATION` object or part of it, follow the procedure "How to Back Up a Novell NetWare System" on page 237, selecting the appropriate item in the Source page of the backup wizard.

Figure 6-19 Backing Up NetWare 5.x Configuration

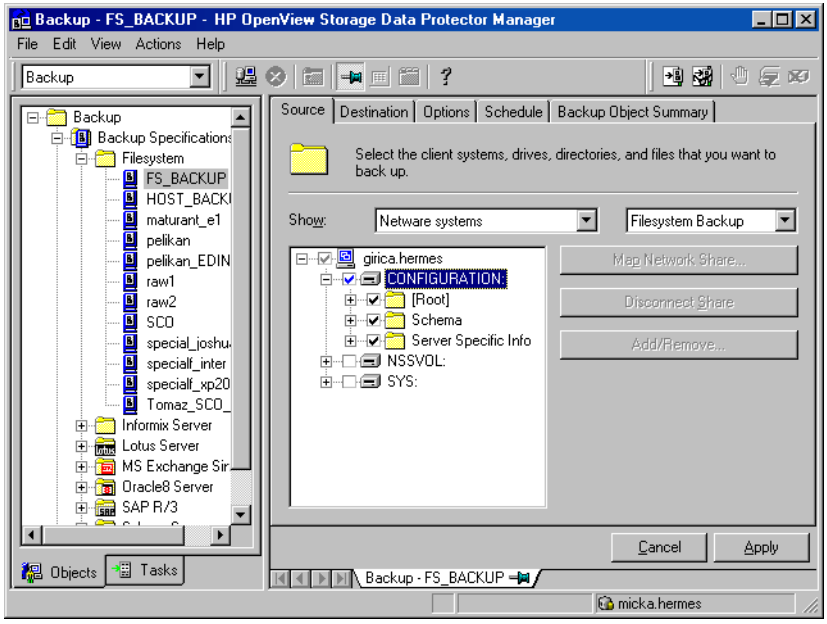

Client Backup with Disk Discovery

You can discover disks (volumes) on NetWare just as you can for UNIX or Windows systems.

How Are Disks Discovered?

If you specify a client backup with disk discovery, Data Protector first contacts the client and discovers all volumes that belong to the client. Then it backs up the `CONFIGURATION` item and each discovered volume as a regular filesystem. The description text for each filesystem object is generated by appending the volume name, in square brackets, to the description of the client backup.

For client backup with disk discovery, it is not possible to select only specific directory trees, because this implies single volume backup. But it is possible to exclude any directory from the backup.

How to Perform a NetWare Client Backup

1. In the `Data Protector Manager`, switch to the `Backup` context.
2. Expand the `Backup` item, and then double-click `Backup Specifications`.
3. In the Results Area, right-click `Filesystem`, and then click `Add Backup`.
4. In the `Create New Backup` dialog box, select one of the available templates.
5. Click `OK` to open the wizard.
6. Click the check box next to the client. This selects the entire client to be backed up, similar to what is shown in Figure 6-14.

See “Using Backup Options” on page 269 for information on using and structuring your backup specifications.

Backing Up NDS/eDirectory

Data Protector backs up NDS/eDirectory using Novell NetWare Storage Management Services (SMS). Data Protector backs up and restores all extensions to the NDS/eDirectory Schema.

NOTE Incremental backup of the NDS/eDirectory database is not possible. A full backup of the NDS/eDirectory database is always performed.

To successfully back up NDS/eDirectory, follow the instructions in the *HP OpenView Storage Data Protector Installation and Licensing Guide* and ensure that:

- `TSANDS.NLM` is loaded
- `HPLOGIN.NLM` is loaded and access information is given to Data Protector

Backing Up NDS/eDirectory

Back up NDS/eDirectory as you would a UNIX or Windows filesystem, except that the mountpoint has to be the `CONFIGURATION` item.

Adding NDS/eDirectory Objects to a Backup Specification

Data Protector offers advanced functionality to back up only a part of NDS/eDirectory. However, unless you understand why some parts can be excluded, it is advisable to back up everything.

Each object in the NDS/eDirectory tree has its own fully distinguished name. For example, leaf object `CN=Admin`, which resides in the container object `O=HSL`, has its fully distinguished name as seen by the SMS (`TSANDS.NLM`):

```
.CN=Admin.O=HSL.[Root]
```

Data Protector uses the fully distinguished name to build the tree structure of NDS/eDirectory as follows:

- The fully distinguished name is reversed.
- The dot-symbol (`.`) separator is replaced with the slash-symbol separator (`/`).

For example, the fully distinguished name

`.CN=Admin.O=HSL.`[Root]

has its counterpart used by Data Protector, containing forward slashes, which are used for Windows as well:

```
/[Root]/O=HSL/CN=Admin
```

Except for this naming rule, the Data Protector backup specification syntax is the same as for Novell NetWare or UNIX filesystem objects.

NOTE

NDS/eDirectory objects (container and leaf objects) are represented and backed up as directories. These objects can be skipped using the `skip` option or backed up using the `only` option. Data Protector views the [Root] object as a non-containment object, so the [Root] object cannot be excluded.

The Mountpoint Configuration File TSANDS.CFG

For the best protection of your NDS/eDirectory data, you should perform a full directory backup of the NDS/eDirectory Schema and all containers in the tree starting with the [Root] object. However, there are situations where you might prefer to begin backing up NDS/eDirectory from a container other than the [Root] object, but a configured user does not have sufficient rights to browse through to the starting container's context.

To facilitate backing up portions of the NDS/eDirectory tree, Novell has provided a text file, `SYS:SYSTEM\TSA\TSANDS.CFG` file, that allows you to specify the names of containers where you want backups to begin. This file is located on the server where `TSANDS.NLM` is loaded.

To begin your NDS/eDirectory backup from the HSL container, create a `TSANDS.CFG` file containing the line:

```
.O=HSL.[Root]
```

An additional mountpoint becomes available to the backup configuration.

Backing Up OpenVMS Systems

This section describes how to back up OpenVMS filesystems.

Backing Up OpenVMS Filesystems

Prerequisites

To back up data on a OpenVMS system, install the OpenVMS Disk Agent on the OpenVMS system. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.

To use backup devices connected to an OpenVMS system with Data Protector, install the General Media Agent on the OpenVMS system. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.

Limitations

- Any file specifications that are entered into the GUI or passed to the CLI must be in UNIX style syntax:

  ```
  /disk/directory1/directory2/filename.ext.n
  ```

  — The string should begin with a slash, followed by the disk, directories, and filename, separated by slashes.

  — Do not place a colon after the disk name.

  — A period should be used before the version number instead of a semi-colon.

  — File specifications for OpenVMS files are case insensitive.

  For example:

  An OpenVMS file specification of:

  ```
  $1$DGA100:[USERS.DOE]LOGIN.COM;1
  ```

  must be specified in the form:

  ```
  /$1$DGA100/Users/Doe/Login.Com.1
  ```

- There is no implicit version number. You always have to specify a version number. Only file versions selected for the backup will be backed up. If you wish to include all versions of the file, select them

all in the GUI window, or, using the CLI, include the file specifications under the `Only (-only)` option, including wildcard characters for the version number, as follows

```
/DKA1/dir1/filename.txt.*
```

- If the `Do not preserve access time attributes (-touch)` option is enabled during a backup, the last accessed date will be updated with the current date and time on ODS-5 disks. On ODS-2 disks, this option has no effect, and all the dates remain unchanged.
- Implementation of incremental backups on OpenVMS systems is based on the revised time. To determine which files have changed, the Data Protector Disk Agent checks when each file was last revised.
- Rawdisk backups are not available on OpenVMS. There is no equivalent to a “BACKUP/PHYSICAL”.
- The `Backup POSIX hard links as files (-hlink)`, `Software compression (-compress)`, and `Encode (-encode)` options are not available on OpenVMS.

  Files with multiple directory entries are only backed up once using the primary pathname. The secondary path entries are saved as soft links. During a restore, these extra path entries will also be restored.

  There is no support for an equivalent to BACKUP/IMAGE. To make a restored copy of an OpenVMS system disk bootable, the OpenVMS WRITEBOOT utility has to be used to write a boot block on to the restored disk.
- Files being backed up are always locked regardless of whether the `Lock files during backup (-lock)` option is enabled or disabled. With the `-lock` option enabled any file opened for write is not backed up. With the `-lock` option disabled any open file is backed up as well.
- The default device and directory for pre- and post-exec command procedures is `/omni$root/bin`. To place the command procedure anywhere else the file specification must contain the device and directory path in UNIX style format. For example: `/SYS$MANAGER/DP_SAVE1.COM`
- When specifying wildcard characters for `Skip (-skip)` or `Only (-only)` filters, use ‘*’ for multiple characters and ‘?’ for single characters.

How to Back Up an OpenVMS System

To back up an OpenVMS filesystem, proceed as follows:

1. In the `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context.
2. Expand `Backup`, right-click `Backup Specifications`, and then click `Add Backup`.
3. In the `Create New Backup` dialog box, select one of the available templates, and then click `OK` to open the wizard.
4. In the drop-down list, select `Filesystem Backup`.
5. Expand the client whose data you want to back up.
6. Select the backup objects. Follow the wizard to select a backup device.
7. Select backup options. See “Using Backup Options” on page 269 for details.
8. Follow the wizard and then save and start your backup.

What Is Backed Up?

The directory structure and the files are backed up, together with the following filesystem information:

- File and directory attributes
- ACL (Access Control List)

Files can be backed up from mounted FILES-11 ODS-2 or ODS-5 volumes only.

Backing Up in a Direct Backup Environment

This section provides the steps for the configuration of a direct backup backup specification. Please refer to *HP OpenView Storage Data Protector Concepts Guide* for a complete information on direct backup concepts.

Prerequisites

- The application and backup systems must be configured for ZDB on HP StorageWorks Disk Array XP. Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.
- If backing up an Oracle9i server on the application system, the application system must be configured for the Oracle9i split mirror or snapshot backup, depending on the disk array used. Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide*.
- The XCopy engine must be configured in the same SAN zone as the source (mirror disk connected to the backup system) and the destination (backup device connected to a Fibre Channel bridge). In other words, the XCopy engine must have SAN access to both the mirror disk connected to the backup system and to backup device connected to a Fibre Channel bridge.
- You need to have HP StorageWorks Disk Array XP agent installed on every system that is to be backed up (application system). Refer to *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- You need to have the General Media Agent and HP StorageWorks Disk Array XP agent installed on every system that controls a backup device (backup system). Refer to *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- You need to have at least one backup device configured for direct backup in the Data Protector cell. Refer to "Configuring Devices for Direct Backup" on page 45.
- You need to have prepared media for your backup. Refer to Chapter 5, "Managing Media," on page 143.
- You need to have appropriate user rights for performing a backup. Refer to Chapter 4, "Configuring Users and User Groups," on page 127.

Limitations

- The systems in the direct backup environment must be HP-UX 11.11.
- The `min` and `max` options for the `Load balancing` option are ignored for direct backup. All devices selected in the backup specification are load balanced, if load balancing is used. Consequently, it is not possible to set the order in which the selected devices are used using the Order devices functionality.
- The pre-exec and post-exec options for backup objects are not available for direct backup of raw logical volumes. They are available for Oracle9i direct backup.
- The backup device must be either attached to an external FC bridge with the XCopy engine, or must have the FC bridge with the XCopy engine embedded internally.
- Backup and restore of an Oracle database installed on raw partitions (rawdisk or raw logical volumes) are not supported.
- Backup and restore of striped logical volumes are not supported.
- The `CRC check` option is ignored with direct backup.
- The disk agent `Concurrency` option is ignored with direct backup.
- The `Block size` option is FC bridge dependent.
- The `Segment size` and `Disk agent buffers` options are ignored with direct backup.
- Object mirroring is not supported with direct backup.

Instant Recovery

Instant Recovery is supported only if:

- Control files and online redo logs do not reside on the same logical volumes as data files.
- A whole database backup was performed. That means that all data files that belong to an Oracle9i Server instance were selected during a backup.

Restore

The data backed up in a direct backup environment can be:

- Restored from a backup medium over the LAN directly to the application system following the Data Protector rawdisk or Oracle restore procedure. Refer to “Restoring Disk Images” on page 342 (rawdisk restore) or *HP OpenView Storage Data Protector Integration Guide* (Oracle restore).

- Restored using the Data Protector instant recovery functionality. Refer to the *HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide*.

Backup Specification Configuration Procedure

A direct backup specification can be configured for the following objects:

- rawdisks
- Oracle9i databases (online)
- Oracle9i databases (offline)

Backing Up Rawdisks

Please refer to the online Help index keyword "configuring direct backup specifications" for detailed information on how to configure a rawdisk direct backup specification.

Backing Up Oracle9i Online and Offline

Please refer to the online Help index keyword "configuring direct backup specifications" for detailed information on how to configure an Oracle9i online or offline direct backup specification.

Starting Direct Backup Using the CLI

After a direct backup specification has been configured, you can start the direct backup session using the GUI as described in the previous section, or using the CLI as follows:

- for rawdisks

  ```
  omnib -datalist <Name>
  ```

- for Oracle9i online and offline

  ```
  omnib -oracle8_list <Name>
  ```

where *<Name>* is the name of the direct backup specification.

Scheduling Unattended Backups

Data Protector allows you to configure unattended backups by scheduling backups of your systems at specific times.

The configuration and your scheduling policies can significantly influence the effectiveness and performance of your backup.

Key Points

- To simplify scheduling, Data Protector provides backup specifications for group clients. All clients configured in one backup specification are backed up at the same time in a single backup session.
- Be sure to have sufficient media and devices to run unattended backups smoothly. See Chapter 9, “Monitoring, Reporting, Notifications, and the Event Log,” on page 379 for details on monitoring running sessions and setting up e-mail or other notifications for a mount request.
- When the scheduled backup is started, Data Protector tries to allocate all the needed resources, such as licenses, devices, and access to IDB. If one of the needed resources is not available, the session is marked as queued. Data Protector will try to find the needed resources for the queued session once every minute until the time-out period is reached.

  Once Data Protector finds the resources, one of the queued sessions is started. The queued sessions are not started in the order they are displayed.
- To prevent Cell Manager overload, a maximum of up to five backup sessions can be started at the same time. If more are scheduled at the same time, the sessions are queued.
- For each individual or periodic scheduled backup, you can specify the following options: `Backup type` (full or incremental), `Network load`, and `Backup protection`. With split mirror or snapshot backup, in the case of ZDB to disk or ZDB to disk+tape (instant recovery enabled), you specify the `Split mirror/snapshot backup` option. For split mirror and snapshot backups, the backup type is ignored (it is set to full).

- Each backup specification can be scheduled multiple times with different option values. Within one backup specification, you can schedule both ZDB to disk and ZDB to disk+tape, and specify a different data protection period for each individual or periodic scheduled backup.
- Data and catalog protection settings determine the period that data is kept on a medium (data protection) and in IDB (catalog protection). See “Data Protection: Specifying How Long Data Is Kept on the Media” on page 272 and “Catalog Protection: How Long Info Is Kept in the Database” on page 274 for details.
- When applying a backup template, the schedule settings of the template override the schedule settings of the backup specification. After applying the template, you can still modify the backup specification and set a different schedule.

NOTE You can schedule backups up to a year in advance. Periodic backups do not have a defined time limit.

Handling Scheduling Conflicts

When scheduling periodic backups, it can happen that the chosen backup start time is already occupied by another scheduled backup in the same backup specification. In that case, Data Protector prompts you that there are scheduling conflicts, and asks if you wish to continue. If you click `Yes`, the new schedule will be applied where possible (on the days when the time slot is still free). If you click `No`, the new schedule will be discarded.

Planning Your Scheduling Policies

See the *HP OpenView Storage Data Protector Concepts Guide* for answers to questions such as:

- How do I plan a scheduling policy for my environment?
- How does the amount of data influence my scheduling policy?
- How long will the backup take?
- How many media do I need for the backup?
- How do I plan for a disaster recovery?

Starting Backups on Specific Dates

Data Protector allows you to define the date and time when you want your unattended backup to start. You usually want to back up on specific dates when configuring exceptions to your regular periodic backups, for example, if you want to back up some data before a specific event.

How to Configure Backups on Specific Dates

To configure a backup on a specific date, you can create a new backup specification, or modify an existing one. For detailed steps, refer to the online Help index keyword “scheduling backups on specific dates and times”.

Starting Periodic Backups

Periodic backups are based on a time period after a specific date. For example, you may configure periodic backups so that a full backup is done on Sunday at 3 a.m. and repeated every two days. The next full backup would be at 3 a.m. the following Tuesday. Periodic backups simplify backup configuration for regularly scheduled backups.

Data Protector provides predefined backup schedules to simplify the configuration.

Predefined Backup Schedules

The predefined backup schedules provided can be used to simplify your configuration. You can modify the schedules later. Schedule types include those described in the following sections:

Daily intensive

Data Protector runs a full backup at midnight and two additional incremental backups at 12:00 (noon) and 18:00 (6 p.m.) every day. This backup type is intended for database transaction servers and other environments with intensive backup requirements.

Daily full

Data Protector runs a full backup every day at 21:00 (9 p.m.). This is intended for backups of single workstations or servers.

Weekly full

Data Protector runs a full backup every Friday and Incr1 backups every day from Monday to Friday at 21:00 (9 p.m.). This is intended for small environments.

Fortnightly full

Data Protector runs a full backup every second Friday. Between these backups, Data Protector runs Incr1 backups every Monday to Thursday, all at 21:00 (9 p.m.).

Monthly full

Data Protector runs a full backup on the first of every month, an Incr1 backup every week, and an incremental backup every other day. This is intended for relatively static environments.

How to Use a Predefined Schedule

To configure a backup using a predefined schedule, you can create a new backup specification, or modify an existing one. For detailed steps, refer to the online Help index keyword “scheduling periodic backups”.

Configuring a Recurring Backup

You can schedule a backup so that it starts at a specific time and date on a set schedule. For example, you could schedule a full backup to take place every Friday at 21:00 (9 p.m.) for the next six months.

How to Configure a Recurring Backup

To configure a recurring backup, you can create a new backup specification, or modify an existing one. For detailed steps, refer to the online Help index keyword “scheduling periodic backups”.

Editing Your Backup Schedule

Clearing a Schedule

To eliminate a schedule that you have already set up, click `Reset` in the `Schedule` property page.

When you clear a schedule, you clear all the schedule settings of a specified mode for the current year.

Undoing the Clear

To undo the schedule clearing, click `Undo` in the `Schedule` property page.

Changing the Start Date

To change the start date, follow the procedure for setting up a backup for a specific date. See “Starting Backups on Specific Dates” on page 252.

Disabling a Schedule

To disable a backup schedule, select the `Disable Schedule` option in the `Schedule` property page. The backup will not be performed until you deselect this option.

Disabling backup schedules does not influence currently running backup sessions.

Skipping Backups During Holidays

By default, Data Protector runs backups on holidays.

If you do not wish to run your backups on holidays, set the `Holidays` option to ON in the `Schedule` page of the Backup wizard. You can identify holidays from the `Holidays` file or dates marked red on the `Schedule Calendar`.

IMPORTANT It is generally not recommended to skip backups on holidays.

To set different holidays, edit the `Holidays` file on the Cell Manager, located in the following directory:

- on a UNIX Cell Manager: `/etc/opt/omni/server/Holidays`
- on a Windows Cell Manager: `<Data_Protector_home>\Config\server\holidays`

Consider the following when editing or adding new entries in the `Holidays` file:

- The first number in each line indicates the consecutive day of the year. The value is ignored by Data Protector, but the value must be set between `0` and `366`. You can set it to 0 to indicate that the number does not correspond to the date that follows it.
- The date is specified as `Mmm d`, where `Mmm` is the three-letter abbreviation of the month and `d` is day of month as a number (for example `Jan 1`). Note that the month must be specified in English, regardless of your locale.
- The description of the holiday is optional and is currently not used by Data Protector.

Regardless of the year specified at the top of the file, the holidays specified in the file *are always used as-is* and must be edited manually if the holidays do not occur on the same dates each year. If you are not using the `Holidays` option for the scheduler, you can remove or comment out the entries in the `Holidays` file to prevent confusion in case of accidental use of a `Holidays` file that is out of date or has not been customized for your country or company specific requirements.

How you configure your scheduling policies strongly influences the effectiveness and performance of your backup. For example, if the date `January 1` is registered as a holiday, Data Protector will not back up on that date. If you have scheduled a full backup for January 1st and an incremental for January 2nd, Data Protector will skip running the full backup on January 1st but will run the incremental backup scheduled for January 2nd. The incremental backup will be based on the last full backup.

Configuring Backup Options When Scheduling Backups

When scheduling a backup, you can set further options. These options are only valid for scheduled backups and not for those started interactively. Data protection that is specified in the `Schedule Backup` dialog overrides protection settings anywhere else in the backup specification.

How to Set Schedule Backup Options

You can set schedule backup options when creating a new backup specification, or when modifying an existing one. For detailed steps, refer to the online Help index keyword “setting schedule backup options”.

Running Consecutive Backups

You can start a backup after one is finished. For example, you can start a backup of an Oracle database after a filesystem backup is finished. For detailed steps, refer to the online Help index keyword “running consecutive backups”.

For details on `pre-` and `post-exec` scripts on UNIX systems, refer to “Examples of Pre-Exec and Post-Exec Commands for UNIX” on page A-21.

Selecting a Backup Type: Full or Incremental

To save time and media during a backup, you can combine full and incremental backups. For example, you can create a second-level incremental backup based on a previous first-level incremental backup, a third-level incremental backup based on a previous second-level incremental backup, and so on.

The backup type (full or incremental) applies to the entire backup specification and only to filesystem objects.

The backup type is ignored for zero downtime backup sessions (split mirror or snapshot backup). It is set to full.

To combine full and incremental backups, make sure that the backup object has exactly the same:

- client name
- drive/mountpoint
- description

  The description can be set for the whole backup specification or for a specific object. Refer to “Backup Specification Options” on page 280 and “Object Options” on page 284.

- owner

  Backup ownership can be set for the whole backup specification. Refer to “Ownership: Who Will Be Able to Restore?” on page 279.

Backup Types

- Full backup

  A full backup consists of all backup objects, even if they have been backed up before. The first backup of an object is always a full backup. Any subsequent backup will be completed as full if no protected full backups with the same ownership are available at the backup time.

- Incr backup

  This backup type is based on any previous, still protected backup chain, either a full or an incremental backup. An incremental backup includes only the files that have changed since the last still protected

backup. Even if the previous backup was an incremental (**Incr** or **Incr1**, **Incr2**, ...) backup, the subsequent incremental backup includes only those files that changed in the interim.

- Incr1 backup

  This backup type refers to the most recent still protected full backup with the same ownership. It does not depend on any previous incremental backups. The files that have changed since the most recent still protected full backup are included in the backup.

- Incr2 backup

  This backup type refers to the most recent still protected full backup, provided that there is no Incr1 done afterwards. If there are several Incr1 backups available, it refers to the most recent one. All files that have changed since the reference backup was done are backed up.

- Incr1-9 backup

  The description above explains the concept of incremental levels, which can be extended up to Incr9.

Table 6-1 shows the relative referencing of backup runs with various backup types. See the text following the table for a full explanation.

Table 6-1 Relative Referencing of Backup Runs

1	Full	<----	Incr1				
2	Full	<----	<----	<----	Incr2		
3	Full	<----	Incr1	<----	Incr2		
4	Full	<----	Incr				
5	Full	<----	Incr1	<----	Incr		
6	Full	<----	Incr1	<----	Incr2	<----	Incr
7	Full	<----	Incr1	<----	Incr	<----	Incr
8	Full	<----	Incr1	<----	Incr3		
9	Full	<----	Incr1	<----	Incr2	<----	Incr3
10	Full	<----	<----	<----	Incr2	<----	Incr3
11	Full	<----	<----	<----	<----	<----	Incr3

How to Read Table 6-1

- The rows in Table 6-1 are independent of each other and show different situations.
- The age of the backups increases from right to left, so that the far left is the oldest and the far right is the most recent backup.
- The full and IncrX represent still-protected objects of the same owner. Any existing IncrX that is not protected can be used for restore, but is not considered for referencing on subsequent backup runs.

Examples:

- In the second row, there is a full, still protected backup and an Incr2 is running. There is no Incr1, so the backup is executed as an Incr1.
- In the fifth row, there is a full backup, an Incr1 and another incremental is running. Data Protector references the currently running backup to the previous incremental, that is Incr1.
- In the eighth row, the Incr3 is executed as Incr2, and in the eleventh row, the Incr3 is executed as Incr1.

How to Select the Backup Type

If you perform an interactive backup, you are prompted to select the backup type. When scheduling a backup, you specify the backup type in the `Schedule Backup` dialog. You can, for example, create a schedule that runs the same backup specification as full on Saturday and as Incr1 on all working days.

Backup Type and the Restore Process

Keep in mind that full backups enable a simple and efficient restore, but require many media that can hold multiple versions of the entire backed up data. The time required to complete a backup is rather long. Incremental backups require fewer media resources, but have a more complex restore algorithm. Compare the following two examples:

1. `full ; Incr ; Incr ; Incr ; Incr  (-> time)`

   This example requires a shorter backup time and the media space required is lower. The restore process is more complex; many media need to be accessed, and the required time is longer if you want to restore to the state of the last Incr.

2. `full ; Incr1 ; Incr1 ; Incr1 ; Incr1  (-> time)`

   This example requires more time for backup and the media space consumption is a bit higher if compared to the first example. The restore process is simple; few media are needed, and the time spent on performing a restore is shorter than in the first example.

Using Backup Templates

Overview

Data Protector backup templates are a powerful tool that can help you simplify your backup configuration. A template has a set of clearly specified options for a backup specification, which you can use as a base for creating and modifying backup specifications. Data Protector enables you to apply a group of options offered by the template.

A template can be used in two ways:

- It can be used to create a new backup specification.
- It can be applied to existing backup specifications to modify these specifications.

Backup templates are created and modified similarly to backup specifications, except that objects and the backup application configuration are not selected within the backup template.

Data Protector Default Backup Templates

Data Protector offers you default templates for different types of data (Filesystem, Oracle/SAP, and so on) to configure a filesystem or an application backup. The templates provide typical settings, which can be used as a basis for your backup specifications.

Blank Backup Templates

In blank backup templates, such as `Blank Filesystem Backup`, `Blank Informix Backup`, and so on, there are no objects or devices selected. Backup specification options and object options have Data Protector default values, and there is no backup schedule. You can separately select the `Load balanced` option, enabling Data Protector to automatically balance the usage of devices selected for the backup.

Options Offered by Templates

When using a backup template for creating or modifying a backup specification, you can select or deselect options offered by the template.

Figure 6-20 Options Offered by Templates

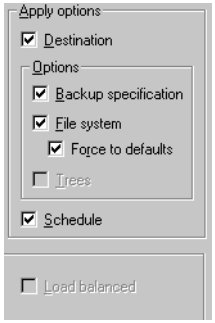

Destination Backup device settings specified in your template apply to your backup specification.

Backup specification Backup specification options specified in the template apply to your backup specification.

Filesystem Filesystem options specified in the template apply to all filesystem objects of your backup specification.

Force to defaults Filesystem object options specified in the template apply to all filesystem objects of your backup specification. These are the options in the `Backup Object Summary` page.

Trees Trees options specified in the template apply to your backup specification.

Schedule Schedule settings specified in the template apply to your backup specification.

Once you have applied the template options, you can still modify your backup specification and change any setting.

For more information on these options, refer to “Using Backup Options” on page 269.

Load balanced This option directs how data is distributed to devices.

See “Load Balancing: Balancing the Usage of Backup Devices” on page 276 for details about the `Load balancing` options.

Using a Backup Template When Creating a New Backup Specification

When creating a new backup specification, Data Protector offers you a set of backup templates, either default templates or templates you have created. Select an appropriate template, or optionally, select or deselect some groups of options, and then proceed with the Backup wizard.

To create a backup specification without predefined settings, select `Blank Filesystem Backup`.

Figure 6-21 Using Templates When Creating New Backup Specifications

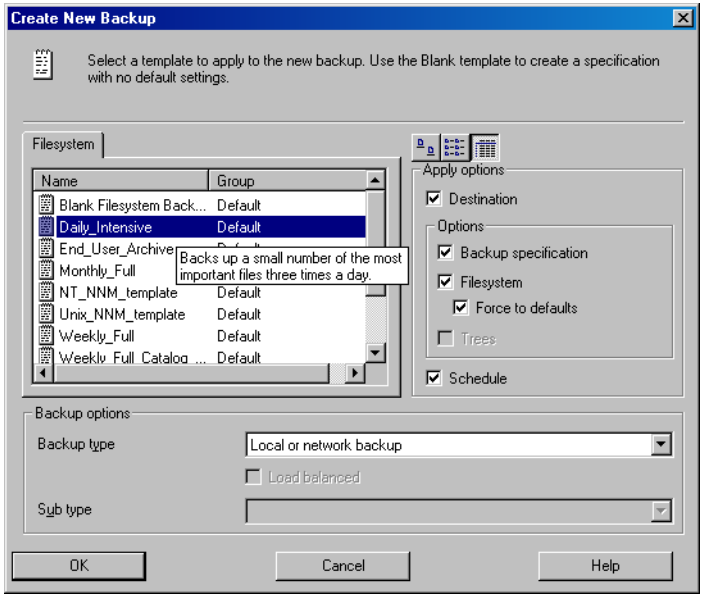

Applying a Backup Template

Data Protector allows you to apply a backup template to saved backup specifications. When applying a template to backup specifications, you can select which option groups should be applied. Refer to “Options Offered by Templates” on page 259.

The result of applying an option group is that all related options in this group are set to the state specified in the template.

To apply a template to backup specifications, right-click the backup specification and click `Apply Template`. The `Apply Template` window appears, in which you apply the desired options. For detailed steps, refer to the online Help index keyword “applying backup templates”.

Figure 6-22 The Apply Template Dialog Box

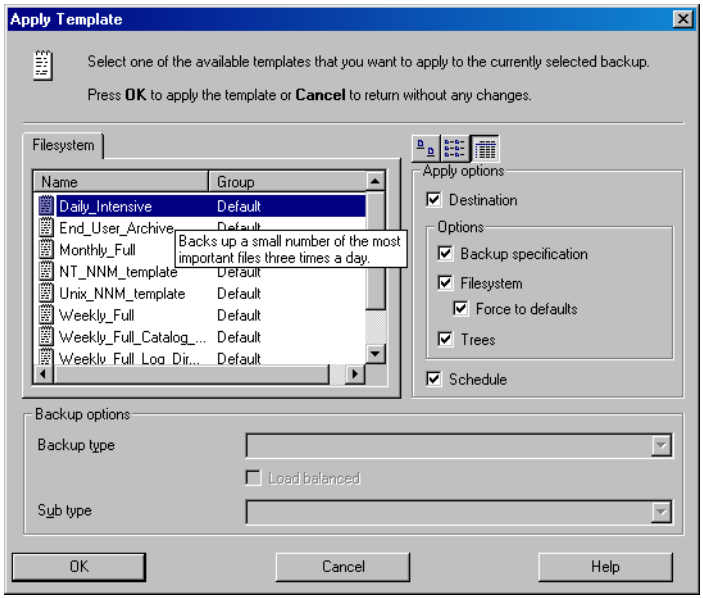

Integration Backup Specification

To apply a template to an integration backup specification, the backup specification you would like to apply should not be opened in the Results Area. If you first click on the backup specification to open it, and then try to apply the template to this backup specification, the `Apply Template` option will not be available.

IMPORTANT

If you select the `Force to defaults` option, the options specified in your template apply to all filesystem objects of your backup specification for which you changed options in the `Backup Object Summary` page.

Creating a New Template

You can create new templates and use them for creating or modifying backup specifications.

To create a new template, use the Backup context. For detailed steps, refer to the online Help index keyword “creating backup templates”.

For more information on individual options, refer to “Using Backup Options” on page 269.

Figure 6-23 Creating a New Template

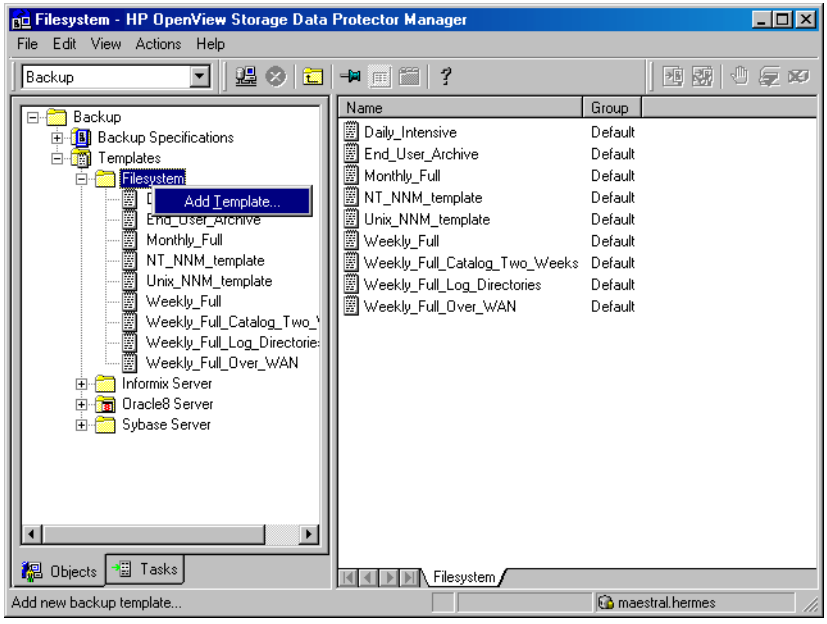

Modifying an Existing Template

You can modify Data Protector default templates, as well as templates that you have created.

To modify an existing template, open the properties of the template. For detailed steps, refer to the online Help index keyword “modifying backup templates”.

For more information on individual options, refer to “Using Backup Options” on page 269.

Handling of Small Reoccurring Backups

When you need to perform reoccurring backups of numerous small objects, you need to run numerous backup sessions. During each backup session, media are loaded and unloaded in the drive. Not only is such backup slow, but it also causes media to deteriorate. To use media more economically and save time, it is recommended to create a file library device and use it to perform small reoccurring backups to disk instead of tape. You can then use the object copy functionality to move the data from the disk to a tape medium.

Using this method, a backup will be performed faster and media will be used more economically because they will be loaded and unloaded only once, during the object copy session.

To perform frequent backups of numerous small objects, perform these tasks:

1. Configure a file library device. Set the block size of each writer to the block size of the device that will be used in the second stage. See “Creation and Configuration of the File Library Device” on page 111.
2. Create one backup specification for all small objects. Use the file library created in the first step for the backup. See “Configuring a Backup” on page 198 for details.
3. Perform or schedule the backup.
4. Use the object copy functionality to move the backed up data to tape. For an example of configuring an object copy specification for this purpose, refer to the online Help index keyword “disk staging”.

Groups of Backup Specifications

Data Protector offers you the ability to organize backup specifications into different groups. The purpose of grouping is to organize the specifications of multiple backups.

For example, backup specifications for *"Corporation X"* can be classified into three different groups:

- **USER_FILES**: This group contains backup specifications that perform weekly, full backups for all users in each of the ten departments. This is the main production backup.
- **SERVERS_DR**: This group contains backup specifications for the company's servers to prepare for disaster recovery. Each time a new server is installed or an existing server is upgraded, a new backup specification is created and added to this group.
- **END_USER**: This group is used to save backup specifications that are made as a result of an end-user request. For example, end users who want to free up some disk space have to archive their own hard disk first.

See Figure 6-24 on page 267.

Such a configuration can result in many backup specifications, often as many as 50, which are hard to manage if they are viewed together. Grouping the backup specifications into meaningful groups can facilitate finding and maintaining single backup specifications. This allows you to apply common options settings from a template to the entire group.

For example, if you want to change the list of devices to all backup specifications in the group, you can selectively apply the device settings of a template.

Figure 6-24 Example of Backup Specification Groups

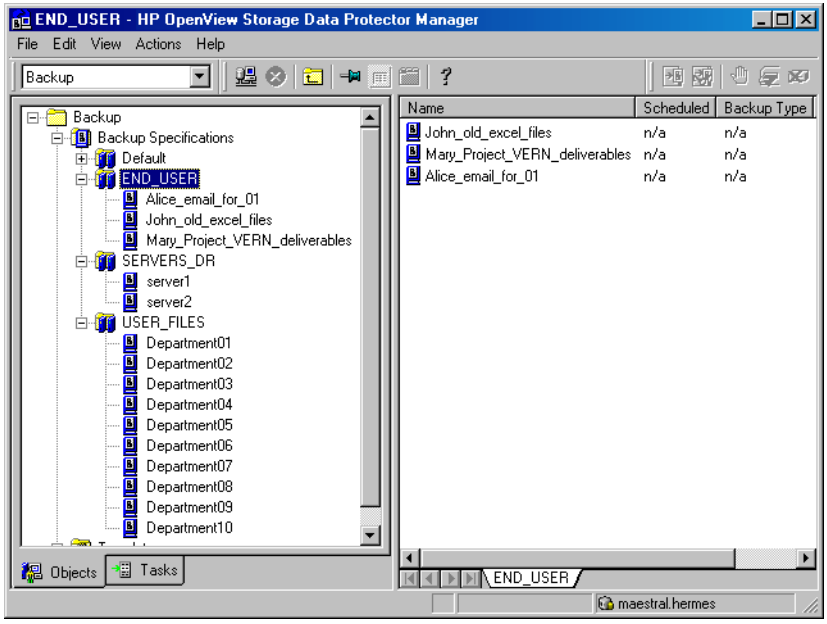

How to View and Create Groups

The following procedure describes how to view the available backup groups and how to create a new one:

1. In the `Data Protector Manager`, switch to the `Backup` context.
2. In the `View` menu, click `By Group`. The list of available backup groups appears under the `Backup Specifications` item. Clicking a group lists the backup specifications within that group.
3. Right-click the `Backup Specifications` item, and then click `Add Group`. The `Add New Group` dialog box appears.
4. In the `Name` text box, enter a name for your new group, and then click `OK`. Your new group will appear under the `Backup Specifications` item.

How to Save a Backup Specification in a Group

While saving the backup specification, you are also adding it to a group of backup specifications. If you do not specify the name of the group, a backup specification will be added to the `Default` group.

How to Delete a Group

Before deleting a group, you have to empty it first. One way of doing this is to move the backup specifications into other groups. See online Help for details.

Using Backup Options

Data Protector offers a comprehensive set of backup options to help you fine-tune your backups. All options have default values that are appropriate for most cases.

The availability of backup options depends on the type of data being backed up. For example, not all backup options available for a filesystem backup are available for a disk image backup. Common and specific application options for Exchange, SQL, and so on, are described in the *HP OpenView Storage Data Protector Integration Guide*.

Additionally, the `User defined variables` function lets you specify a variable name and its value for flexible operation on some platforms and integrations, for example, for backing up MPE platforms.

The backup options can be grouped as follows:

- Backup specification options, such as **Ownership** and **pre-** and **post-exec** options for the whole backup specification.
- Object options specifying how different backup objects, such as filesystems or disk images, are backed up.

  It is important to understand that object options can be set on two levels. First, you can set the *default object options* for all filesystems and for all disk image objects in the backup specification separately. Then you can set them differently *for a specific object*. These settings will override the defaults. For example, to compress data from all clients except for one with a slow CPU, set `compression` to `ON` when setting filesystem options. Then, select the slow client and set `compression` to `OFF` for this client.
- Device options define the behavior of backup devices. If you do not set the device options, the values are read from the device definition.
- Schedule options define the backup type, network load, and data protection for each individual or periodic scheduled backup. With split mirror or snapshot backup, in the case of ZDB to disk or ZDB to disk+tape (instant recovery enabled), you specify also the `Split mirror/snapshot backup` option.

For split mirror and snapshot backups, the backup type is ignored (it is set to full). Data protection that is specified in the `Schedule Backup` dialog overrides protection settings anywhere else in the backup specification.

See Figure 6-25 for a graphic scheme of how some of these options work. You can use backup templates to apply the same group of options to a number of backup specifications. Applying a template changes the backup specification according to the template. If you later change the template, you have to apply it again if you want the changes to take effect.

You can selectively apply schedule, device, and object options and the private exclusion list.

See also “Using Backup Templates” on page 259 for details about the backup templates.

Figure 6-25 **Backup Options**

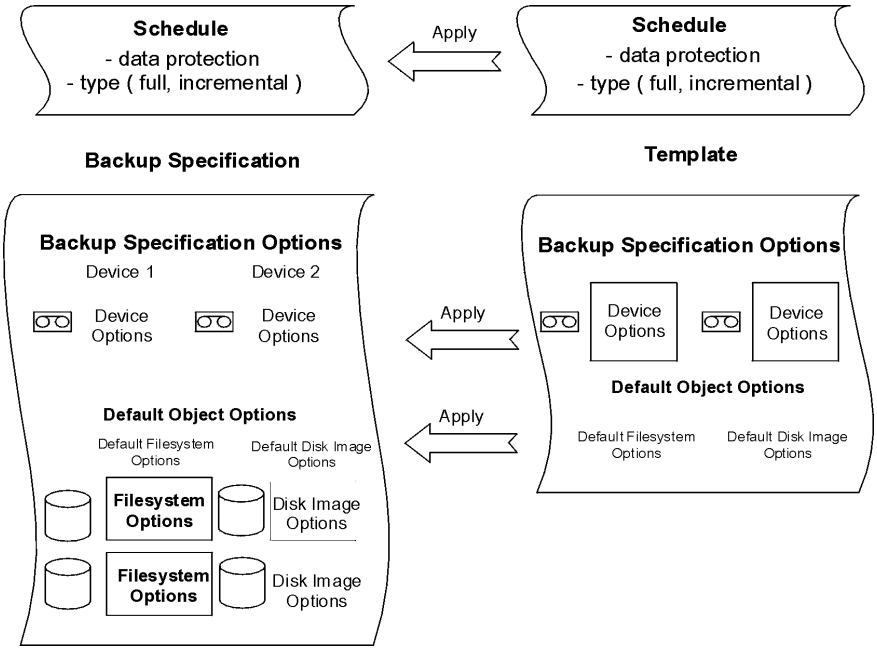

Most Frequently Used Backup Options

This section describes the options that are most likely to be modified according to specific backup policies. These are the following:

- “Data Protection: Specifying How Long Data Is Kept on the Media” on page 272
- “Catalog Protection: How Long Info Is Kept in the Database” on page 274
- “Logging: Changing Details About Data Stored in the Database” on page 275
- “Load Balancing: Balancing the Usage of Backup Devices” on page 276
- “Ownership: Who Will Be Able to Restore?” on page 279

Data Protection: Specifying How Long Data Is Kept on the Media

Configuring protection policies is extremely important for the safety of your data and for successful management of your environment. See the *HP OpenView Storage Data Protector Concepts Guide* for more detailed information on how to define these policies.

Based on your company data protection policies, you have to specify how long your backed up data is kept on the medium. For example, you may decide that data is out of date after three weeks and can be overwritten during a subsequent backup.

The `Protection` option can be specified for backup and object copy operations.

Limitation

Due to operating system limitations, the latest protection date that can be set is Jan 18th, 2038.

You can specify data protection in different places. Different combinations are available, depending on whether you are running an interactive backup, starting a saved backup specification, or scheduling a backup. The default value is `Permanent`.

- Interactive backups

  When configuring an interactive backup, you can change the default data protection for the entire backup. See Figure 6-26 on page 273. Additionally, you can specify different data protection periods for individual backup objects. The protection that is specified on the backup object level overrides the default protection setting. See Figure 6-27 on page 274.

- Backups using a saved backup specification

  When starting saved backups using the GUI, the data protection is applied as described for interactive backups.

  When starting saved backups using the CLI, you can also specify data protection. This will override all data protection settings in the backup specification.

- Scheduled backups

  You can specify a different period of protection for each individual or periodic scheduled backup. The data protection specified in the `Schedule Backup` dialog overrides all other data protection settings in the backup specification. If you leave the default protection, data protection is applied as described for interactive backups.

On how to specify data protection, refer to the online Help index keyword “specifying data protection”.

NOTE

If you apply a backup template to an existing backup specification and select the `Filesystem` and/or `Schedule` options, the protection settings from the template will replace the previous data protection settings in the respective parts of the backup specification. For more information, refer to “Options Offered by Templates” on page 259.

Figure 6-26 Backup Options: Protection

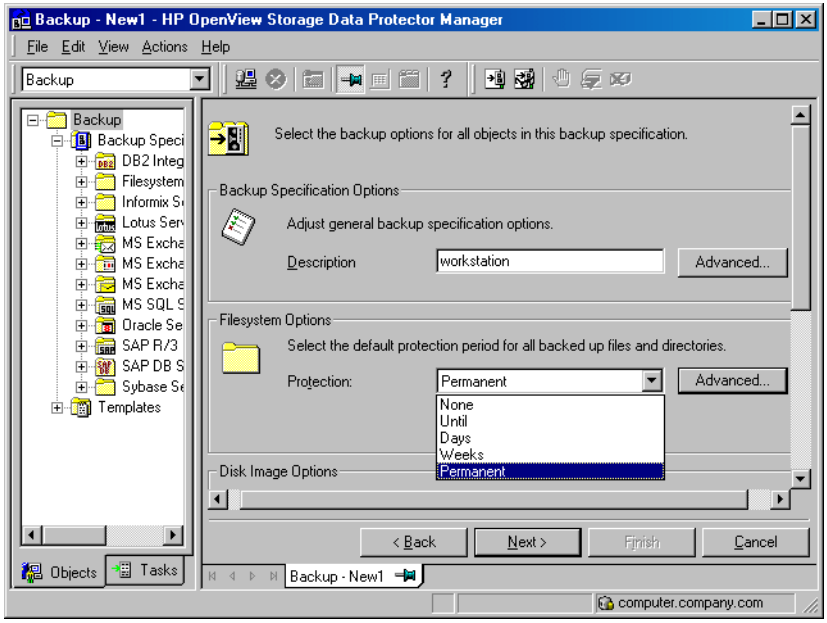

Figure 6-27 **Backup Object Properties - Options: Protection**

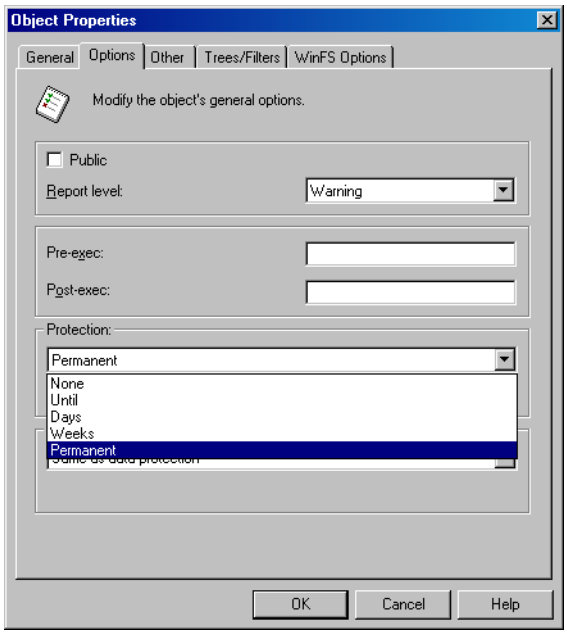

Catalog Protection: How Long Info Is Kept in the Database

Besides the `Protection` option, which controls how long data is protected on media, you can set the `Catalog Protection` option, which controls the time for which information about backed up files and directories is kept in IDB. Catalog protection and data protection can be set independently. Catalog protection has no effect if the log level is `Log None`.

The `Catalog Protection` option can be specified for backup and object copy operations.

The default value for catalog protection is `Same as data protection`. This means that you can browse and select files or directories as long as the media are available for restore.

NOTE If data protection expires, the catalog protection is cancelled. That is, when the data protection ends and a medium is overwritten, the catalogs for the objects are removed regardless of the catalog protection.

Even when catalog protection expires, you are still able to restore, but you must specify filenames manually.

Be aware that catalog protection, together with logging level, has a very big impact on the growth of the IDB. Therefore, it is very important to define a catalog protection policy appropriate to your environment. Refer to the IDB section in the *HP OpenView Storage Data Protector Concepts Guide* for more information on catalog protection and usage recommendations.

Limitation Due to operating system limitations, the latest protection date that can be set is Jan 18th, 2038.

Logging: Changing Details About Data Stored in the Database

The logging level determines the volume of detail on the backed up files and directories written to the IDB during a backup session. Note that you can restore your data regardless of the logging level used during a backup session.

Logging level can be specified for backup and object copy operations.

Data Protector provides the following four logging levels:

Table 6-2

Log All	This is the default logging level. All detailed information about backed up files and directories (names, versions, and attributes) is logged to the IDB. You can browse directories and files before restoring and in addition look at file attributes. Data Protector can fast position on the tape when restoring a specific file or directory.

Table 6-2

Log Files	When this logging level is selected, detailed information about backed up files and directories (names and versions) is logged to the IDB. You can browse directories and files before restoring, and Data Protector can fast position on the tape when restoring a specific file or directory. The information does not occupy much space, since not all file details (file attributes) are logged to the database.
Log Directories	When this logging level is selected, all detailed information about backed up directories (names, versions, and attributes) is logged to the IDB. You can browse only directories before restoring. However, during the restore Data Protector still performs fast positioning because a file is located on the tape near the directory where it actually resides. This option is suitable for filesystems with many auto-generated files, such as news and mail systems.
No Log	When this logging level is selected, no information about backed up files and directories is logged to the IDB. You will not be able to search and browse files and directories before restoring.

The different logging level settings influence the IDB growth, backup speed, and the convenience of browsing data for restore.

To be able to select the most appropriate logging level setting, it is important to understand the consequences. Refer to the *HP OpenView Storage Data Protector Concepts Guide* for more information on logging level and usage recommendations.

Load Balancing: Balancing the Usage of Backup Devices

What Is Load Balancing?

By default, Data Protector automatically balances the usage of backup devices specified for backup. This is also called load balancing, and it ensures equal usage of the devices. When you run backup with the `Load Balancing` option, Data Protector uses devices in the order they are specified in the load balanced backup specification.

NOTE

If you disable the `Load Balancing` option, you have to select the backup device which is used to back up each object in the backup specification. If a device becomes unavailable, then the objects that should be backed up to the device will not be backed up.

See the *HP OpenView Storage Data Protector Concepts Guide* for more information on load balancing.

When to Use Load Balancing

It is recommended that you use the `Load Balancing` option when you want to back up a large number of objects to a number of available devices, and you would like Data Protector to keep all the devices busy all of the time. You should use `Load Balancing` to minimize the impact of unavailable devices on the backup. A device may become unavailable because it:

- failed during a backup
- stopped during a backup
- is in use by another session
- cannot be started at all

When Not to Use Load Balancing

Deselecting the `Load Balancing` option is recommended when

- you want to back up a small number of objects
- objects are backed up on simple devices, such as DDS
- you want to manually select the devices to which objects will be backed up
- you want to know on which medium/media objects will be backed up

How Are the Parameters Used?

The `Load Balancing` option has `MIN` and `MAX` parameters:

`MIN` specifies the minimum number of backup devices out of the list of devices in the backup specification that must be available for the session to start. Available means that they are not used by some other backup session and that you have enough licenses.

`MAX` specifies the maximum number of devices that are used at the same time, even if there are more devices defined in the backup specification. The rest are used if needed.

For example, assume that there are four devices in the backup specification and `MIN` and `MAX` are both configured at two. The backup session will queue until any of those two devices can be used. If any of them fail, one of the two devices in reserve will be used.

How Are Objects Assigned to an Available Device?

The first device from the list of devices is started. The number of selected objects for a device is defined by its `concurrency`. The next device is started and objects are selected until there are no more objects in the list or the maximum number of devices are running.

Objects to be backed up are assigned according to the following criteria:

- Objects that reside on the client connected to the backup device have a higher priority.
- Objects are selected so that the number of Disk Agents per client is kept as low as possible.

The size of objects does not play a role in assigning an object to a device.

If a device becomes unavailable, the following happens:

- All objects backed up to the device before the failure time are actually backed up.
- All objects that are being backed up to the device at failure time are aborted.
- All objects pending to be backed up to the device will be backed up to some other available device specified in the backup specification, if the maximum number of devices has not been used.

Example

For example, assume that there are 100 objects configured for backup to four devices with concurrency set to three and with load balancing parameters `MIN` and `MAX` both configured at two. If at least two devices are available, the session will start with three objects being backed up in parallel to each of the first two available devices. The other 94 objects will be pending and will not be assigned to a particular device at that time.

Once a backup of a particular object is done, the next pending object is started and assigned to the device that has less than three concurrent objects being backed up. Load balancing ensures that the two devices are running in parallel as long as there are still pending objects to be backed up. If a device fails during backup, one of the two devices in reserve is used. The objects that were being backed up to the failed device are

aborted, while the next three pending objects are assigned to the new device. This means that each failure of a device can cause a maximum of three objects to be aborted, provided that other devices are available for the backup session to continue.

The following rules should be considered when applying device options from a template:

- If the load balancing option is not selected in the template, the device options are not used with the backup specification.
- If the load balancing option is selected in both the template and the backup specification, the device options are applied.
- If load balancing is only selected in the template, the device options are applied only if the backup specification has no devices.

For more information on failed backups, refer to “Managing Failed Backups” on page 311.

Ownership: Who Will Be Able to Restore?

Who Is a Backup Session Owner?

A user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the backup session is not considered interactive.

If a modified backup specification is started by a user, the user is the owner unless the following conditions apply:

- The user has the `Switch Session Ownership` user right.
- The backup session owner is explicitly defined in the backup specification, where the username, group or domain name, and the system name are specified. In that case, the backup session owner is the user specified in the backup specification.

If a backup is scheduled on a UNIX Cell Manager, the session owner is `root:sys`, unless the above conditions apply.

If a backup is scheduled on a Windows Cell Manager, the session owner is the user specified at installation time, unless the above conditions apply.

Who Can Restore a Private Object?

The following users can restore a private object:

- Members of the `Admin` and `Operator` user group.

- The backup session owner who has the Start Restore user right. Other user rights may be required, such as Restore to Another Client.
- Users who have the See Private Objects user right.

Why Change the Backup Owner?

Sometimes, you may want to change the backup owner. For example, if the administrator configures and schedules a backup specification, operators are allowed to run it, but they cannot modify or save it. If the Private backup option is set for all objects, the operators are not able to restore anything, but can still manage backups and restart failed sessions.

Changing the owner works only for saved backup specifications. If the backup configuration is changed and not saved, the backup is treated as an interactive backup and the owner is not changed. This could result in a different kind of backup than expected. For example, if you interactively start an incremental backup and you are not the owner of the full backup, you will get another full backup instead of an incremental one.

List of Data Protector Backup Options

This section describes three sets of backup options. The options are ordered alphabetically within each set.

Backup Specification Options

Setting Options for a Backup Specification

1. Select the backup specification whose options you want to set.
2. Click the Options tab.
3. Under Backup Specification Options, click Advanced. The Backup Options window appears.
4. Select the options for General and Clustering. The EMC and StorageWorks tabs appear only if you have the respective devices connected and configured.

   Ignore Clustering if you do not have the MC/Service Guard or the Microsoft Cluster Server installed and configured.
5. Click OK to confirm and exit the Backup Options window. Refer to online Help for details.

Figure 6-28 Backup Specification Options - General

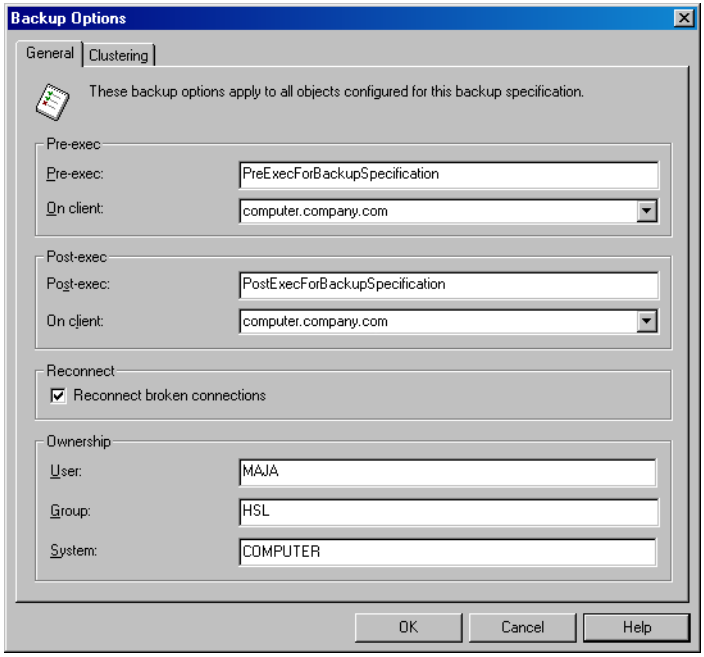

Available Backup Specification Options

Description

You can type in any text to describe the purpose or contents of the backup specification. This text has no effect on the backup session.

Load Balancing

By default, this option is enabled in the `Create New Backup` dialog. If you disabled it there, you can select it later in the `Destination` property page of the backup specification, in the `Backup` tab.

If this option is selected, Data Protector dynamically assigns backup objects to available devices. This means that devices are evenly used, and if one fails, a backup continues on other available devices. If it is not selected, the backup objects are backed up to devices assigned to them in the exact order specified.

The default value is `ON`.

See “Load Balancing: Balancing the Usage of Backup Devices” on page 276 for more information.

Ownership

The session owner is the user who started the interactive backup, unless the owner is specified in the backup specification. Otherwise, the owner is:

- `root` on UNIX Cell Managers
- the user specified at installation time on Windows Cell Managers

The default value is `not specified`.

See “Ownership: Who Will Be Able to Restore?” on page 279 for more information.

You can change the session owner by using the following steps:

1. In the `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context.
2. Double-click `Backup Specifications`, then right-click on the backup specification you want to modify.
3. Choose `Properties`, `Options`, then, under `Backup Specification Options`, choose `Advanced`. Choose the `General` tab.
4. Modify session ownership as necessary. Use uppercase on Windows systems.

NOTE Make sure to specify the information as it was specified when the user was configured.

Pre-Exec

The command specified in this field is run on a specified client before any object is backed up. If the client is not defined, the command is run on the Cell Manager.

See “Pre- and Post- Exec Commands on Windows Systems” on page 298 for details of specifying pre-exec commands on Windows.

See “Pre- and Post- Exec Commands on UNIX Systems” on page 304 for details of specifying pre-exec commands on UNIX.

See “Examples of Pre-Exec and Post-Exec Commands for UNIX” on page A-21 for some sample scripts on UNIX.

The default value is `not specified`.

Post-Exec

The command specified in this field is run on a specified client after all objects have been backed up. If the client is not defined, the command runs on the Cell Manager.

See “Pre- and Post- Exec Commands on Windows Systems” on page 298 for details of specifying post-exec commands on Windows.

See “Pre- and Post- Exec Commands on UNIX Systems” on page 304 for details of specifying pre-exec commands on UNIX.

See “Examples of Pre-Exec and Post-Exec Commands for UNIX” on page A-21 for some sample scripts on UNIX.

Default value is `not specified`.

Reconnect Broken Connections

When this option is set, Data Protector reconnects the

- Backup Session Manager and Disk Agents or Media Agents (control connections) or
- Disk Agent and Media Agent during backup, and Media Agents if object mirroring is enabled (data connections)

in the event of short-term network problems. Otherwise, the session is aborted.

This setting is useful if you have, for example, the Cell Manager on one LAN and Disk Agents or Media Agents on another. Assuming that the connection between these two LANs is unreliable (WAN connections), Data Protector tries to reconnect for 600 seconds by default. This time-out period can be set in the `omnirc` variable `OB2RECONNECT_RETRY`.

The default value is `OFF`.

Object Options

Setting the Filesystem Options

1. Select the backup specification and from the `Options` property page, under `Filesystem Options`, click `Advanced`.
2. Select the options to be set from the `Options`, `Other`, `WinFS Options`, or `Netware Options` tabs.

NOTE

On the `Options` tab, if specifying Pre- and/or Post- exec command names, you may or may not have to specify the full paths for the commands.

See “Pre- and Post- Exec Commands on Windows Systems” on page 298 for details of specifying pre-exec commands on Windows.

See “Pre- and Post- Exec Commands on UNIX Systems” on page 304 for details of specifying pre-exec commands on UNIX.

3. Click `OK` to confirm and exit this dialog box.

   See online Help for specific help on each option.

Figure 6-29 Filesystem Options

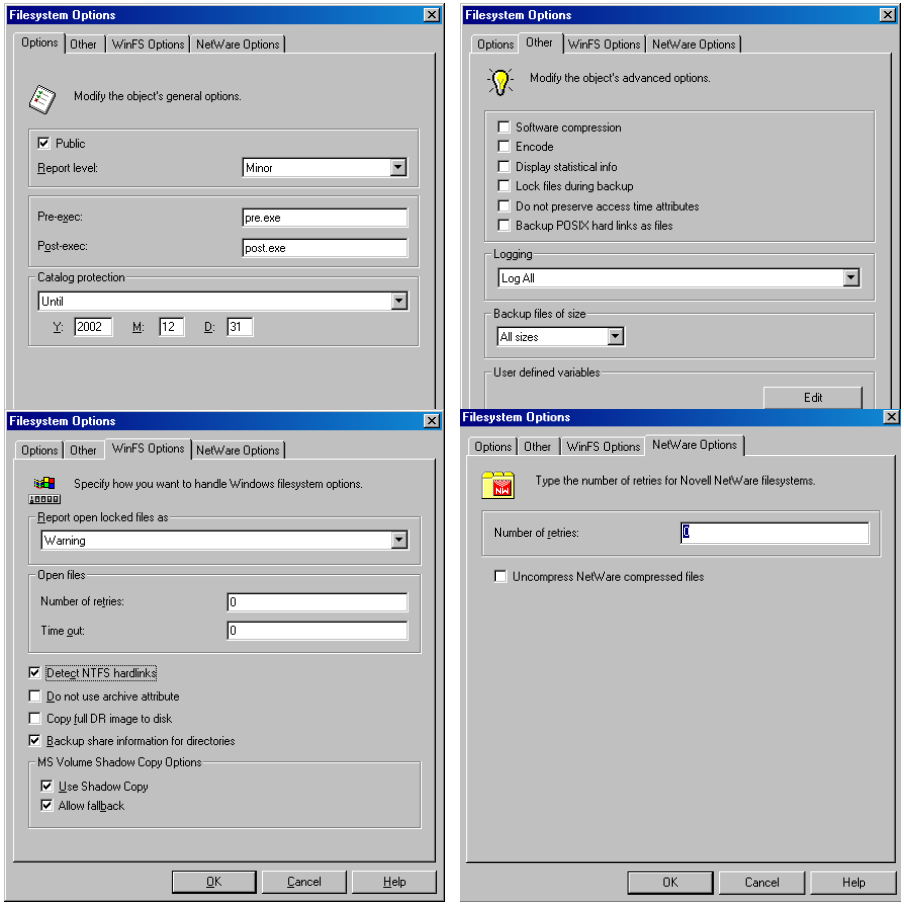

Setting the Disk Image Options

1. Select the backup specification.
2. Select the `Options` property page.
3. Under `Disk Image Options`, click `Advanced`.
4. Click either the `Options` or the `Other` tab, and specify the options as desired. For a description of each option, click `Help` in the dialog box.

NOTE

On the Options tab, if specifying Pre- and/or Post- exec command names, you may or may not have to specify the full paths for the commands.

See “Pre- and Post- Exec Commands on Windows Systems” on page 298 for details of specifying pre-exec commands on Windows.

See “Pre- and Post- Exec Commands on UNIX Systems” on page 304 for details of specifying pre-exec commands on UNIX.

Figure 6-30 Disk Image Options

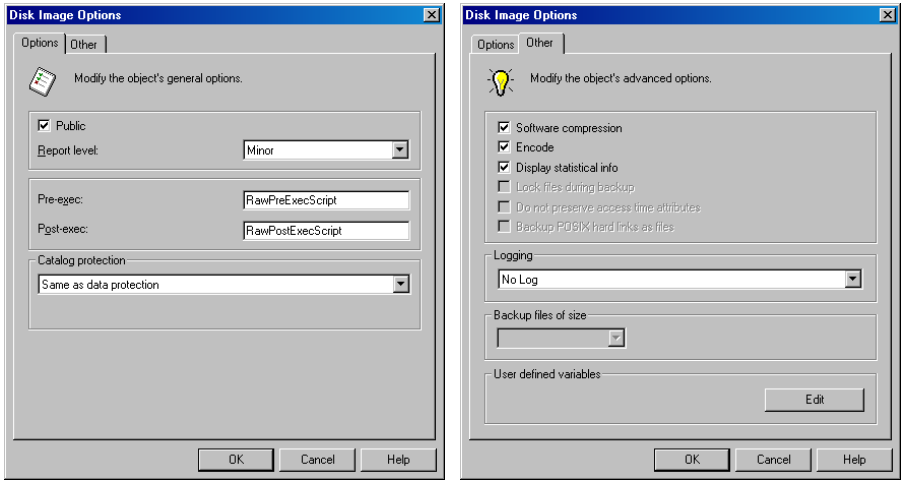

5. Click OK to confirm and exit the dialog box.

Setting the Object Specific Options

1. Select the backup specification whose options you want to set.
2. Select the Backup Object Summary property page.
3. Right-click the backup object, then select Properties. The contents of the Object Properties window depend on the type of backup object you selected. This can be a UNIX filesystem, a Windows filesystem, or a UNIX disk image.

The Object Properties window for a **Windows filesystem** contains the General, Options, Other, Trees/Filters, and the WinFS Options tabs. Options, Other, and WinFS Options are the same as shown in Figure 6-29, while General and Trees/Filters are shown in Figure 6-31.

Figure 6-31 Object Properties - General and Trees/Filters

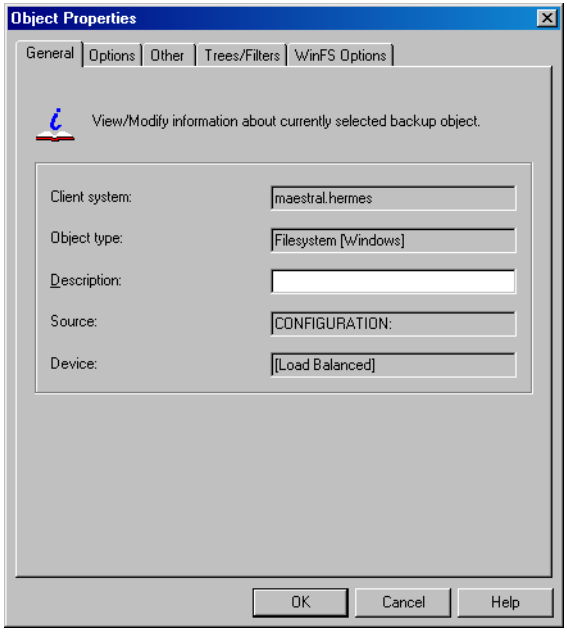

The Object Properties window for a **UNIX filesystem** contains the General, Options, Other, and the Trees/Filters tabs. Options and Other are the same as shown in Figure 6-29, while General and Trees/Filters are the same as in Figure 6-31, except that Object type is described as Filesystem [UNIX].

The Object Properties window for a **disk image** object contains the General, Options, Other, and the Disk Image Options tabs. Options and Other are the same as for the backup specification shown in Figure 6-30, while General and Disk Image Options are shown in Figure 6-32.

Figure 6-32 **Object Properties - General and Disk Image Options**

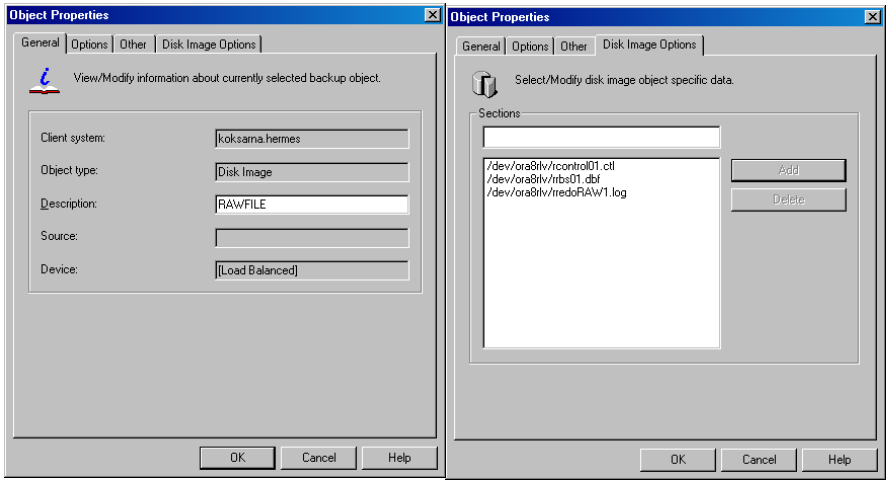

4. Set the options and click `OK` to confirm the selection. See below for details on a particular option.

Available Object Options

Allow Fallback (Windows-Specific Options)

If the `Use Shadow Copy` option is specified, but the shadow copy creation on the system where the VSS filesystem backup is running fails, the backup also fails by default. However, you can avoid backup failure by specifying the `Fallback to legacy` option. In this case, the backup will continue as a non-VSS backup.

Backup Files of Size

Use this option to specify the size of backed up files. You can back up `All Files` (default), `Files Bigger Than`, `Smaller Than`, or files within a specified size range in kilobytes.

Backup POSIX Hard Links as Files

This option is only relevant for UNIX filesystems.

A hard link is a directory entry that actually points to a physical file. If this option is not set, Data Protector traverses the directory trees twice. In the first traverse,

a table of all hard links that point to the same files is made. In the second traverse, only one hard link is backed up with the file contents, while all the others are backed up as hard links. The first traverse also allows Data Protector to estimate the size of the backup.

If set, Data Protector backs up the entire file contents for each hard link. Data Protector traverses the filesystem tree only once, thus significantly speeding up the backup process.

Use this option when there are no hard links in your directory. When this option is set, Data Protector cannot estimate the size of the backup or display the percentage of the backup finished.

The default value is `OFF`.

Catalog Protection

The default `Catalog Protection` value is **Same as data protection**. It can be changed by specifying the **None**, **Until**, **Days**, and **Weeks** values.

Refer to “Catalog Protection: How Long Info Is Kept in the Database” on page 274 for more information.

Do Not Preserve Access Time Attributes

When this option is not set, the access time attributes remain as they were before the backup: they are reset to their original values after each file is backed up. When this option is set, the access time values are set to the moment of backup.

See also “Backing Up UNIX Systems” on page 206.

The default value is `OFF`.

This option is not supported on Novell NetWare.

Do Not Use Archive Attribute (Windows-Specific Option)

Data Protector clears the archive attribute after each backup (after the file has been read). If you have other applications that make specific use of this attribute, you should use this option.

The default value is OFF.

Detect NTFS Hardlinks (Windows-Specific Option)

This option is similar to Backup hardlinks as files except that it is only valid for NTFS and the default value is OFF, meaning that hardlinks are backed up as ordinary files. The default value is OFF because the NTFS hardlinks are not often used and setting this option decreases backup performance.

Encode

Open Systems and public networking make data security in large enterprises essential. Data Protector lets you encode file and disk image data so that it becomes unreadable. Data is encoded before it is transferred over the network and written to the media. Data Protector uses a fixed, built-in algorithm for this purpose.

The default value is OFF.

Lock Files During Backup

If set, files are locked while being backed up, preventing them from being modified during the backup. Mandatory locking is used.

The default value is OFF.

This option is not supported on Novell NetWare.

Logging

The default logging level is **Log All**. It can be changed to **No Log**, **Log Directories**, or **Log Files**.

For more information on each logging level, see “Logging: Changing Details About Data Stored in the Database” on page 275.

Number of Retries (Novell NetWare Specific)

This option defines the number of times Data Protector attempts to back up a file. If a backup cannot be made within this number of retries, Data Protector issues an

error message. If you use applications that open and release files, you can use this option to increase the probability that the files are backed up.

The default value is 1.

Open Files (Windows-Specific Option)

This option controls what Data Protector does when it encounters open Windows files. If the **Number of retries** value is specified, this number defines how many times Data Protector tries to back up an open or busy file. The **Time out** value is the amount of time in seconds during which Data Protector waits before retrying to back up an open or busy file.

Protection (Data Protection)

This option enables you to set the protection level for backed up data. In this way, you prevent the backup media from being overwritten for the specified period. The `Protection` values are **None**, **Until**, **Days**, **Weeks**, and **Permanent**.

The default value is **Permanent**.

Public/Private

This option lets you set the access rights for restoring data that you back up. If a filesystem is backed up with the **Private** setting, it can be restored only by you or users who are part of the Data Protector `Admin` group.

Setting the value to **Public** lets anyone with the `Start Restore` user right restore the data.

The default value for filesystem backups is **Private**, and **Public** for integration backups.

Report Level

This option defines the level of errors that are reported for an object during a backup session. Setting a level means errors of this level and higher are reported. You can choose from **Warning**, **Minor**, **Major**, and **Critical** report level.

For example, when the value `Minor` is set, only errors graded as **Minor**, **Major**, and **Critical** are reported in the `Messages` field. Messages keyed as **Normal** always appear in the `Messages` field. The default value is **Warning**.

NOTE The number of messages per backup system stored in the IDB is limited to 3000.

Report Open Locked Files As (Windows-Specific Option)

This option sets the report level for files that are opened and locked at the time Data Protector attempts to back them up. Data Protector reports such files as per the regard to the **Report Level** setting. The default value is **Warning**.

Software Compression

Data Protector can compress data on a Disk Agent client before sending it to a Media Agent client. This feature is known as software compression. Select `Software compression` in the `Other` property page of the `Object Properties` window to enable software compression. In this way, you reduce traffic over the network, as well as the number of media needed for backup, thus improving overall backup performance. Depending on the data type, compression ranges from 30% to 70% and is based on the Lempel-Ziv 4.3 compression algorithm, which is compatible with the standard UNIX `compress` utility. Note that the progress indication on the monitor is not accurate if this option is used.

The default value is `OFF`.

This option is not supported on Novell NetWare. However, it is possible to uncompress files that were compressed with this option using older versions of Data Protector.

NOTE

Most modern backup devices provide built-in hardware compression that can be set when you create a device file or SCSI address in the device configuration procedure. Do not use software and hardware compression at the same time, since double compression decreases performance without giving better compression results. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details on how to enable hardware compression.

HP Ultrium LTO devices do not let you disable automatic hardware compression. Do not set the Keep the default software compression when you configure an HP Ultrium LTO drive with Data Protector.

Uncompress NetWare Compressed Files (Novell Netware Specific Option)

By default, Data Protector backs up Novell NetWare compressed files in their compressed format. Though this approach speeds up the backup process, it makes it impossible to restore the Novell NetWare compressed files to a non-compressed Novell NetWare volume. When this option is selected, Novell NetWare compressed files are uncompressed before being backed up. Files backed up in this form can be restored to a non-compressed Novell NetWare volume.

Use Shadow Copy (Windows-Specific Option)

When performing filesystem backup on Windows Server 2003 systems, Data Protector uses MS Volume Shadow Copy service (VSS) for coordinating the point-in-time backup. VSS allows you to create shadow copy backups of volumes and exact point-in-time copies of files, including all open files. This means that the VSS mechanism commits all pending I/O operations and holds incoming writing requests during the preparation of a shadow copy volume. In this way all files on the filesystem are closed and unlocked during the shadow copy creation.

User Defined Backup Variables

Set user defined backup variables (a variable name and its value) to enable flexible operation on some platforms and integrations with Data Protector. For detailed steps, refer to the online Help index keyword “setting user definable backup variables”.

The list of variables and their values that are configurable with Data Protector is dynamic and comes with Data Protector patches.

Figure 6-33 User Defined Variables

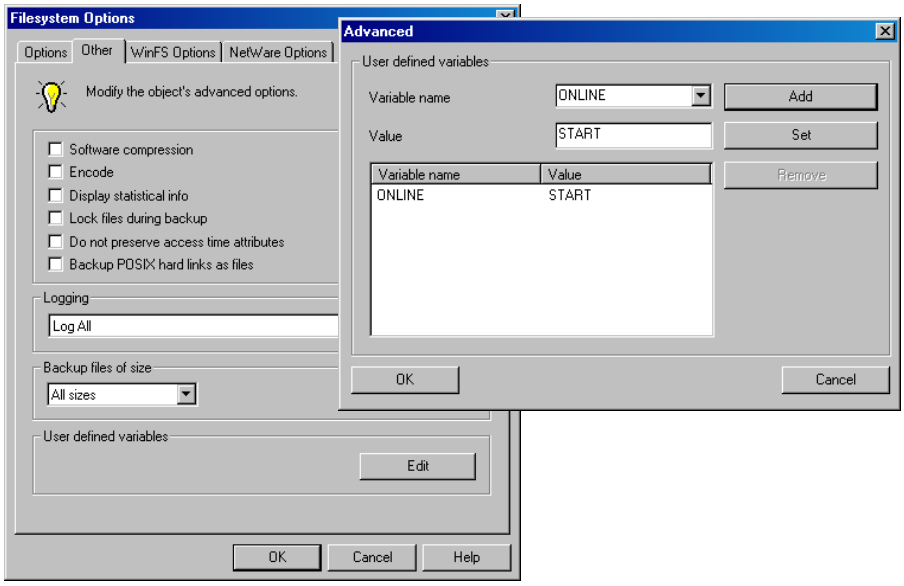

Device Backup Options

You can select the options listed below for each backup device in use. None of the settings are needed, because `CRC Check`, `Concurrency`, and `Media Pool` use the default values that are set when the device is configured. The `Prealloc List` value is specified along with the media pool settings.

Available Device Backup Options

CRC Check

Set this option to have Data Protector calculate the CRC (Cyclic Redundancy Check) when a backup runs. CRC is an enhanced `checksum` function that lets you later confirm using the `Verify` option whether or not data has been written correctly to the medium.

This option can be specified for backup and object copy operations.

The default value is `OFF`.

Concurrency

Concurrency allows more than one Disk Agent to write to one backup device. Data Protector can then keep the devices streaming if data can be accepted faster than a Disk Agent can send it. The maximum concurrency value is 32.

Data Protector provides default values for all supported devices.

This option can be specified for backup and object copy operations.

Media Pool

This option selects the media pool with the media you will use for a backup. If not defined, a default pool, which is a part of device specification, is used.

This option can be specified for backup and object copy operations.

Prealloc List

The **Prealloc List** is a subset of media in the media pool used for a backup. It specifies the order in which the media will be used. When using the **Prealloc List** and the `Strict` media allocation policy with the backup device, Data Protector expects the sequence of the media in the device to correspond with that specified in the **Prealloc List**. If the media are not available in this sequence, Data Protector issues a mount request. If no media are specified in this list, then the Data Protector allocation procedure is used to allocate media.

This option can be specified for backup and object copy operations.

Use preferred MultiPath host

This option is available only for multipath devices. To set a preferred host, select this option and select the host from the drop down list. During a backup session, Data Protector will try to use this host first, regardless of the predefined order.

Pre- and Post-Exec Commands

Before a backup or restore session begins, an additional action is sometimes necessary. For example, you may want to check the number of files to back up, stop some transaction processing, or shut down a database. Such actions are performed using `pre-` and `post-exec` commands. `Pre-` and `post-exec` commands are not supplied by Data Protector. Depending on your needs, you have to write your own executables to perform the required actions.

IMPORTANT

`Pre-` and `post-exec` commands are potentially dangerous because they enable numerous possible exploits if they are used by unauthorized personnel. For security considerations, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

For backup, `pre-` and `post-exec` commands can be configured on two levels:

Backup Specification

The `pre-exec` command is executed before the backup session starts. The `post-exec` command is executed when the backup session stops. You specify these commands as backup options for the entire backup specification. By default, `pre-` and `post-exec` commands for the session are executed on the Cell Manager, but you can choose another system.

Specific Backup Object

The `pre-exec` command for a specific backup object starts before the object is backed up. The `post-exec` command for the backup object is executed after the object is backed up. You specify these commands as backup options that apply for all objects, or for individual objects. `Pre-` and `post-exec` commands for the object are executed on the system where the Disk Agent that backs up the object is running.

`Pre-` and `post-exec` commands are run in the following order:

1. The `pre-exec` command for the entire backup specification starts and completes.
2. For each object in the backup specification:
   a. The `pre-exec` starts and completes.
   b. The object is backed up.
   c. The `post-exec` (for each object in the backup specification) starts and completes.
3. The `post-exec` command for the entire backup specification starts and completes.

Figure 6-34 Pre- and Post-Exec Control Flow

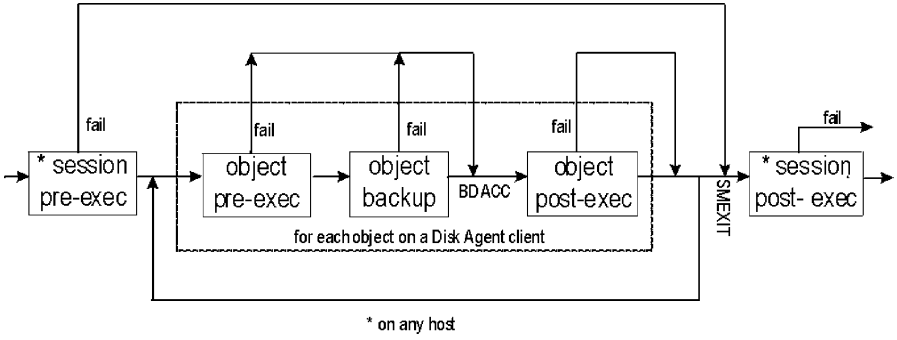

Pre- and Post- Exec Commands on Windows Systems

This section describes how to implement `pre-` and `post-exec` commands on Windows Cell Managers and clients.

How to Write the Commands

Pre- and post-exec commands can be written as executables or batch files. The supported extensions for pre- and post-exec commands on Windows systems are `.bat`, `.exe`, and `.cmd`. To run a pre-/post-exec script with an unsupported extension (for example `.vbs`), create a batch file (`.bat`) that starts the script. Then configure Data Protector to run the batch file as a pre-/post-exec command which then starts the script with the unsupported extension.

All the commands that run within the batch file must return an exit code 0 to signify success or greater than 0 to signify a failure.

Carefully follow the implementation guidelines provided in this section.

Pre- and Post-Exec Commands for a Backup Specification

`Pre-` and `post-exec` commands for a backup session are started before and after the session. These commands are usually executed on the Cell Manager, but you can choose another system.

Where to Locate the Commands

`Pre-` and `post-exec` scripts are started by the Data Protector CRS when executed on the Cell Manager; and under the Data Protector Inet Service account (by default, Local System account) when executed remotely.

On the Cell Manager, the scripts can be located in any directory. On the systems other than the Cell Manager, the scripts must be located in the *<Data_Protector_home>*\bin directory.

For the scripts located in the *<Data_Protector_home>*\bin directory, specify only the filename, otherwise, specify the full pathname of the script.

How to Specify the Filename or Pathname

In the backup specification, click the `Options` tab. Under `Backup Specification Options`, click `Advanced`. Write the filename or pathname in the `Pre-exec` and/or `Post-exec` text box.

When entering a full pathname, if your directory names are longer than 8 characters, write the pathname either in quotes or in the short 8.3 MS-DOS compatible form.

IMPORTANT

If you use quotes ("") to specify a pathname, do not use the combination of backslash and quotes (\"). If you need to use a trailing backslash at the end of the pathname, use the double backslash (\\).

Environment Variables

The following environment variables are set by Data Protector upon starting a backup session, and can be used only in `pre-` and `post-exec` scripts for a backup specification on the Cell Manager:

DATALIST	The name of a backup specification.
MODE	Backup operation type, such as full, incremental, incremental1, incremental2.
OWNER	Owner of the session.

The contents of this variable are in the same format as in the database (case-sensitive):
<user>.<group>@<hostname> for UNIX
<DOMAIN>\<user>@<hostname> for Windows

PREVIEW The value is 1 if a preview is running and 0 if a backup is running. Use this variable to modify your commands so that they are executed only during a backup and not during a preview. By default, the `pre-` or `post-exec` commands are not executed for preview. You can enable them by setting the global option `ExecScriptOnPreview`.

RESTARTED Set to 1 if this is a restarted backup session, otherwise set to 0. The `post-exec` can use this variable to prevent an additional restart in the case that **SMEXIT** equals 0.

SESSIONID Is used to identify a finished session and is recorded in the database. You cannot use this to preview a session (use SESSIONKEY).

SESSIONKEY Is used to identify a running session. You may, for example, abort a backup session before it is started if something is wrong.

SMEXIT The exit code of the Session Manager is the same as the exit code of the `omnib` command. You can only use this variable with the `post-exec` command. Agents can refer to Disk Agents, Media Agents, Application Agents, Symmetrix Agents, and so on.

Table 6-3 SMEXIT VALUES

Value	Description
0	All files were successfully backed up.
10	All agents completed successfully, but not all files were saved.
11	One or more agents failed or there was a database error.
12	None of the agents completed the operation; session was aborted by Data Protector.

Table 6-3 SMEXIT VALUES

Value	Description
13	Session was aborted by a user.

Key Points

- ✓ The `pre-` and `post-exec` commands for a backup specification have to be executables or batch files. It is important to specify a filename extension on Windows.
- ✓ The `pre-` and `post-exec` commands can be located in any directory on the Cell Manager, or on any other system where the Disk Agent is running in the *<Data_Protector_home>*`\bin` directory. If they are located in a directory other than *<Data_Protector_home>*`\bin` the full pathname must be specified.
- ✓ The execution of `pre-` and `post-exec` commands is implemented using the Windows pipe mechanism. All processes started in the `pre-` or `post-exec` functions must finish before processing continues.
- ✓ A `pre-` or `post-exec` command must return a non-negative value upon successful completion.
- ✓ If a `pre-exec` command fails (returns a value less than 0), the status of the backup session is set to `Failed` and the session is aborted. A `post-exec` command is not executed.
- ✓ If a `post-exec` command fails (returns a value less than 0), the backup session status is set to `Completed with errors`.
- ✓ `Post-exec` command is always executed, unless the session is aborted and the `pre-exec` command is *not* executed or not set. If the `OB2FORCEPOSTEXEC` `omnirc` variable is set, the `post-exec` command is always executed.
- ✓ The `pre-` and `post-exec` commands for a backup specification are by default NOT executed during a preview of the backup. This behavior is defined by the `ExecScriptOnPreview` variable in the global options file. See “Global Options File” on page 613 for details on how to modify these values.
- ✓ `Pre-` and `post-exec` commands are handled in the same way as commands entered at the DOS prompt. Therefore, special characters, such as the pipe (|) and the redirect symbols (>, <) are not allowed.
- ✓ While `pre-` or `post-exec` commands are running, the backup session cannot be aborted.

- ✓ The pre- and post-exec commands run in the background mode. Therefore, do not use any commands that require user interaction.
- ✓ Standard output of the pre- and post-exec commands is written to the IDB as messages and shown on the monitor screen of the Data Protector GUI.
- ✓ You can disable a session's pre- and post-exec command execution on the Cell Manager by setting SmDisableScript global option to 1.
- ✓ You can disable remote session pre- and post-exec command execution on any client by adding OB2REXECOFF=1 into the omnirc file on the specific client.
- ✓ You can secure the client by specifying which Cell Managers are allowed to access the client. Only permitted Cell Managers will be able to execute pre- and post-exec commands on the client. For more information on securing a client, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Pre- and Post-Exec Commands for a Specific Backup Object

Pre- and post-exec commands for an object are executed before and after the backup of the object, respectively. You can specify these commands for all objects in a backup specification, or for each individual object. When backing up integrations, for example Oracle, the database is considered as an object, so the commands are executed before and after the database backup. These commands are executed on the system where the Disk Agent is running.

Where to Locate the Commands

Pre- and post-exec scripts for a backup object are started under the Data Protector Inet Service account (by default, Local System account).

The exec scripts for backup objects can reside in any directory on the system where the Disk Agent is running. However, for client backups, they must reside in *<Data_Protector_home>*\bin. If the scripts are located in the *<Data_Protector_home>*\bin, specify only the filename, otherwise the full pathname must be specified.

How to Specify the Filename or Pathname

To apply pre- and post-exec commands to all objects in the backup specification, click the Options tab in the backup specification. Under Filesystem Options (Disk Image Options in a saved backup specification for disk image backup), click Advanced.

To apply pre- and post-exec commands to individual objects only, click the Backup Object Summary tab in the backup specification. Right-click an object and click Properties. In the Object Properties dialog box, click the Options tab.

To apply pre- and post-exec commands to an integration object, click the Options tab in the backup specification. Under Application Specific Options, click Advanced.

Write the filename or pathname in the Pre-exec and/or Post-exec text box.

When entering a full pathname, if your directory names are longer than 8 characters, write the pathname either in quotes or in the short 8.3 MS-DOS compatible form.

Environment Variables

BDACC

The Disk Agent sets its exit code (0 is successful) to the **BDACC** environment variable. This variable can be checked in the post- exec command, thus making the post-exec command dependent upon successful termination of the Disk Agent.

NOTE

If you perform a host backup, the pre-exec script is started once, before the first filesystem backup for the particular system, while the post-exec script is started after the backup. In this case, **BDACC** cannot be exported because the variable is related to a single filesystem object, not to a whole client.

Key Points

- ✓ The pre- and post-exec commands for a backup object have to be executable or batch files. It is important to specify the filename extension on Windows.
- ✓ The pre- and post-exec commands can be located in any directory on the system where the Disk Agent is running except for client backups. If they are located in a directory other than *<Data_Protector_home>*\bin the full pathname must be specified.
- ✓ If a pre-exec command fails (returns a non-zero value), the backup of this object is aborted. The status of the object is set to Aborted and the backup Disk Agent stops processing but the post-exec command is executed (unless the post-exec command is dependant on the BDACC environment variable). No backup of the object exists.

- ✓ If a post-exec command fails (returns a non-zero value), the backup object status is set to Aborted. The backup of the object exists and data can be restored.
- ✓ The pre- and post-exec commands are handled in the same way as commands entered at the DOS prompt. Therefore, special batch characters such as the pipe (|) and the redirect symbols (>, <) are not allowed.
- ✓ While pre- or post-exec commands are running, the backup session cannot be aborted.
- ✓ The pre- and post-exec processes run in the background mode. Therefore, do not use any commands that require user interaction.
- ✓ Standard output of the pre- and post-exec commands is written to the IDB as messages and shown on the monitor screen of the Data Protector GUI.
- ✓ The pre- and post-exec scripts have to send some output at least every 15 minutes by default, or the sessions waiting for the scripts are aborted. You can change this time interval by modifying the ScriptOutputTimeout variable in the global options file.
- ✓ Time-out is provided. If no message is received within the specified time-out in seconds, the session is aborted.
- ✓ You can disable a pre- and post-exec script by adding the line OB2OEXECOFF=1 in the omnirc file on any client.
- ✓ You can secure the client by specifying which Cell Managers are allowed to access the client. Only permitted Cell Managers will be able to execute pre- and post-exec commands on the client. For more information on securing a client, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Pre- and Post- Exec Commands on UNIX Systems

This section describes how to implement pre- and post-exec commands on UNIX Cell Managers and clients.

How to Write the Commands

Pre- and post-exec commands can be written as shell scripts.

See “Examples of Pre-Exec and Post-Exec Commands for UNIX” on page A-21.

Pre- and Post-Exec Commands for a Backup Specification

`Pre-` and `post-exec` commands for a backup session are started before and after the backup session, respectively. These commands are usually executed on the Cell Manager, but you can choose another system as well.

Where to Locate the Commands

`Pre-` and `post-exec` commands for backup specifications on UNIX systems are started by the backup session owner, unless the backup session owner has the `Back up as root` permission and the commands are then started under root.

On the Cell Manager, the `exec` commands for backup specifications can reside in any directory.

On a remote UNIX client, the `exec` commands for backup specifications must be located as follows:

- Solaris 7/8/9, HP-UX: `/opt/omni/lbin`
- Solaris 2.6, other UNIX systems: `/usr/omni/bin`

For the commands located in the `/opt/omni/lbin` or in the `/usr/omni/bin` directory, specify only the filename, otherwise, specify the full pathname.

How to Specify the Filename or Pathname?

For information on how to specify the commands, refer to the online Help index keyword “pre- and post-exec commands for backup specifications”.

Environment Variables

The following environment variables are exported upon starting a backup session, and can be used in `pre-` and `post-exec` scripts for a backup specification session on any host:

DATALIST The name of the backup specification.

MODE Backup operation type, such as full, incremental, incremental1, and so on.

OWNER Owner of the session.

The content of this variable is in the same format as in the database (case-sensitive):
<user>.<group>@<hostname> for UNIX and
<DOMAIN>\<user>@<hostname> for Windows.

PREVIEW Set to 1, if the preview is running. Set to 0, if a backup is running. Use this variable to modify your commands so that they are executed only during a backup and not during a preview. By default, `pre-` and `post-exec` commands are not executed for preview. You can enable this with global option `ExecScriptOnPreview`.

RESTARTED Set to 1 if this is a restarted Backup session, otherwise set to 0. The `post-exec` can use this variable to prevent an additional restart if **SMEXIT** equals 0.

SESSIONID Is used to identify a finished session and is recorded in the database. You cannot use this to preview a session (use SESSIONKEY).

SESSIONKEY Is used to identify a running session. You may, for example, abort a backup session before it is started if something is wrong.

SMEXIT The exit code of the Session Manager is the same as the exit code of the `omnib` command. You can only use this variable with the `post-exec` command. Agents can refer to Disk Agents, Media Agents, Application Agents, and Symmetrix Agents. Refer to Table 6-3 on page 300 for details on SMEXIT values.

Key Points

Check the following before configuring `pre-` and `post- exec` commands for a backup specification on a local or remote host:

- ✓ If a `pre-exec` command fails (returns a non-zero value), the backup status of the session is set to `Failed` and the session is aborted. A `post-exec` command is not executed.
- ✓ If a `post-exec` command fails (returns a non-zero value), the backup of the session is set to `Completed with errors`.
- ✓ `Post-exec` command is always executed, unless the session is aborted and the `pre-exec` command is *not* executed or not set. If the `OB2FORCEPOSTEXEC` `omnirc` variable is set, the `post-exec` command is always executed.

- ✓ The pre- and post-exec commands for a backup specification are by default NOT executed during a preview of a backup. This behavior is defined by the ExecScriptOnPreview variable in the global options file. See "Global Options File" on page 613 for details.
- ✓ While the pre- or post-exec commands are running, the backup session cannot be aborted.
- ✓ The pre- and post-exec processes operate in the background mode. Therefore, do not use any interactive commands for pre- and post-exec processing.
- ✓ The pre- and post-exec scripts have to send some output at least every 15 minutes by default, or the sessions waiting for the scripts are aborted. You can change this time interval by modifying the ScriptOutputTimeout variable in the global options file.
- ✓ Time-out is provided. If no message is received within the specified time-out in seconds, the session is aborted.
- ✓ If there is no executable script on the host or if the path of the script is wrong, Data Protector displays an error message that the script failed and the session is aborted.
- ✓ If a command writes any text to stdout, this text is sent to the Session Manager and written to the database. A stderr is redirected to /dev/null. You can redirect it to stdout to get error messages logged to the database.

NOTE

A pre- or post-exec script may hang because it did not close all file descriptors before forking a new process. If the new process runs in the background and does not exit, such as, for example, the database server process (dbstart), the scripts hang. You can use the detach command. The source of the detach command is provided in the detach.c file, but is officially unsupported. For example:

```
/opt/omni/lbin/utilns/detach pre_script [arguments...]
```

- You can disable a session's pre- and post-exec command execution on the Cell Manager by setting the SmDisableScript global option to 1.

- You can disable remote session `pre-` and `post-exec` command execution on any client by adding `OB2REXECOFF=1` into the `omnirc` file on the specific client.
- You can secure the client by specifying which Cell Managers are allowed to access the client. Only permitted Cell Managers will be able to execute `pre-` and `post-exec` commands on the client. For more information on securing a client, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Pre- and Post-Exec Commands for a Specific Backup Object

`Pre-` and `post-exec` commands for an object are executed before and after the backup of the object, respectively. You can specify these commands for all objects in a backup specification, or for each individual object. When backing up integrations, for example Oracle, the database is considered as an object, so the commands are executed before and after the database backup. These commands are executed on the system where the Disk Agent is running.

Where to Locate the Commands

`Pre-` and `post-exec` commands for backup objects on UNIX systems are started by the backup session owner, unless the backup session owner has the `Back up as root` permission and the commands are then started under root.

The `exec` commands for backup objects can reside in any directory on the system where the Disk Agent is running. However, for client backups, they must reside in `/opt/omni/lbin` on HP-UX or in `/usr/omni/bin` on other UNIX systems. If the commands are located in the `/opt/omni/lbin` or in `/usr/omni/bin` directory, specify only the filename, otherwise the full pathname must be specified.

How to Specify the Filename or Pathname

For information on how to specify the commands, refer to the online Help index keyword “pre- and post-exec commands for backup objects”.

Environment Variables

The following environment variables are exported upon starting a backup session, and can be used in the `pre-` and `post-exec` scripts for an object on the system where the Disk Agent is running:

BDACC

The Disk Agent sets its exit code (0 is successful) to the BDACC environment variable. This variable can be checked in the `post-exec` script, thus making the `post-exec` command dependent on the successful termination of the Disk Agent.

NOTE

If you perform a host backup, the `pre-exec` script is started once, before the first filesystem backup for the particular system, while the `post-exec` script is started after the backup. In this case, BDCACC cannot be exported because the variable is related to a single filesystem object, not to a whole client.

Key Points

Check the following key points before configuring the `pre-` and `post-exec` commands:

- ✓ The `pre-` and `post-exec` commands for an object are executed during the preview of a backup. Therefore, you may want to preview your backup first and then add the `pre-` and `post-exec` commands, or check the PREVIEW environment variable in your scripts.
- ✓ If a `pre-exec` command for an object fails (returns a non-zero value), the backup status of the object is set to `Aborted` and the Disk Agent stops processing but the `post-exec` command is executed (unless the `post-exec` command is dependant on the `BDACC` environment variable). No backup of the object exists.
- ✓ If a `post-exec` command fails (returns a non-zero value), the backup status of the object is set to `Aborted`. A backup of the object exists and data can be restored.
- ✓ The `pre-` and `post-exec` commands should send some output to the Disk Agent at least every 120 minutes by default, or the backup of the object is aborted. This time period can be changed by modifying the `SmDaIdleTimeout` variable in the global options file.
- ✓ `Pre-` and `post-exec` commands are handled in the same way as commands entered at the shell prompt. Special shell characters, such as the pipe (`|`) and the redirect symbols (`>`,`<`) are not allowed.
- ✓ While the `pre-` and `post-exec` commands are running, the backup session cannot be aborted.

- ✓ The `pre-` and `post-exec` processes operate in background mode. Therefore, do not use any interactive commands for the `pre-` and `post-exec` processing.
- ✓ If a command writes any text to `stdout`, this text is received by the Disk Agent, sent to the Session Manager, and written to the database. A `stderr` is redirected to `/dev/null`. You can redirect it to `stdout` to get error messages logged to the database.
- ✓ The `pre-` and `post-exec` commands for an object have to be located on the client where the Disk Agent is running.
- ✓ The `pre-` and `post-exec` commands must be executable and specified with the full pathname.
- ✓ You can disable `pre-` and `post-exec` scripts by adding the line `OB2OEXECOFF=1` into the `omnirc` file on any client.
- ✓ You can secure the client by specifying which Cell Managers are allowed to access the client. Only permitted Cell Managers will be able to execute `pre-` and `post-exec` commands on the client. For more information on securing a client, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Managing Failed Backups

During a backup, some systems may not be available because they were shut down, there were some networking problems, or similar occurrences. This results in some systems not being backed up entirely.

Setup Notification

Data Protector lets you configure a notification so that you are informed about unexpected events, such as a mount request or a device error during a backup session. You can choose among the methods that most suit your needs, for example e-mail, or a broadcast message to your Windows display.

See Chapter 9, “Monitoring, Reporting, Notifications, and the Event Log,” on page 379 for details.

Checking Failed Backups

One of the most important aspects of managing backups is the regular checking of the backup status. Data Protector provides a comprehensive reporting functionality that allows you to view reports on the backup status. See “Monitoring Sessions” on page 381 for details on the reporting functionality.

Warnings When Backing Up System Disks

Data Protector issues warnings when backing up the system disk on Windows systems. This is because certain files on the system disk are always busy and cannot be opened by any application, including the Disk Agent. The contents of these files can only be backed up as a part of `CONFIGURATION`.

When these files are accessed by a filesystem backup, such as when the whole system disk is backed up, Data Protector fails to open them and reports warnings or errors, depending on the backup options. See “Using Backup Options” on page 269.

While this behavior is correct from the filesystem backup point of view, it can create a manageability problem. Due to the large number of warnings that are always reported, it is likely that a failure of another file may be overlooked.

These specific files can only be backed up through a `CONFIGURATION` backup. Knowing this, you can exclude them from a filesystem backup to avoid warnings.

The following example is a list of files that cannot be opened on an active Windows system with the Windows software installed on the C: drive:

<%SystemRoot%>\system32\config\default

<%SystemRoot%>\system32\config\default.LOG

<%SystemRoot%>\system32\config\SAM

<%SystemRoot%>\system32\config\SAM.LOG

<%SystemRoot%>\system32\config\SECURITY

<%SystemRoot%>\system32\config\SECURITY.LOG

<%SystemRoot%>\system32\config\software

<%SystemRoot%>\system32\config\software.LOG

<%SystemRoot%>\system32\config\system

<%SystemRoot%>\system32\config\SYSTEM.ALT

For each user who is logged on, the following files also cannot be opened:

<%SystemRoot%>\Profiles*<user>*\NTUSER.DAT

<%SystemRoot%>\Profiles*<user>*\ntuser.dat.LOG

IMPORTANT

When performing a filesystem backup of a system disk, the previously listed files are not backed up. Excluding them only solves the problem of managing the session reports. You should perform a CONFIGURATION backup to back up the contents of these specific files.

When backing up an inactive system disk (for example in a dual-boot situation) the previously listed files are not a part of the currently active CONFIGURATION. These files can be backed up in a filesystem backup, and should not be excluded.

Preventing Backup Failure

Data Protector provides a set of features that improve backup robustness, thus lessening the chance that a backup could fail.

If a backup of an object fails to start, Data Protector tries to back up this object again at the end of the backup session. If it fails again, the object is not backed up, and the status of the object and the session is set to `Failed`. A backup is repeated when it is scheduled. If some objects finish properly, the session status is completed with failures.

Clients that are not up and running when they are scheduled to be backed up are retried after the rest of the objects are completed. Before the first failed object is retried, the backup session is suspended for 30 seconds. This waiting time can be changed using the `WaitBeforeRetry` global option. See “Global Options File” on page 613 for information on how to change global options.

IMPORTANT If you have an infrequent backup schedule, this may result in a period of time when there is no recent backup of your data.

NOTE Data Protector always needs one full backup of data. If no protected full backup is available, a full backup will be done next time, even though an incremental backup was scheduled. To avoid this, run a full backup of the failed system interactively before you schedule a backup.

For details on full and incremental backup behavior, see the *HP OpenView Storage Data Protector Concepts Guide*.

When the `Reconnect Broken Connection` backup option is set, Data Protector reconnects the

- Backup Session Managers and Disk or Media Agents (control connections) or
- Disk Agent and Media Agent during backup (data connections)

in case of short-term network problems during a backup session (this often happens on unreliable LAN networks). The time-out value can be defined by the `OB2RECONNECT_RETRY` `omnirc` variable. Refer to “Using Omnirc Options” on page 615 for information on using the `omnirc` file.

Enabling Wake ONLAN Support

If you have any machines that support remote power-up (**Wake ONLAN**), you can use the Data Protector Wake ONLAN support. When a Backup Session Manager fails to connect to a client that is configured to use Wake ONLAN support, it sends a wake-up request according to the Wake ONLAN protocol, and retries connecting to the client. This allows the full use of the power-saving features of desktop systems, which would otherwise interfere with the backup process.

NOTE

You can enable Wake ONLAN support for computers equipped with a Wake ONLAN-compatible LAN interface, such as the HP NightDIRECTOR series. The Wake ONLAN (WOL) option is available in the BIOS setup.

When you install a Disk Agent on a Windows client and add it to a cell, the client's Mac address is automatically discovered. You can also manually change the Mac address in the same section where you enable the Wake ONLAN (WOL) option, as shown below.

Use the following steps to enable Wake ONLAN support for the Windows client:

1. In the `Data Protector Manager`, switch to the `Clients` context.
2. In the Scoping Pane, right-click the client whose WOL option you want to enable, and then click `Properties`.
3. Click the `Advanced` tab.
4. Under the `Magic Packet` section, select the `Enable Magic Packet` check box, and then click `Apply`.

Restarting Failed Backups

Data Protector provides a simple way of restarting the backup of failed objects only. This can be done as follows:

1. In the `Data Protector Manager`, switch to the `Internal Database` context.
2. Under `Internal Database`, expand the `Sessions` item.
3. In the `Results Area`, search for your backup.

You can sort your sessions using the buttons on the top of each of the columns.

4. Right-click the failed session, and then select `Restart Session`.

   A dialog box appears asking you to confirm that you want to restart the session. Click `Yes`.

7 Copying Data

In This Chapter

This chapter describes Data Protector functionalities that enable you to duplicate backed up data during or after a backup. The chapter is organized as follows:

Overview

Duplicating backed up data brings several benefits. You can copy data to improve its security and availability, or for operational reasons.

Data Protector provides the following methods of duplicating backed up data: object copy, object mirror, and media copy. These methods are described in subsequent sections of this chapter. For a comparison of the duplication methods, refer to the *HP OpenView Storage Data Protector Concepts Guide*.

You can also use a combination of the duplication methods. For example, you can create object copies or media copies of data that is the result of object mirroring. Or, you can copy entire media containing object copies.

Copying Objects

What Is Object Copy

The Data Protector object copy functionality enables you to copy selected object versions to a specific media set. You can select object versions from one or several backup sessions. During the object copy session, Data Protector reads the backed up data from the source media, transfers the data, and writes it to the target media.

The result of an object copy session is a media set that contains copies of the object versions you specified.

The following is characteristic of the object copy functionality:

- Start of session

  An object copy session can be started interactively or automatically.

- Selection of media

  As source media, you can use original media sets containing backups, media sets containing object copies, or media sets that are media copies.

  However, the selection of the media set is not possible after the start of the object copy session. In case of a mount request, you need to provide the specific medium that is requested by Data Protector, or its identical copy (created using the media copy functionality).

- Media type

  You can copy objects to media of a different type. Furthermore, the block size of the destination device can be larger than the block size of the source device.

- Media policy

  You can append data to media already containing backups or object copies.

Why Use Object Copy?

Additional copies of backed up data are created for multiple purposes:

- Vaulting

You can make copies of backed up objects and keep them in several locations.

- Freeing media

  To keep only protected object versions on media, you can copy such object versions, and then leave the medium for overwriting.

- Demultiplexing of media

  You can copy objects to eliminate interleaving of data.

- Consolidating a restore chain

  You can copy all object versions needed for a restore to one media set.

- Migration to another media type

  You can copy your backups to media of a different type.

- Support of advanced backup concepts

  You can use backup concepts such as disk staging.

Using Object Copy

Below are the prerequisites and limitations of the object copy functionality:

Prerequisites

- You need to have a Media Agent installed on every system that will participate in an object copy session.
- You need to have at least two backup devices configured in the Data Protector cell.
- You need to have media prepared for the object copy session.
- You need to have appropriate user rights for performing an object copy session.

Limitations

- It is not possible to copy objects backed up using the ZDB to disk or NDMP backup functionality.
- It is not possible to create multiple copies of one object version in one object copy session.
- The destination devices must have the same or a larger block size than the source device.

- The same medium cannot be used as a source medium (to copy objects from) and as a target medium (to copy objects to) in the same object copy session.
- During object copying, the media used as sources are unavailable for restore.
- The `Reconnect` functionality between agents is not available in an object copy session.
- It is not possible to demultiplex SAP DB, DB2, or SQL integration objects.

How to Use the Object Copy Functionality

First, create an object copy specification. In the specification, select the objects you want to copy, the media and devices you want to use, session options, and the media location priority that influences how Data Protector selects the media set in case the same object resides on several media sets.

IMPORTANT

The Data Protector SAP DB, DB2, and Microsoft SQL Server integrations have interdependent data streams. Hence the object copy operation must preserve the layout of objects on media to enable a restore. To ensure this, select all objects of these integrations with the same backup ID for copying. Otherwise, a restore from the copy will not be possible.

The minimum number of devices required for copying objects of these integrations equals the number of devices used for backup. The concurrency of the devices used for backing up and copying the objects must be the same.

The media from which you will copy objects are referred to as **source media**, while the media to which you will copy the objects are referred to as **target media**. The source media and target media can be of a different media type.

Selection of Devices

You need separate devices to be used with the source media and the target media. The destination devices can have a larger block size than the source devices. However, to avoid impact on performance, it is recommended that the devices have the same block size and are connected to the same system or to a SAN environment.

Object copying is load balanced by default. Data Protector makes optimum use of the available devices by utilizing as many devices as possible.

If you do not specify the source devices to be used in the object copy specification, Data Protector uses the default devices. By default, the devices that were used for writing the objects are used as source devices. If destination devices are not specified per object, Data Protector selects the most suitable devices automatically from those you selected in the object copy specification.

Devices are locked at the beginning of the session. Devices that are not available at that time cannot be used in the session, as device locking after the beginning of the session is not possible. If a media error occurs, the device with errors will be avoided within that copy session.

What Is the Result?

The result of a successful object copy session is an additional copy of the selected backup objects on a different media set.

Configuring Object Copy

An object copy session is based on an object copy specification. Configuring an object copy specification consists of the following steps:

1. Selecting the type of session - interactive or automated (post-backup or scheduled).
2. Selecting what to copy.
3. Selecting which devices to use for copying.
4. Selecting how to copy - the session options.

Figure 7-1 **Selecting Objects for Interactive Copying**

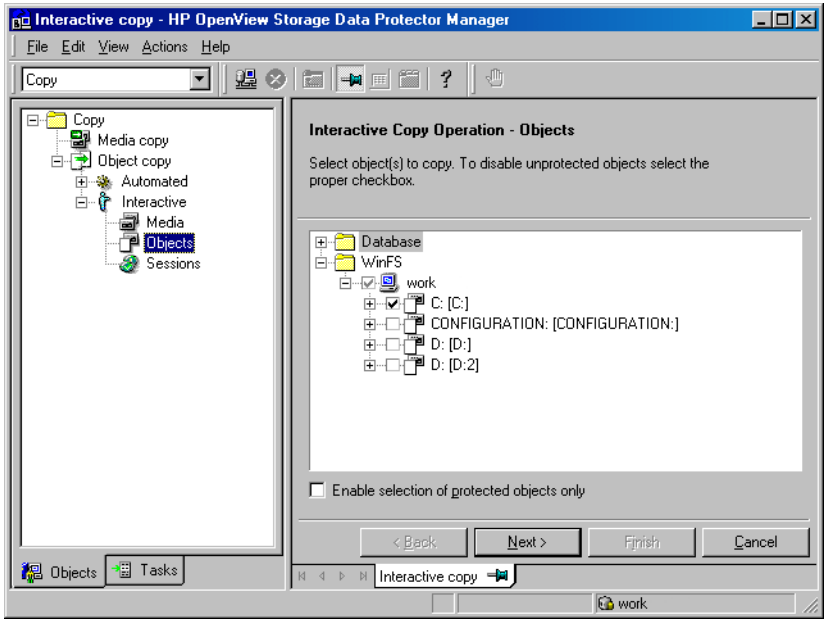

Interactive Object Copying

An interactive object copy session is started using an interactive object copy specification. You cannot save an interactive object copy specification; you can only start an object copy session.

In the `Copy` context, expand `Copy`, then expand `Interactive`, and click `Media`, `Objects`, or `Sessions`. For detailed steps, refer to the online Help index keyword “interactive object copying”.

Automated Object Copying

Data Protector offers two types of automated object copying: **post-backup object copying** and **scheduled object copying**.

In an automated object copy specification, you can specify one or more criteria for the selection of object versions that will be copied:

- Backup specifications - to copy only object versions backed up using specific backup specifications.
- Data protection - to copy only protected object versions.
- Number of existing copies - to copy only object versions that do not have more than the specified number of successful copies.

- Libraries - to copy only object versions located on the media in the specified libraries.
- Time frame (only in a scheduled object copy specification) - to copy only object versions backed up in the specified period of time.

Post-Backup Object Copying

Post-backup object copying takes place after the completion of a backup session that is specified in the automated object copy specification. It copies objects selected according to the automated object copy specification that were backed up in that particular backup session.

In the `Copy` context, expand `Copy`, then expand `Automated`, right-click `Post Backup`, and click `Add`. For detailed steps, refer to the online Help index keyword “post-backup object copying”.

Scheduled Object Copying

Scheduled object copying takes place at a user-defined time. Objects backed up during different backup sessions can be copied in a single scheduled object copy session.

You can configure scheduled object copying to run on specific dates at specific times, or to run periodically. You can reset, disable, or enable a schedule, and disable or enable object copying on holidays.

In the `Copy` context, expand `Copy`, then expand `Automated`, right-click `Scheduled`, and click `Add`. For detailed steps, refer to the online Help index keyword “scheduled object copying”.

Object Copy Options

You can specify data protection, catalog protection, and logging level for object copies in the object copy specification. Equivalent options are used for backup as well. For more information on these options, see “Most Frequently Used Backup Options” on page 271.

Depending on your policy, backed up objects and their copies may have the same or different option values specified. For example, you can specify the `No Log` value for a backup object to increase the backup performance, and then specify the `Log All` value for the same object in a subsequent object copy session.

To create identical copies of backed up objects, specify the same logging level for object copies. Consider that each object copy with a logging level higher than `No Log` has an impact on the IDB size. For information on the impact of these options on the IDB size, refer to the *HP OpenView Storage Data Protector Concepts Guide*.

IMPORTANT

Do not select the `Recycle data and catalog protection after successful copy` option if you are copying objects from a ZDB to disk+tape session, otherwise instant recovery from this backup using the GUI will no longer be possible after the media are overwritten.

Selecting the Media Set to Copy From

If an object version that you want to copy exists on more than one media set, which has been created using one of the Data Protector data duplication methods, any of the media sets can be used as a source for copying. By default, Data Protector automatically selects the media set that will be used. You can influence the media set selection by specifying the media location priority.

The overall process of media selection is the same as for restore. For details, see "Selecting the Media Set to Restore From" on page 376. When copying objects interactively, you can manually select the media set to copy from when your starting point is `Objects` or `Sessions`. You cannot select media when configuring automated object copying, as the backup of the objects is often performed at a later time.

Object Copy Completion Status

You can copy objects that have the status `Completed` or `Completed/Errors`, provided that all the media on which they reside are logged in the IDB. If the copy operation is successful, the status of the copied object is the same as the status of the corresponding backed up object.

A post-backup object copy session does not start if the backup session failed. If the backup session has been aborted, but contains completed objects, a post-backup object copy session copies the completed objects by default. To disable the copying of aborted sessions, set the global variable `CopyStartPostBackupOnAbortedSession` to 0.

If you have aborted an object copy session, or if it failed for other reasons, the object copies that are results of such a session have the status `Failed`. An object copy with the status `Failed` cannot be copied again; its data and catalog protection are set to `None`.

Tasks Based on Object Copy

The object copy functionality also enables you to perform the following tasks:

Vaulting

Vaulting is a process of storing media in a safe place, often called a vault, where they are kept for a specific period of time. For details, refer to “Vaulting Media” on page 186.

It is recommended to keep a copy of the backed up data on site for restore purposes. To obtain additional copies, you can use the object copy, object mirror, or media copy functionality, depending on your needs.

For an example of using the object copy functionality to obtain copies for vaulting, refer to the online Help index keyword “copying objects for vaulting purposes”.

Freeing Media

You can minimize the media space consumption by keeping only protected backups and overwriting unprotected ones. As a single medium may contain both, you can copy protected objects to a new media set and leave the medium for overwriting.

On how to free media, refer to the online Help index keyword “freeing media”.

Demultiplexing of Media

Multiplexed media contain interleaved data of multiple objects. Such media may arise from backup sessions with the device concurrency more than 1. Multiplexed media may compromise the privacy of backups and require more time for restore.

Data Protector offers a possibility of demultiplexing of media. Objects from a multiplexed medium are copied to several media that you specify.

On how to demultiplex media, refer to the online Help index keyword “demultiplexing media”.

Consolidating a Restore Chain

You can copy a restore chain (all backups that are necessary for a restore) of an object version to a new media set. A restore from such a media set is faster and more convenient, as there is no need to load several media and seek for the needed object versions.

On how to consolidate a restore chain, refer to the online Help index keyword “consolidating a restore chain”.

Migration to Another Media Type

You can migrate backed up data to another media type. For example, you can copy objects from file devices to LTO devices or from DLT devices to LTO devices.

On how to migrate to another media type, refer to the online Help index keyword “migrating to another media type”.

Disk Staging

You can use the object copy functionality to implement disk staging. For information on the disk staging concept, refer to *HP OpenView Storage Data Protector Concepts Guide*.

You can perform disk staging using the object copy specification configured specifically for this purpose. For an example of disk staging, refer to the online Help index keyword “disk staging”.

Object Mirroring

What Is Object Mirroring?

The Data Protector object mirror functionality enables writing the same data to several media sets simultaneously during a backup session. You can mirror all or some backup objects to one or more additional media sets.

Benefits of Object Mirroring

The use of the object mirror functionality serves the following purposes:

- It increases the availability of backed up data due to the existence of multiple copies.
- It enables easy multi-site vaulting, as the backed up data can be mirrored to remote sites.
- It improves the fault tolerance of backups, as the same data is written to several media. A media failure on one medium does not affect the creation of the other mirrors.

Using Object Mirroring

Below are the limitations of the object mirror functionality:

Limitations

- It is not possible to mirror objects backed up using the ZDB to disk, direct backup, or NDMP backup functionality.
- It is not possible to mirror an object to the same device more than once in a single session.
- Block size of the devices must not decrease within a mirror chain. This means the following:
  - The devices used for writing mirror 1 must have the same or a larger block size than the devices used for backup.
  - The devices used for writing mirror 2 must have the same or a larger block size than the devices used for writing mirror 1, and so on.

How to Use Object Mirroring

You specify object mirroring when configuring a backup specification. On how to configure a backup specification, see “Creating a Backup Specification” on page 199.

In the backup specification, select the objects you want to mirror, and then specify the number of mirrors. To be able to specify more than 5 mirrors, increase the `MaxNumberOfMirrors` variable value in the global options file.

Selection of Devices

Specify separate devices for the backup and for each mirror. To avoid impact on performance, it is recommended that the devices have the same block size and are connected to the same system or to a SAN environment.

IMPORTANT

The minimum number of devices required for mirroring SAP DB, DB2, or Microsoft SQL Server integration objects equals the number of devices used for backup.

When a backup session with object mirroring starts, Data Protector selects the devices from those you specified in the backup specification.

Object mirroring is load balanced by default. Data Protector makes optimum use of the available devices by utilizing as many devices as possible. When you perform an object mirror operation from the command line, load balancing is not available.

What Is the Result?

The result of a successful backup session with object mirroring is one media set containing the backed up objects and additional media sets containing the mirrored objects. The mirrored objects on these media sets are treated as object copies.

Copying Media

What Is Media Copying?

The Data Protector media copy functionality enables you to copy media after a backup has been performed. Media copying is a process that creates an exact copy of a medium containing a backup. You can use it to duplicate media for archiving or vaulting purposes.

After the media have been copied, you can move either the original media or the copies to an off-site vault, and keep the other set of media on site for restore purposes. On how to configure Data Protector for vaulting, see “Vaulting Media” on page 186.

Besides manually started media copying, Data Protector also offers automated media copying. For more information, see “Automated Media Copying” on page 333.

How to Copy Media

In the `Devices & Media` context, browse for a medium, right-click it and click `Copy`. For detailed steps, refer to the online Help index keyword “copying media”.

You need two devices with the same media type, one for the **source medium**, one for the **target medium**. A source medium is the medium being copied, while a target medium is the medium to which data is copied.

You can specify the protection period for the target medium, during which the data on the medium cannot be overwritten. The default protection is the same as for the original. Other options are Permanent and Until (specified date). A medium is protected until the protection of the object with the longest protection period on the medium expires.

You need to start the copying of each medium separately, as only one medium can be copied in a copy session. The copy operation is not available for media in free pools and for NDMP media.

What Is the Result?

The result of copying media is that you have two identical sets of media, the original media set and the copy. Either of them can be used for restore.

After the source medium has been copied, Data Protector marks it as non-appendable to prevent appending new backups. (This would result in the original being different from its copy.) The copy is also marked as non-appendable.

You can make multiple copies of the original media. You cannot, however, make copies of copies, also known as second generation copies.

NOTE

When copying media, it is possible that the target medium reaches the end of the tape before the source medium. This may happen if the source medium was written in streaming mode and you make a copy on a busy system or through a loaded network, which can create blank space where the tape has stopped and started again. You can prevent this by enabling tape padding when you format media. See “Formatting Media” on page 155.

Moving Copies

Typically, you want to move the copies of the media to a safe place. See “Vaulting Media” on page 186 and “Ejecting a Medium from a Device” on page 183 for more information.

Exporting Copies

Exporting a medium removes all information regarding this medium from the IDB. If you export the original medium, but one or more copies of the medium exist, one of the copies becomes the original.

If you try to import the removed copy, but the original media are not in the IDB, you have to import these media using the `force` option. See “Importing Media” on page 159 for instructions.

Restoring from a Copy

By default, Data Protector restores data from the original media set. However, if the original media set is not available, but a copy is available, the copy is used for the restore.

If neither the original nor a copy is available in the device during restore, Data Protector issues a mount request, displaying both the original and the copy as the media required for restore. You can use any one of these.

If you perform a restore using a standalone device, you can choose to restore from the copy rather than from the original. To do this, insert the copy in the device that will be used for the restore, or select the device containing the copy. However, if you perform a restore using a library device and the original is in the library, Data Protector will use it for the restore.

For detailed instructions on how to restore data from the media archive, see “Vaulting Media” on page 186.

Automated Media Copying

What Is Automated Media Copying?

Automated media copying is an automated process that creates copies of the media containing backups.

Data Protector offers two types of automated media copying: **post-backup media copying** and **scheduled media copying**.

Post-Backup Media Copying

Post-backup media copying takes place after the completion of a backup session. It copies all media used in that particular session.

In the `Devices & Media` context, right-click `Automated Operations` and click `Add Post-Backup Media Operation`. For detailed steps, refer to the online Help index keyword “post-backup media copying”.

Scheduled Media Copying

Scheduled media copying takes place at a user-defined time. Media used in different backup specifications can be copied in a single session. You create an automated media copy specification to define which media will be copied.

In the `Devices & Media` context, right-click `Automated Operations` and click `Add Scheduled Media Operation`. For detailed steps, refer to the online Help index keyword “scheduled media copying”.

You can configure scheduled media copying to run on specific dates at specific times, or to run periodically. You can reset, disable, or enable a schedule, and disable or enable automated media copying on holidays. For details, refer to the online Help index keyword “automated media copying”.

Limitations

- You cannot use standalone devices for automated media copying; only library devices can be used.
- The source medium and the target medium must be of the same type.
- You cannot copy NDMP media.

How Does Automated Media Copying Operate?

First you create an automated media copy specification. When the automated media copy session begins, Data Protector generates a list of media, referred to as **source media**, based on the parameters specified in the automated media copy specification. For each source medium, a **target medium** is selected to which the data will be copied. The target media are selected from the same media pool as the source media, from a free pool, or from the blank media in a library.

Selection and Use of Devices

For each source medium, Data Protector selects a pair of devices from the devices that you specified in the automated media copy specification. The automated media copy functionality provides its own load balancing. Data Protector tries to make optimum use of the available devices by utilizing as many devices as possible and selecting local devices if they are available.

Devices are locked at the beginning of the session. The devices that are not available at that time cannot be used in the session, as device locking after the beginning of the session is not possible. Note that at least a pair of devices must be available for each media type for the entire session to complete successfully. If the minimum number of devices necessary for the session cannot be locked, the session fails.

For devices with multiple configured paths, the local paths are preferred. If no local path is available, any available path in the predefined order is used.

If a media error occurs, the device with errors will be avoided within that automated media copy session. However, if there are no other devices available, it will be reused.

Destination Pool of the Copies

The source medium defines the destination pool of the target medium. This means that the copied media will belong to the same pool as the original media.

Data Protection of the Copies

The default protection period for the copy is the same as the protection for the original. You can set a different protection period when creating or modifying the automated media copy specification.

Mount and Cleanme Request Handling

The automated media copy functionality does not handle mount or cleanme requests. If a mount request is received, the media pair concerned is aborted, but the session continues. You can manually copy the media that were not copied after the automated media copy session finishes.

For examples of use, refer to the *HP OpenView Storage Data Protector Concepts Guide*.

8 Restore

In This Chapter

This chapter describes restore topics, such as how to restore specific data and how to use restore options to achieve a desired restore behavior.

For information on how to restore database applications such as Oracle, SAP R/3, MS Exchange, MS SQL, Informix, IBM DB2 UDB or Sybase, refer to the *HP OpenView Storage Data Protector Integration Guide*.

For information on how to restore the IDB, refer to Chapter 11, "Managing the Data Protector Internal Database," on page 457 and "Recovering the IDB" on page 494.

Restoring Your Data

A restore is a process that recreates the original data from a backup copy on a disk. This process consists of the preparation and actual restore of the data, and optionally some post-restore actions that make the data ready for use.

The Data Protector Internal Database (IDB) keeps track of data, including what files from which system are kept on a particular medium. The IDB provides fast and convenient access to the data to be restored.

Data Protector offers you some special restore features:

- The ability to restore on different levels: session, client, object, directory, specific file, or specific file version
- The option to specify an alternative location to restore your data
- Cross-platform restore
- Parallel restore of multiple objects from a session, on a client, or in a cell
- The possibility of automatic or manual selection of the media set to restore from

Depending on the platform, the way you specify these features and available options can vary.

Standard Restore Procedure

Prerequisite

In order to perform a restore, you need to have the appropriate user rights. These rights are defined according to the user group.

What You Need to Do to Perform a Restore

As part of the standard restore procedure, you need to do the following:

- Select the data to be restored
- Find the media needed
- Start the restore session

Other Settings

Other settings are already predefined according to the backup process, but can be modified. If you want to change these predefined settings, you can specify the following:

- The backup version you want to restore
- The location you want to restore data to
- The device to restore from
- How to handle file conflicts with existing files
- Restore options, such as locking files during restore
- The priority of selecting media containing the same object version according to their locations
- Manual selection of the media set to restore from if more than one media set contains the same object version

For detailed steps of standard restore tasks, refer to the online Help index keyword “standard restore procedure”.

Selecting Your Data for Restore

The Data Protector `Restore` context offers two possible ways of browsing objects for restore:

- `Restore Objects` with a list of backed up objects classified by client systems in the cell and by different data types, such as Filesystem, Disk Image, Internal Database, and so on.
- `Restore Sessions` with a list of filesystem sessions with all objects backed up in these sessions. You can choose to view only sessions from the last year, last month, or last week. By default, all filesystem sessions are listed. You cannot perform restore of the online database integrations from a specific backup session.

You can select either one object to perform a single restore, or multiple objects to perform a parallel restore. For more information on parallel restore, refer to “Restoring Files in Parallel” on page 371.

You can also specify a `Search Interval` and browse only objects backed up within a specific timeframe.

Data Protector offers the `Restore by Query` task, which searches for files and directories you want to restore and restores them. Refer to “Restoring by Query” on page 373.

Figure 8-1 Selecting Data for Restore

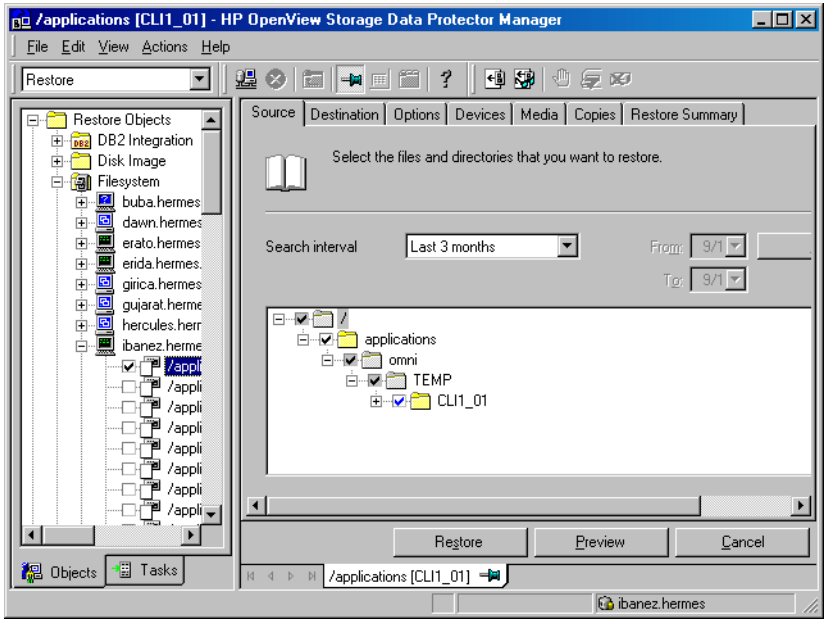

Selecting a Backup Version

When selecting data that you want to restore, the last backup version is selected by default. This means that only directories and/or files from the last backup session are selected for restore. Directories and files in the same tree structure that have not been backed up in the same backup session are shaded.

If you want to restore the data from any other backup session, browse for the file or directory that you want to restore, right-click it, and click `Restore Version`.

In the `Version` tab, click "`. . .`" to get additional information about the backup versions. The "`. . .`" button is available only if the backup was performed using a logging level that logs attributes.

Handling File Conflicts

In the `Description` property page of your restore, you can specify how to handle conflicts between the version currently on the disk and the backup version of a file. `File Conflict Handling` offers you three

possible options: Keep most recent, No overwrite, and Overwrite. For more information on these options, refer to “Restore Options” on page 365.

Specifying Restore Location

By default, Data Protector restores data to the same client and directory from which it was backed up. You can change these default settings in the Destination property page by specifying where to restore your data to:

- With the appropriate user rights, you can restore to another client.
- You can restore to another directory.

This specification can be set on a per-object basis.

Additionally, Data Protector offers the Restore As/Into option for specifying a different location for individual files and directories from the same backup object. This specification can be set on a per-object basis or for the individual files.

For more information on specifying restore location, refer to “Restoring Files to Different Paths” on page 370.

Setting Restore Options

Set restore options in the Options property page of your restore. These are available according to the type of data being restored. For example, not all restore options available for a filesystem restore are available for a disk image restore. For more information on restore options, refer to “Restore Options” on page 365.

Restoring Under Another Device

By default, the device used for restore is the same device as the one the backup was made to. You can restore your data from any device configured in the same Data Protector cell. To specify a new device, click the Change button in the Devices property page of your restore. The new device will be used for this session only.

NOTE

With *some* database integrations, you can set the changed device as a default restore device for *all* Data Protector integration restore sessions (regardless of the type of integration), by clicking the `Save as default` button.

Finding Needed Media

To get a list of media on which your data is stored, go to the `Media` property page after you select data for restore.

You can also find the media needed for the restore by clicking the `Needed Media` button in the `Start Restore Session` dialog box. This dialog box appears when you start the restore.

If an object version exists on more than one media set, you can influence the selection of the media set that will be used for the restore by setting the media location priority, or manually select the media set that will be used. For more information, refer to “Selecting the Media Set to Restore From” on page 376.

Previewing and Starting a Restore

Ensure that the media are loaded properly before starting the restore. Otherwise, the media will not be detected.

If restoring objects selected in the `Source` property page of your restore, click `Start` to start the restore process or `Preview` if you want to preview it.

If also restoring objects selected in the Scoping Pane, in the `Actions` menu click `Preview Restore` to preview the restore process, or click `Start Restore` to actually start it.

Aborting a Restore

Aborting a restore session stops the restore. Data processed before the session was aborted is restored to the specified location.

To abort a restore session, click `Abort` in the `Actions` menu.

You can also abort restore sessions in the Data Protector `Monitor` context.

Restoring Disk Images

A disk image restore is a sector-by-sector restore of a disk image backup. Data Protector restores a complete image of the disk that was backed up (as a disk image) at a certain point in time. This method is particularly fast. It is available for Windows and UNIX systems.

Prerequisites

You need to meet the following prerequisites in order to perform a disk image restore:

- The disk must have been backed up using the disk image backup.
- To restore a disk image on a disk other than the disk from which you backed it up, the new disk must be of the same size or larger.
- On UNIX systems, unmount the disk before a disk image restore and then mount it back afterwards.

Limitation

On Windows systems, disk image restore fails if a file or section is in use.

Procedure

To restore a disk image backup, expand the disk image object under the `Restore` context as shown in Figure 8-2 on page 343, and then use the standard restore procedure. Refer to “Standard Restore Procedure” on page 337.

Figure 8-2 Disk Image Objects

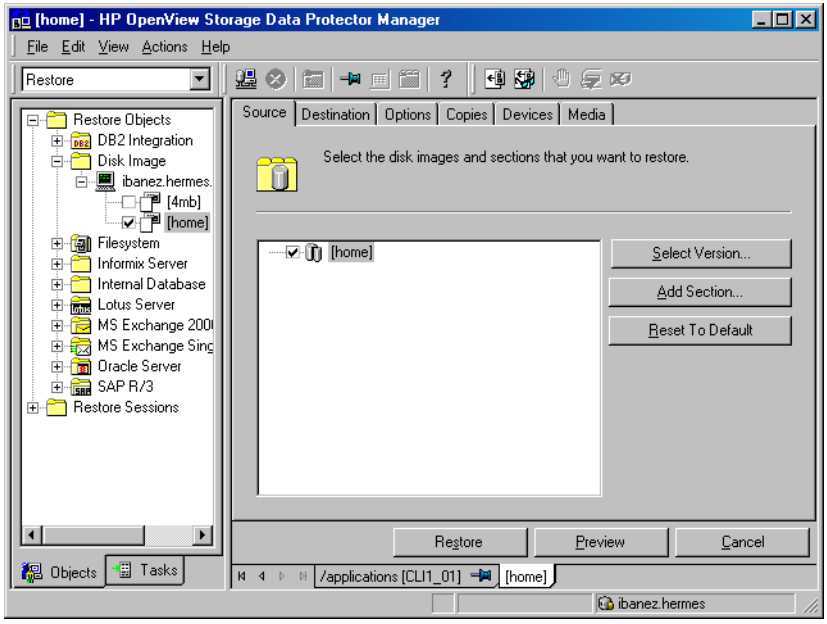

Restoring Your Data to a Shared Disk

Data Protector allows you to restore UNIX and Windows data to a Windows shared disk, even if the data was not originally backed up from the shared disk. The Data Protector user account and its Inet service must have permission to access the remote computer and permission on the Disk Agent client. Refer to “Setting the User Account for the Data Protector Inet Service” on page 232 for more information on how to use the appropriate logon account.

Here are some cases in which one would restore a UNIX or Windows filesystem to a Windows shared disk:

- If the system is not part of the Data Protector cell and does not have the Data Protector Disk Agent installed.
- If you are restoring to platforms not directly supported by Data Protector, such as Windows for Workgroups or Windows 3.1 systems.
- If you want to make the data available from several systems.

NOTE

When you restore your data to a different filesystem type than it was backed up from (UNIX to Windows), filesystem-specific attributes may be lost.

How to Restore to a Shared Disk

In the `Destination` property page of your restore, you can specify the target client and a Windows shared disk as a new location for the data you want to restore. For detailed steps, refer to the online Help index keyword “shared disks, restoring to”.

Restoring UNIX Systems

What Is Restored?

When restoring files to the original location from which the backup was performed, Data Protector restores the files, including file attributes.

System-specific data, such as ACL (Access Control List) on UNIX, is restored only on the same filesystem type and operating system from which the backup was made.

Restoring Regular UNIX Files

Use the standard restore procedure to restore UNIX files and directories. Refer to “Standard Restore Procedure” on page 337.

Restoring VxFS

When restoring VxFS data backed up to a temporary directory, use the `Restore As` option and restore it to the desired location. Refer to “Restoring Files to Different Paths” on page 370 for information on how to use the `Restore As` option.

Restoring OmniStorage Backups

Beside restoring backed up data into an OmniStorage controlled file system (MFS), Data Protector A.05.50 offers the possibility to restore OmniStorage filesystem data backed up with OmniBack II or Data Protector, using a normal filesystem restore on HP-UX 11.x. In this case, the “migration attributes” of OmniStorage, like migration policies, will be lost.

OmniStorage files can be restored to any filesystem on HP-UX, but in order to retain the VxFS specific file attributes it is recommended that the target filesystem is of JFS type with a VxFS3 or later layout.

Restoring Disk Images

Refer to “Restoring Disk Images” on page 342.

Restoring to a Shared Disk

Refer to “Restoring Your Data to a Shared Disk” on page 344.

Restoring Windows Systems

What Is Restored?

When restoring a Windows filesystem, Data Protector restores the data within the files and directories, as well as Windows-specific information about the files and directories.

Consider the filesystem restore limitations when restoring to a different filesystem from the one where the backup was performed. See “Filesystem Limitations” on page 348.

The following Windows-specific information is restored:

- Full Unicode filenames
- FAT16, FAT32, VFAT, and NTFS attributes
- Directory share information

  If a directory is shared on a network during backup, the share information is stored on the backup medium. The directory will be shared on the network after the restore by default (unless a shared directory with the same share name already exists). To prevent restoring share information for directories that are being restored, deselect the `Restore share information for directories` option.

Figure 8-3 Restoring Share Information for Directories

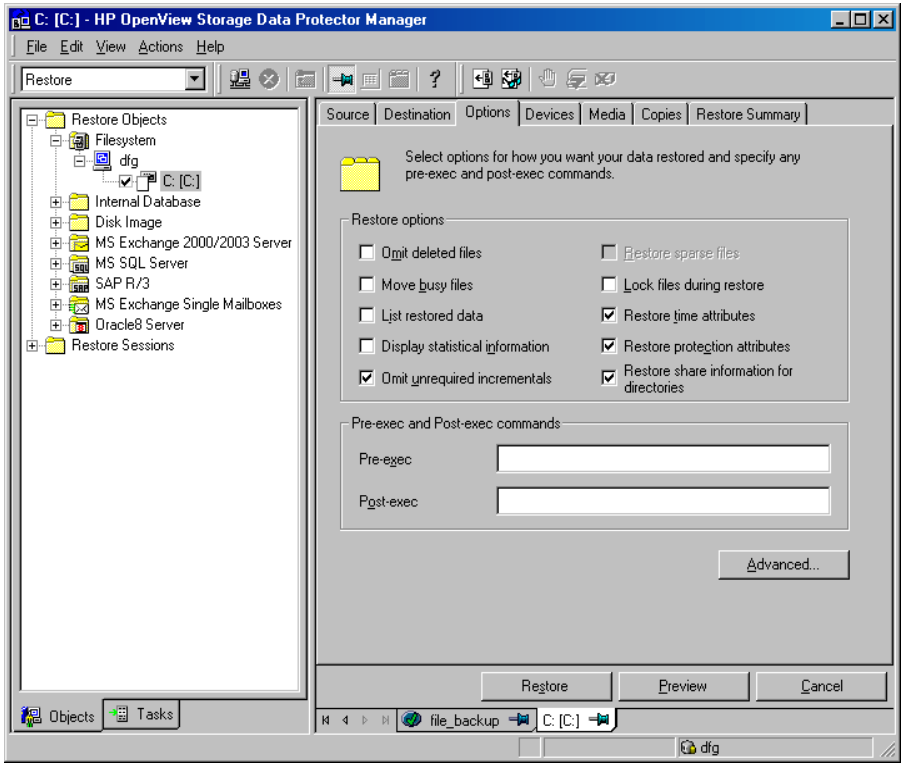

Windows directory share information can only be restored to a Windows system with a Data Protector A.05.50 Disk Agent or newer. If this requirement is not met, the directory will still be restored, but the Disk Agent will ignore the directory share information.

NOTE

File conflict handling options apply also for the restore of the directory share information. For example, if the `No overwrite` restore option is used for the restore, the directory share information for directories that exist on the disk, is preserved. Refer to “Restore Options” on page 365 for more information on how file conflict handling options affect your restore.

- NTFS alternate data streams

For example, Object IDs on Windows 2000 are backed up as sets of alternate data streams.

- NTFS security data

Additionally, the following applies on Windows systems, using NTFS 3.x:

- The NTFS filesystem supports reparse points.

  The volume mount points, Single Instance Storage (SIS), and directory junctions are based on the reparse point concept. These reparse points are selected like any other filesystem object.

- The NTFS filesystem supports sparse files as an efficient way of reducing the amount of allocated disk space.

  These files are backed up as sparse to save tape space. Sparse files are backed up and restored as sparse to the NTFS 3.x filesystem only.

- Some of the NTFS filesystem specific features are controlled by the system services, which maintain their own data records. These data structures are backed up as a part of `CONFIGURATION`.

- Encrypted files.

Filesystem Limitations

You can select a different target filesystem from the one where the backup was performed. This functionality has limitations that should be taken into consideration. See Table 8-1 on page 348.

Table 8-1 Windows Filesystem Restore Limitations

	TO					
FROM	FAT32	FAT16	CDFS	UDF	NTFS 3.0[a]	NTFS 3.1[b]
FAT32	FC	FC	N/A	N/A	FC	FC
FAT16	FC	FC	N/A	N/A	FC	FC
CDFS	FC	FC	N/A	N/A	FC	FC
UDF	FC	FC	N/A	N/A	FC	FC
NTFS 3.0[a]	*	*	N/A	N/A	FC	FC
NTFS 3.1[b]	*	*	N/A	N/A	FC	FC

How to Read This Table

a	Also called NTFS 5.0. It is used by Windows 2000.
b	Also called NTFS 5.1. It is used by Windows XP/Server 2003.
FC	Full Compatibility, meaning that the file attributes are entirely preserved.
*	Reparse points, sparse files and encrypted files are not restored. Files are restored without security information and alternate data streams.

Table 8-1 shows that NTFS 3.x filesystem objects can only be adequately restored to the NTFS 3.x filesystem. The filesystem-specific attributes and alternate data streams are lost when restoring into a different or older filesystem version.

- A Windows reparse point, such as a directory junction or a volume mountpoint, can only be restored to an NTFS 3.x filesystem. UNIX reparse points cannot be restored to an NTFS 3.x filesystem.

NOTE

When you restore an NTFS 3.x filesystem that contains SIS reparse points, a full disk condition may occur. This happens if the original file is restored into multiple target files, which can take up more space than available.

- Sparse files are restored as sparse to the NTFS 3.x filesystem only.
- User Disk Quotas cannot be restored using Data Protector.
- If a user attempts to restore a sparse file to a non-NTFS 3.x filesystem, Data Protector will issue a warning. A sparse file restored to a filesystem other than NTFS 3.x will not include zero sections.
- Microsoft encrypted NTFS 3.x files can only be restored to the NTFS 3.x filesystem, because other filesystem drivers cannot decrypt them.

Restoring Regular Windows Files and Directories

Use the standard restore procedure to restore Windows files and directories. Refer to “Standard Restore Procedure” on page 337.

Restoring Shared Disks

Objects that were backed up as shared disks are associated with the Disk Agent client that was used to back them up. If the environment has not changed, you can restore the shared disk as you would a local Windows filesystem. By default, the same Disk Agent client that was used to back up the shared disk is used to restore the data to the original location.

For information on how to choose and configure the Disk Agent client that restores the shared disks, refer to “Backing Up Windows Shared Disks” on page 230.

For information on restoring a UNIX or Windows filesystem to a shared disk, refer to “Restoring Your Data to a Shared Disk” on page 344.

Restoring Disk Images

Refer to “Restoring Disk Images” on page 342.

Restoring the Windows CONFIGURATION

To restore the Windows `CONFIGURATION`, select the `CONFIGURATION` object and follow the standard restore procedure. See Figure 8-4.

Figure 8-4 Restoring Windows CONFIGURATION

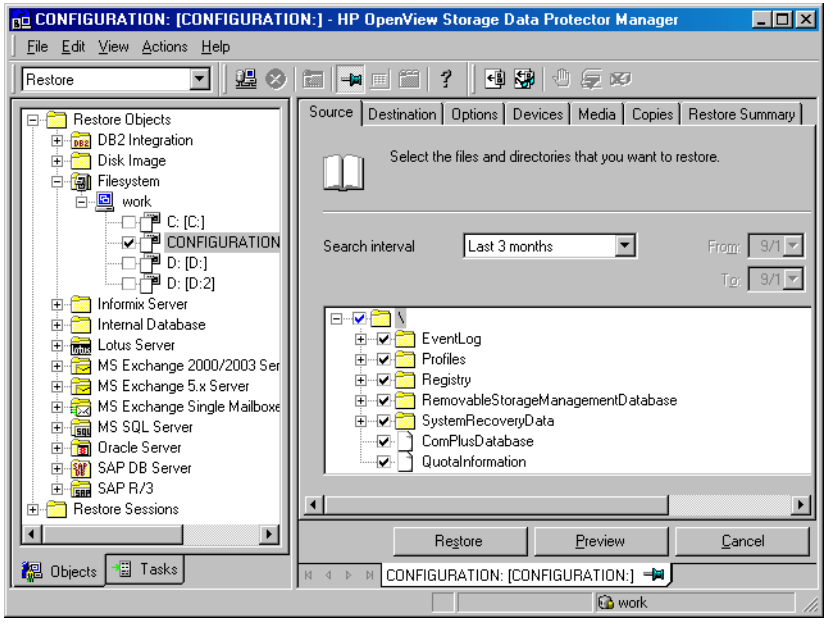

Prerequisites

The CONFIGURATION consists of data structures that influence system operation. Therefore, the system must be prepared for such a restore. The prerequisites depend on the contents of the CONFIGURATION item and the Windows operating system version. Refer to “Backing Up CONFIGURATION” on page 218. They can be summarized as follows:

- User profiles that are currently being used cannot be restored. The login account has to be changed or the relevant service has to be stopped.

  Refer to “Restoring Windows User Profiles and Event Logs” on page 356 for details.

- You have to boot the system in the Active Directory restore mode to restore the Active Directory.

  Refer to “Restoring Windows Services” on page 354 for details.

When the whole CONFIGURATION is restored, restart the system to read the restored data in the Registry. Refer to “Restoring the Windows Registry” on page 353 for details.

Restoring the SysVol

You can perform a restore of SysVol directory in one of three modes:

- Nonauthoritative restore

  If at least one domain controller in the domain is available and working, files are restored to their original location. The restored data is not propagated to other domain controllers.

- Authoritative restore

  Perform an authoritative restore if critical SysVol data is deleted from the local domain controller and the deletion is propagated to other domain controllers.

- Primary restore

  If all domain controllers in the domain are lost and you want to rebuild a domain controller from backup, the FRS is informed that you are restoring primary files, and files are restored to their original location.

Restoring the Windows System State

Prerequisites

If you use Active Directory, which is always a part of the System State, you have to start the system in the Active Directory restore mode.

Refer to “Restoring Windows Services” on page 354 for details on Active Directory modes.

You restore the System State by selecting the following objects in the Restore wizard:

1. System State objects that belong to `CONFIGURATION`. Refer to “Backing Up the Windows System State” on page 220 for a list of these objects.
2. The `SystemVolumeInformation` folder and the boot files. These are located on the system drive.

Figure 8-5 Selecting System State Items

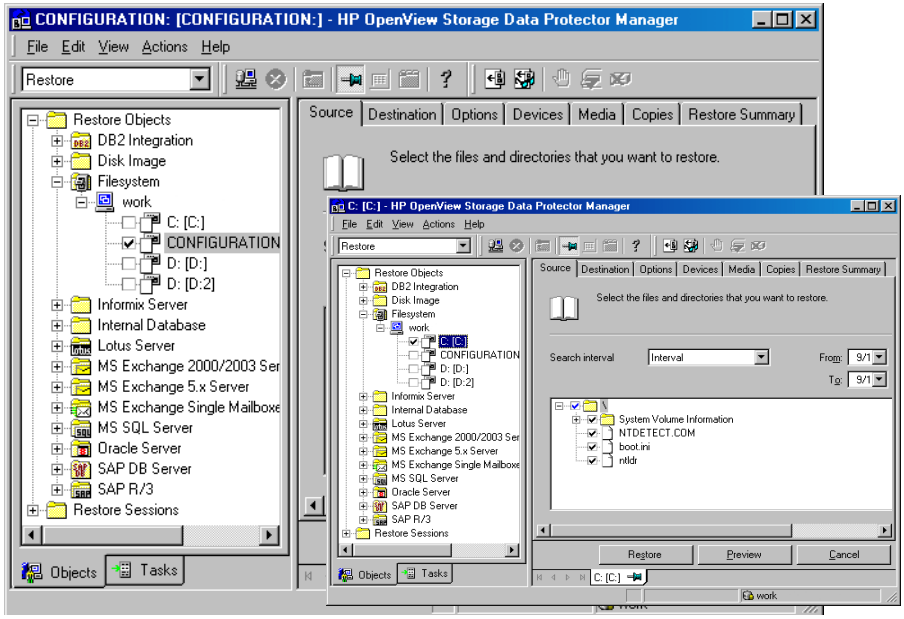

NOTE From the Data Protector point of view, the System State consists of ordinary filesystem objects and `CONFIGURATION` objects. As opposed to selecting objects in the Backup wizard, different objects for restore are selected in separate Restore wizards.

Once the restore session is completed, restart the system.

Restoring the Windows Registry

To restore the Windows `Registry`, expand the `CONFIGURATION` item and select only the `Registry` item.

Once the restore session is completed, restart the system.

NOTE

If you select the whole Windows `Registry` for a restore, some of the Registry keys are not restored, and others are treated in a special way during a restore. This is because certain keys are being used by the operating system. You can find them under the following `Registry` key:

`\HKEY_LOCAL_MACHINE\SYSTEM\CurentControlSet\Control\BackupRestore\KeysNotToRestore`

Restoring Windows Services

To restore Windows services, expand `CONFIGURATION` and select the service you want to restore.

Prerequisites

The following information that belongs to Windows services can be selected under `CONFIGURATION`:

- `COMPlusDatabase`
- `FileReplicationService`
- `RemovableStorageManagementDatabase`
- `ActiveDirectoryService`
- `TerminalServiceDatabase`
- `CertificateServer`
- `DHCP`, `WINS`, and `DNSServerDatabase`

For a detailed explanation of these terms, refer to the “Glossary”.

The list below describes specifics related to restoring a particular Windows service.

Active Directory Restore

If you want to restore the Active Directory service, restart the system using the `Directory Services Restore Mode` start-up option.

When the system is started in the `Directory Services Restore Mode`, the domain user accounts cannot be used. Configure the `Data Protector Inet` and the `crs` service (for a Cell Manager) to log on using the local system account and then restart the services. Refer to “Setting the User Account for the Data Protector Inet Service” on page 232 for more details.

Select `Active Directory`, and set a replication mode by choosing among the `Windows specific options`: `Primary`, `Nonauthoritative`, `Authoritative`. For information on these options, refer to “Active Directory Specific Options” on page 368.

NOTE

To perform an `Authoritative` restore, you also need to run `ntdsutil.exe` after the restore session has finished. For example, to perform a typical `authoritative` restore, at a command prompt enter `ntdsutil`, then `authoritative restore`, then `restore database`. Restart the server and wait for replication to take place.

TIP

You can also create a post-exec command to perform the additional action needed for the Active Directory authoritative restore. For example, to perform an authoritative restore of an entire directory, use the following line:

```
ntdsutil "popups off" "authoritative restore" "restore
database" quit quit
```

Figure 8-6 Active Directory Restore Modes

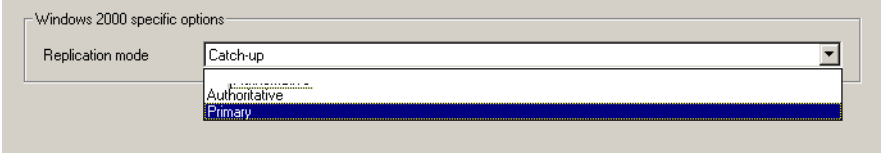

Certificate Services Restore

Certificate Server Services are restored offline. You have stop them before you can start a restore. `Authoritative` is the only possible replication mode.

Once the restore has finished, restart the system.

Remote Storage Service Restore

Although the RSS databases are part of System State data, you restore them manually. The RSS database must be restored offline. You can provide pre- and post-exec scripts to stop and restart the service, or you can stop and restart it manually before and after the restore, respectively.

Select the following directories for restore:

- `<%SystemRoot%>\System32\RemoteStorage`
- `<%SystemRoot%>\System32\NtmsData`

Restoring DFS

Data Protector restores the configuration of the Windows Distributed File System (DFS) as part of one of the following:

- Windows Registry, if the DFS is configured in a standalone mode.
- Windows Active Directory, if the DFS is configured in a domain mode.

Restoring Windows User Profiles and Event Logs

To restore the Windows User Profiles and Event Logs, expand the `CONFIGURATION` object and select the items you want to restore.

User Profiles

Data Protector will not restore any files that are currently accessed. You have to log off the system and stop all the services that are running under the user account whose profiles you want to restore.

The restore session can be started from another system or by logging on the restore target system as a different user.

Deleted User Profiles

A user profile can only be restored when its location is already defined on the system. Individual files of existing user profiles or deleted profiles can be still restored as long as they exist among the system's profiles. Otherwise, you need to recreate them before restoring the files. Proceed as follows:

1. Log on as the user whose profile you want to restore in order to create a default user profile.
2. To keep the restored files unmerged, you can delete the files in the newly created profile before running a restore session.
3. Log off and start the restore session by logging on as a different user or by using another system.

The system may assign a different name to the user. In this case, use the `Restore As` option to restore the files to the newly assigned location.

When the restore has finished, restart the system.

User Disk Quotas

User Disk Quotas cannot be restored using Data Protector. The backed up information can be restored using Microsoft utilities.

Restoring Windows TCP/ IP Services

WINS, DHCP, DNS Servers

On a Windows Server that runs a Microsoft TCP/IP protocol and is configured as a **WINS Server**, a **DHCP Server,** or a **DNS Server,** you can restore the services that manage network communication.

To restore Windows TCP/IP services, expand the `CONFIGURATION` item and select `WNS`, `DHCP`, or `DNSServerDatabase`.

Each of these services is automatically stopped before the restore.

When the restore has finished, restart the system.

Restoring Novell Netware Filesystems

Use the standard restore procedure to restore Novell NetWare filesystems. Refer to “Standard Restore Procedure” on page 337.

Restoring Namespace Information and Volume Space Restrictions

To restore only volume space restrictions, specify the `Volume space restrictions only restore` option in the `Destination` page. The object selected for the restore must be a volume.

Data Protector restores Novell NetWare volume namespace information during a regular filesystem restore session. Namespace information is restored on a per-file/directory basis for the following Name Spaces: DOS, Mac, NFS, OS/2.

To restore files or directories, note the following:

- Backed up namespace information will be successfully restored only if the same Name Spaces are installed on the volume where you are attempting to restore the data.
- DOS namespace exists on each installed Novell NetWare volume and is always restored.
- A Mac's resource fork can only be restored to a volume that has the Mac namespace installed.
- Specific namespace information depends on the existence of NDS/eDirectory objects, such as user and group IDs in NFS namespace.
- After restoring the Queue objects, manually create a queue directory in the `SYS:SYSTEM` directory with the proper name *<queue_ID>*`.qdr`. Use the appropriate utility (`NWADMIN.EXE` or `SYSCON.EXE`) to retrieve the *<queue_ID>*`.qdr` from NDS/eDirectory.
- NSS volumes on Novell NetWare 5.1 support files up to 4 GB. Other Novell NetWare systems do not have file size limitations.
- You cannot restore Novell NetWare files to non-NetWare platforms.

Restoring File Ownerships and Trustees

Data Protector restores owner and trustee information on a per-file/directory basis. The owner and trustees of the file or directory are restored correctly if the relevant objects exist in the NDS/eDirectory database.

At restore time, select `Trustee only restore` and the appropriate Trustee Conflict Handling option in the `Destination` page of the `Restore` context.

Restoring the Novell NetWare CONFIGURATION

Data Protector enables you to restore the special data structure known as `CONFIGURATION`, which consists of the following components:

CONFIGURATION Components

- Server Specific Info
- Schema
- Root

To restore only the `Server Specific Info` object of `CONFIGURATION`, refer to “Restoring Server Specific Info Object” on page 359. For restore of other specific components of `CONFIGURATION`, refer to the Novell NetWare documentation.

Restoring Server Specific Info Object

If you want to restore only the `Server Specific Info` object, select only the `Server Specific Info` object for restore in the Data Protector GUI. Refer to Figure 8-7 on page 360.

Figure 8-7 **Restoring Server Specific Info**

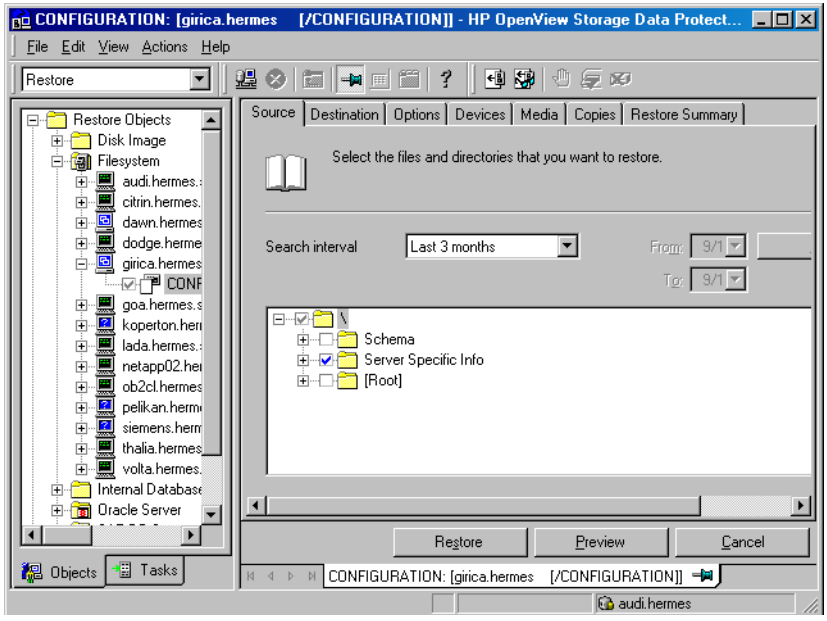

Then follow the standard restore procedure and set the following restore options:

- Select the `Overwrite` option.
- Deselect the `Omit unrequired incrementals` option.

Refer to “Standard Restore Procedure” on page 337. For more information on these restore options, refer to “Restore Options” on page 365.

`Server Specific Info` is restored into the `SYS:SYSTEM\`*`<server_name>`* directory.

Restoring Novell NDS/eDirectory

Prerequisites

The prerequisites for performing a successful restore are the same as for a backup of the NDS/eDirectory database. Data Protector restores NDS/eDirectory objects in the same way as Novell NetWare filesystem data, except in the following cases:

- NDS/eDirectory objects cannot be restored to other Novell NetWare volumes.
- Container and leaf objects (treated as directories by Data Protector) cannot be restored into other container objects or as other container objects.

Restoring the NDS/eDirectory does not affect the current partitioning and replication in the NDS/eDirectory tree. If partitions and replicas exist when NDS/eDirectory information is restored, those partitions and replicas are fully utilized. If partition information does not exist at restore time, the entire tree structure is placed in one partition.

NOTE Data Protector does not restore the NDS/eDirectory partitions and replica information. Partitions and replicas have to be manually reestablished.

For Novell NDS/eDirectory restore you can specify how to handle conflicts between the version currently on the disk and the backup version of a file. `File Conflict Handling` offers you three possible options: `Keep most recent`, `No overwrite`, and `Overwrite`. For more information on these options, refer to "Restore Options" on page 365.

Restoring NDS/eDirectory Schema and NDS/eDirectory Objects

Data Protector allows single NDS/eDirectory object restore. Within a Data Protector restore session, it is possible to:

- Restore the trees of the NDS/eDirectory using the `-trees` option
- Exclude a subtree of the NDS/eDirectory using the `-exclude` option
- Skip NDS/eDirectory objects using the `-skip` option
- Overwrite existing NDS/eDirectory objects using the `-overwrite` option

Troubleshooting

Sometimes an NDS/eDirectory restore session is completed successfully but some of the objects are not correctly restored and are marked as unknown. This happens when the NDS/eDirectory container object is deleted from NDS/eDirectory after the backup session. To solve this problem, restore this object again using the `-overwrite` option.

Restoring OpenVMS Filesystems

Use the standard restore procedure to restore OpenVMS filesystems. Refer to “Standard Restore Procedure” on page 337.

What Is Restored?

The directory structure and the files are restored, together with the following filesystem information:

- The directory and file attributes.
- ACL (Access Control List) if available (see Limitations below).
- Secondary file entries.

  Files with multiple directory entries are backed up once using the primary path name. Secondary path entries are saved as soft links. During a restore, these extra path entries are restored. Refer to “Limitations” in “Backup Specification Configuration Procedure” on page 249.

Files can be restored to mounted FILES-11 ODS-2 or ODS-5 volumes only.

Limitations

- For files and directories saved on any other operating system platform not all file attributes are restored and no ACL will be restored in this case.
- Directories that are created during a restore but have not been included in a save will get the attributes of the first file restored in the directory unless disabled by the `-no_protection` option.
- Any file specifications that are entered into the GUI or passed to the CLI must be in UNIX style syntax:

  ```
  /disk/directory1/directory2/filename.ext.n
  ```

  — The string should begin with a slash, followed by the disk, directories, and filename, separated by slashes.

  — Do not place a colon after the disk name.

  — A period should be used before the version number instead of a semi-colon.

— File specifications for OpenVMS files are case insensitive.

For example:

An OpenVMS file specification of:

```
$1$DGA100:[USERS.DOE]LOGIN.COM;1
```

must be specified in the form:

```
/$1$DGA100/Users/Doe/Login.Com.1
```

- There is no implicit version number. You always have to specify a version number. Only file versions selected for the backup will be restored. If you wish to include all versions of the file, select them all in the GUI window, or, using the CLI, include the file specifications under the `Only (-only)` option, including wildcard characters for the version number, as follows

  ```
  /DKA1/dir1/filename.txt.*
  ```

- If you restore to a location other that the original location, only the disk device and starting directory are changed. The original directory path is added to the destination path to form the new restore location.
- If the `Restore Time Attributes (-notouch)` option is disabled during a restore, the last accessed date will be updated with the current date and time on ODS-5 disks. On ODS-2 disks, the original dates will be set on the files.
- A file saved as a soft link will be restored using the equivalent of a "DCL SET FILE/ENTER" command. No data will be restored in this case. The soft link entered points to the primary path/filename for this file from the time the file was saved. If the primary path/filename does not exist or was not restored, the creation of the soft link will fail.

  To make a restored copy of an OpenVMS system disk bootable, the OpenVMS WRITEBOOT utility has to be used to write a boot block after the disk has been restored.

- The `Move Busy Files (-move)` and `Restore Sparse Files (-sparse)` options are not available on OpenVMS.
- Files backed up from an ODS-5 disk on an OpenVMS system that have extended filesystem names (i.e. upper and lower case letters, Unicode characters, etc.) may not be restored to an ODS-2 disk.
- Files being restored are always locked regardless of whether the `Lock`

`Files during Restore (-lock)` option is enabled or disabled.

- The default device and directory for pre- and post-exec command procedures is `/omni$root/bin`. To place the command procedure anywhere else the file specification must contain the device and directory path in UNIX style format: For example:`/SYS$MANAGER/DP_SAVE1.COM`
- If the `Restore Protection Attributes (-no_protection)` option is disabled, the files are created with the default owner, protection and ACL.
- When specifying wildcard characters for `Skip (-skip)` or `Only (-only)` filters, use '*' for multiple characters and '?' for single characters.

Restore Options

Data Protector offers a set of comprehensive restore options that allow fine-tuning of a restore. All these options have default values which are appropriate in most cases.

Restore options depend on the data being restored. For example, restore options for a filesystem are different from those for a disk image restore.

List of Restore Options

The following list of restore options can be set for a particular object. They apply to all the data restored from the backed up object.

General Restore Options

Target Client By default, you restore to the same client system from which the data was backed up. You can select another system in your cell from the drop-down list. The Disk Agent is started on the selected client system and the data is restored there. You need to have the `Restore to other clients` user right to be able to restore to another client system.

Omit Deleted Files By default, this option is disabled. When restoring a directory from which files were deleted between a full and an incremental backup, these files are also restored.

If this option is enabled, Data Protector attempts to recreate the state of the restored directory tree as it was when the last incremental backup was run, while preserving files that were created or modified after the last incremental backup. The files that were deleted between a full and incremental backup are restored and then deleted.

CAUTION

When this option is enabled, if the directory contains files that did not exist there at the time of the last incremental backup, but their modification time is older than the time of the incremental backup, Data Protector will delete these files as well.

When using the `Restore As` or `Restore Into` functionality, be careful when selecting the new location to prevent accidental deletion of existing files.

The time on the Cell Manager and Data Protector clients must be synchronized for this option to function properly.

Move Busy Files This option is relevant if a file on the disk is being used by an application when a restore wants to replace this file. The option is used with the `Keep most recent` or `Overwrite` options. By default, this option is disabled.

On UNIX systems, Data Protector moves the busy file *`filename`* to `#`*`filename`* (adds a hash in front of the filename). The application will keep using the busy file until it closes the file. Subsequently, the restored file is used.

On Windows systems, the file is restored as *`filename.001`*. All applications keep using the old file. When the system is rebooted, the old file is replaced with the restored file.

List Restored Data When this option is enabled, Data Protector displays the names of the files and directories in the monitor window as the objects are being restored. By default, this option is disabled.

Display Statistical Information When this option is enabled, Data Protector reports statistical information (such as size and performance) for each object that is restored. You can view the information in the monitor window. By default, this option is disabled.

Omit Unrequired Incrementals This option enables repositioning within a medium when restoring individual files of a specific object. A Media Agent restores a specific item, repositions itself directly on the next requested item, and continues the restore. This improves restore performance when restoring multiple single files. Note that several Disk Agents may be started per object. Disable this option if you intend to restore empty directories. By default, this option is enabled.

Restore Sparse Files This option restores sparse files in their original compressed form. This is important because sparse files can consume additional disk space unless they are restored in their original form. By default, this option is disabled.

This option applies to UNIX sparse files only. Windows sparse files are always restored as sparse.

Lock Files During Restore This option denies access to files during the restore. By default, this option is disabled.

Restore Time Attributes This option preserves the time attribute values of each restored file. When this option is disabled, Data Protector sets the time attributes of the restored objects to the current date and time. By default, this option is enabled.

Restore Protection Attributes This option preserves the original protection attributes of each restored file. If this option is disabled, Data Protector applies the protection attributes of the current restore session. By default, this option is enabled.

On Windows systems, this option applies to file attributes only. Security information is always restored, even when this option is disabled.

Pre- and Post-Exec Commands

For general information on `pre-` and `post-exec` commands, refer to "Pre- and Post-Exec Commands" on page 297. For examples of these commands on UNIX, refer to "Examples of Pre-Exec and Post-Exec Commands for UNIX" on page A-21. Note that `pre-` and `post-exec` commands are executed before and after the restore of each object, and not the entire restore session.

Pre-Exec This option allows you to enter a command to be executed before the restore of each object is initiated. This command must return success for Data Protector to proceed with the restore. The `pre-exec` command is executed on the client system where the Disk Agent is running. On how to specify the command, refer to online Help.

Post-Exec This option allows you to enter a command to be executed after the restore of each object is completed. The `post-exec` command is executed on the client system where the Disk Agent is running. On how to specify the command, refer to online Help.

File Conflict Handling Options

Keep Most Recent If this option is selected, the most recent versions of files are kept. If a file on the disk is newer than the backed up version, the file is not restored. If a file on the disk is older than the backed up version, the file is overwritten with the newer version from the backup. By default, this option is enabled.

No Overwrite If this option is selected, files that exist on the disk are preserved. This means that they are not overwritten by other versions of these files from the backup. Only non-existing files are restored from the backup. By default, this option is disabled.

Overwrite If this option is selected, existing files on the disk are replaced with files from the backup. By default, this option is disabled.

Active Directory Specific Options

Authoritative The Active Directory database is *not* updated after the restore, and the restored data overwrites the existing data in the target destination. An authoritative restore can only be performed by running `ntdsutil.exe` from the command prompt after the restore session has finished.

Nonauthoritative The `Nonauthoritative` replication mode is the default option. The Active Directory database is updated after the restore using standard replication techniques.

Primary The `Primary` replication mode allows you to keep the Directory Service online, and is used when you restore `FileReplicationService` along with the Active Directory service. This option must be used when all replication partners for a replicated share have been lost. With regard to the Certificate Server and the Active Directory Server, `Primary` is the same as `Authoritative`.

User Defined Restore Variables

You can use variables (a variable name and its value) for flexible operations on some platforms and integrations with Data Protector. For detailed steps, refer to the online Help index keyword “setting user definable restore variables”.

Figure 8-8 **User Defined Restore Variables**

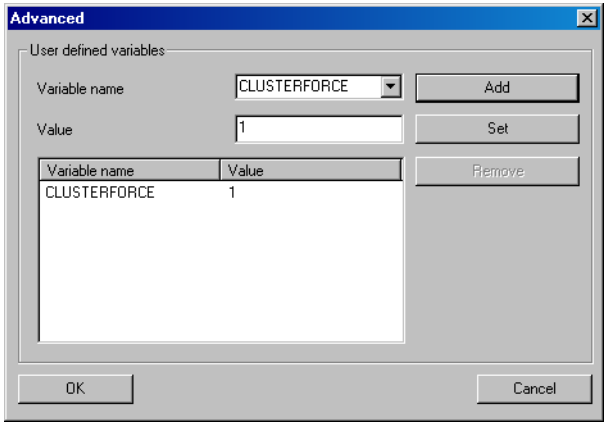

Restore Techniques

The following restore techniques apply to the UNIX and Windows platforms.

Restoring Files to Different Paths

By default, Data Protector restores data to the same client and directory from which it was backed up. You can restore your data to a different client system and directory. For individual files and directories, you can specify a different path and different name.

Different Location for an Object

In the `Destination` page of your restore, you can specify a different restore location for an object selected for restore:

- With appropriate user rights you can restore to a different client system by selecting the client system in the `Target client` drop-down list. By default, Data Protector restores the object using the same directory structure. For example, if the object was originally backed up from the `C:\temp` directory on system A, it will restore the data to the `C:\temp` directory on system B.
- You can restore to a different directory by selecting the `Restore to new location` option, and then entering or browsing for a new path in the text box. The original path is appended to the new one. For example, if data was backed up from the `C:\sound\songs` directory and you enter `\users\bing` as the new path, the data is restored to the `C:\users\bing\sound\songs` directory.

For detailed steps, refer to the online Help index keyword “location options for restore”.

Different Location for Individual Files

The individual location specified under the `Restore As/Into` option overrides the default destination specified in the `Destination` property page.

You can restore individual files and directories to different paths and under a different name using the `Restore As/Into` option available from the `Source` property page of your restore.

This capability is available for the initially selected tree node (directory) and for tree nodes that are not hierarchically dependent on any already selected tree nodes. A selected tree node is indicated by a blue check mark, and a dependent tree node is indicated by a black check mark.

`Restore Into` appends the source path to the new one entered under `Location`. For example, if the `colors.mp3` file was backed up from the `C:\sound\songs` directory and you enter `\users\bing` as the new path, the file is restored to the `C:\users\bing\sound\songs` directory.

`Restore As` replaces the source path with the one entered under `Location`. The destination path can be a new directory or an existing one. You can rename the files and directories as you restore them. For example, if the `colors.mp3` file was backed up from the `C:\sound\songs` directory and you enter `\users\bing\colors.mp3` as the new path, the file is restored to the `C:\users\bing` directory.

CAUTION

Consider the risk of deleting data with the `Overwrite` option enabled when:

- Specifying restore under a name that already exists
- Entering an existing path without specifying the file or directory name

For example, when you enter the new path `\users\bing` in the `Location` text box when restoring the file `colors.mp3`, but you do not enter the name of the file, then the `colors.mp3` file will be restored as `bing`. What used to be the `bing` directory is deleted and substituted with the restored file.

Restoring Files in Parallel

What Is Parallel Restore?

Parallel restore lets you restore files to multiple disks at the same time, assuming that the disks have been backed up to the same device using a concurrency higher than 1. This improves the speed of the restore. This behavior is complementary to a parallel backup, where files from multiple disks are backed up concurrently to the same device.

Figure 8-9 Restoring Files in Parallel

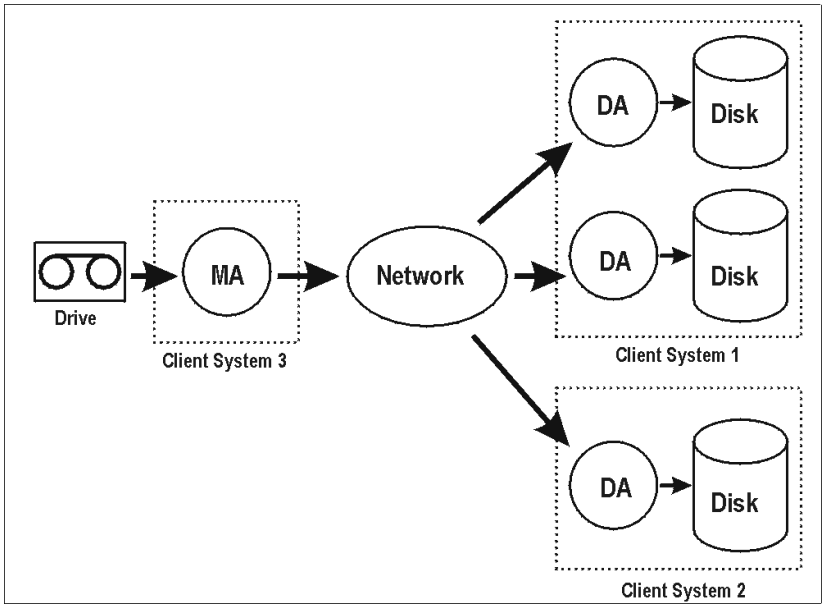

The figure shows an example of restoring files in parallel from one medium. Each object uses a different DA.

How to Run a Parallel Restore

Select the data that you want to restore to different disks and start the restore. Data Protector asks you if you want to perform a parallel or single restore. Choosing parallel restore enables multiple Data Protector Disk Agents to run in parallel. Refer to See “Selecting Your Data for Restore” on page 338.

Viewing Files Not in the IDB

Data Protector allows you to view and restore data directly from backup media even though the information about this data is no longer in the IDB.

When to Restore Directly from Media

The following must apply in this case:

- You have removed information about backed up data or media from the IDB.

- The catalog protection has expired. Refer to “Most Frequently Used Backup Options” on page 271 for more information about data and catalog protection.
- The media are not from the same Data Protector cell and, as such, are not recognized in the IDB of the cell. In this case, you need to import it first.

Prerequisite

A large amount of memory on the Cell Manager is required. The amount of memory needed can be estimated by using the following formula: *number_of_files* multiplied by 200 bytes.

Limitations

- You cannot list database application objects from the media.
- Files that span several media cannot be restored directly fro media. All media needed to restore the file have to be imported, and then the file can be restored using the `List From Database` option.

How to Restore Directly from Media

To restore data directly from media, click `List From Media` in the `Actions` menu of your restore context, and follow the `Restore from media` wizard. For detailed steps, refer to the online Help index keyword “restoring directly from media”.

Restoring Files in Use

Data Protector allows you to back up and restore files, such as databases or word processing documents, that are in use (opened) by other applications.

Data Protector provides restore options that allow you to specify the behavior when files being restored are in use by setting the **Lock Files During Restore** and the **Move Busy Files** options. Refer to “Restore Options” on page 365.

Restoring by Query

If you do not know the full path of a file that you want to restore, you can search for the file in the IDB, provided that the logging level at backup time was set to `Log files` or `Log all`. You can search for files and directories using the `Restore by Query` task if you know at least a part of the file name.

Figure 8-10 Restore by Query

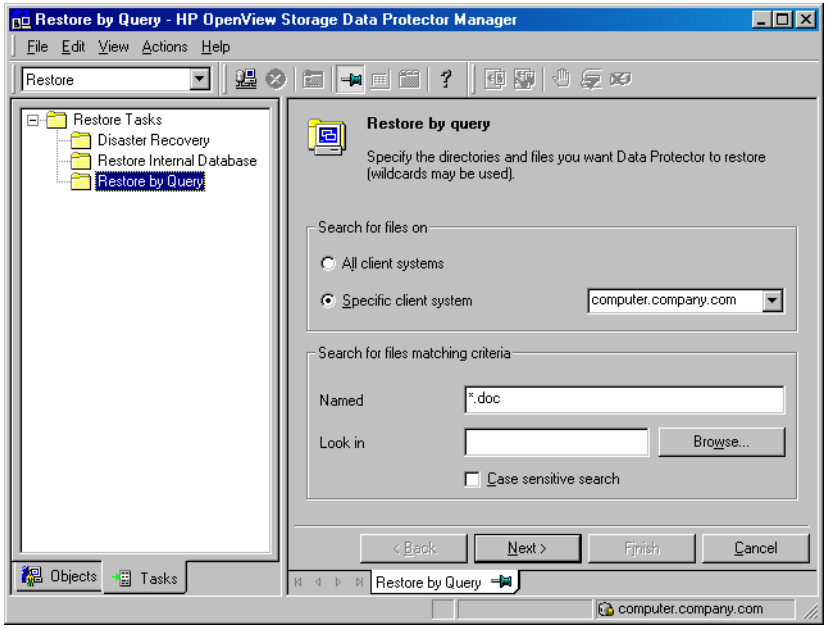

How to Restore by Query

Start the `Restore by Query` task from the `Restore` context of the `Data Protector Manager`. Use the `Tasks` navigation tab. See Figure 8-10. For detailed steps, refer to the online Help index keyword “restore by query”.

When specifying filenames containing non-ASCII characters using the `Restore by Query` feature, replace non-ASCII characters with wildcard characters, unless you operate in a pure Windows environment.

Skipping Files for Restore

Data Protector allows you to skip certain files during restore. By using wildcard characters (`*` or `?`), you can skip files matching specific criteria. For example, entering `*.exe` skips the files that end in `.exe`.

How to Skip Files for Restore

In the `Source` property page of your restore, select the tree node to be restored and right-click it to open its properties. In the `Skip` property page, specify the criteria to match the files to be skipped. For detailed steps, refer to the online Help index keyword “skipping files”.

Figure 8-11 Skipping Files for Restore

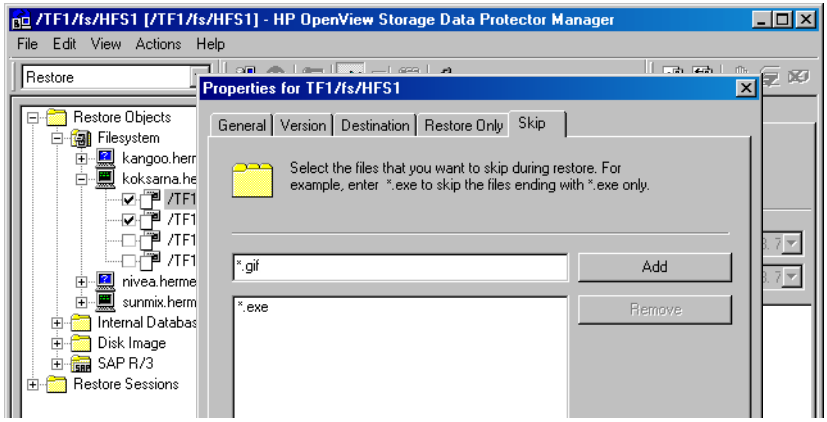

Selecting Only Specific Files (Matching) for Restore

Data Protector allows you to restore only specific files. By using wildcard characters (`*` or `?`), you can restore files matching specific criteria. For example, entering `*.exe` restores only the files that end in `.exe`.

How to Match Files for Restore

In the `Source` property page of your restore, select the tree node to be restored and right-click it to open its properties. In the `Restore Only` property page, specify the criteria to match the files to be restored. For detailed steps, refer to the online Help index keyword “selecting only specific files for restore”.

Figure 8-12 Matching Files for Restore

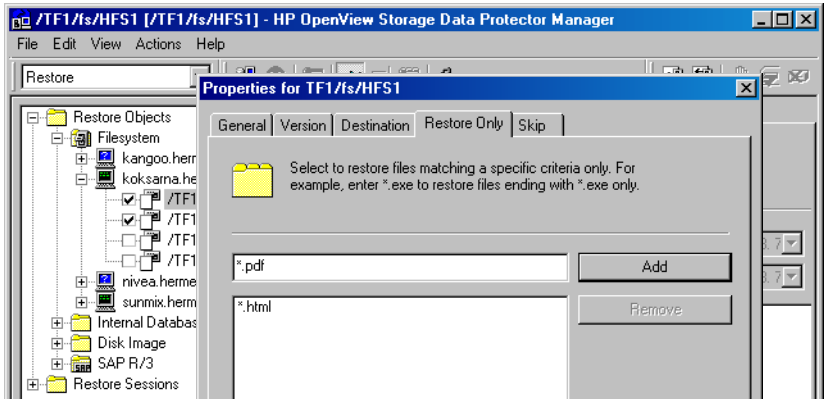

Restoring Files and Directories Manually

You need to restore a file or a directory manually when you can no longer browse for the file or directory. This happens when the catalog protection for your data has expired, or when backup was done using the `No log` option.

Prerequisite

To add a file or a directory manually, you need to know the exact path and the name of the file or the directory. The file and path names are case sensitive.

How to Add Files and Directories Manually

In the `Restore Summary` page of your restore, write the exact path and name of the file or the directory, and then click `Add`. For detailed steps, refer to the online Help index keyword “manually restoring files or directories”.

Selecting the Media Set to Restore From

If an object version that you want to restore exists on more than one media set, which has been created using one of the Data Protector data duplication methods, any of the media sets can be used for the restore. By default, Data Protector automatically selects the media set that will be used. You can influence the media set selection by specifying the media location priority.

NOTE For media location priority to take effect, you need to specify the location of each medium. This can be done for individual or multiple media.

You can also manually select the media set you want to use for the restore, except when restoring integration objects. If a medium needed for the restore is unavailable, a mount request will be issued. For detailed steps, refer to the online Help index keyword “finding media for restore”.

NOTE When restoring data, copies obtained using the media copy functionality are not listed as needed media in the `Media` tab and their location priority is not considered. Such a copy is used only if the original medium (the medium that was used as a source for copying) is unavailable or unusable.

Media Location Priority

To influence the selection of the media set, set the media location priority. Data Protector will use the media set with the highest priority (priority `1` is the highest, priority `None` is the lowest) if more than one media set matches the conditions of the media set selection algorithm. For more information on the algorithm, refer to the *HP OpenView Storage Data Protector Concepts Guide*.

There are two levels of specifying the media location priority:

- You can set it globally, in the `Devices & Media` context. The priority will be considered in all restore sessions, unless it is overridden on the restore session level. See Figure 8-13 on page 378.

  For detailed steps, refer to the online Help index keyword “setting media location priority”.

- You can set it for a specific restore session in the `Restore` context. It will override the priority setting on the global level. See Figure 8-14 on page 378.

  For detailed steps, refer to the online Help index keyword “finding media for restore”.

Figure 8-13 Specifying the Media Location Priority Globally

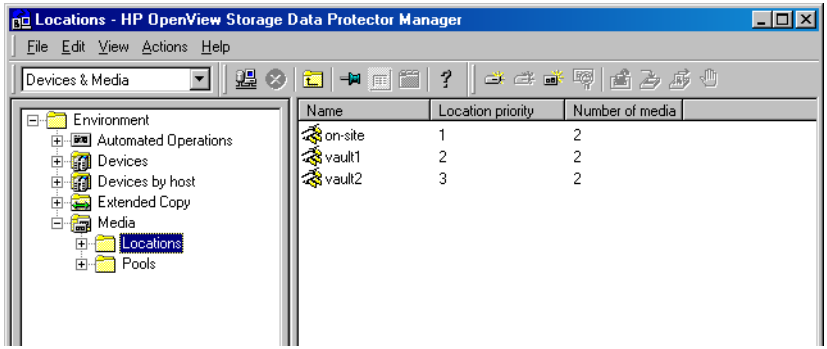

Figure 8-14 Specifying the Media Location Priority per Restore Session

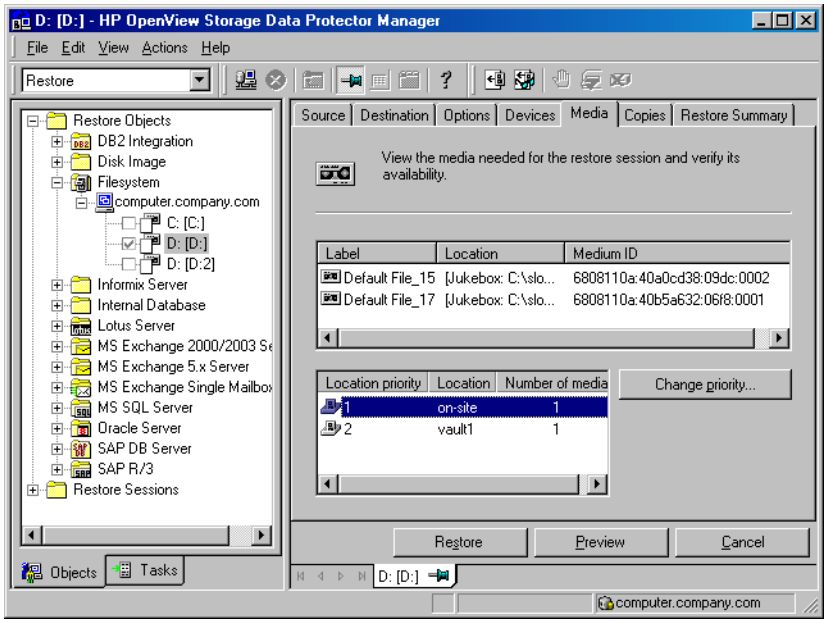

9 Monitoring, Reporting, Notifications, and the Event Log

In This Chapter

This chapter consists of the following sections:

You can monitor several cells at the same time using the Manager-of-Managers functionality. See Chapter 10, "Manager-of-Managers Environment," on page 433 for more information.

If you do not have access to the Data Protector user interface, you can still view reports and set notifications using your Web browser. See "Configuring Reports and Notifications on the Web" on page 426 for information on how to do this.

Monitoring Sessions

Data Protector enables you to manage running sessions and to respond to mount requests. You can view the status of sessions, their type, owner, session ID, and start time, as well as the names of the corresponding backup specifications.

When you run an interactive backup, restore, or media management session, a monitor window is displayed, showing the objects and backup devices used, and the messages generated during the session. Note that even if the user interface is closed, the session continues.

You can change the level of reported messages during a backup or restore session by changing the Report level option when configuring a backup specification or when starting a restore session.

NOTE Only the Data Protector users in the `Admin` group and those granted the `Monitor` user rights are given access to the Data Protector monitoring functionality.

Monitoring Current Sessions

You can monitor currently running sessions in the `Monitor` context of the Data Protector GUI.

NOTE A currently running session is displayed in the `Monitor` context after the pre-exec script has finished.

Refresh Interval

At refresh intervals (by default, every 5 seconds), the list of currently running sessions is automatically updated with new sessions if there are any. To change the default refresh interval using the Data Protector GUI, proceed as follows:

1. In the `File` menu, click `Preferences`.

   The `Preferences` dialog box is displayed.

2. Click the `Monitor` tab.
3. Specify the refresh interval in seconds for the Cell Manager and for the Manager-of-Managers (MoM).

To monitor a currently running session using the Data Protector GUI, proceed as follows:

1. In the Context List, click `Monitor`.

   In the Results Area, all currently running sessions are listed.
2. Double-click the session you want to monitor. See Figure 9-1 on page 382.

Figure 9-1 Monitoring a Current Session

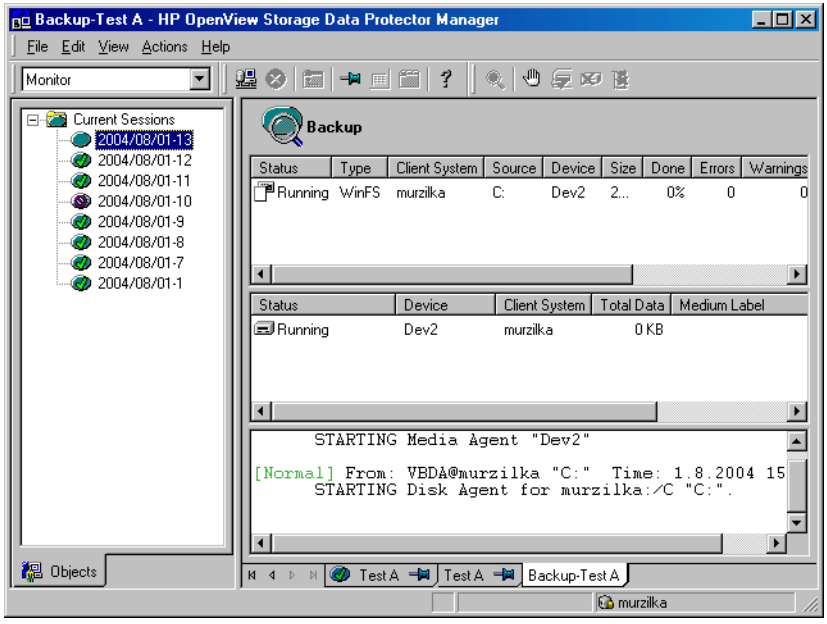

Clearing Sessions

To remove all completed or aborted sessions from the Results Area of the Monitor context, proceed as follows:

1. In the Scoping Pane, click `Current Sessions`.
2. In the `Actions` menu, select `Clear Sessions`. Or click the `Clear Sessions` icon on the toolbar.

To remove a particular completed or aborted session from the current sessions list, right-click the session and select `Remove From List`.

NOTE

All completed or aborted sessions are automatically removed from the Results Area of the `Monitor` context if you restart the Data Protector GUI.

For detailed information on a completed or aborted session, see "Viewing Previous Sessions".

Viewing Previous Sessions

To view a previous session using the Data Protector GUI, proceed as follows:

1. In the Context List, click `Internal Database`.

   MoM

   If you are running the Manager-of-Managers, select `Monitor` in the Context List, and then select a Cell Manager of your choice. From the Tools menu, select `Database Administration` to open a new Data Protector GUI with the `Internal Database` context selected.

2. In the Scoping Pane, expand `Sessions` to display all the sessions stored in the IDB.

   The sessions are sorted by date. Each session is identified by a session ID consisting of a date in the YY/MM/DD format and a unique number.

3. Right-click the session and select `Properties` to view details on the session.

4. Click the `General`, `Messages` or `Media` tab to display general information on the session, session messages, or information on the media used for this session, respectively. See Figure 9-2 on page 384.

Figure 9-2 Viewing a Finished Session

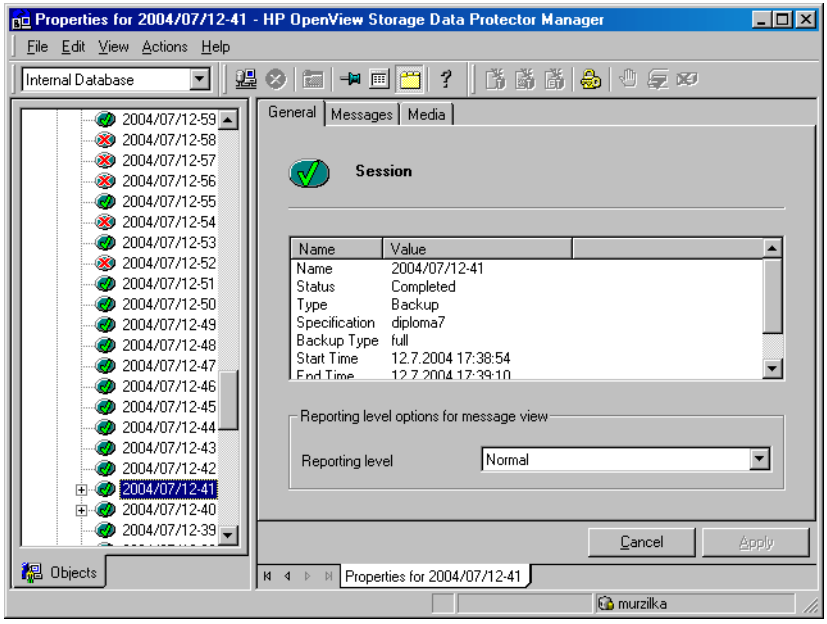

Responding to Mount Requests

Data Protector issues a mount request in the following cases:

- The end of the currently used medium has been reached and Data Protector needs a free medium.
- A mail slot is open. In this case, shut the mail slot.

You respond to a mount request to confirm that the needed medium is in a device. Use the following procedure to respond to the mount request while monitoring the session:

1. In the Context List, select `Monitor`.
2. Insert the needed medium into the device. If you have a library device, it is not necessary to use the slot requested by the mount request.
3. In the Results Area, double-click the session with the mount request status to display details about the session.

4. Select the device with the mount request status.
5. In the Actions menu, click Confirm Mount Request. The status of the session and device changes to Running.

TIP You can also right-click the device with the mount request status and select Confirm Mount Request.

Restarting Failed Backups

During backup, some systems may not be available because they were shut down, there were some networking problems, and so on. This results in some systems not being backed up or being backed up just partially - some objects failed.

This section gives you detailed instructions on how to restart failed backup sessions. For more information on how to manage failed backups, see “Managing Failed Backups” on page 311.

You cannot restart failed sessions that are the result of an unsaved backup specification.

After you have resolved the related problems, restart a failed session, as follows:

1. In the Data Protector Manager, switch to the Internal Database context.

   If you are running the Manager-of-Managers, select Clients in the Context List, and then expand Enterprise Clients. Select a Cell Manager with the failed backup. From the Tools menu, select Database Administration to open a new Data Protector window with the Internal Database context displayed.
2. Under the Internal Database item, expand the Sessions item.
3. In the Results Area, search for your backup.

   You can sort your sessions using the buttons on the top of each of the columns.
4. Right-click on your failed session, and then select Restart Failed Object.
5. Click Yes to confirm.

Aborting Running Sessions

You can abort a session if you want to stop a backup, restore, or media management operation. A backup copy or restored data will exist only for data that was backed up or restored before you aborted the session.

1. In the Context List, click `Monitor`. The progress and status of current sessions appear in the Results Area.

   If you are running the Manager-of-Managers, expand the `Enterprise Monitor` in the Scoping Pane, and then select the Cell Manager you want to monitor. The progress and status of current sessions appear in the Results Area.

2. Click the column headings to sort the sessions.
3. Right-click the session that you wish to abort and select `Abort`.

If you abort a backup session while it is still determining the sizes of the disks that you have selected for the backup, it does not abort immediately. The backup is aborted once the size determination (treewalk) is completed.

TIP

If you started a backup, restore, or media management session interactively, you can also abort the session in the Data Protector `Backup`, `Restore`, or `Devices & Media` context respectively.

Changing the Amount of Messages Shown

You can change the level of reported messages for backup and restore sessions by changing the Backup and Restore options.

See “Using Backup Options” on page 269 for information on which backup options affect your displayed messages.

See “Restore Options” on page 365 for information on which restore options affect your displayed messages.

Monitoring Several Cells Simultaneously

You can monitor several cells at the same time using the Manager-of-Managers functionality.

See Chapter 10, “Manager-of-Managers Environment,” on page 433 for more information.

Data Protector Reporting

What Is Reporting?

Data Protector reports provide various information on your backup environment. For example, you can check the status of the last backup, check which systems in your network are not configured for backup, check the status of devices, and more.

Data Protector reporting represents a powerful, customizable, and flexible tool for managing and planning your backup environment.

You can configure reports and report groups using the Data Protector GUI or any Web browser with Java support.

NOTE

Only the Data Protector users in the Admin group and those granted the `Reporting and notifications` user rights are given access to the Data Protector reporting functionality.

Prerequisite

The Data Protector user under whose account the CRS service is running should not be removed. This user is configured by default at installation time. On a Windows Cell Manager, this is the user under whose account the installation was performed. On a UNIX Cell Manager, this is the root user of the Cell Manager.

Report Groups

You can gather various reports in a report group, which can be scheduled, started interactively, or triggered by a notification.

Starting Reports

Reports can be started using the Data Protector GUI, the Data Protector command-line interface, the Data Protector Web reporting interface, the Data Protector scheduler, a notification event, or a post-exec script that includes a Data Protector command-line interface command.

Reports on Multiple Cells

Reporting is also available for a multiple cell configuration when you use the Manager-of-Managers functionality.

Report Parameters

Reports can be customized by configuring optional input parameters (optional selections). Some input parameters allow multiple selections.

If no optional input parameters (optional selections) are specified when configuring a report, a default value is set, which is *<all>* in case of objects and *<no time limit>* in case of time frames.

To configure a report or report group, you need to provide the following information:

- name of the report
- type of report
- send method
- recipient(s)
- format

All other input parameters (selections) depend on the type of the report.

Report Formats

Output of the reports is provided in various formats and optionally displays input parameters (selections), too. Refer to “Report Formats” on page 402.

Report Send Methods

Reports can be sent using various methods. Refer to “Report Send Methods” on page 404.

Report Types

Data Protector provides various types of reports, as shown in Table 9-1:

Table 9-1

Backup Specifications	Provides information on backups, such as average size of backed up objects, schedule of backups, filesystems not configured for backup, and so on.
Configuration	Provides information on the configuration of the Data Protector cell, on devices not configured for backup, on systems not configured for backup, and so on.
IDB	Provides information on the size of the IDB and on the results of the IDB purge sessions.
Pools and Media	Provides information on media pools and used media.
Sessions in Time Frame	Provides information on backup sessions that have run in a specified period of time.
Single Session	Provides detailed information on a specific session.

Backup Specification Reports

The following table lists the Backup specification reports. Backup specification reports provide information on backups, such as average size of backed up objects, schedule of backups, filesystems not configured for backup, and so on.

For supported formats, refer to “Report Formats” on page 402.

Table 9-2 Backup Specification Reports

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Trees in Backup Specification `dl_trees`	Lists all trees in the specified backup specification. It also shows names of drives and the name of a tree.	none	• Backup Specifications • Backup Specification Group	all formats
Objects without Backup `obj_nobackup`	Lists all objects that are part of a backup specification and do not have a valid backup (successfully completed backup, the protection has not yet expired).[a]	none	• Backup Specifications • Backup Specification Group • Number of Days[b]	all formats
Object’s Latest Backup `obj_lastbackup`	Lists all objects for each specified backup specification, together with the last full and the last incremental backup time.	none	• Backup Specifications • Backup Specification Group • Number of Days[b]	all formats
Average Object Size `obj_avesize`	Displays the average size of an object in the specified backup specification. It displays the size of the full and the incremental backup of the object.	none	• Backup Specifications • Backup Specification Group • Number of Days[b]	all formats

Table 9-2 Backup Specification Reports

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Not Configured Filesystems `fs_not_conf`	Lists all disks (filesystems) that are not configured in any of the selected backup specifications.	none	• Backup Specifications • Backup Specification Group	all formats
Backup Specification Information `dl_info`	Shows the following information: backup specification name, type, group, owner, and pre & post exec commands for all specified backup specifications.	none	• Backup Specifications • Backup Specification Group	all formats
Backup Specification Schedule `dl_sched`	Lists the next backup time for each specified backup specification.	none	• Backup Specifications • Backup Specification Group	all formats

a. N/A for backup specifications for integrations.
b. Counted from the moment of starting the report backwards.

Configuration Reports

The following table lists the Configuration reports. Configuration reports provide information on the configuration of the Data Protector cell, devices, systems not configured for backup, and so on. For supported formats, refer to “Report Formats” on page 402.

Table 9-3 Configuration Reports

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Cell Information `cell_info`	Lists Data Protector cell related information (number of clients, backup specifications, media management server, licensing server).	none	none	all formats
Configured Clients Not Used by Data Protector `hosts_unused`	Lists all configured clients that are not used for backup and do not have any device configured.	none	none	all formats
Configured Devices Not Used by Data Protector `dev_unused`	Lists configured devices that are not used for backup at all.	none	none	all formats
Look up Schedule `lookup_sched`	Lists backup specifications that are scheduled to start in the next specified number of days.	Number of Days	none	all formats
Clients Not Configured for Data Protector `hosts_not_conf`	Lists clients in the selected domains that are not part of the current cell.	Network Ranges	none	all formats
Licensing `licensing`	Lists all licenses with their total and available amount.	none	none	all formats

Table 9-3 Configuration Reports

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Client Backup `host`	Lists information about the specified clients such as: filesystems not configured, all objects, and all objects with a valid backup. Reports also list times and average sizes.	Host Name	none	all formats

IDB Reports

The following table lists the IDB reports. IDB reports provide information on the size of the IDB and on the results of the IDB purge sessions. For supported formats, refer to “Report Formats” on page 402.

Table 9-4 IDB Reports

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
IDB Size `db_size`	Provides a table that contains information about the Media Management DB, Catalog DB, DB extension files, statistics for DC binary files, SMBF, and SIBF and low DB disk space.	none	none	all formats
IDB Purge `db_purge`	Lists all purged sessions together with the following information: start time, end time, duration, inactivity time, and number of the file name records and the amount of Mb read.	none	none	all formats

Table 9-4 IDB Reports

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Purge Preview `db_purge_preview`	Lists the following information: overall number of filenames in database (in thousands), estimated number of obsolete filenames in database (in thousands) and estimated duration of database purge (in seconds).	none	none	all formats

Table 9-4 IDB Reports

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
System Dynamics `db_system`	Lists for each Data Protector client in the cell: the number of filenames (in thousands) in the IDB, the number of active filenames (in thousands) in the IDB, the IDB filenames growing ratio (new filenames per day), the number of deleted filenames in the IDB per day, active growth per year, and a dynamics indicator (medium/high/low/critical). The filenames that are not active are filenames of the backed up files in the IDB that have no associated file versions in the IDB. The active growth per year is calculated in two ways: If there is no Data Protector database purge session recorded in the Data Protector database, the active growth per year is calculated on the basis of data in last 11 days and then extrapolated to one year. If there is a Data Protector database purge session recorded in the Data Protector database, the active growth per year is calculated on the basis of data in the time span since the last Data Protector database purge session and then extrapolated to one year.	none	none	all formats

Pools and Media Reports

The following table lists the Pools and Media reports. Pools and media reports provide information on media pools and used media. For supported formats, refer to “Report Formats” on page 402.

Table 9-5 Pools and Media Reports

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Extended List of Media `media_list_extended`	Lists all media matching the specified search criteria. For each medium, the following information is provided: medium ID, medium label, medium location, medium condition, medium protection, used and total space (MB), time when medium was last accessed, media pool and media type, and the backup specifications that have used this medium during the backup.	none	• Description • Locations • Pool names • Media Types (DDS, DLT and so forth) • Condition • Expiration[1] • Timeframe[2] • Library Devices	all formats
List of Pools `pool_list`	Lists all pools matching the specified search criteria. For each pool the following information is provided: pool name, description, media type, total number of media, number of full and appendable media containing protected data, number of free media containing no protected data, number of poor, fair and good media.	none	• Pool Names • Locations • Media Types (DDS, DLT, and so forth) • Library Devices • Timeframe[2]	all formats

Table 9-5 Pools and Media Reports

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Media Statistics `media_statistics`	Reports the statistics on the media matching the search criteria. The following information is provided: number of media; number of scratch media; number of protected, good, fair and poor media; number of appendable media; and total, used, and free space on media.	none	• Description • Locations • Poolnames • Media Types (DDS, DLT and so forth) • Status • Expiration[1] • Timeframe[2] • Library Devices	all formats
List of Media `media_list`	Lists all media matching the specified search criteria. For each medium, the following information is provided: medium ID, medium label, medium location, medium condition, medium protection, used and total space (MB), time when medium was last accessed, and media pool and media type.	none	• Description • Locations • Poolnames • Media Types (DDS, DLT and so forth) • Condition • Expiration[1] • Time frame[2] • Library Devices	all formats

1. The following are possible:

Don't care / Unprotected / Protected; the last with the following suboptions:

Number of remaining days in which the data protection will expire, counted from the moment of starting the report / Never

2. Timeframe in which the medium was used for a backup.

Relative time: the first parameter sets the starting point of the timeframe (number of hours counted from the moment of starting the report backwards), the second parameter sets the end point of the timeframe (number of hours counted from the starting point).

Absolute time: the first parameter sets the starting point of the timeframe (date), the second parameter sets the end point of the timeframe (date).

Sessions in Timeframe Reports

The following table lists the Data Protector Sessions in Timeframe reports. Sessions in Timeframe reports provide information on backup sessions that have run in a specific period of time. For supported formats, refer to "Report Formats" on page 402.

Table 9-6 Sessions in Timeframe Reports

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
List of Backup Sessions `list_sessions`	Lists all sessions in the specified timeframe.	TimeFrame[1]	• Backup Specifications • Backup Specification Group	all formats
Session Flow `session_flow`	Graphically presents the duration of each session for the specified timeframe. A flow chart of the backup sessions matching the search criteria is shown.	TimeFrame[1]	• Backup Specifications • Backup Specification Group	HTML

Table 9-6 Sessions in Timeframe Reports

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Device Flow `device_flow`	Graphically presents the usage of each medium. A flow chart of the backup sessions matching the search criteria is shown.	TimeFrame[1]	• Backup Specifications • Backup Specification Group	HTML
Used Media `used_media`	Lists media that have been used during the backup sessions in the specified timeframe, together with their statistics.	TimeFrame[1]	• Backup Specifications • Backup Specification Group	all formats
Client Statistics `host_statistics`	Lists clients and their backup status statistics. Only the clients that match the search criteria are listed.	TimeFrame[1]	• Backup Specifications • Backup Specification Group • Hostnames	all formats
Backup Statistics `backup_statistics`	Shows statistics about backup status in the selected timeframe.	TimeFrame[1]	• Backup Specifications • Backup Specification Group	all formats
Backup Errors `backup_errors`	Displays a list of messages that occurred during backup. The messages are grouped by client.	TimeFrame[1]	• Backup Specifications • Backup Specification Group • Hostnames • Message Level	all formats

Table 9-6 Sessions in Timeframe Reports

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Extended Report on Used Media `used_media_extended`	Provides extended information about all media that were used in the selected session.	TimeFrame[1]	• Backup Specifications • Backup Specification Group	all formats

1. Timeframe in which the medium was used for a backup.

Relative time: the first parameter sets the starting point of the timeframe (number of hours counted from the moment of starting the report backwards), the second parameter sets the end point of the timeframe (number of hours counted from the starting point).

Absolute time: the first parameter sets the starting point of the timeframe (date), the second parameter sets the end point of the timeframe (date).

Single Session Report

The following table lists the Data Protector Single Session Reports. For supported formats, refer to “Report Formats” on page 402.

Table 9-7 Single Session Reports

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Single Session `single_session`	Displays all relevant information about a single Data Protector backup session.	Session ID	Message Level	all formats
Session Objects `session_objects`	Lists all backup objects and their statistics that took part in a selected session.	Session ID	none	all formats

Table 9-7 Single Session Reports

Report and `omnirpt` Option	Description	Required Selections	Optional Selections	Supported Formats
Session per Client `session_hosts`	Provides information about each client that took part in the selected session. Using the `Generate multiple reports` option, this report can be split into smaller reports, one for each client.	Session ID	Message Level	all formats
Session Devices `session_devices`	Provides information about all devices that were used in the selected session.	Session ID	none	all formats
Session Media `session_media`	Provides information about all media that were used in the selected session.	Session ID	none	all formats

Report Formats

Data Protector reports can be produced in various formats.

If you view each report individually, the report is displayed in the Data Protector Manager and you do not have to choose the report format.

If you group reports into report groups so that you can send reports on a specific event or schedule the reports, you also need to specify the format and the recipients of each report.

The following is a list of report formats:

ASCII — A report is generated as plain text.

HTML — A report is generated in HTML format. This format is useful for viewing using a Web browser. For example, you can check if your systems have been backed up by clicking a link and viewing the report on the intranet.

Short	A report is generated as plain text, but in a short, summary form, showing the most important information. This is the suggested format for broadcast messages.
Tab	A report is generated with fields separated with tabs.

TIP

The Tab format is useful to import the reports into some other applications or scripts for further analysis, such as Microsoft Excel.

The following command creates a list of media used in the last 24 hours in a Microsoft Excel spreadsheet:

```
omnirpt -report used_media -timeframe 24 24 -log
used_media.xls -tab
```

Report Send Methods

Report Send Methods

Reports can be sent using various methods:

- Email send method
- Broadcast message send method
- SNMP send method
- External send method
- Log to file send method

The following sections describe specifics of each method.

E-mail Send Method

The e-mail send method enables you to send or receive an e-mail with the output of a report.

IMPORTANT

Due to security features of Microsoft Outlook, using the e-mail send method may cause the CRS service to stop responding. For details and solutions, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

Limitation

Due to the operating system limitations, international characters in localized E-mail notifications and reporting can be displayed incorrectly on UNIX if they are passed between systems using a different locale.

The display of HTML E-mail report on Windows depends on the E-mail client settings. Many email clients display the report as plain ASCII text. To ensure the report displays correctly as HTML, open it in a Web browser.

Prerequisite for Windows

To send an e-mail with the output of a report from a Windows system, you need to have a mail profile created. You can either use an existing profile or create a new mail profile, named `OmniBack`.

To use an existing mail profile, add the following line to the Data Protector `omnirc` file:

`OB2_MAPIPROFILE=<existing_MAPI_profile_name>`

For more information on the `omnirc` file, see "Using Omnirc Options" on page 615.

Creating a Data Protector Mail Profile

To create a new mail profile, named `OmniBack`, on the Windows 2000 system with Microsoft Outlook 2002 installed, perform the following steps:

1. In the Windows `Control Panel`, double-click the `Mail` icon.

   The `Mail Setup - Outlook` dialog box is displayed.

2. Click `Show Profiles`.

   The `Mail` dialog box is displayed.

3. Click `Add`.

   The `New Profile` dialog box is displayed.

4. Type `OmniBack` in the `Profile Name` text box and click `OK` to start the `E-Mail Accounts` wizard.

5. Select `Add a new e-mail account` and click `Next`.

6. In the `Server Type` page, select `Microsoft Exchange Server` and click `Next`.

7. In the `Exchange Server settings` page, type the name of the local Microsoft Exchange Server system, your username and click `Next`.

8. Click `Finish` to end the wizard.

   In the `Mail` dialog box, the `OmniBack` profile is listed among other profiles set on your system.

On UNIX systems, no additional configuration is needed.

Broadcast Message Send Method

The broadcast message send method allows you to send a broadcast message with the output of the report to specified systems.

Broadcast messages can be sent to Windows systems only, by specifying the system to which the broadcast message should be sent. Broadcast messages are limited in length, so the `short` format is preferred. The reports are limited to 1000 characters.

Log to File Send Method

The log to file send method allows you to post a log file with the output of the report to a specified file.

The log file is posted to the Cell Manager system. Specify the name of the file to which you want to post the report. The file will be overwritten if it exists.

SNMP Send Method

SNMP send method allows you to send an SNMP trap with the output of the report. The SNMP trap can be further processed by applications using SNMP traps.

NOTE

On a UNIX Cell Manager, SNMP traps are sent to the systems configured in the notification.

On a Windows Cell Manager, SNMP traps are sent to the systems configured in the Windows SNMP traps configuration.

Windows

To configure Windows SNMP traps, proceed as follows:

1. Run `omnisnmp.exe` command from `<Data_Protector_home>\bin` directory.
   It will create the appropriate Data Protector entry in the System registry under `CurrentControlSet\Services\SNMP\Parameters\ExtensionAgents`
2. On the Cell Manager, click `Settings`, and then `Network and Dial-up Connections`.
3. In the `Advanced` menu, select `Optional Networking Components` to start the wizard.
4. In the wizard, select `Management and Monitoring` tools and click

Next.

5. Follow the wizard to install the Management and Monitoring tools.
6. Open Control Panel, Administrative Tools and then Services.
7. Right-click SNMP Service and select Properties.
   a. Select the Traps tab and enter public in the Community name text box and the hostname of the VPO Management Server in the Trap Destinations text box.
   b. Select the Security tab. Under Accepted community names, select the community public (by default), click Edit and set Community rights to READ CREATE.
   c. Confirm your settings.
8. Start omnisnmp.

External Send Method

The external send method allows you to process the output of the report in your own script. The script receives the output as standard input (STDIN). The recommended format for script processing is the tab format.

The script, which is located on the Cell Manager system, must reside in the /opt/omni/lbin (UNIX systems) or *<Data_Protector_home>*\bin (Windows systems) directory. You need to provide only the name of the script, not the entire path.

Note that only .bat, .exe, and .cmd are supported extensions for external scripts on Windows systems. To run an external script with an unsupported extension (for example .vbs), create a batch file (.bat) that starts the script. Then configure Data Protector to run the batch file as an external script which then starts the script with the unsupported extension.

TIP You can use this delivery method to perform a scheduled eject of the specified media. Refer to “Scheduled Eject of Media” on page 184.

Configuring Reports Using the Data Protector GUI

This section describes how to configure Data Protector reports using the Data Protector GUI.

NOTE

To display the input parameters (selections) in the output of a report, select the `Show selection criteria in report` option in the `Report Wizard`. The `Show selection criteria in report` is not available for the reports that have no required or optional input parameters (selections). The output of the report displays only required parameters and optional parameters with the changed default values.

Configuring Report Groups and Adding Reports

Report Groups

You can start Data Protector reports individually (interactively) or you can group them into report groups and then start the report group. You can add individual reports to an already configured report group.

Using the Data Protector GUI, a report group allows you to:

- Start all the reports at once (interactively).
- Schedule the group to start the reports at a specified time.
- Start the group when triggered by a notification.

Examples

These are some examples of the use of reports:

- A backup operator wants to receive an email with the status of the backup performed on the previous night.
- Administrators of specific departments want to receive a broadcast message with information on the backup of the systems they are responsible for.
- A full report with tab delimited data is posted as a log file and is used by an application that records backup statistics.

Administrators can configure a report group and add a separate report for each of the requirements. They can schedule the report group to be executed early enough in the morning, so that all recipients receive the reports before coming to work.

NOTE The Mount Request Report and Device Error Report can only be used in a report group and are not available as interactive reports.

To configure a report group, do the following:

1. In the `Data Protector Manager`, switch to the `Reporting` context.
2. Click the `Objects` tab below the Scoping Pane to switch to the Objects view.
3. Right-click `Reports` and then select `Add Report Group`. The `Add Report Group` wizard appears.

   Follow the wizard. You will go through the following steps:

   a. Name the report group.

   b. Optionally schedule when the group should be started. For more information on how to use the Scheduler, see “Scheduling Unattended Backups” on page 250.

   c. Choose and configure a report for the group. For each report, you must configure a format used to deliver a report, recipients for each report, and a send method. See “Report Formats” on page 402 for more information on report formats. See “Data Protector Notifications” on page 414 for more information on various send methods.

NOTE To trigger a report group by a notification, you first need to configure a report group and then configure the notification to use the `Use Report Group` send method.

4. The report group is created and displayed in the Scoping Pane.
5. To add multiple reports to the group, right-click the group and then select `Add Report`.

Running Reports and Report Groups Using the Data Protector GUI

Data Protector reports can be run individually, or they can be grouped into report groups and then run.

Running Individual Reports

To run each report individually, do the following:

1. In the `Data Protector Manager`, switch to the `Reporting` context.
2. Click the `Tasks` tab below the Scoping Pane to switch to the tasks context. Browse the provided reports and select the one that you want.
3. Follow the `Report Wizard` to configure and run the report.

Running Report Groups

To run a configured report group, do the following:

1. In the `Data Protector Manager`, switch to the `Reporting` context.
2. In the Scoping Pane, browse for and right-click the report group you want to run and then click `Start`.
3. Click `Yes` to confirm.

Running Reports and Report Groups Using the Command-Line Interface

Data Protector reports can be generated using the command-line interface. The command-line interface allows you to include Data Protector reports in some other configuration scripts you are using. It allows you to generate individual reports, run report groups, and define report formats and send methods.

The `omnirpt` command is used to generate reports. For a detailed description of the command, see the `omnirpt` man page.

Here are some examples of `omnirpt` usage:

`omnirpt -rptgroup <ReportGroup>`

Runs the report group named *<ReportGroup>*.

NOTE You first need to configure a report group using the Data Protector GUI or Web reporting interface before running it using the Data Protector command-line interface.

`omnirpt -report host -host <Hostname> -html`

This generates a Client Backup Report for system *<System_Name>* in the HTML format.

Example 1 The following command creates a Session Flow Report for the last 24 hours and logs it to the file in HTML format, as shown in Figure 9-3 on page 412:

```
omnirpt -report session_flow -timeframe 24 24 -log
session_flow.html -html
```

Figure 9-3 A Session Flow Report

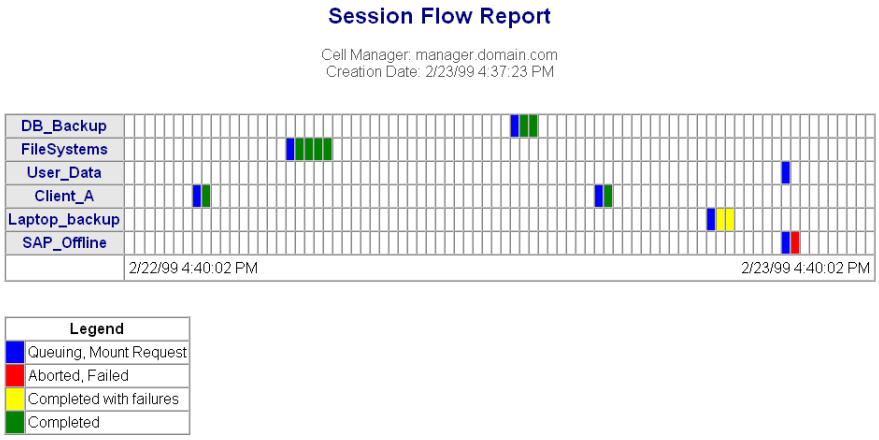

Example 2

The following command creates a Media Statistics Report on media in poor condition and logs it in the file in the ASCII format, as shown in Figure 9-4 on page 412:

```
omnirpt -report media_statistics -status poor -log
media_statistics.txt -ascii
```

Figure 9-4 A Media Statistics Report

```
Media Statistics

Cell Manager: popsicle.bbn.hp.com
Creation Date: 2/23/99 4:47:24 PM

# Media: 2
# Scratch: 2
# Protected: 0
# Good : 0
# Fair: 0
# Poor: 2
Total [MB]: 7.79
Used [MB]: 0.75
Free [MB]: 7.04
```

Example 3

The following command creates a Device Flow Report for the last 24 hours and sends it via email in HTML format, as shown in Figure 9-5 on page 413:

```
omnirpt -report device_flow -timeframe 24 24 -email
ulmo@outersea.ea -html
```

Figure 9-5 A Device Flow Report

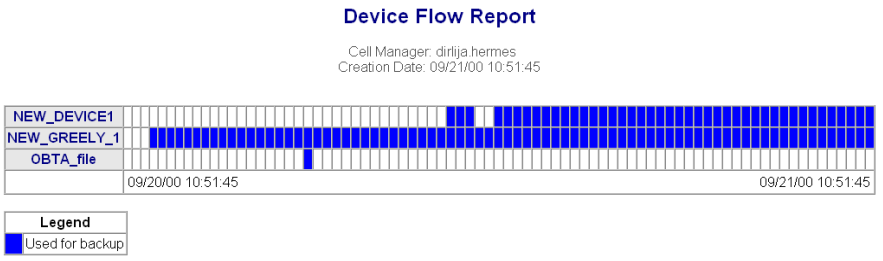

Data Protector Notifications

What Are Notifications?

The Data Protector notification functionality allows you to receive notifications when specific events occur. For example, when a backup session is completed, you can receive an email with a status of the session.

You can set up a notification so that it triggers a report. For more information about Data Protector reports, refer to “Data Protector Reporting” on page 388.

NOTE

Only the Data Protector users in the Admin group and those granted the `Reporting and notifications` user rights are given access to Data Protector notification functionality. For more information on Data Protector users and users groups, refer to the “Configuring Users and User Groups” on page 127.

Configuring Notifications

Notifications can be configured using the Data Protector user interface or any Web browser with Java support.

Notifications can be customized by configuring input parameters.

All notifications have the following common input parameters:

- Name (a name for the notification)
- Message Level (the default value depends on the notification and is listed for each notification in the table below)
- Send Method (the default value is Data Protector Event Log)

Notification Types

There are two main types of notifications:

- Notifications that are triggered when an event occurs:
  - — Alarm
  - — Backup Error
  - — Device Error

  - Start of Session
  - End of Session
  - IDB Corrupted
  - Mail Slots Full
  - File Library Disk Usage
  - Mount Request
- Notifications that are scheduled and started by the Data Protector checking and maintenance mechanism:
  - Health Check Failed
  - IDB Purge Needed
  - IDB Space Low
  - IDB Tablespace Space Low
  - License Warning
  - License Will Expire
  - Not Enough Free Media
  - Unexpected Events
  - User Check Failed

For more information on the Data Protector checking and maintenance mechanism, refer to “Data Protector Checking and Maintenance Mechanism” on page 725.

Table 9-8 Data Protector Notifications

Name	Optional Input Parameters	Default Message Level and Optional Input Parameter Default Values	Message Displayed
IDB Corrupted	none	• Critical	Corruption in the *<DB_part>* part of the internal database has been detected *<error_message>*

Table 9-8 Data Protector Notifications

Name	Optional Input Parameters	Default Message Level and Optional Input Parameter Default Values	Message Displayed
Backup Error	Single Message Level (<Any>/Warning/Minor/Major/Critical/Normal - only the Data Protector messages of the specified level of messages and above trigger this notification)	• Major • Major	Backup session *<session_ID>* of the backup specification *<backup_spec>* has errors: *<number_of_errors>*
Unexpected Events	Number of Events (threshold value for the number of events in the Data Protector Event Log that triggers this notification)	• Warning • 20	Data Protector Event log increased for *<Number of Events>* unexpected events in last day
Health Check Failed	none	Critical	Health check message: *<healthcheck_command>* failed, check HealthCheck.log file.
User Check Failed	Command Path	• Major • none	User check failed with exit code *<error_code>*:*<error_description>*
Start of Session	Data list	• Warning • Any	Session *<Session_ID>* started for backup specification *<Datalist>*, backup group *<BackupGroup>*.
End of Session	• Datalist • Session Status	• Warning • All • Completed with Errors	Session *<session_ID>* of backup specifications *<Datalist>* completed with overall status *<Session Status>*

Table 9-8 Data Protector Notifications

Name	Optional Input Parameters	Default Message Level and Optional Input Parameter Default Values	Message Displayed
Device Error	Device	• Critical • *<Any>*	Error on device *<Device>* occurred
IDB Space Low	• Maximum Size of filenames.dat [MB] • Disk Free for Internal Database [MB] • DCBF Size Limit [MB]	• Major • 250 MB • 50 MB • 250 MB	Internal database is running out of space
IDB Tablespace Space Low	Tablespace Used Threshold [%]	• Major • 85%	Tablespace *<Tablespace_name>* is running out of space.
IDB Purge Needed	• Days Last Purge [days]) • Num. Estimated Filenames [mio] • Estimated Time Purge [min] • Num. Filenames [mio]	• Warning • 180 days • 6 million • 120 minutes • 100 million	Filename purge should be run for Internal Database
Mount request	Device	• Warning • *<Any>*	Mount request on device *<Device>*

Table 9-8 Data Protector Notifications

Name	Optional Input Parameters	Default Message Level and Optional Input Parameter Default Values	Message Displayed
Not Enough Free Media	• Media Pool • Number of Free Media (threshold value for the lowest number of free media that triggers this notification)	• Warning • *<Any>* • 2	Media pool *<Media Pool>* contains only *<number_of_media>* free media
Mail Slots Full	• Device	• Warning • *<Any>*	All mail slots of library *<Device>* are full. Please remove them immediately
File Library Disk Usage	• File Library Device Name	• Warning • <All>	The *<File Library Device Name>* is low in disk space in the directory *<File Library Device Directory>*.
License Warning	none	• Warning	*<n>* licenses need to be purchased for category *<name of the license>*. Run `omnicc -check_licenses -detail` for more information.
License Will Expire	License expires in days	• Warning • 10	The first license will expire in *<License expires in days>* days
Alarm	none	• Warning	Alarm: *<Alarm_message>*

Explanation of Some Notifications

Alarm The Alarm notification is used to display critical Data Protector messages triggered by Data Protector internal conditions.

IDB Purge Needed By default, once per day Data Protector will check the IDB Purge Needed condition as a part of Data Protector checking and maintenance mechanism and trigger the notification if:

- For any Data Protector client in the cell, the number of days since the last IDB filename purge is larger than the *<Days Last Purge [days]>* input parameter and at least one of the following two conditions is true:
  - The number of filename records likely to be purged is larger than the *<Num. Estimated Filenames [mio]>* input parameter.
  - It is estimated that more than *<Estimated Time Purge [sec]>* seconds will be needed to finish the purge.
- The number of filenames in the IDB is larger than the *<Num. Filenames [mio]>* input parameter.

For more information on the Data Protector checking and maintenance mechanism, refer to “Data Protector Checking and Maintenance Mechanism” on page 725.

IDB Space Low By default, once per day Data Protector will check the IDB Space Low condition, and will trigger notification if the allocated space for CDB extension files is running low, if any of the disks containing the IDB are running out of space, or if the allocated space for all DC directories is running low. In other words, the notification will be triggered if any of the following is true:

- The difference between the maximum size of *all* CDB extension files, (the sum of all CDB extension files maximum sizes) and the current size of all CDB extension files drops below the *<Maximum Size of filenames.dat [MB]>* input parameter.
- The free disk space on *any* of the disks containing the IDB drops below the *<Disk Free for Internal Database [Mb]>* input parameter.
- The difference between the maximum size of *all* DC directories and the current size of all DC directories drops below the *<DCBF Size Limit [MB]>* input parameter.

For more information on the Data Protector checking and maintenance mechanism, refer to “Data Protector Checking and Maintenance Mechanism” on page 725.

IDB Tablespace Space Low By default, once per day Data Protector checks the IDB Tablespace Space Low condition as part of the Data Protector checking and maintenance mechanism, and triggers the notification if the allocated space for any of the tablespaces is running low. By default, this notification is triggered when 85% of allocated space is used.

Health Check Failed As a part of the Data Protector checking and maintenance mechanism, Data Protector will by default once per day start the Health Check, which starts the `omnihealthcheck` command and triggers the notification if the `omnihealthcheck` command fails. For more information on the `omnihealthcheck` command, refer to the `omnihealthcheck` man page. The `omnihealthcheck` command checks:

- whether the Data Protector services (`rds`, `crs`, `mmd`, `omnitrig`, and `OmniInet`) are active
- whether the Media Management database is consistent
- whether at least one backup of the IDB exists

The exit code of the command is 0 (OK) only if all three checks completed successfully (exit code for every check was 0). Exit values other than 0 indicate that one of the checks failed. For more information on exit codes, refer to the `omnihealthcheck` man page.

User Check Failed By default, once per day Data Protector will start the User Check, which executes the script/command specified as the *<script/command pathname>* input parameter. Create the command/script in the `/opt/omni/lbin` (HP-UX and Solaris) or *<Data_Protector_home>*`\bin` (Windows) directory of the application system. Enter the filename here. The notification is triggered if the script/command exits with the return value other than 0.

For more information on the User Check Failed notification, refer to “The User Check Failed Notification” on page 726.

Start of Session The Start of Session notification is triggered when a Data Protector session for the backup specification(s) specified by the *<Datalist>* input parameter starts with the status specified by the *<Session Status>* input parameter. The default value is Session *<Session_ID>* started for backup specification *<Datalist>*, backup group *<BackupGroup>*.

End of Session The End of Session notification is triggered when a Data Protector session for the backup specification(s) specified by the *<Datalist>* input parameter ends with the status specified by the *<Session Status>* input parameter. The default value is Completed with Errors.

Notification Send Methods

Notifications can be sent using various methods:

- Email send method
- Broadcast message send method
- SNMP send method
- External send method
- Log to file send method
- Use Report Group send method
- Data Protector Event Log send method

NOTE By default, all notifications are configured to be sent to the Data Protector Event Log. In order to send an additional notification using some other send method, an additional notification has to be configured.

E-mail Send Method

E-mail notifications enable you to receive an e-mail with desired information when a specified event occurs.

IMPORTANT Due to security features of Microsoft Outlook, using the e-mail send method may cause the CRS service to stop responding. For details and solutions, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

Prerequisite for Windows To send an e-mail notification from a Windows system, you need to have a mail profile created. You can either use an existing profile or create a new mail profile, named `OmniBack`.

To use an existing mail profile, add the following line to the Data Protector `omnirc` file:

`OB2_MAPIPROFILE=<existing_MAPI_profile_name>`

For more information on the `omnirc` file, see “Using Omnirc Options” on page 615.

Creating a Data Protector Mail Profile

To create a new mail profile, named `OmniBack`, on the Windows 2000 system with Microsoft Outlook 2002 installed, perform the following steps:

1. In the Windows `Control Panel`, double-click the `Mail` icon.

   The `Mail Setup - Outlook` dialog box is displayed.

2. Click `Show Profiles`.

   The `Mail` dialog box is displayed.

3. Click `Add`.

   The `New Profile` dialog box is displayed.

4. Type `OmniBack` in the `Profile Name` text box and click `OK` to start the `E-Mail Accounts` wizard.

5. Select `Add a new e-mail account` and click `Next`.

6. In the `Server Type` page, select `Microsoft Exchange Server` and click `Next`.

7. In the `Exchange Server settings` page, type the name of the local Microsoft Exchange Server system, your username and click `Next`.

8. Click `Finish` to end the wizard.

   In the `Mail` dialog box, the `OmniBack` profile is listed among other profiles set on your system.

On UNIX systems, no additional configuration is needed.

Broadcast Message Send Method

Broadcast message notifications allow you to send a broadcast message to systems when a specified event occurs.

Broadcast messages can be sent to Windows systems only, by specifying the system to which the broadcast message should be sent. Broadcast messages are limited in length, so the `short` format is preferred. The reports are limited to 1000 characters.

Log to File Send Method

Log to file notifications allow you to post a log file with desired information when a specified event occurs.

The log file is posted to the Cell Manager system. Specify the name of the file to which you want to post the report.

SNMP Send Method

SNMP traps notifications allow you to send an SNMP trap with desired information when a specified event occurs. The SNMP trap can be further processed by applications using SNMP traps.

NOTE

On a UNIX Cell Manager, SNMP traps are sent to the systems configured in the notification.

On a Windows Cell Manager, SNMP traps are sent to the systems configured in the Windows SNMP traps configuration.

Windows

To configure Windows SNMP traps, proceed as follows:

1. On the Cell Manager, open `Settings`, `Network and Dial-up Connections`.
2. In the `Advanced` menu, select `Optional Networking Components` to start the wizard.
3. In the wizard, select `Management and Monitoring` tools and click `Next`.
4. Follow the wizard to install the `Management` and `Monitoring` tools.
5. Open `Control Panel`, `Administrative Tools`, `Services`.
6. Right-click `SNMP Service` and select `Properties`.
   a. Select the `Traps` tab and enter `public` in the `Community name` text box and the hostname of the VPO Management Server in the `Trap Destinations` text box.

b. Select the Security tab. Under Accepted community names, select the community public (by default), click Edit and set Community rights to READ CREATE.

c. Confirm your settings.

7. Start omnisnmp.

External Send Method

External script notification allows you to process the output of the report in your own script. The script receives the output as standard input (STDIN). The recommended format for script processing is the tab format.

The script, which is located on the Cell Manager, must reside in the /opt/omni/lbin (HP-UX and Solaris systems) or *<Data_Protector_home>*\bin (Windows systems) directories. You need to provide only the name of the script, not the whole path.

Note that only .bat, .exe, and .cmd are supported extensions for external scripts on Windows systems. To run an external script with an unsupported extension (for example .vbs), create a batch file (.bat) that starts the script. Then configure Data Protector to run the batch file as an external script which then starts the script with the unsupported extension.

TIP You can use this delivery method to perform a scheduled eject of the specified media. Refer to “Scheduled Eject of Media” on page 184.

Use Report Group Send Method

Report group notification allows you to start a report group when a specified event occurs. See “Configuring Reports Using the Data Protector GUI” on page 408 for more information on report groups.

Data Protector Event Log Send Method

By default, all notifications are sent to the Data Protector Event Log. The Data Protector Event Log is accessible only for Data Protector users in the Admin group and to Data Protector users that are granted the

Reporting and notifications user rights. You can view or delete all events in the Data Protector Event Log. Refer to “Data Protector Event Log” on page 430.

Configuring Notifications

To configure a notification, do the following:

1. In the Data Protector Manager, switch to the Reporting context.
2. Click the Objects tab below the Scoping Pane to switch to the Objects view.
3. Right-click Notifications and then select Add Notification. The Add Notification wizard appears. Follow the wizard.

TIP To trigger a report group by a notification, configure a report group and then configure the notification to use the Use Report Group send method.

4. The notification is created and displayed in the Scoping Pane.

Configuring Reports and Notifications on the Web

You can use your Web browser to view Data Protector reports and notifications.

Using the web reporting and notifications interface, you can view, configure, and start Data Protector reports and notifications from any system on your network. You can configure reports and notifications that are delivered using various reporting methods and formats.

All reporting and notifications functionality accessible using the Data Protector GUI is also accessible using Data Protector web reporting and notifications. See below for the limitations.

When you install the Data Protector Cell Manager, the web reporting user (called Java) is automatically created. By default, no password is needed to use the Data Protector web reporting and notifications. By configuring a Web user password you restrict the access to the Data Protector web reporting and notifications functionality.

Requirement

- A supported Web browser must be installed. Refer to the *HP OpenView Storage Data Protector Software Release Notes* for a list of Web browsers supported for web reporting and notifications.
- Java VM (versions 1.1 or later) must be installed on the system and enabled in the Web browser.

Limitations

The following is a list of Data Protector web reporting and notifications interface limitations:

- You cannot edit, view, or delete the saved reports using the web reporting and notifications interface.
- You cannot start a report group using the web reporting and notifications interface.
- Whenever multiple input parameters (selections) are to be *typed* in the web reporting and notifications interface, every parameter (selection) has to be enclosed in double quotes if it contains spaces.

To use Data Protector web reporting and notifications, do the following:

1. Have a system with a configured and running web server. Data Protector works with all popular web servers.
2. Copy Data Protector Java programs to the web server. The system does not have to be a Data Protector client. The steps are described below.
3. Optionally, configure a password to limit access to Web reports. The steps are described below.

Copying Data Protector Java Programs to the Web Server

To allow access to Data Protector Web reporting and notifications interface from a browser from any system, copy Data Protector Java reporting programs to the web server.

From the system with the Data Protector user interface installed, copy the following directory with all subdirectories:

- On Windows: `<Data_Protector_home>\java`
- On UNIX:`/opt/omni/java`

Access the `\bin\WebReporting.html` (Windows systems) or the `/bin/webreporting.html` (UNIX systems) file from the copied `java` folder in a browser to display the Data Protector reporting. Make this file available to the users of the web reporting in the full URL form. For example, you can put a link to this file from your intranet site.

Restricting Access to Web Reporting

When you install the Data Protector Cell Manager, the web reporting and notifications user (called `java`) is automatically created. By default, no password is needed to use the Data Protector web reporting and notifications. By configuring a web user password, you restrict the access to Data Protector web reporting and notifications functionality. Any user using web reporting and notifications will have to provide this password to browse the Data Protector reports on the web.

To change the password for the Data Protector web reporting and notifications interface, do the following:

1. In the `Data Protector Manager`, switch to the `Users` context.

2. Choose `Action, Set Web User Password`. A dialog box appears, where you change the password.

   Any user using web reporting and notifications interface will have to provide this password to browse the Data Protector reports on the web.

Generating the Reports

To generate reports using the Data Protector Web reporting and notifications interface, you have to access this interface. The actual steps depend on your configuration. Once you are logged on the Cell Manager, you can generate various types of reports. See “Data Protector Reporting” on page 388 for more information on report types.

To view a report, click the report and provide the needed information.

When the report is displayed, you can print the report or save it. When you save the report, you can also add this report to an existing or a new report group. See the next section for more information.

Configuring Notifications

To configure notifications using the Data Protector Web reporting and notifications interface, you have to access this interface. The actual steps depend on your configuration. Once you are logged on the Cell Manager, you can configure notifications. See “Data Protector Notifications” on page 414 for more information on notifications.

To configure a notification, select `Notifications` and click `Add Notification`. Provide the needed information and save the notification.

Configuring Report Groups

Report Groups

See “Configuring Report Groups and Adding Reports” on page 408 for more information on report groups.

In the web reporting and notifications interface, you can create a new report group when you save the report:

1. Choose the report you want to generate.
2. Enter the needed information.

3. Once the report is displayed, click `Save`. Enter the report name and a new or an existing report group to which you want to add the report.

Data Protector Event Log

The Data Protector Event Log represents a centralized event management mechanism, dealing with specific events that occurred during the Data Protector operation. The events are logged in *<Data_Protector_home>*\log\server\Ob2EventLog.txt (Windows Cell Manager) or in /var/opt/omni/server/log/Ob2EventLog.txt (UNIX Cell Manager). Viewing the Data Protector Event Log using the Data Protector GUI helps you troubleshoot possible problems.

The events are logged by the notifications functionality. Refer to "Data Protector Notifications" on page 414 for more information on notifications.

NOTE Only the Data Protector users in the Admin group and those granted the Reporting and notifications user rights are given access to Data Protector Event Log functionality.

Event Log

To access the Event Log, select the Reporting context in the Data Protector GUI and expand Reporting. Select Event Log to display events.

NOTE The Data Protector Event Log is not refreshed automatically. If you want to view new messages, refresh it manually by pressing F5.

Deleting Event Log Contents

Right-click Event Log and select Empty Event Log. This will delete all entries in the Event Log.

NOTE Deleting the Event Log contents will not delete the *<Data_Protector_home>*\log\server\Ob2EventLog.txt file (Windows Cell Manager) or the /var/opt/omni/server/log/Ob2EventLog.txt file (UNIX Cell Manager).

When the Data Protector graphical user interface is started by a user, if there are new notifications that have not been seen by this user in the Data Protector Event Log, the following message is displayed:

Figure 9-6 The Event Log Message

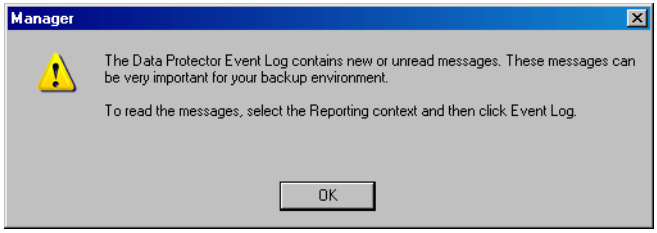

10 Manager-of-Managers Environment

In This Chapter

This chapter shows you how to configure and use the Data Protector Manager-of-Managers, which is used to control an enterprise backup environment. It consists of the following sections:

NOTE MoM is subject to specific Data Protector licenses. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.

Manager-of-Managers

The Data Protector Manager-of-Managers (MoM) allows administrators to centrally manage a large environment consisting of several Data Protector cells, also known as MoM clients, from a single point. Refer to the *HP OpenView Storage Data Protector Concepts Guide* for further details about the enterprise environment.

NOTE Each MoM client and the MoM Manager need to run the same version of Data Protector.

The Data Protector MoM is flexible enough to expand the backup environment as the enterprise grows. It provides the following features:

Centralized management of all tasks

Enables configuration, management, and control over the enterprise environment from a single point. This includes configuring backup, media management, restoring, and monitoring; and reporting about the status of the whole backup environment.

Centralized Media Management Database

Optionally, all the cells in the environment can share a common, central database to manage devices and media within the enterprise. The Centralized Media Management Database (CMMDB) enables you to share high-end devices between cells. This means that any device in a cell using the CMMDB is available to all cells using the CMMDB.

Centralized licensing

Data Protector enables you to configure centralized licensing for the whole MoM environment. All Data Protector licenses are installed and kept on the MoM Manager and can be allocated to specific cells as needed.

Configuring the Manager-of-Managers

To configure the MoM environment, you need to do the following:

1. Set up the MoM Manager. See “Setting Up MoM Manager” on page 437.
2. Import Data Protector cells into MoM environment. See “Importing Data Protector Cells” on page 438.
3. Create a Data Protector user in the Admin user group on every cell in the environment (MoM administrator). See “Adding a MoM Administrator” on page 438.
4. Restart Data Protector services. See “Restarting Data Protector Services” on page 439.

Optionally, you can also:

- Configure the Centralized Media Management Database. See “Configuring a Centralized Media Management Database” on page 443.
- Configure centralized licensing. See “Centralized Licensing” on page 447.
- Distribute the MoM configuration. See “Distributing the MoM Configuration” on page 453.

Prerequisites

Choose the system you will configure as your MoM Manager. Follow the guidelines below:

- The MoM Manager system should be highly reliable.
- The system has to already be a Data Protector Cell Manager with the software installed. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information on how to configure the Data Protector Cell Manager system.

Install the required licenses on the MoM cell and every prospective MoM client cell.

Setting Up MoM Manager

To set up an enterprise environment, configure one of your Cell Managers as a Manager-of-Managers.

1. In the `Data Protector Manager`, click `Clients` in the Context List.
2. In the `Actions` menu, click `Configure CM as Manager-of-Managers Server`.
3. Stop and restart Data Protector services. Refer to the section "Restarting Data Protector Services" on page 439.
4. Run the MoM graphical user interface:
   - On Windows platforms, do one of the following:
     — Click `Start` on the Windows taskbar and click `Manager-of-Managers` from the `HP OpenView Storage Data Protector` program group to start the MoM GUI for the complete Data Protector functionality.
     — Use the `mom` command to start the GUI for the complete Data Protector functionality.

     Context-specific options for this command enable you to start one or more Data Protector contexts. For example, the command
     `mom -backup -restore`
     starts the Data Protector `Backup` and `Restore` contexts, as well as the monitor context.

     To specify the Cell Manager you want to connect to, use the following command: `mom -server <Cell Manager_name>`.
   - On UNIX: run the `/opt/omni/bin/xomnimom` command to start the Data Protector Manager-of-Managers GUI with all Data Protector contexts activated (with the exception of Internal Database and Devices & Media contexts). If additional context options are specified, this command starts only the specified Data Protector context(s).

   For more information on these commands, refer to the *omnigui* man page.

Importing Data Protector Cells

Once you have configured the MoM Manager, you can start adding (importing) the Data Protector cells to the MoM environment. To import a Data Protector cell to the MoM environment, proceed as follows:

1. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.
2. Right-click `Enterprise Clients`, and then click `Import Cell Manager`.

IMPORTANT In order to import a Cell Manager into the MoM as an Enterprise Client, you must be a member of the admin user group on that Cell Manager. If you are not, the import will fail.

3. Enter, or browse for, the name of the Cell Manager that you want to import, and then click `Finish`. The selected Cell Manager is now a part of your MoM environment.

NOTE If you are adding a Cell Manager installed on a cluster to the MoM cell, ensure that you enter its **virtual server name**.

Adding a MoM Administrator

A MoM administrator can perform administration tasks in all cells in the enterprise environment.

You need to have a certain user that is in the Admin user group on every Cell Manager in the MoM environment. For example, you may have a user called *MoM_Admin*. This user will be the MoM administrator.

1. Using the Data Protector Manager, connect to each Cell Manager in the MoM environment as a member of the `Admin` user group (the `User configuration` user right is required).
2. Add the user that will be the MoM Administrator to the Data Protector `Admin` user group.

On how to add users, see “Adding or Deleting a User” on page 137.

Restarting Data Protector Services

After you have configured the MoM environment, you will be notified to restart the Data Protector services.

If you use the `Windows Service Control Manager` to start and stop services on the Cell Manager, only the current and previous copies of the database log are kept. Using the `omnisv -stop` and the `omnisv -start` commands will save all previous database logs.

1. Stop all Data Protector services by entering the following command:
   - on Windows: `<Data_Protector_home>\bin\omnisv -stop`
   - on UNIX: `/opt/omni/sbin/omnisv -stop`

NOTE `omnisv` command is not supported in the cluster.

MC/ServiceGuard If the Cell Manager is configured on MC/SG, stop the Data Protector package using the following command:

`cmhaltpkg <pkg_name>` , where `<pkg_name>` is the name of the Data Protector cluster package.

This command stops the Data Protector package and dismounts the Data Protector shared volume group.

Microsoft Cluster Server If the Cell Manager is configured on Microsoft Cluster Server, take the `OBVS_VELOCIS` cluster group offline (using the `Cluster Administrator` utility on the active node).

2. Restart the Data Protector services by entering the following command:
   - on Windows: `<Data_Protector_home>\bin\omnisv -start`
   - on UNIX: `/opt/omni/sbin/omnisv -start`

MC/ServiceGuard If the Cell Manager is configured on MC/SG, restart the Data Protector package using the following command:

`cmrunpkg -n <node_name> <pkg_name>`

Microsoft Cluster Server

If the Cell Manager is configured on MSCS, bring the `OBVS_VELOCIS` and `OBVS_MCRS` cluster groups online using the `Cluster Administrator` utility.

Centralized Media Management Database (CMMDB)

The IDB is an embedded database that keeps information about backup, restore, and media management sessions, devices, and media. It consists of five parts that are located on the Cell Manager.

- MMDB - Media Management Database
- CDB - Catalog Database
- DCBF - Detail Catalog Binary Files
- SMBF - Session Messages Binary Files
- SIBF - Serverless Integrations Binary Files

In a typical cell-oriented environment, all parts are located on the Cell Manager system and each keeps information on devices, media, and backup information for that cell. For security reasons, it is impossible to access and use this data from another Data Protector cell. Therefore, media and devices used in that cell cannot be accessed and used in some other cell without moving them to that cell.

In larger multi-cell environments with high-end backup devices, you may want to share these devices and media among several cells. This can be achieved by having one centralized MMDB database for all the cells and keeping an individual CDB for each cell. This allows media and device sharing while preserving the security capabilities of the multi-cell structure.

With the CMMDB, media are owned by the Data Protector cell that performed the first backup on the media. The media owner is displayed in the media view. While media are protected, only backups from that cell can be appended on the media. Therefore, media can only be owned by one cell at a time. Once the protection expires, the media become available to other cells again.

NOTE A backup anywhere in the enterprise environment will not run if the cell running the backup does not have access to the CMMDB. For example, this happens if a network failure occurs between the cell and the MoM cell.

Figure 10-1 The Central Media Management Database

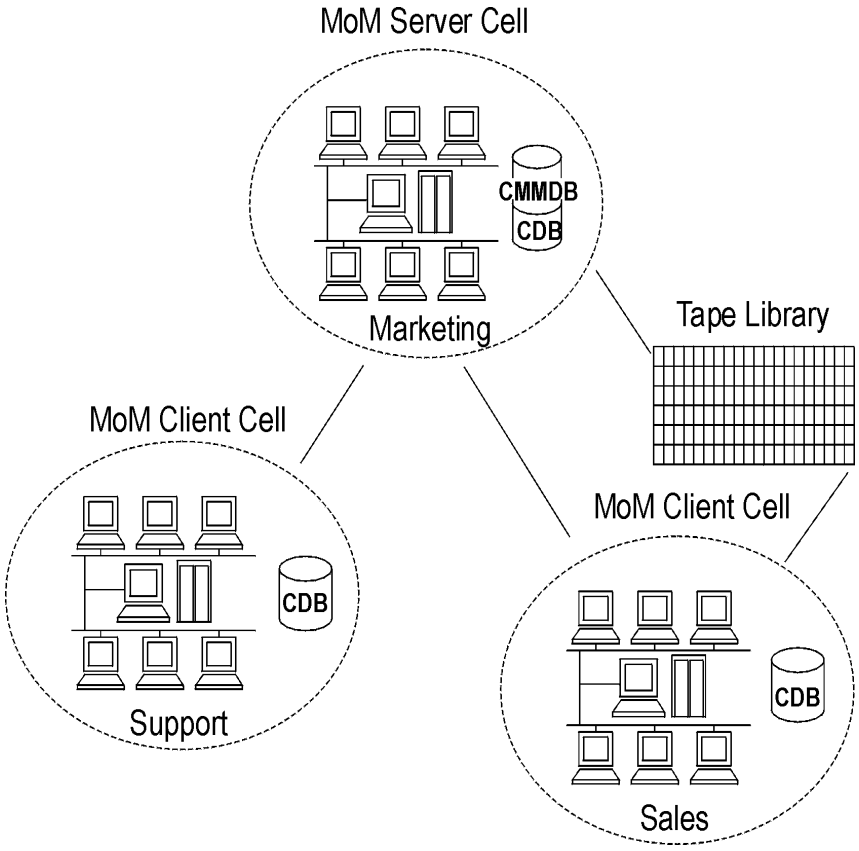

Configuring a Centralized Media Management Database

It is not required to set up a Centralized Media Management Database (CMMDB). If you do not set up a CMMDB, Data Protector will work in a multi-cell environment, but each cell will have its own IDB. See "Centralized Media Management Database (CMMDB)" on page 441 for more information on this functionality.

This section describes how to configure a Centralized Media Management Database for the whole multi-cell environment. If it is needed, this process will merge the local Media Management Database into the CMMDB. You can decide for each cell if it will use the CMMDB or its own local MMDB.

IMPORTANT The CMMDB has a major effect on licensing. Immediately after the MMDB is changed from local to remote, all the licenses associated with libraries and devices are taken (validated) from the MoM Manager and can be removed from client cells.

When the CMMDB is used, it does not have to reside on the MoM Manager system. The CMMDB can reside on any Cell Manager in the MoM environment. The Cell Manager on which the CMMDB is located is specified in the file `mmdb_server` in the following directory:

- On Windows: `<Data_Protector_home>\Config\server\cell`
- On UNIX: `/etc/opt/omni/server/cell`

Each medium with protected data on it has information showing which cell currently owns the data. Once this protection has expired, any cell can reuse the medium. If a tape has been initialized by one cell, any other cell can use it, as long as it does not have any protected data on it. If a tape is loaded in a library and not yet initialized, any cell can initialize it, assuming that there is a `loose` media allocation policy and no other tapes are available.

The media allocation rules apply in exactly the same way to shared tapes, except that appendable media can only be appended by the cell that owns it.

On the MoM, add one cell at a time to the CMMDB.

Prerequisites

- Data Protector Cell Managers in all cells have to have the same version of Data Protector installed and running.
- Check that there are no backup, restore, or media management sessions running on any of the cells to be added to the multi-cell environment.

How to Configure the CMMDB

To configure the CMMDB in the MoM environment, two phases are required:

1. Configuration of the CMMDB on the MoM Manager. See “Configuring the CMMDB on the MoM Manager” on page 444.
2. Configuration of the CMMDB on the client cell. See “Configuring the CMMDB on the Client Cell” on page 445.

NOTE

Once you have configured the CMMDB and start using it, it is not possible to split it back into local MMDBs. It is not recommended to recover the old state of a MMDB. Instead, you should create a new MMDB from scratch.

Configuring the CMMDB on the MoM Manager

Log on to the MoM Manager and perform the following steps:

1. Copy the following directory to a temporary location for safety reasons:
   - On Windows: `<Data_Protector_home>\db40\datafiles\mmdb`
   - On UNIX: `/var/opt/omni/server/db40/datafiles/mmdb`
2. Run the following command to merge the local MMDB into the CMMDB:
   - On Windows: `<Data_Protector_home>\bin\omnidbutil -mergemmdb <Cell_Server_Hostname>`
   - On UNIX: `/opt/omni/sbin/omnidbutil -mergemmdb <Cell_Server_Hostname>`

TIP If you are configuring a new cell, (and you do not yet have devices and media configured) there is no need to merge the database. You only want to merge cells with the CMMDB that already have devices and media configured.

3. Run the following command to synchronize the local CDB:
   - On Windows: *<Data_Protector_home>*\bin\omnidbutil -cdbsync *<Cell_Server_Hostname>*
   - On UNIX: /opt/omni/sbin/omnidbutil -cdbsync *<Cell_Server_Hostname>*
4. On the MoM Server, edit the duplicated names of media pools and devices (in the user interface). The duplicated names have a "_N" appended to their name, where N represents a number. This always happens to default pools if they exist on both cells. In this case, manually change the backup specifications that use these devices to use the new device names. It would be a good idea to add a line to the media pool's description to say from which cell the pool has come.

Repeat the steps 2 to 4 for all client cells that you want to add to the CMMDB.

Configuring the CMMDB on the Client Cell

On each MOM client cell, perform the following:

1. Log on to the Cell Manager of the client cell as a member of the Admin user group.
2. Create the file containing the name of the MMDB Server (fully qualified):
   - On Windows: *<Data_Protector_home>*\Config\server\cell\mmdb_server Save the file as Unicode.
   - On UNIX: /etc/opt/omni/server/cell/mmdb_server
3. Stop and restart the Data Protector services. See "Restarting Data Protector Services" on page 439.

4. Update configuration files by running the following command:
   - On Windows: `<Data_Protector_home>\bin\omnicc -update_mom_server`
   - On UNIX: `/opt/omni/bin/omnicc -update_mom_server`

Centralized Licensing

It is not required to set up centralized licensing. Individual licenses can be installed on each Cell Manager. Without centralized licensing, these individual licenses are restricted to the cell on which they are installed, and all licensing administration tasks have to be performed locally.

NOTE If you have clusters configured in the MoM cell, make sure you identify a cluster client with its virtual hostname.

Why Use Centralized Licensing?

Data Protector allows you to configure centralized licensing for the whole MoM environment. All licenses are installed and kept on the MoM Manager system and can be allocated to specific cells as needed.

Centralized licensing simplifies license management. Licensing administration is performed by the MoM administrator for all cells in the MoM environment. This also includes the distribution and moving of the licenses.

When licenses are installed locally on the Cell Managers, they cannot be moved among the cells without the approval of the *HP Password Delivery Center*. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions on how to move licenses.

Setting Up Centralized Licensing

Prerequisite

If you are consolidating existing Data Protector cells into a MoM environment, send a request to *HP Password Delivery Center* to move the licenses from the existing Cell Managers to the new MoM Manager.

Configuring Centralized Licensing

1. Log on to the MoM Manager and create the `licdistrib.dat` file:
   - On Windows: *<Data_Protector_home>*`\Config\server\cell\licdistrib.dat`
   - On UNIX: `/etc/opt/omni/server/cell/licdistrib.dat`

2. Log on to each client Cell Manager in the MoM environment and create the `lic_server` file with the name of the MoM Manager:

   On Windows:
   <Data_Protector_home>`\Config\server\cell\lic_server`

   On UNIX: `/etc/opt/omni/server/cell/lic_server`

3. Stop and restart Data Protector services on each Cell Manager where you made the changes. See “Restarting Data Protector Services” on page 439.

4. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.

5. In the Scoping Pane, right-click the Cell Manager that has the licensing information you want to change, and then click `Configure Licensing` to open the wizard. The types and numbers of licenses available to your selected Cell Manager are displayed.

   The **USED** column shows the number of licenses assigned to that particular Cell Manager. Increasing the number in this column will correspondingly decrease the number of available licenses, and vice-versa.

   The **AVAILABLE** column shows the number of licenses available to the entire enterprise. This is the number of licenses not taken by any cell within the enterprise environment.

   The **TOTAL** column shows the total number of licenses both used and available in the entire enterprise.

6. Click the `Remote` option to change the licensing from local to remote. Note that **USED** column is changed into **ALLOCATED**.

7. Modify the license configuration. Note that only **ALLOCATED** column is available during the modification process.

Releasing Licenses

To release (give up) a license type, thus increasing the number available, reduce its corresponding number in the **ALLOCATED** column.

Assigning Licenses

To assign a license type, increase its corresponding number in the **ALLOCATED** column by double-clicking it.

8. Click `Finish` to apply the configuration.

9. Repeat the steps for all Cell Managers for which you want to set up the centralized licensing.

NOTE

Data Protector checks the license configuration with the MoM Manager every hour. The licensing status is kept for 72 hours. In case of a communication problem, after this 72 hour period, local licenses are used.

Moving Licenses in the MoM Environment

If you have not configured centralized licensing, you cannot move licenses between cells without the approval of the *HP Password Delivery Center*. This is, however, possible in the MoM environment with configured centralized licensing, where the MoM administrator allocates licenses as needed.

In the example below, assume that the clients from one cell were moved to another. This resulted in the need to move the licenses.

Enterprise Environment Before the Reorganization

Assume that two Cell Managers, Aztec and Mayan, are configured in the enterprise environment with centralized licensing. Aztec is an HP-UX Cell Manager with a Cell Manager for UNIX - Single Drive license. There is also an NDMP server connected in the cell that requires an NDMP Server Backup Extension license. Mayan is also an HP-UX Cell Manager with one Cell Manager for UNIX - Single Drive license.

Reorganization of the Enterprise Environment

The Aztec cell needs to be reorganized, with most of the clients and the NDMP server being transferred to the Mayan cell. Mayan now needs the NDMP Server Backup Extension license. Follow the procedure described below to move the license:

1. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.
2. Right-click the Aztec Cell Manager and then click `Configure Licensing`. The types and numbers of licenses available to the Aztec Cell Manager are displayed. Remove the NDMP Server Backup Extension license.
3. Click `Finish` to apply the configuration.
4. Right-click the Mayan Cell Manager and then click `Configure Licensing`. Add the NDMP Server Backup Extension license.
5. Click `Finish` to apply the configuration.

Enterprise Environment After the Reorganization

The Aztec Cell Manager now has one Cell Manager for UNIX - Single Drive license and the Mayan Cell Manager has a Cell Manager for UNIX - Single Drive license and an NDMP Server Backup Extension license for the NDMP server.

For more information on Data Protector licensing policies, see the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Deactivating Centralized Licensing

Centralized licensing can be deactivated and changed back to local licensing.

Deactivation Procedure

1. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.
2. In the Scoping Pane, right-click the Cell Manager for which you want to deactivate centralized licensing, then click `Configure Licensing` to open the wizard. The types and numbers of licenses available to your selected Cell Manager are displayed.
3. Click the `Local` option to change licensing from remote to local.
4. Click `Finish` to apply the configuration.
5. Repeat the steps for all Cell Managers for which you want to deactivate centralized licensing.
6. Log on to the MoM Manager and mount the following directory:
   - On Windows: *<Data_Protector_home>*`\Config\server\cell`
   - On UNIX: `/etc/opt/omni/server/cell`
7. Rename the `licdistrib.dat` file, for example, to `licdistrib.old`.

The changes will take effect after you stop and restart Data Protector services on the MoM Manager and each Cell Manager where you made the changes. See “Restarting Data Protector Services” on page 439.

Working with a MoM Environment

The Manager-of-Managers interface enables you to configure, manage, and control an enterprise backup environment from a single point.

In the MoM user interface, you can import and export cells, move clients among cells, and distribute the MoM configuration to other cells in the environment.

Other tasks are performed on the MoM Manager in the same way as if you were a local administrator. Follow the standard procedure to configure backup and restore, manage devices and media for a specific cell, configure Data Protector users and user groups, add clients, monitor running sessions and the status of the backup environment, and configure reporting and notifications.

Importing and Exporting Data Protector Cells

Importing a cell into a MoM environment allows it to be centrally managed using the MoM Manager. Exporting a cell will remove it from the enterprise environment.

NOTE

Cluster clients identify themselves to the MoM Manager with their virtual server names. If you import or export a cluster in a MoM environment, use only its virtual server name.

Importing a Cell Manager

1. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.
2. Right-click `Enterprise Clients`, and then click `Import Cell Manager`.
3. Select a Cell Manager you want to import and click `Finish`.

Exporting a Cell Manager

1. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.
2. In the `Scoping Pane`, right-click the Cell Manager you want to export, and then click `Export Cell Manager`.
3. Confirm your choice.

Moving Client Systems Among Cells

Data Protector allows you to move systems among cells. During the process, Data Protector:

- Checks whether the client to be moved is configured in any backup specification and removes all backup objects belonging to this client from backup specifications configured on the initial Cell Manager, while backup objects of other clients remain intact. Data Protector thus ensures no orphan backup objects remain in backup specifications after the client is moved to another cell.
- Checks whether there are any devices configured on the system and leads you through the steps to move devices to another system.
- Checks whether there are media used in the devices on this system and leads you through the steps to move media.

Moving Clients

1. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.
2. Expand the Cell Manager that has the system that you want to move to another cell.
3. Right-click the client system and then click `Move Client System to Other Cell` to open the wizard.
4. Select the target Cell Manager and click `Finish` to move the client.

Distributing the MoM Configuration

Data Protector allows you to create a common user class specification, holidays file settings, global options file settings, and vaulting on all Cell Managers in a MoM environment.

How to Distribute the MoM Configuration

To distribute the `MoM` configuration, follow these steps:

1. In the `Data Protector Manager-of-Managers` click `Clients` in the Context List, right-click `Enterprise Clients`, and then click `Distribute Configuration`.
2. In the `Distribute Configuration` dialog box, select the type of configuration and the Cell Managers to which you want to distribute the selected configuration.
3. Click `Finish` to distribute the configuration.

Configuring Data Protector Users

You can add users or user groups to a MoM environment as you would for a single Cell Manager. This procedure updates all Cell Managers with the new users. See Chapter 4, “Configuring Users and User Groups,” on page 127 for more information about users and user groups.

To configure Data Protector users or user groups, follow these steps:

1. In the `Data Protector Manager-of-Managers`, click `Users` in the Context List.
2. Select a Cell Manager to which you want to add users.
3. In the `Edit` menu, click `Add` and select `Users` if you want to add a new user, or `User Group` if you want to add a new user group.
4. Enter the required information and click `Finish`.

Managing Devices and Media for a Specific Cell

You can configure devices and media for specific devices and media anywhere within your enterprise environment. To do so, follow these steps:

1. In the `Data Protector Manager-of-Managers`, click `Clients` in the Context List.
2. Select the cell that has the devices or media that you want to manage.
3. In the `Tools` menu, click `Device & Media Administration`. In the `Device and Media` context, configure devices and media as if you were a local administrator.

Restoring, Monitoring, and Reporting in an Enterprise Environment

Restoring data in an enterprise environment is the same as restoring data in a single cell environment.

Select data from the appropriate source and restore as described in Chapter 8, “Restore,” on page 335.

Data Protector allows you to monitor currently running or previously run sessions for any cell in the enterprise environment. When you use Web Reporting, you can also get reports on the entire enterprise environment using the MULTICELL item in the Scoping Pane.

See Chapter 9, “Monitoring, Reporting, Notifications, and the Event Log,” on page 379 for more information on how to use these features in an enterprise environment.

11 Managing the Data Protector Internal Database

In This Chapter

This chapter provides information about the Data Protector internal database (IDB) and tasks related to managing the database. It is organized as follows:

About the Data Protector Internal Database

What Is the Data Protector Internal Database (IDB)?

The Data Protector internal database (IDB) is an embedded database, located on the Cell Manager, which keeps information regarding what data is backed up; on which media it resides; the result of backup, restore, copy, and media management sessions; and what devices and libraries are configured.

Why Is the IDB Used?

There are three key reasons for using the IDB:

- Fast and convenient restore

  The information stored in the IDB enables you to browse the files and directories to be restored. You can quickly find the media required for a restore and therefore make the restore much faster.

- Backup management

  The information stored in the IDB enables you to verify the result of backup sessions.

- Media management

  The information stored in the IDB enables you to allocate media during a backup and copy sessions, track media management operations and media attributes, group media in different media pools, and track media locations in tape libraries.

How to Manage the IDB

One of the important steps in setting up your Data Protector backup environment is to configure the IDB. Once the IDB is configured as described in “Configuring the IDB” on page 464, you will be notified if you need to perform any of the IDB maintenance tasks.

The IDB maintenance tasks, and the cases when they need to be performed, are described in “Maintaining the IDB” on page 479.

If you receive error messages, refer to “Troubleshooting the IDB” on page 711 and “Recovering the IDB” on page 494.

For information on IDB limitations, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

The IDB Architecture

The IDB consists of the following parts:

- MMDB (Media Management Database)
- CDB (Catalog Database)
- DCBF (Detail Catalog Binary Files)
- SMBF (Session Messages Binary Files)
- SIBF (Serverless Integrations Binary Files).

Figure 11-1 IDB Architecture

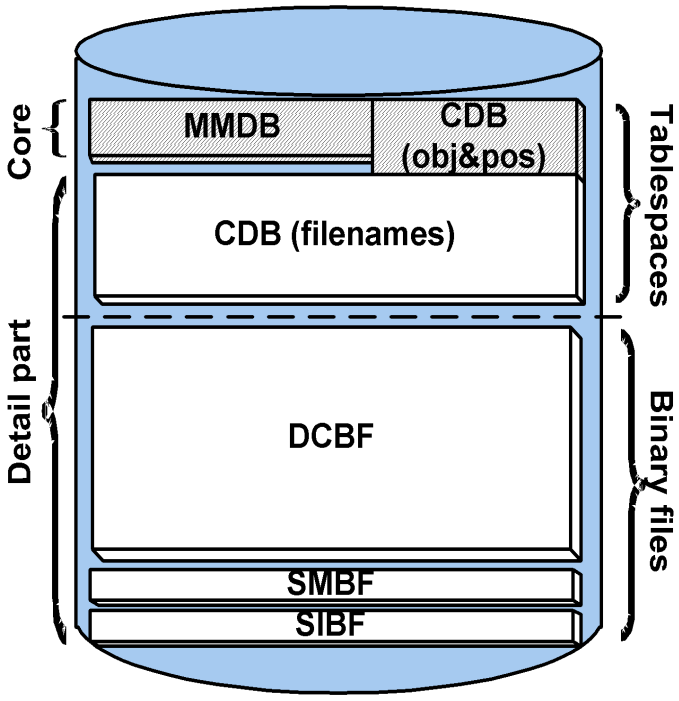

Each of the IDB parts stores specific Data Protector information (records), influences the IDB size and growth in different ways, and is located in a separate directory on the Cell Manager.

MMDB — The Media Management Database stores information about the following:

- configured devices, libraries, library drives, and slots
- Data Protector media
- configured media pools and media magazines

CDB — The Catalog Database stores information about the following:

- Backup, restore, copy, and media management sessions. This is the copy of the information sent to the Data Protector Monitor window.
- Backed up objects, their versions and object copies.
- Pathnames of backed up files (filenames) together with client system names. Filenames are stored only once per client system. The filenames created between backups are added to the CDB.
- Positions of backed up objects on media. For each backed up object, Data Protector stores information about the media and data segments used for the backup. The same is done for object copies and object mirrors.

DCBF — The Detail Catalog Binary Files part stores file version information. This is information about backed up files, such as file size, modification time, attributes/protection, and so on.

One DC (Detail Catalog) binary file is created for each Data Protector medium used for backup. When the medium is overwritten, the old binary file is removed and a new one is created.

SMBF — The Session Messages Binary Files part stores session messages generated during backup, restore, copy, and media management sessions. One binary file is created per session. The files are grouped by year and month.

SIBF — The Serverless Integrations Binary Files part stores raw NDMP restore data. This data is necessary for restore of NDMP objects.

The MMDB and CDB parts are implemented using an embedded database consisting of tablespaces. This database is controlled by the `rds` database server process. All changes to the MMDB and CDB are updated using transaction logs. CDB (objects and positions) and MMDB present the core part of IDB.

The DCBF, SMBF, and SIBF parts of the IDB consist of binary files. Updates are direct (no transactions).

In the Manager-of-Managers (MoM) environment, the MMDB can be moved to a central system to create the Central Media Management Database (CMMDB).

For additional information on each of the IDB parts, refer to the *HP OpenView Storage Data Protector Concepts Guide*.

Configuring the IDB

The IDB configuration helps to manage the following:

- the size of the IDB and available disk space
- the location of the IDB directories
- transaction log usage
- the IDB backup necessary in case of IDB corruption or a disaster
- configuration of the IDB reports and notifications

Once the IDB is configured, it should be maintained only when you are notified about the need.

General Procedure

This is the general procedure for IDB configuration:

1. Allocate disk space for future needs.
   Refer to “Allocating Disk Space for Future Use” on page 464.
2. Prepare for the IDB recovery.
   Refer to “Preparing for IDB Recovery” on page 466.
3. Set the appropriate reports and notifications about the IDB.
   Refer to “Configuring the Database Reports and Notifications” on page 477.

Allocating Disk Space for Future Use

Over time, the IDB can occupy a considerable amount of disk space on the Cell Manager. You need to plan in advance and consider the allocation of the disk space for future IDB needs.

Prerequisites

- You need to understand the key factors influencing the IDB growth, such as number of files, file dynamics, environment growth, and so on. Refer to the *HP OpenView Storage Data Protector Concepts Guide* for additional information.
- You need to set logging level and catalog protection polices according to your environment requirements and available disk space. To get this information, together with the usage recommendations for logging level and catalog protection settings, refer to the *HP OpenView Storage Data Protector Concepts Guide*.

- You need to estimate future IDB size (disk space necessary for future IDB needs). Refer to the *HP OpenView Storage Data Protector Concepts Guide* for the IDB size estimation.

How Much Disk Space Is Needed?

The disk space needed to accommodate the IDB varies significantly as a function of many configuration aspects and policies used in defining and operating backups.

The following simplified scenario of an environment requires about 900 MB of disk space for the IDB after 3 months, with very little growth afterwards:

- 100 systems to be backed up (10,000 files each; without mail-servers)
- 350 GB total data volume
- filesystem backups with typical dynamics of 3% of new files per month
- one full backup and four incremental backups per week
- logging level set to Log all (to allow convenient browsing of filenames before restore). This is the most demanding logging option.
- catalog protection setting of three months for the full backups and two weeks for the incremental backups.

Note that large configurations or long catalog protection periods in the IDB can require more than 20 GB for the IDB.

A detailed estimation can be performed using the IDB Capacity Planning Tool located on the Cell Manager:

- On UNIX: `/opt/omni/doc/C/IDB_capacity_planning.xls`
- On Windows: `<Data_Protector_home>\docs\IDB_capacity_planning.xls`

What to Plan for in Advance

Typically the IDB grows rapidly in the beginning, until the catalog retention periods have been reached. After that, the growth of the IDB is mainly determined by the dynamics of systems that have a large percentage of new files per month and the growth of the environment itself (new systems to be backed up).

It is important to understand the various IDB growth functions:

- The filenames part of the IDB is proportional to the total number of filenames in the cell (but not the data volume and the number of backups). Typically the filename growth is moderate, with the exception of some mail servers or other systems with a large amount of automatically generated files.
- The file versions part of the IDB grows with the number of backups and object copies, the number of files in the cell, and the duration of the catalog protection.
- Using the IDB transaction log files requires additional disk space. Size prediction is not simple. Dominating factors influencing the size are the number of new filenames being backed up and the total backup activities (or weeks, if scheduled backups are the main operation) between IDB backups.

Preparing for IDB Recovery

You need to make advance preparations in order to be able to recover the IDB at any point in time. The IDB recovery restores information stored in the IDB and is essential for the restore of backed up data in case the Cell Manager crashes.

Prepare for IDB recovery by:

- Considering recommendations for optimizing robustness. Refer to “Robustness Considerations” on page 466.
- Relocating IDB directories. Refer to “The IDB Directories” on page 467.
- Enabling of transaction logs. Refer to “Enabling Transaction Logs” on page 473.
- Configuring the IDB backup and backing it up regularly. Refer to “Configuring the Database Backup” on page 474.

Robustness Considerations

This section outlines some aspects and recommendations you should consider to optimize robustness and reliability of the IDB.

- The core part of the IDB, which contains CDB (objects & positions) and MMDB, is essential for the operation of Data Protector.

- The DCBF and SMBF parts of the IDB are not required for basic operation of Data Protector, such as backup and restore. However, if they are not present, restore becomes less convenient (no filename browsing) and the session messages are lost.
- If the IDB recovery file and the IDB transaction logs are lost, normal operation would not be affected, but IDB restore would be considerably more difficult, and replaying the IDB data generated since the last IDB backup would not be possible. Instead, the used media would need to be reimported.

Recommendations to Optimize Robustness

- Ensure that the IDB recovery file and the transaction logs do not reside on the same physical disk as the core part of the IDB.

  This is to ensure a fast and simple restore of the IDB in case the physical disk A crashes. It also for the replay of the transactions that happened since the last IDB backup. Refer to Figure 11-2.
- Relocating the DCBF, SMBF, and SIBF parts to a disk other than the one that holds the core part of the IDB is also recommended, but less important. If this is done, the load on disk A is reduced significantly and IDB space management is easier, because these parts are usually the largest part of the IDB.

TIP

Following the recommendations to optimize robustness will also increase performance, allowing for more backup activities on the Cell Manager system.

The IDB Directories

The IDB is located on the Cell Manager. In order to improve space management, you may want to relocate some IDB directories.

Limitations

- The IDB files can be located only on locally attached disks (not using NFS or on shared disks).

- If the IDB is installed in a cluster, it must be installed on disks in the cluster group (Microsoft Cluster Server) or cluster package (MC/ServiceGuard).

Table 11-1 Location of IDB Directories on Windows

IDB	Location on Windows
Tablespaces (CDB and MMDB)	*<Data_Protector_home>*\db40\datafiles
Binary files (DCBF, SMBF, SIBF)	• *<Data_Protector_home>*\db40\dcbf • *<Data_Protector_home>*\db40\msg • *<Data_Protector_home>*\db40\meta
Transaction logs	*<Data_Protector_home>*\db40\logfiles\syslog
IDB recovery file	*<Data_Protector_home>*\db40\logfiles\rlog

Table 11-2 Location of IDB Directories on UNIX

IDB	Location on UNIX
Tablespaces (CDB and MMDB)	/var/opt/omni/server/db40/datafiles
Binary files (DCBF, SMBF, SIBF)	• /var/opt/omni/server/db40/dcbf • /var/opt/omni/server/db40/msg • /var/opt/omni/server/db40/meta
Transaction logs	/var/opt/omni/server/db40/logfiles/syslog
IDB recovery file	/var/opt/omni/server/db40/logfiles/rlog

Figure 11-2 Recommended Location of IDB Directories

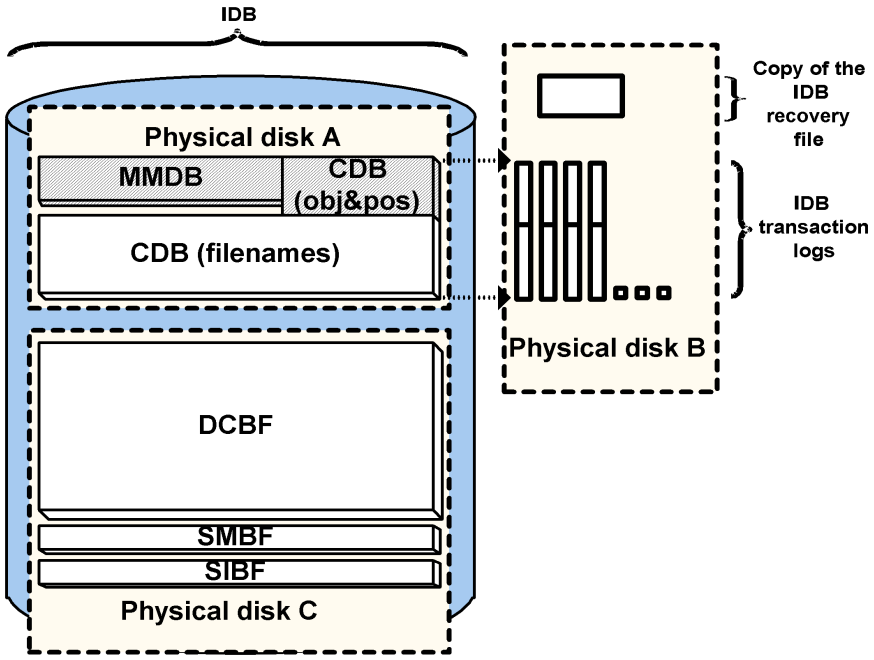

Relocating the IDB Directories

You can change the location of any of the following IDB directories:

- the `datafiles` directory, containing CDB (objects, positions, and filenames) and MMDB parts of the IDB
- the `logfiles` directory, containing transaction logs and the IDB recovery file
- the `dcbf` directory, containing the DCBF part of the IDB
- the `msg` directory, containing the SMBF part of the IDB
- the `meta` directory, containing the SIBF part of the IDB

You can also modify the directory path for the `dcbf` directory (using the Data Protector user interface) and for the `msg` and `meta` directories (using the global options file).

NOTE

On UNIX, you can use symbolic links to relocate the directories, but the links are not allowed beneath the `/var/opt/omni/server/db40/datafiles` directory.

Follow the described below to relocate the IDB directories:

1. Stop all backups and other Data Protector activities and run the `omnisv -stop` command to stop the Data Protector services:
   - On Windows: *<Data_Protector_home>*`\bin\omnisv -stop`
   - On UNIX: `/opt/omni/sbin/omnisv -stop`

   If the IDB is installed on MC/ServiceGuard, run the `cmhaltpkg` *<pkg_name>* command on the active node to stop the Data Protector package, where *<pkg_name>* is the name of the Data Protector cluster package.
2. Rename the *<IDB_dir>* directory that you want to move to *<IDB_dir>*`.save`. For example, to relocate the transaction logs and the IDB recovery file, rename *<Data_Protector_home>*`\db40\logfiles` to *<Data_Protector_home>*`\db40\logfiles.save` (on Windows), or `/var/opt/omni/server/db40/logfiles` to `/var/opt/omni/server/db40/logfiles.save` (on UNIX).
3. Create a new empty directory with the same relative path, for example *<Data_Protector_home>*`\db40\logfiles` on Windows systems, or `/var/opt/omni/server/db40/logfiles` on UNIX systems.
4. On Windows, add a new disk or mount a new volume at an NTFS folder as *<Data_Protector_home>*`\db40\`*<IDB_dir>*. For example, mount it as *<Data_Protector_home>*`\db40\logfiles`.

   On UNIX, add a new disk or create a new logical volume and mount it as `/var/opt/omni/server/db40/`*<IDB_dir>*. For example, mount it as `/var/opt/omni/server/db40/logfiles`.
5. Copy the contents of *<IDB_dir>*`.save` into *<IDB_dir>* on the new disk or new volume.

6. Run the `omnisv -start` command to start the Data Protector services:

   - On Windows: `<Data_Protector_home>\bin\omnisv -start`
   - On UNIX: `/opt/omni/sbin/omnisv -start`

   If the IDB is installed on MC/ServiceGuard, run the `cmrunpkg <pkg_name>` command on the active node to start the Data Protector package.

Creating an Additional Copy of the IDB Recovery File

Creating an additional copy of the IDB recovery file prevents you from losing important data for IDB recovery.

Use the following steps to make another copy of the IDB recovery file:

1. Stop all backups and other Data Protector activities and run the `omnisv -stop` command to stop the Data Protector services.

   - On Windows: `<Data_Protector_home>\bin\omnisv -stop`
   - On UNIX: `/opt/omni/sbin/omnisv -stop`

   If the IDB is installed on MC/ServiceGuard, run the `cmhaltpkg <pkg_name>` command on the active node to stop the Data Protector package, where `<pkg_name>` is the name of the Data Protector cluster package.

   If the IDB is installed on Microsoft Cluster Server, take the `OBVS_VELOCIS` cluster group offline using the Cluster Administrator utility and stop the `Inet` service on the active node.

2. Edit the global options file by setting the value for the `RecoveryIndexDir` variable: specify an additional location where Data Protector makes a copy of the IDB recovery file, `obrindex.dat`. It is recommended to specify a different physical disk.

3. Run the `omnisv -start` command (on UNIX, located in the `/opt/omni/sbin` directory) to start the Data Protector services.

   - On Windows: `<Data_Protector_home>\bin\omnisv -start`
   - On UNIX: `/opt/omni/sbin/omnisv -start`

   If the IDB is installed on MC/ServiceGuard, run the `cmrunpkg <pkg_name>` command on the active node to start the Data Protector package.

If the IDB is installed on Microsoft Cluster Server, bring the `OBVS_VELOCIS` and `OBVS_MCRS` cluster groups online using the Cluster Administrator utility and restart the `Inet` service.

Creating or Relocating DC Directories

Creating a DC Directory

Create a DC directory using the `Database` context in the `Data Protector Manager`. See Figure 11-3. For detailed steps, refer to the online Help index keyword “creating DC directories”.

Relocating a DC Directory

To change the location of a DC directory, proceed as follows:

1. Create a new DC directory on a new location, using the Data Protector user interface. See Figure 11-3.
2. Verify that the new DC directory has been created and has enough disk space.
3. Move DC binary files from the source DC directory to the new DC directory.
4. Run the `omnidbutil -remap_dcdir` command to update the pathnames of DC binary files.
5. Remove the old DC directory from the list of configured DC directories.

Figure 11-3 Creating a DC Directory

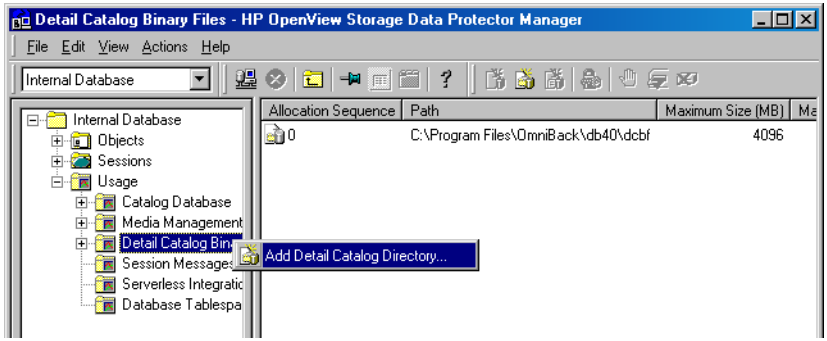

Enabling Transaction Logs

Transaction logs used by the MMDB and CDB parts of the IDB are created in the following directory:

- On Windows: `<Data_Protector_home>\db40\logfiles\syslog`
- On UNIX: `/var/opt/omni/server/db40/logfiles/syslog`

By default, transaction logging is disabled. If enabled, transaction logs from the latest IDB backup are kept until the next backup. If a transaction log file reaches 2 MB, a new one is created. An IDB backup removes all existing transaction logs, except for the currently active one, and starts to create new ones.

Why Enable Transaction Logs?

In order to perform the most convenient IDB recovery method, **guided autorecovery**, with replaying logs, you need to have available the transaction log files created after the last IDB backup.

Disk Space Considerations

The disk space used for the transaction logs depends on the amount of backups done between two IDB backups. If the filenames are already in the IDB, the amount is fairly small and the reserved space of 100 MB should be enough for most cases. If new filenames are backed up, the disk space usage is considerable (estimation is 200 bytes per filename). It is recommended to enable transaction logs *after* the first full backup of the environment (when all filenames are stored in the IDB).

How to Enable the Transaction Logs

To enable transaction logs, proceed as follows:

1. Stop all backups and other Data Protector activities and run the `omnisv -stop` command to stop Data Protector services:
   - On Windows: `<Data_Protector_home>\bin\omnisv -stop`
   - On UNIX: `/opt/omni/sbin/omnisv -stop`
2. Ensure that there is enough disk space in the following directory:
   - On Windows: `<Data_Protector_home>\db40\logfiles\syslog`
   - On UNIX: `/var/opt/omni/server/db40/logfiles/syslog`
3. Edit the `velocis.ini` file (`rdmserver.ini` on HP-UX 11.23) and set the value of the Archiving parameter to 1.
   - On Windows: `<Data_Protector_home>\db40\datafiles\catalog\velocis.ini`
   - On UNIX: `/var/opt/omni/server/db40/datafiles/catalog/velocis.ini` (`rdmserver.ini` on HP-UX 11.23)
4. Start the Data Protector services using the `omnisv -start` command:
   - On Windows: `<Data_Protector_home>\bin\omnisv -start`
   - On UNIX: `/opt/omni/sbin/omnisv -start`

Configuring the Database Backup

An essential part in the IDB configuration is to configure the backup of the IDB itself. Once the IDB backup is performed regularly, the most important preparation for recovery in case of a disaster is done. The IDB recovery is essential for restore of other backed up data in the event that the Cell Manager crashes.

How to Configure the IDB Backup

Configure the IDB backup like any standard backup, but be sure to select the `Internal Database` object and specify the object options in the Backup Object Summary page of the IDB backup specification. For detailed steps, see the Data Protector online Help index keyword “configuring IDB backups”.

Figure 11-4 Selecting the Internal Database Object

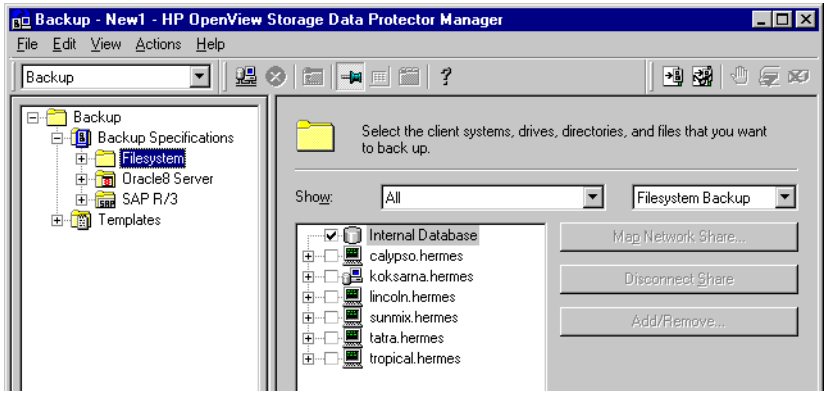

Recommended IDB Backup Configuration

We recommend the following when configuring the IDB backup:

- Create a separate backup specification for the IDB. This simplifies scheduling and restoring in case of a disk crash. To create an IDB backup specification, follow the standard backup procedure, but select the `Internal Database` object.
- Schedule the IDB backup to be performed once per day. This ensures that you always have an almost up-to-date backup of the IDB.
- Perform the IDB backup using a separate media pool on separate media, on a specific device. Make sure you know which media you use for the IDB backup. You can configure a `Session Media Report` to be informed about the media used for the backup. This greatly simplifies eventual restore. If possible, use a device locally connected to the Cell Manager. Refer to “Data Protector Reporting” on page 388.
- Set data protection and catalog protection to a few days only. Set these options such that you have at least the last two IDB backup versions protected.
- Always have the `Check Internal database` option enabled (default). See Figure 11-5.
- Do not overwrite the previous IDB backup with the new one (keeping several copies is suggested).

What Happens During the IDB Backup

During the IDB backup, Data Protector does the following:

- Checks the consistency of the IDB, thus preventing the backing up and later restoring a corrupted IDB. For this check to happen, you need to have the `Check Internal database` option enabled (default).

  The check operation takes approximately 1.5 hours for a 10 GB database with a `fnames.dat` file size of 1 GB.

- Backs up the IDB online (while the IDB is in use). Therefore, other backup or restore sessions can run while the IDB backup runs. But, if possible, back up the IDB when no other backup and restore activities are in progress.

- Backs up all Data Protector configuration data, including the data on devices, backup specifications, and schedules. This simplifies recovery in case of a disaster.

NOTE Only one IDB backup can run at a time.

Disabling the Automatic Check Before Backup

By default, Data Protector automatically checks the consistency of the IDB before the database is backed up. You can enable or disable the automatic consistency check. It is strongly recommended that you keep the automatic IDB check enabled.

In environments where the Cell Manager is used heavily and the time needed to perform the check of the IDB creates a problem, you may need to disable the `Check Internal database` option. In such cases, consider the following suggestions:

- Schedule the IDB backup with the IDB check option enabled to be performed when the automatic check activity is acceptable.
- Schedule the daily IDB backup with the IDB check option disabled.
- Keep at least the most recent checked IDB backup.

For detailed steps, refer to the online Help index keyword “disabling automatic IDB checks”.

Figure 11-5 The Check Internal database Option (enabled by default)

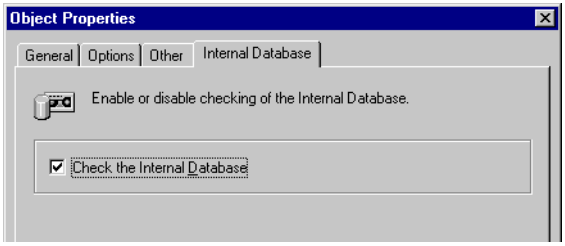

Configuring the Database Reports and Notifications

Configure the IDB reports and notifications so that you are notified if you need to perform IDB maintenance tasks such as purging the IDB, extending the size of the IDB, and so forth.

IDB Reports

The following list presents the IDB reports:

IDB Purge Preview Report — Lists the number of filenames per client, the estimated number of obsolete filenames per client, and the estimated duration of the filename purge session per client.

Report on System Dynamics — Reports on the dynamics of the growth of filenames on a particular client.

IDB Purge Report — Lists the filenames that have been removed from the IDB.

IDB Size Report — Lists the sizes of the individual parts of the IDB.

There are also other Data Protector reports to be considered. For example, the `List of Sessions` report shows the number of files backed up in one session. Refer to “Data Protector Reporting” on page 388 for more information.

IDB Notifications

The following list presents IDB notifications:

IDB Space Low — Informs you if the IDB is running out of space.

IDB Tablespace Space Low	Informs you if a tablespace in the IDB is running out of space.
IDB Purge Needed	Informs you if you need to run the filename purge of the IDB.
IDB Corrupted	Informs you if any kind of IDB corruption is detected.

For detailed information on each report and notification, refer to “Report Types” on page 390.

Procedure for Configuring IDB Reports and Notifications

Configure the IDB reports and notifications using the `Reporting` context in the `Data Protector Manager`. For detailed steps, refer to the online Help index keywords “configuring IDB reports” and “configuring IDB notifications”.

What’s Next?

Once you have configured the IDB reports and notifications, you have completed the last step in IDB configuration. If you need to perform any IDB maintenance task, you will be notified by Data Protector. Now, you can continue to set up your environment.

Maintaining the IDB

Once you have configured the IDB, you need to perform IDB maintenance tasks in the following cases:

- the IDB is running out of space

  This requires you to provide more space for the IDB or decrease the volume of data written to the IDB. If configured, the `IDB Space Low` or `IDB Tablespace Space Low` notification informs you about this.

- the IDB needs a file version purge

  The granularity of the purge is the complete medium. This means that the catalog protection for all object versions on the medium must expire before the file version records are purged. Then, the related medium binary file containing the detail catalog is removed. This purges many file versions in a very short time. This happens automatically on a daily basis. Obsolete sessions and messages are also purged automatically.

- the IDB needs a filename purge

  This requires you to perform a filename purge operation. In an environment that can generate 100,000 obsolete filenames per day, the frequency is once per year. You will be notified automatically if the filenames purge is needed. Filenames purge can be executed selectively on a per host basis. The operation must run exclusively, so no backups can run concurrently. This purge takes more time to execute than in the previous version of Data Protector.

- the dynamics of the client system are high or critical

  This means that there is a large number of new or changed filenames over time. This has an impact on the IDB if the filenames are logged to the IDB. If configured, the `System Dynamics` report informs you about this.

- you want to move the IDB to a different Cell Manager

- you want to check the size of the IDB

  The `IDB Size` report informs you of the size of the IDB.

- the IDB does not work properly (might be corrupted) and you want to check its consistency

The `IDB Corrupted` notification informs you about IDB corruption.

Refer to Table 11-3 for information on which of the maintenance tasks you can perform in which cases.

Table 11-3 IDB Maintenance Tasks

Situation	Which Task Can You Perform?	Reference
The IDB is running out of space	• Extend the size of the IDB • Purge the IDB filenames • Reduce the growth of the IDB • Reduce the current size of the IDB	• "Extending the IDB Size" on page 485 • "Purging Obsolete Filenames" on page 485 • "Reducing the IDB Growth" on page 482 • "Reducing the IDB Size" on page 483
Obsolete filenames in the IDB	• Purge IDB filenames	• "Purging Obsolete Filenames" on page 485
The dynamics of a client system are high or critical	• Reduce the growth of the IDB • Extend the size of the IDB	• "Reducing the IDB Growth" on page 482 • "Extending the IDB Size" on page 485
You want to check the size of the IDB	• Check the size of the IDB	• "Checking the IDB Size" on page 487
The IDB does not work properly (might be corrupted)	• Check the consistency of the IDB	• "Checking the Consistency of the IDB" on page 488
You want to move the IDB to a different Cell Manager	• Move the IDB to a different Cell Manager on the same platform	• "Moving the Database to a Different Cell Manager" on page 489

Reducing the IDB Growth

You can reduce the growth of the IDB by reducing the logging level and catalog protection settings of your backup and object copy specifications. These actions do not influence the current size of the IDB, but they do influence its future growth.

The effect of reducing the logging level is a reduction in browse comfort at restore time.

The effect of reducing the catalog protection is that browsing is not possible for some restores (namely of those backups that have exceeded the catalog protection).

Refer to the *HP OpenView Storage Data Protector Concepts Guide* for information on key factors and tunable parameters for IDB growth and performance, as well as for usage recommendations.

How to Reduce the IDB Growth

Modify the backup specifications by changing the logging level and catalog protection settings using the `Data Protector Backup` context in the `Data Protector Manager`. See Figure 11-6. For detailed steps, refer to the online Help index keyword “reducing IDB growth”.

By reducing the logging level settings for a backup specification, you reduce the amount of data (files/directories) that will be stored in the IDB (`Log all -> Log files -> Log directories -> No log`).

By reducing the catalog protection, you reduce the protection for the restore browse information in the IDB only. The information is still stored on media.

Figure 11-6 **Changing Logging Level and Catalog Protection Settings for Backup**

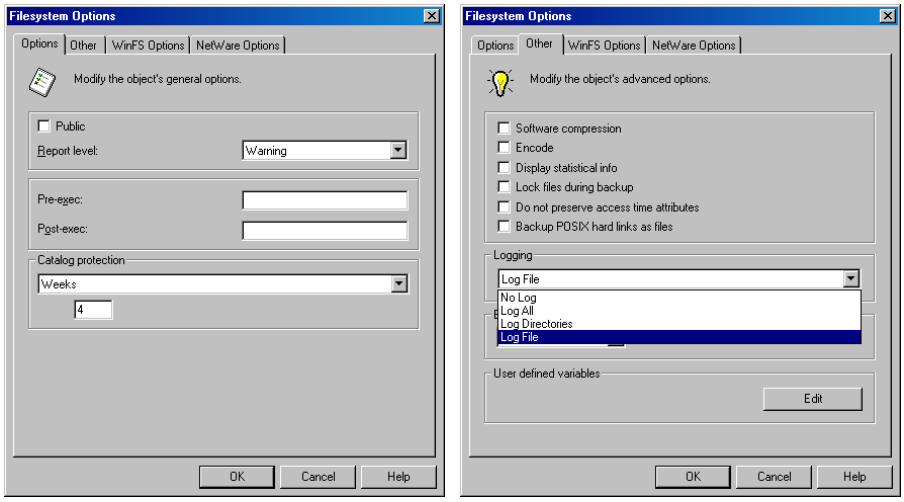

Reducing the IDB Size

You can reduce the IDB size by changing the catalog protection settings for a complete backup or object copy session (all objects in the session) or for specific objects only.

The effect of reducing the catalog protection is that browsing is not possible for some restores (namely of those backups that have exceeded the catalog protection).

This action does not influence the future growth of the IDB.

When Does the Change Take Effect?

The change takes effect:

- If catalog protection is removed from all objects on the medium.
- Once per day (by default, at noon) when Data Protector automatically removes obsolete data from the IDB. The time can be specified in the `DailyMaintenanceTime` global options variable, using the twenty-four hour clock notation. Refer to “Global Options File” on page 613.

You can start the purge immediately by running the `omnidbutil -purge -dcbf` command. Refer to the `omnidbutil` man page for information on removing other obsolete items from the IDB.

By changing the catalog protection, you change protection in the IDB only. The information is still stored on media. Therefore, if you export a medium and import it back, Data Protector rereads information about catalog protection from the media.

How to Reduce the IDB Size

Change the catalog protection setting using the `Internal Database` context in the `Data Protector Manager`. See Figure 11-7 and Figure 11-8. For detailed steps, refer to the online Help index keyword “reducing IDB current size”.

Figure 11-7 Changing Catalog Protection for a Session

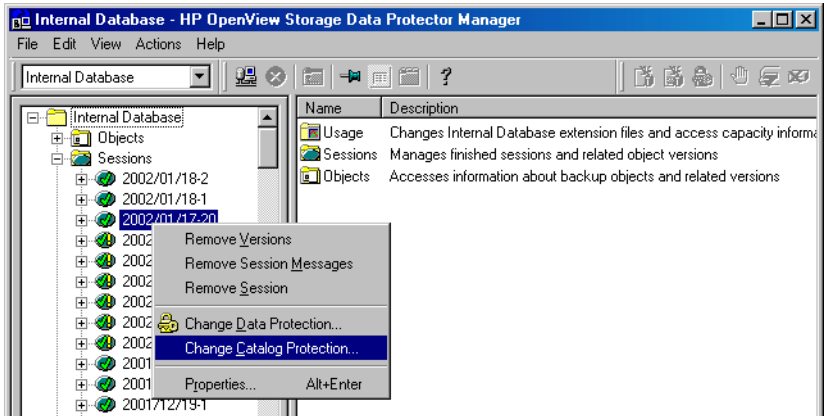

Figure 11-8 Changing Catalog Protection for an Object

Purging Obsolete Filenames

During the purge process, Data Protector automatically checks for and purges obsolete filenames from the IDB to free up space for new information. A filename becomes obsolete when there are no file versions for the filename in the IDB.

Use the `Internal Database Purge Preview Report` and `Internal Database Purge Report` to get more information about the purge. Refer to "Configuring the Database Reports and Notifications" on page 477.

How to Purge Obsolete IDB Filenames

Purge the IDB when no other backups are running on the Cell Manager. Run the following command:

```
omnidbutil -purge -filenames
```

You can limit the purge to one or more clients by running the following command:

```
omnidbutil -purge -filenames <host_1 ... host_n>
```

Data Protector skips purging filenames on the clients that have fewer than 1,000,000 obsolete filenames. In order to purge filenames on these clients as well, use the `-force` subcommand.

Extending the IDB Size

It is required to extend the IDB size for the following reasons:

- The space for the filenames is consumed and another `fnames.dat` file is needed, or some other tablespace needs to be extended.
- More disk space is needed for the detail part of the IDB (file versions and attributes).

You can extend the size of the IDB in either of two ways:

- By creating new DC (Detail Catalog) directories and, possibly, locating them on different disks.
- By creating additional `fnames.dat` files.
- By extending other tablespaces.

Creating New DC Directories

You create a new DC directory using the `Internal Database` context in the `Data Protector Manager`. See Figure 11-3 on page 473. For detailed steps, refer to the online Help index keyword “creating DC directories”.

Creating New fnames.dat Files

What Are fnames.dat Files?

The `fnames.dat` files contain information on the names of backed up files. Typically, these files occupy about 20% of the IDB. The default size of a `fnames.dat` file is 2 GB; the maximum size is 32 GB.

How to Create fnames.dat Files

You add new `fnames.dat` files using the `Internal Database` context in the `Data Protector Manager`. See Figure 11-9. For detailed steps, refer to the online Help index keyword “creating fnames.dat files”.

On Windows Cell Managers, it is recommended that the extension files are created on the same logical disk as the IDB.

The IDB extension files are backed up as a part of the IDB backup and are restored using the IDB recovery.

Figure 11-9 Creating a New `fnames.dat` File

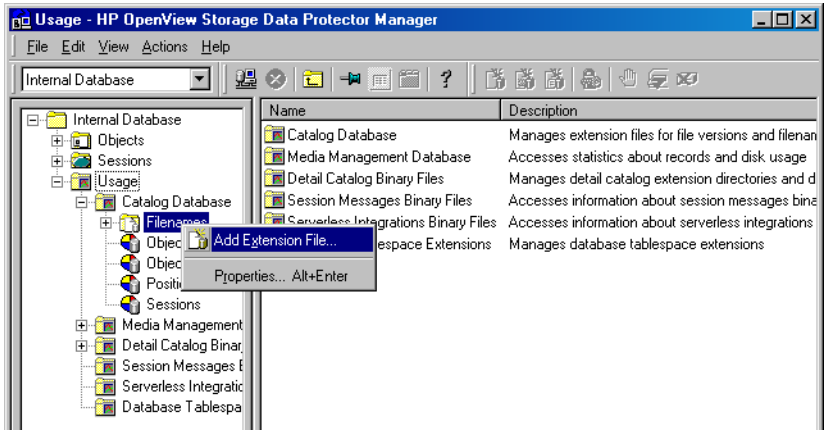

Extending Other Tablespaces

By default, the IDB Tablespace Space Low notification is triggered when 85% of the space allocated for a specific tablespace is used. On how to extend the specified tablespace, refer to the online Help index keyword “extending IDB size”.

Checking the IDB Size

You can check the current size of the IDB parts using the Data Protector GUI.

Also, if configured, the `IDB Size Report` as well as the `IDB Space Low` and `IDB Tablespace Space Low` notifications inform you about the IDB size.

How to Check IDB Size

Check the size of the IDB parts, CDB, MMDB, DCBF, SMBF, and SIBF using the `Internal Database` context in the `Data Protector Manager`. See Figure 11-10. For detailed steps, refer to the online Help index keyword “checking, IDB size”.

Figure 11-10 Checking the Size of the fnames.dat File (CDB Part)

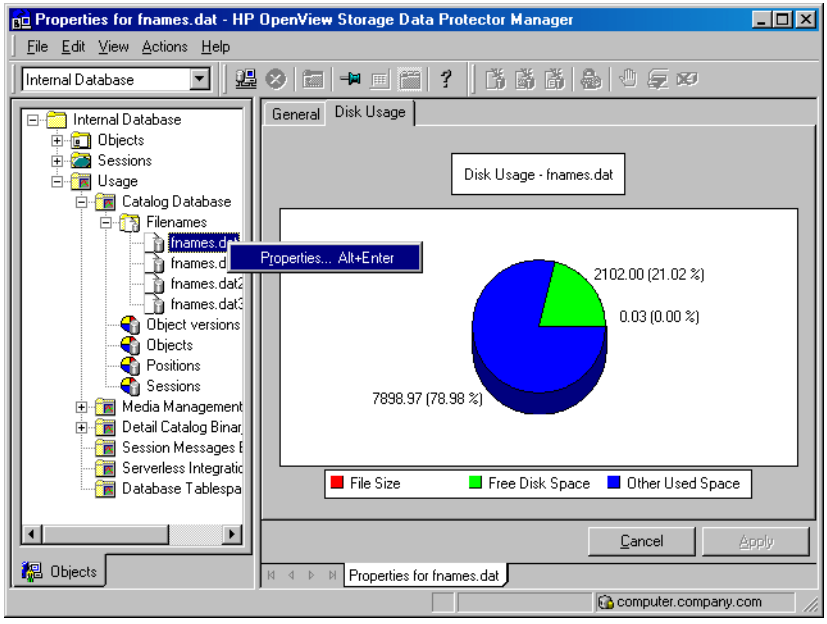

Checking the Consistency of the IDB

Data Protector by default checks the consistency of the IDB before the IDB is backed up. This is extremely important for recovering the IDB and backed up data in case of a disaster.

Additionally, you can manually perform the following IDB checks:

Check of the core part of the IDB	Checks the MMDB (Media Management Database) and CDB (Catalog Database) parts without information about filenames. It takes approximately 5-10 minutes for a medium size IDB. To perform it, run the `omnidbcheck -core` command.
Filenames check	Checks IDB information about filenames. It takes approximately one hour for a medium size IDB. To perform it, run the `omnidbcheck -filename` command.
Simple check of the DCBF part	Checks if the DC binary files exist and what their size is. It takes approximately 10-30 seconds for a medium size IDB. To perform it, run the `omnidbcheck -bf` command.
Complete check of the DCBF part	Checks the consistency of media positions and the DC binary files. It takes approximately 10 minutes for each GB of the DCBF part. To perform it, run the `omnidbcheck -dc` command.
Check of the SMBF part	Checks for the presence of session messages binary files. It takes approximately 5-10 minutes. To perform it, run the `omnidbcheck -smbf` command.
Check of the SIBF part	Checks the consistency of object versions and Serverless Integrations Binary Files. It takes approximately 10 minutes for each GB of the SIBF part. To perform it, run the `omnidbcheck -sibf` command.
Quick check	Checks the core part (MMDB and CDB), filenames, and the DCBF part. It takes approximately two and a half hours for a medium size IDB. To perform it, run the `omnidbcheck -quick` command.

Extended check	Checks the critical part (MMDB and CDB), filenames, the DCBF part, and the DC part. To perform it, run the `omnidbcheck -extended` command.

If you run into problems using the IDB, refer to the troubleshooting section “Troubleshooting the IDB” on page 711 and “Recovering the IDB” on page 494.

Moving the Database to a Different Cell Manager

You can move the IDB to a different Cell Manager that runs on the same operating system by following the steps below:

1. Stop all Data Protector services on the source and target systems using the `omnisv -stop` command:

   - On Windows: `<Data_Protector_home>\bin\omnisv -stop`
   - On UNIX: `/opt/omni/sbin/omnisv -stop`

   If the IDB is installed on MC/ServiceGuard, run the `cmhaltpkg <pkg_name>` command on the active node to stop the Data Protector package, where `<pkg_name>` is the name of the Data Protector cluster package.

   If the IDB is installed on Microsoft Cluster Server, take the `OBVS_VELOCIS` cluster group offline using the Cluster Administrator utility and stop the `Inet` service on the active node.

2. Copy the following IDB files to the target system:

   - Tablespaces to the same relative pathname:

     On Windows systems: `<Data_Protector_home>\db40\datafiles` to `<Data_Protector_home>\db40\datafiles`

     On UNIX systems: `/var/opt/omni/server/db40/datafiles` to `/var/opt/omni/server/db40/datafiles`

   - Extension files to the same full pathname as they were on the source system. You can get a list of the files by using the `omnidbutil -extendinfo` command.

- SMBF files to the same relative pathname:

  On Windows systems: *<Data_Protector_home>*\db40\msg to *<Data_Protector_home>*\db40\msg

  On UNIX systems: /var/opt/omni/server/db40/msg to /var/opt/omni/server/db40/msg

- SIBF files to the same relative pathname:

  On Windows systems: *<Data_Protector_home>*\db40\meta to *<Data_Protector_home>*\db40\meta

  On UNIX systems: /var/opt/omni/server/db40/meta to /var/opt/omni/server/db40/meta

- DC directories to the same or other locations. You can get the list of DC directories using the omnidbutil -list_dcdir command.

3. Start Data Protector services on the target system using the omnisv -start command:

   - On Windows: *<Data_Protector_home>*\bin\omnisv -start
   - On UNIX: /opt/omni/sbin/omnisv -start

   If the IDB is installed on MC/ServiceGuard, run the cmrunpkg *<pkg_name>* command on the active node to start the Data Protector package.

   If the IDB is installed on Microsoft Cluster Server, bring the OBVS_VELOCIS and OBVS_MCRS cluster groups online using the Cluster Administrator utility and restart the Inet service.

4. Run the omnidbutil -change_cell_name command.
5. Relocate DC directories on the target system.
6. Run the omnidbutil -remap_dcdir command for Data Protector to refresh the new locations of the DC directories.

Restoring the IDB

If you have backed up the IDB using the standard procedure, you can restore it using the methods described in this section.

For a detailed description of how to handle the IDB recovery in case of a disaster, refer to “Recovering the IDB” on page 494.

Restoring the IDB consists of two phases:

1. Restoring the IDB to a temporary location.

IMPORTANT This step is necessary because the IDB is in use during the restore. If you try to restore the IDB to the original location, you will corrupt the IDB.

2. Moving the IDB to the original location.

Ensure that you have enough disk space before you begin.

Restoring the IDB to a Temporary Directory

To restore the IDB files to a temporary location, proceed as follows:

1. In the `Data Protector Manager`, switch to the `Restore` context.
2. Expand the `Internal Database` item.
3. Expand the client system with the IDB backup and then click the database object to open the `Source` property page.
4. In the `Source` property page, select the IDB directories that you want to restore. By default, the last backup version is selected for restore. If you want to restore any other version, right-click the selected directory and click `Restore Version`. From the `Backup version` drop-down list, select the backup version that you want to be restored. Click `OK`.
5. In the `Destination` property page select `Restore to new location` option and select the temporary directory for IDB files (for example, the `temp` directory).

NOTE You should not select the *<Data_Protector_home>* directory, as this directory is the original location of the IDB.

If you want to restore to a different system, specify the new Cell Manager's name.

6. Click `Restore`.

Moving the IDB to the Original Location

After you have restored the IDB to a temporary location, you need to move the IDB directories to their original location. Proceed as follows:

On a UNIX Cell Manager

1. Stop all running Data Protector sessions and close the Data Protector GUI. This prevents access to the IDB.
2. Stop all Data Protector processes by running:

   `/opt/omni/sbin/omnisv -stop`
3. Move the existing IDB directories:

   `/var/opt/omni/server/db40` and `/etc/opt/omni/server`

   This prevents merging of old and new files.
4. Copy the IDB directories from the temporary directory to the original directories

   `/var/opt/omni/server` and `/etc/opt/omni/server`

   If your extension files were located on some other directory, be sure to copy them to the original disk and directory as well.
5. Restart the Data Protector processes unless you are moving the IDB to original location as a part of migration of your existing Cell Manager to an IA-64 based HP-UX 11.23 system. Run the following command:

   `/opt/omni/sbin/omnisv -start`

On a Windows Cell Manager

1. Stop all running Data Protector sessions and close the Data Protector GUI. This prevents access to the IDB.

2. Stop all Data Protector services by running:

   ```
   <Data_Protector_home>\bin\omnisv -stop
   ```

3. Move the existing IDB directories (`db40` and `config`) from the *<Data_Protector_home>* directory. This prevents merging of old and new files.

4. Copy the IDB directories from the temporary directory to the original directory *<Data_Protector_home>*.

   If your extension files were located on some other directory, be sure to copy them to the original disk and directory as well.

5. Restart the Data Protector services by running:

   ```
   <Data_Protector_home>\bin\omnisv -start
   ```

TIP You can check the consistency of the IDB after the restore. See “Checking the Consistency of the IDB” on page 488 for more information.

Recovering the IDB

When Is Recovery Needed?

IDB recovery is needed if all or some of the IDB files are not available or are corrupted.

There are three levels of IDB issues, each with its own techniques for repair:

- Troubleshoot the IDB problems that are caused by OS configuration issues, such as not mounted filesystems, naming service problems, and so on. Refer to the troubleshooting section “Troubleshooting the IDB” on page 711.
- Omit or remove non-core parts (binary files or filenames part) of the IDB that contain problems. This is possible if the identified level of IDB corruption is minor or major (meaning the corruption is not in the core part of the IDB).
- Perform a complete recovery. This consists of restoring the IDB and updating information that has been modified since the last IDB backup. This is a must if the identified level of IDB corruption is critical (meaning the corruption is in the core part).

Complete Recovery

Complete recovery consists of two phases:

1. IDB restore, which gets the IDB to the last (available) consistent state.
2. Updating the IDB from the last consistent state up to the last moment when the IDB was still operational.

Depending on how well you prepared for IDB recovery before problems occurred (availability of IDB recovery file, IDB backup, original device and transaction logs), the recovery procedure can differ. If all these are available, you can use a very convenient IDB recovery method, guided autorecovery.

Identifying the Level of Database Corruption

IDB Corruption Levels

There are three levels of IDB corruption: critical, major, and minor. The level depends on the part of the IDB where the corruption occurs.

You can use the IDB consistency check to determine which part of the IDB is corrupted. Depending on the level of corruption, the IDB recovery procedure differs.

Figure 11-11 IDB Corruption Levels

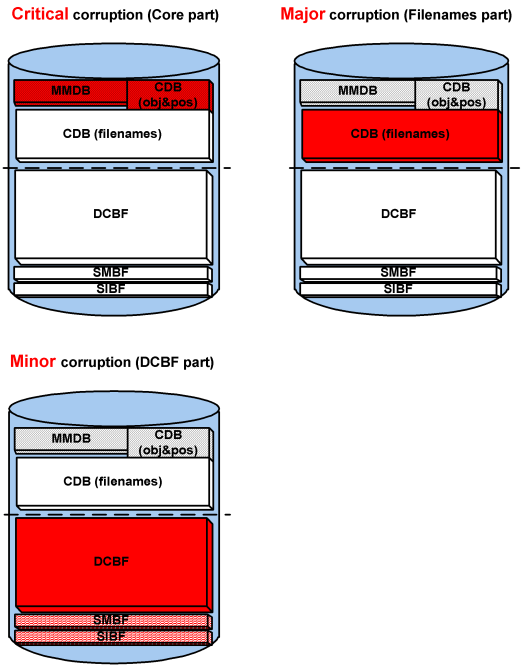

How to Identify the Corruption Level

Identify the level of IDB corruption using the `omnidbcheck -extended` command.

NOTE

The extended check may take several hours. To avoid an extended period of system downtime, you can run subparts of the `omnidbcheck` command instead. For example, run the `omnidbcheck -core` to determine whether the core part of the IDB is corrupted.

After identifying the level of corruption, perform the appropriate recovery procedure. Refer to “Overview of IDB Recovery Methods” on page 496.

Overview of IDB Recovery Methods

Several recovery methods are available for recovering the IDB. Depending on the identified level of corruption, your requirements, and the availability of the IDB recovery file and the original device and transaction logs, the recovery procedure can differ.

The Most Convenient Complete Recovery

When the complete IDB is missing or the core part is corrupted, the corruption level is critical. If the IDB recovery file and the original device used for the IDB backup are available, you can perform the Guided Autorecovery (IDB Restore and Replay Logs). Refer to “Performing Guided Autorecovery” on page 501. Otherwise, follow one of the methods given under “More Recovery Methods” on page 497.

The guided autorecovery method guides you through restoring the IDB and replaying transaction logs. If transaction logs are not available, you can still update the IDB by importing all media since the last IDB backup.

Omitting (Removing) Corrupted IDB Parts

If the identified level of corruption is major or minor (corruption is not in the core part), you can consider omitting (removing) the missing or corrupted parts of the IDB or perform the complete IDB recovery instead.

When the filename tablespace is corrupted, the corruption level is major. Refer to “Handling Major Database Corruption in the Filenames Part” on page 499.

When the DC binary files are missing or corrupted, the corruption level is minor. Refer to “Handling Minor Database Corruption in the DCBF Part” on page 498.

More Recovery Methods

These recovery procedures are adapted to specific situations. They assume that you want to recover the complete IDB, but for some reason you cannot perform the guided autorecovery method. The recovery consists of restoring the IDB and updating the IDB.

Table 11-4 Restoring the IDB

Current situation	Remark	Recovery Procedure
The IDB recovery file is available but the original device used for the IDB backup has changed.	The method is essentially the same as the guided autorecovery method, but less guided, and more complex and time consuming.	“Recovering the IDB Using IDB Recovery File and Changed Device” on page 502.
The IDB recovery file is not available.	The method is essentially the same as the guided autorecovery method, but less guided, and more complex and time consuming.	“Recovering the IDB Without the IDB Recovery File” on page 504.
You want to recover the IDB from a specific IDB backup (not the latest one).	This method does not provide the latest state of the IDB as a result.	“Recovering the IDB from a Specific IDB Session” on page 506.

Table 11-4 Restoring the IDB

Current situation	Remark	Recovery Procedure
You want to recover to a different disk layout.	This method is equivalent to disaster recovery from a Data Protector configuration where you lost the IDB transaction logs, the IDB recovery file, and the `media.log` file. It is far more complex than the guided autorecovery and does not provide the latest state of the IDB as a result.	"Recovering the IDB to a Different Disk Layout" on page 506.

If the transaction logs are available, the recovery procedures in Table 11-4 guide you through replaying the IDB transaction logs. Refer to "Replaying IDB Transaction Logs" on page 508.

If the transaction logs are not available, you can update the IDB by importing media. Refer to "Updating the IDB by Importing Media" on page 509.

Handling Minor Database Corruption in the DCBF Part

If you detect that the IDB corruption is of minor severity, it means that some DC binary files are missing or corrupted. If this is the case, there is no need for complete IDB recovery. You can easily recreate the binary files by importing catalog from media. Choose the recovery procedure depending on the corruption type.

Recovering if DC Binary Files Are Missing

DC binary files are organized so that one binary file exists for each medium. If some DC binary files are missing, media positions of some media point to the non-existent files. An error message is displayed when browsing the relevant filesystems. Proceed as follows:

1. From the `omnidbcheck -bf` output, identify the Medium ID of the missing binary file. Run the `omnimm -media_info <Medium>` command to get other attributes of the medium, such as medium

label and media pool.

2. Run the `omnidbutil -fixmpos` command to establish consistency between media positions (mpos) and binary files.
3. Import the catalog from the media to recreate the binary files. Refer to “Importing the Catalog from Media” on page 160.

Recovering if DC Binary Files Are Corrupted

If some DC binary files are corrupted, you can remove the DC binary files and recreate them. The only effect of removing the files is that some media positions point to the non-existent binary files, and thus an error message is displayed when browsing the relevant filesystems. Proceed as follows:

1. From the `omnidbcheck -dc` output, identify the Medium ID of the corrupted DC binary file. Run the `omnimm -media_info <Medium>` command to get other attributes of the medium, such as medium label and media pool.
2. Identify the DC binary file for the affected medium. DC binary files are named: *<Medium>*`_`*<TimeStamp>*`.dat` (in the *<Medium>*, and colons ":" are replaced with underscores "_").
3. Remove the corrupted DC binary files.
4. Run the `omnidbutil -fixmpos` command to establish consistency between media positions (mpos) and binary files.
5. Import the catalog from the media to recreate the binary files. Refer to “Importing the Catalog from Media” on page 160.

Handling Major Database Corruption in the Filenames Part

If you detect that the corruption is of major severity, which means that a filename tablespace is corrupted, you can remove the detail catalogs (filenames and DC binary files) instead of recovering the whole IDB.

The procedure is fast and results in an IDB without detail catalogs (as though all backups were done with the `No log` option). The IDB is still fully operational in terms of all backups, restores, and media management operations, except that browsing is not possible (information about backed up data should be read from media).

Since all detail catalogs are lost, this method of recovery is only applicable if:

- The catalogs created by subsequent backups are good enough.
- There is no IDB backup available.

Recovery Procedure

Proceed as follows:

1. Run the command:

   `omnidbutil -writedb -no_detail -cdb <Directory> -mmdb <Directory>`

   to write the IDB without detail catalogs to ASCII files.

2. Run the command:

   `omnidbutil -readdb -cdb <Directory> -mmdb <Directory>`

   to read the IDB from the ASCII files.

The operation lasts approximately 5-20 minutes.

After the detail catalogs are removed, all DC binary files can be deleted, although the DC directories are still registered. Subsequent backups will store the file versions in the DC binary files.

Prerequisites for IDB Recovery

- Mount a disk of the same size as before the disaster on the same directories as at the IDB backup time (on Windows systems, the same drive letters must be assigned). If this cannot be ensured, follow the procedure for recovering the IDB to a different disk/volume layout. You can use the `-preview` option of the `omnidbrestore` command to see where the files will be restored.
- Verify that Data Protector is installed on the Cell Manager and the system where a device is attached (preferably, the device used for the IDB backup).
- If possible, move the `media.log` file from the previous installation to a safe place. It will provide you with the information about the media used since the last IDB backup. This is very helpful for updating the IDB if transaction logs are not available.
- If the IDB is installed on MC/ServiceGuard, the following commands

have to be run on the active node before performing the recovery:

1. `cmhaltpkg <pkg_name>` , where `<pkg_name>` is the name of the Data Protector cluster package.

   This command stops the Data Protector package and dismounts the Data Protector shared volume group.

2. `vgchange -a e /dev/<vg_name>` , where `<vg_name>` is the name of Data Protector shared volume group.

   This command activates the Data Protector shared volume group. To list volume groups on your system, run `ll /dev/*/group`.

3. `mount /dev/<vg_name>/<lv_name>/<MountPoint>` , where `<MountPoint>` is the name of the mount point for the Data Protector shared volume group.

   This command mounts the Data Protector shared volume group.

When the guided autorecovery has finished, run the `cmrunpkg <pkg_name>` command on the active node to start the Data Protector package.

- If the IDB is installed on Microsoft Cluster Server, take the `OBVS_VELOCIS` cluster group offline (using the Cluster Administrator utility on the active node) and stop the `Inet` service before performing the recovery.

  When the recovery has finished, bring the `OBVS_VELOCIS` and `OBVS_MCRS` cluster groups online using the Cluster Administrator utility and restart the `Inet` service.

Performing Guided Autorecovery

Guided autorecovery is the most convenient IDB recovery method. You can perform it if the IDB recovery file and the original device used for the IDB backup together with the IDB backup medium are available.

This method guides you through restoring the IDB and replaying transaction logs since the last IDB backup. If the transaction logs are not available, you can still update the IDB since the last IDB backup by importing media.

Transaction replay updates the core part of the IDB. Binary files are not updated and changes to binary files are lost.

The following are not available for the backups that were running from the last IDB backup before the IDB corruption:

- Session messages
- Browsing of file versions (restores of complete objects are possible). Import the catalog on the media used by the backups to recover the changes.
- SIBF updates. Export and import the media used by the backups to recover the changes.

Prerequisites

Refer to “Prerequisites for IDB Recovery” on page 500.

Cluster

Additional steps are required if you are performing restore of the IDB installed in a cluster. “Prerequisites for IDB Recovery” on page 500.

Recovery Procedure

To recover the IDB, run the `omnidbrestore -autorecover` command.

The command reads the IDB recovery file and if IDB backups are logged to the file, it stops the services and starts restore of the IDB back in place. All the options are generated automatically using data from the IDB recovery file.

Once the restore is complete, the `omnidbrestore` checks if transaction logs are available to be replayed. If logs are available, you are asked to confirm the replay of the logs. If this step is cancelled or transaction logs are not available, output describes how to update the IDB since the last IDB backup by:

- importing media
- finding the transaction logs and replaying them later

Once you replay logs or import media to update the IDB, the full IDB should be successfully recovered.

Recovering the IDB Using IDB Recovery File and Changed Device

Use this procedure to recover the IDB if the IDB recovery file (`obrindex.dat`) is available but the original device used for the IDB backup is different from the one to be used for recovery, or the medium is located in a different slot.

Prerequisites

Refer to “Prerequisites for IDB Recovery” on page 500.

Cluster

Additional steps are required if you are performing restore of the IDB installed in a cluster. “Prerequisites for IDB Recovery” on page 500.

Recovery Procedure

1. Run the following command to create a text file with the restore job options:

   ```
   omnidbrestore -logview -autorecover -skiprestore -save
   C:\TEMP\restjob.txt
   ```

IMPORTANT

The specified `-logview` command lists first transaction logs, next to the session IDs. Remember the first transaction log for the session you want to restore, because you will need it in order to update the IDB after the restore. For example, from the output `2001/02/09-2 AAAAAAH`, you would remember the first transaction log `AAAAAAH` in order to restore the `2001/02/09-2` session.

The created `restjob.txt` file has the information on original devices and on slots in which media were originally located (at IDB backup time).

For example, if the IDB backup was done on a DDS drive with the SCSI address `scsi0:0:0:0`, a file like this is created:

```
-name LDEV
-policy 1
-type 1
-dev scsi0:0:0:0
-mahost goedl.hermes
-maid 0100007f:3a486bd7:0410:0001
-position 3:0
-daid 977824764
```

2. Modify the `restjob.txt` file to specify the current device or the slot in which the media are currently located.

For example, if the DDS drive that had the SCSI address `scsi0:0:0:0` at backup time has the SCSI address `scsi0:0:1:0` at restore time, the `restjob.txt` file should be modified accordingly:

```
-name LDEV
-policy 1
-type 1
-dev scsi0:0:1:0
-mahost cm.dom.com
-maid 0100007f:3a486bd7:0410:0001
-position 3:0
-daid 977824764
```

3. Run the restore with the `omnidbrestore -read C:\TEMP\restjob.txt` command.

The command guides you through restoring the IDB and replaying transaction logs since the last IDB backup.

If the transaction logs are not available, you can still update the IDB by importing all media used since the last IDB backup. In this case, refer to “Updating the IDB by Importing Media” on page 509.

Recovering the IDB Without the IDB Recovery File

Use this procedure to recover the IDB if the IDB recovery file (`obrindex.dat`) is not available.

Prerequisites

Refer to “Prerequisites for IDB Recovery” on page 500.

Cluster

Additional steps are required if you are performing restore of the IDB installed in a cluster. “Prerequisites for IDB Recovery” on page 500.

Recovery Procedure

1. Configure the device using the `Data Protector Manager`.
2. Find the medium with the latest IDB backup.
3. Insert the medium into the device and use the following command to display the contents of the medium:

   `omnimlist -dev <LogicalDevice>`

The information you need for the IDB restore is the Medium ID and Disk Agent ID for the backup session you want to restore.

4. Use the following command to display the information on the device configuration:

   ```
   omnidownload -dev <LogicalDevice>
   ```

   The information you need for the IDB restore is the following:

   - Mahost (a Media Agent host)
   - Policy (number)

     A policy number can be obtained using the following translation: 1 for Standalone devices, 3 for Stacker devices, 10 for SCSI Libraries, and 5 for Jukebox devices.
   - Media type (number)

     A media type number can be obtained using the following translation: 1 for DDS, 3 for ExaByte, 10 for DLT, or 7 for File.
   - SCSI address
   - Robotics SCSI address (only if using Exchanger library devices)

5. Run the `omnidbrestore` command using the obtained information:

   ```
   omnidbrestore -policy <log. device policy> -type <log.
   device_type> [-ioctl <RoboticsDevice>] -dev <PhysicalDevice>
   -mahost <DeviceHostname> -maid <mediumID> -daid <DAID>
   ```

   For example, you would use the following command to restore the IDB from a backup session with the medium ID 0100007f:3a486bd7:0410:0001 and the Disk Agent ID 977824764, performed using a standalone device of the type DLT, connected to the system `cm.dot.com` and with the SCSI address `scsi0:1:2:0`:

   ```
   omnidbrestore -policy 1 -type 10 -dev scsi0:1:2:0 -mahost
   cm.dom.com -maid 0100007f:3a486bd7:0410:0001 -daid  977824764
   ```

The command guides you through restoring the IDB and replaying transaction logs since the last IDB backup.

If the transaction logs are not available, you can still update the IDB by importing all media used since the last IDB backup. In this case, refer to “Updating the IDB by Importing Media” on page 509.

Recovering the IDB from a Specific IDB Session

Use this procedure to recover the IDB from a backup other than the latest one if the IDB recovery file (`obrindex.dat`) is available.

Prerequisites

Refer to "Prerequisites for IDB Recovery" on page 500.

Cluster

Additional steps are required if you are performing restore of the IDB installed in a cluster. "Prerequisites for IDB Recovery" on page 500.

Recovery Procedure

1. Check all backups using the following command:

   ```
   omnidbrestore -autorecover -logview -skiprestore
   ```

2. Choose the backup session you want to restore from and perform the restore by running the `omnidbrestore -autorecover -session <sessionID>` command.

   For example, if you choose to restore from the backup session `2000/12/26-1` and the original device used for the IDB backup exists, run:

   ```
   omnidbrestore -autorecover -session 2000/12/26-1
   ```

   The command guides you through restoring the IDB and replaying transaction logs since the last IDB backup. If the transaction logs are not available, you can still update the IDB by importing all media used since the last IDB backup. In this case, refer to "Updating the IDB by Importing Media" on page 509.

3. Bring the `OBVS_VELOCIS` and `OBVS_MCRS` cluster groups online using the Cluster Administrator utility and restart the `Inet` service.

4. Run `omnidbutil -fixmpos`.

Recovering the IDB to a Different Disk Layout

You can restore the IDB to a disk of a different size than before the disaster, and to different directories than at the backup time.

Prerequisites

Refer to "Prerequisites for IDB Recovery" on page 500.

The following prerequisite must also be met before recovering the IDB to a different disk layout:

- Import the media with the IDB backup.

Cluster

Additional steps are required if you are performing restore of the IDB installed in a cluster. “Prerequisites for IDB Recovery” on page 500.

Recovery Procedure

After you meet the prerequisites, proceed as follows to recover the IDB:

1. In the `Data Protector Manager`, browse the `Internal Database` backup object and select it for restore. Refer to “Selecting Your Data for Restore” on page 338.
2. For the `db40/datafiles` directory, use the `Restore As/Into` option to specify a restore location other than the default one. Refer to “Restoring Files to Different Paths” on page 370.
   You may want to restore the Detail Catalog and Session Messages Binary Files to a different restore location. In this case, also use the `Restore As/Into` option.
3. Start the IDB restore. Refer to “Previewing and Starting a Restore” on page 341.
4. Move the `db40/datafiles` directory back in place and start the Data Protector services using the `omnisv -start` command.
   - On Windows: `<Data_Protector_home>\bin\omnisv -start`
   - On UNIX: `/opt/omni/sbin/omnisv -start`

   If the IDB is installed on MC/ServiceGuard, run the `cmrunpkg <pkg_name>` command on the active node to start the Data Protector package, where `<pkg_name>` is the name of the Data Protector cluster package.

   If the IDB is installed on Microsoft Cluster Server, bring the `OBVS_VELOCIS` and `OBVS_MCRS` cluster groups online using the Cluster Administrator utility and restart the `Inet` service.
5. If you restored the Detail Catalog and Session Messages Binary Files to a different restore location, you need to do the following:
   a. Create a new DC directory and remove the old one. Refer to “Creating a DC Directory” on page 472.
   b. Run the `omnidbutil -remap_dcdir` command to update the pathnames of DC binary files.
6. Verify that you have all files back by running the `omnidbcheck` command.

What's Next?

After you have restored the IDB, you need to update the IDB by importing media if the `media.log` file is available. Refer to "Updating the IDB by Importing Media" on page 509.

Replaying IDB Transaction Logs

In a successful `omnidbrestore -autorecover`, transaction logs are already replayed. Use this procedure only if you need to retry replaying of transaction logs or you postponed it before.

Replaying transaction logs after the IDB restore is completed recovers the IDB to the same state as before the crash, except that binary files are not updated and changes to binary files are lost.

The following are not available for the backups that were running from the last IDB backup until the IDB corruption:

- Session messages.
- Browsing of file versions (restores of complete objects are possible). Perform the import catalog on the media used by the backups, to recover the changes.
- SIBF updates. Export and import the media used by the backups to recover the changes.

Limitation

Replay of the transaction logs can only be done if archiving of the transaction logs is enabled. The archiving parameter in the `velocis.ini` (`rdmserver.ini` on HP-UX 11.23) file must be set to 1.

Prerequisites

- Transaction logs must be available. For more information on transaction logs, refer to "Preparing for IDB Recovery" on page 466. You can verify that the transaction logs are available by listing the directory: `/db40/logfiles/syslog`

  If transaction logs are not available, refer to "Updating the IDB by Importing Media" on page 509.
- If the IDB is installed on MC/ServiceGuard, run the `cmhaltpkg <pkg_name>` command on the active node before running the `omnidbrestore` command in the procedure below, to stop the Data Protector package. Before running the `omnidbcheck` command in the procedure below, run the `cmrunpkg <pkg_name>` command on the active node to start the Data Protector package, where `<pkg_name>` is the name of the Data Protector cluster package.

- If the IDB is installed on Microsoft Cluster Server, take the `OBVS_VELOCIS` cluster group offline using the Cluster Administrator utility and stop the `Inet` service on the active node before running the `omnidbrestore` command in the procedure below. Before running the `omnidbcheck` command in the procedure below, bring the `OBVS_VELOCIS` and `OBVS_MCRS` cluster groups online using the Cluster Administrator utility and restart the `Inet` service.

How to Replay Transaction Logs

Proceed as follows:

1. Run the following command to replay the transaction logs:

   ```
   omnidbrestore -replay_only -firstlog
   <FirstTransactionLog>
   ```

   where *<first_trans_log>* is the first transaction log that was created just after the IDB backup was started.

   At the end of the `omnidbrestore -autorecover` output, Data Protector displays the exact command you should use to replay the transaction logs, giving you the name of the first transaction log.

   For example, the command could be:

   ```
   omnidbrestore -replay_only -firstlog AAAAAC
   ```

   where `AAAAAC` is the first transaction log created after the IDB backup was started.

2. Run the `omnidbcheck` command.

This completes the recovery procedure.

Updating the IDB by Importing Media

To successfully complete the IDB recovery, you need to update the IDB changes after the IDB is restored.

If transaction logs are not available, update the changes by importing all media since the last IDB backup. Do this once the IDB restore has finished.

To verify that transaction logs are available, or to update the changes using transaction logs, refer to “Replaying IDB Transaction Logs” on page 508.

To update the changes by importing media, proceed as follows:

1. Start the Data Protector processes and services using the `omnisv -start` command:
   - On Windows: `<Data_Protector_home>\bin\omnisv -start`
   - On UNIX: `/opt/omni/sbin/omnisv -start`

2. Increase the session counter to 200 using the following command:

   ```
   omnidbutil -set_session_counter 200
   ```

   If necessary, you can now start with backups.

3. Export and import the media with the last IDB backup. This creates consistent information about the last IDB backup.
4. Import (export if already in IDB) the media used between the last IDB backup and the time of the IDB recovery. See the `/var/opt/omni/server/log/media.log` (on UNIX systems) or `<Data_Protector_home>\log\server\media.log` (on Windows systems) file for a list of media.
5. Run the `omnidbcheck` command.

The complete IDB should be successfully recovered.

NOTE If recovering an IDB that encompasses a CMMDB or a remote MMDB to a different disk layout, you need to run the `omnidbutil -cdbsync` command after updating the IDB.

12 Disaster Recovery

In This Chapter

This chapter provides an overview of disaster recovery on Windows UNIX clients and Cell Managers. The following sections are included:

Introduction

This section explains the basic terms used in the Disaster Recovery chapter. For overview and concepts of the available disaster recovery methods as well as table outlining the possible combinations of disaster recovery methods and operating system, please see the Disaster Recovery section in the *HP OpenView Storage Data Protector Concepts Guide*.

For a list of supported disaster recovery methods for a particular operating system, refer to the support matrices in the *HP OpenView Storage Data Protector Software Release Notes*.

What Is a Computer Disaster?

A **computer disaster** refers to any event that renders a computer system unbootable, whether due to human error, hardware or software failure, virus, natural disaster, etc. In these cases it is most likely that the boot or system partition of the system is not available and the environment needs to be recovered before the standard restore operation can begin. This includes repartitioning and/or reformatting the boot partition and recovery of the operating system with all the configuration information that defines the environment. *This has to be completed in order to recover other user data.*

What Is an Original System?

Original system refers to the system configuration backed up by Data Protector before a computer disaster hit the system.

What Is a Target System?

Target system refers to the system after the computer disaster has occurred. The target system is typically in a non-bootable state and the goal of Data Protector disaster recovery is to restore this system to the original system configuration. The difference between the crashed and the target system is that the target system has all faulty hardware replaced.

What Are Boot and System Disks/Partitions/Volumes?

A **boot disk/partition/volume** refers to the disk/partition/volume that contains the files required for the initial step of the boot process, whereas the **system disk/partition/volume** refers to the disk/partition/volume that contains the operating system files.

NOTE Microsoft defines the boot partition as the partition that contains the operating system files and the system partition as one that contains the files required for the initial step of the boot process.

What Is a Hosting System?

Hosting system is a working Data Protector client used for Disk Delivery Disaster Recovery with Disk Agent installed.

What Is Auxiliary Disk?

Auxiliary disk is a bootable disk that has a minimal operating system with networking and Data Protector Disk Agent installed. It can be carried around and used to boot the target system in Phase 1 of Disk Delivery Disaster Recovery of UNIX clients.

What Is a Disaster Recovery Operating System (DR OS)?

Disaster recovery operating system (DR OS) is operating system environment where the process of disaster recovery is running. It provides Data Protector a basic runtime environment (disk, network, tape and filesystem access). It has to be installed and configured before the Data Protector disaster recovery can be performed.

DR OS can be either temporary or active. **Temporary DR OS** is used exclusively as a host environment for some other operating system restore along with the target operating system configuration data. It is deleted after the target system is restored to the original system configuration. **Active DR OS** not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces it's own configuration data with the original configuration data.

What Are Critical Volumes?

Critical volumes are volumes required to boot the system and Data Protector files. Regardless of the operating system, these volumes are:

- boot volume
- system volume
- Data Protector executables
- IDB (Cell Manager only)

NOTE If IDB is located on different volumes than all volumes where IDB resides, are critical.

Apart from the critical volumes stated above, CONFIGURATION is also a part of the critical volumes set for Windows systems. Services are backed up as a part of the CONFIGURATION backup.

Some items included in the CONFIGURATION can be located on volumes other than system, boot, Data Protector or IDB volumes. In this case these volumes are also part of critical volumes set:

- user profiles volume
- Certificate Server database volume on Windows Server
- Active Directory Service volume on domain controller on Windows Server
- quorum volume on Microsoft Cluster Server.

What is Online Recovery?

Online recovery is performed when Cell Manager is accessible. In this case most of Data Protector functionalities are available (Cell Manager runs the session, restore sessions are logged in the IDB, you can monitor the restore progress using GUI, etc.).

What is Offline Recovery?

Offline recovery is performed if the Cell Manager is not accessible (for example, due to network problems, Cell Manager has experienced a disaster, online recovery has failed, etc.). Only standalone and SCSI Library devices can be used for offline recovery. Note that recovery of Cell Manager is always offline.

What is Local/Remote Recovery?

Remote recovery is performed if all Media Agent hosts specified in SRD file are accessible. If any of them fails, disaster recovery process fails over to local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise Data Protector prompts you to select the device which will be used for restore. Note that offline OBDR is always local.

Disaster is always serious, however the following factors can exacerbate the situation:

- The system has to be returned to online status as quickly and efficiently as possible.
- Administrators are not familiar with the required steps to perform the disaster recovery procedure.
- The available personnel to perform the recovery may only have fundamental system knowledge.

Disaster recovery is a complex task that involves extensive planning and preparation before execution. You have to have a well-defined, step-by-step process in place to prepare for, and recover from, disastrous situations.

The Recovery Process

The disaster recovery process consists of 4 phases with the *Phase 0* (preparation) being the prerequisite for a successful disaster recovery. In *Phase 1*, DR OS is installed and configured, which usually includes repartitioning and reformatting of the boot partition, since the boot or system partition of the system are not always available and the environment needs to be recovered before normal restore operations can resume. Operating system with all the configuration information that defines the environment with Data Protector (as it was) is restored in *Phase 2*. Only after this step is completed, is the restore of applications and user data possible (*Phase 3*). A well-defined, step-by-step process has to be followed to ensure fast and efficient restore.

Preparing for a Disaster Recovery

Carefully follow the instructions in this section to prepare for a disaster recovery and to ensure fast and efficient restore. Preparation does not depend on the disaster recovery method, however, it does include developing a detailed disaster recovery plan, performing consistent and relevant backups and updating the SRD file on Windows.

This section contains the general preparation procedure for disaster recovery for all disaster recovery methods. Additional preparation is required for each particular disaster recovery method. Refer to corresponding sections for additional preparation steps.

Planning

Developing a detailed disaster recovery plan has a major impact on the success of a disaster recovery. To deploy disaster recovery in a large environment with many different systems, proceed as follows:

1. **Plan**

   Planning must be prepared by IT administration and should include the following:

   - Determine the systems that need to be recovered as well as the time and level of recovery. Critical systems are all systems required for network to function properly (DNS servers, domain controllers, gateways, etc.), Cell Managers and Media Agent clients.
   - Determine a recovery method to be used (impacts the required preparations).
   - Determine a method to obtain the required information at recovery time, such as the media that holds the IDB, location of updated SRD file and location and labels of the Cell Manager backup media.
   - Create a step-by-step detailed checklist to guide you through the process.
   - Create and execute a test plan to confirm that the recovery will actually work.

2. **Prepare for recovery**

   Depending on the recovery method to be used, the preparation should include:

   On UNIX systems:

   - Creation of tools, such as the auxiliary disk with the minimum operating system, network resources, and the Data Protector Disk Agent installed.
   - Creation of pre-execution scripts, which collect the storage structure and other client-specific preparations.

   On Windows systems:

   - Updating **System Recovery Data (SRD)** and storing it to a safe place. You should restrict access to SRD files due to security reasons.

   On all systems:

   - Performing regular and consistent backups.

3. **Perform recovery procedures**

   Follow the procedures and checklists you have tested to recover the crashed system.

Consistent and Relevant Backup

In the case of a disaster, the target system should be put back into the state it was at the time of the last valid known backup. Additionally, the system should function as it had functioned just before the last valid backup performance.

NOTE

On UNIX systems, some daemons or processes are active as soon as the system finishes booting, for various reasons (HP-UX example: License server at run level-2). Such an early process may even read the data into memory and write a “dirty flag” into some file while it runs. A backup taken at the standard operating stage (the standard run level-4) cannot be expected to yield a problem-free restart of such an application. To follow the example, the license server, if started after such a pseudo recovery, will realize that the data read from the file is inconsistent and will refuse to run the service as expected.

On Windows, while the system is up and running, many system files cannot be replaced because the system keeps them locked. For example, the user profiles that are currently being used cannot be restored. The login account has to be changed or the relevant service has to be stopped.

Data consistency of an application can be violated depending on what is active on the system when the backup runs, thereby causing re-start and execution issues after recovery.

How to Create a Consistent and Relevant Backup?

- ✓ Ideally, you would perform a backup with the relevant partition(s) set off-line, which is usually not possible.
- ✓ Examine the activity on the system during the backup. Only operating system related processes and database services which are backed up online can remain active during the backup execution.
- ✓ None of the low-level (UNIX) or background-level (Windows) application specific services should be running.

What should be included in the consistent and relevant backup depends on the disaster recovery method you plan to use and other system specifics (for example, disaster recovery of Microsoft Cluster). See the sections pertaining to particular disaster recovery methods.

Updating and Editing the System Recovery Data (SRD)

What Is SRD?

System recovery data (SRD) is a Unicode text file that contains information required for the configuration and restore of the Windows target system. A SRD file is generated when `CONFIGURATION` backup is performed on a Windows client and then stored in:

- On a Windows Cell Manager: *<Data_Protector_home>*`\Config\server\dr\srd`
- On a UNIX Cell Manager: `/etc/opt/omni/server/dr/srd/`

IMPORTANT

When IDB is not available, information about objects and media is stored only in SRD file.

The SRD filename on the Cell Manager is identical to the hostname of the computer where it was generated - for example `computer.company.com`.

After the `CONFIGURATION` backup, the SRD contains only system information required for installation of the DR OS. In order to perform a disaster recovery, additional information about backup objects and corresponding media must be added to the SRD. The SRD can be updated only on a Windows client. The name of the updated SRD file is `recovery.srd`.

How to Update SRD?

There are three different methods possible for updating the SRD file:

- Update SRD File Wizard
- `omnisrdupdate` command as a standalone utility
- `omnisrdupdate` command as a backup session post-exec script

Using SRD Update Wizard

To update the SRD file using the Update SRD File Wizard, proceed as follows:

1. In the `Data Protector Manager` switch to the `Restore` context and then click the `Tasks` Navigation tab.
2. In the `Scoping Pane` of the `Tasks` Navigation tab, check the `Disaster Recovery`.
3. In the `Results Area`, check the `SRD File Update` option button, select the client and click `Next`.
4. For each of the critical objects, select an object version and click `Next`.
5. Type the destination directory where the updated SRD file is to be placed and click `Finish`.

IMPORTANT

Because the SRD file is saved on the Cell Manager system, it is not accessible if the Cell Manager fails. As a result, you need an additional copy of the Cell Manager's SRD which should be stored in a vault. In addition to the Cell Manager, you should save the updated SRD file to several secure locations as a part of the disaster recovery preparation policy. See “Preparation” on page 527.

Using omnisrdupdate

It is also possible to update the SRD file using the `omnisrdupdate` command as a standalone command. The `omnisrdupdate` command is located in the *<Data_Protector_home>*`\bin` directory.

`Omnisrdupdate` requires a `session_ID` to update an existing SRD file with backup object information belonging to the given session. Using this value, `omnisrdupdate` will update the SRD file with the backup object information which belongs to the passed `session_ID` value. After the SRD is updated it will be saved back on the Cell Manager.

This procedure will only succeed if all critical backup objects (as specified in the SRD file) were actually backed up during the specified session. To view which objects are considered as critical for the SRD update, open the SRD file in a text editor and find the objects section. All critical objects for the SRD update are listed there. Note that the database is represented as "/".

Here is an example of an objects section of the SRD file:

```
-section objects
-objcount 3
-object /C -objtype 6 -objpurpose 283
-endobject /C
-object / -objtype 3 -objpurpose 32
-endobject /
-object /CONFIGURATION -objtype 6 -objpurpose 4
-endobject /CONFIGURATION
-endsection objects
```

In this case, there are 3 critical objects: `/C`, `/` (database) and `/CONFIGURATION`.

TIP

To obtain the session ID, execute the `omnidb` command with the option `-session`. To obtain the latest session ID, at the command prompt type **`omnidb -session -latest`**.

The updated SRD file should be kept in a safe place so that it is not lost in the case of disaster. To locate where the updated SRD file will be saved, use the `-location` option with the `omnisrdupdate` command.

There can be more than one `-location` parameters specified (including network shares on which you have write permission), each of which will receive an updated copy of the SRD file. See “Preparation” on page 527.

To determine for which hostname the SRD file from the Cell Manager should be updated, use the option `-host` with the command `omnisrdupdate`. If you don’t specify the hostname, the local host is assumed. SRD file on the Cell Manager is not updated.

Example

To update the SRD file with the backup object information which belongs to a session `2002/05/02-5` for the client with the hostname `computer.company.com` and to store an updated copy of the SRD file on the floppy disk and in the `SRDfiles` share on computer with the hostname `computer2`, type **`omnisrdupdate -session 2002/05/02-5 -host computer.company.com -location a: -location \\computer2\SRDfiles`**

Make sure that you have the write permission on that share.

Using a Post-Exec Script

Another method to update the SRD is using the `omnisrdupdate` command as a backup post-exec script. To do so, either modify an existing backup specification or create a new one. Perform the following steps to modify a backup specification so that the SRD file is updated with information about backed up objects when the backup session stops:

1. In the `Backup` context, expand the `Backup Specifications` item and then `Filesystem`.
2. Select the backup specification that you would like to modify (it must include all backup objects marked as critical in the SRD file, otherwise the update will fail. It is recommended to perform the client backup with disk discovery) and click `Options` in the `Results Area`.
3. Click the `Advanced` button under the `Backup Specification Options`.
4. Type **`omnisrdupdate.exe`** in the `post-exec` text box.
5. In the `On client` drop down list, select the client on which this post-exec script will be executed and confirm with OK. This should be the client that was marked for backup on the source page.

When `omnisrdupdate` command is executed as a post-exec utility, the session ID is obtained automatically from the environment and the user is not required to specify the session ID.

All other options can be specified the same way as with the standalone utility (`-location` <*path*>, `-host` <*name*>).

Editing the SRD File

It is possible, that the information about backup devices or media stored in the SRD file is out of date at the time disaster recovery is being performed. In this case edit the SRD file to replace the incorrect information with the relevant information before performing the disaster recovery. See “Recovery Using an Edited SRD File” on page 580.

IMPORTANT You should restrict access to the SRD files due to security reasons.

Assisted Manual Disaster Recovery of a Windows System

The following sections explain how to prepare and execute an Assisted Manual Disaster Recovery on Windows systems. For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

Assisted Manual Disaster Recovery is an elementary method that consists of the following steps:

1. Installing the Windows system to its original location. This includes the creation and formatting of the boot and system partition, needed for the Windows installation.
2. Creating and formatting additional partitions as they existed on the crashed system, including original drive letter assignments.
3. Executing the Data Protector `drstart.exe` command, which will install a temporary Data Protector suite and start the restore of the system critical volumes.
4. Booting the system.
5. Recovering the vendor-specific partition, if it existed before the disaster.

NOTE The preparation and recovery procedure are different for the recovery of a Data Protector client and of a Data Protector Cell Manager. The differences are marked in the text.

Note that Windows provide additional possibilities to recover a system before deciding on a disaster recovery. This can be done by booting the system in the safe mode or from the recovery floppy disks and trying to resolve problems. Another option is to start the computer using the last known good configuration.

Requirements

- The partitions have to be the same size or larger than the partitions on the failed disk. This way the information stored on the failed disk can be restored to the new one. Also, the type of filesystem and compression attributes of the volumes must match (FAT, NTFS).
- The hardware configuration of the target system must be the same as of the original system. This includes SCSI BIOS settings (sector remapping).

Limitation

- Internet Information Server (IIS) Database, Terminal Services Database and Certificate Server Database are not restored automatically during Phase 2. They can be restored on the target system using the standard Data Protector restore procedure.

Preparation

To prepare for a successful disaster recovery, you should follow the instructions related to the general preparation procedure together with the specific method requirements. Advance preparation is essential to perform the disaster recovery fast and efficiently. You should also give special attention to the disaster recovery preparation of the Cell Manager and Microsoft Cluster Server.

CAUTION It is too late to prepare for a disaster recovery once a disaster has occurred.

See also “Preparing for a Disaster Recovery” on page 519, for the general preparation procedure for all disaster recovery methods before completing the steps listed in this section. To recover from a disaster quickly and efficiently, consider the following steps and prepare your environment accordingly:

1. You need a Windows bootable installation CD-ROM to enable your system to start from the CD-ROM. If you do not have a bootable CD-ROM, use the standard procedure for booting the computer from diskettes.

2. Ensure that you have drivers for the system you want to recover. You may need to install some drivers, such as network, HBA and SCSI drivers during Windows Setup.
3. To recover the crashed system, you need the following information about the system before the disaster (stored also in the SRD file):
   - If DHCP was not used before the disaster, the TCP/IP properties (IP address, Default gateway, Subnet mask and DNS order)
   - Client properties (Hostname)
4. Ensure that the following is true:
   - You should have a successful full client backup. See “Backing Up Filesystems (Logical Disk Drives)” on page 213 and “Backing Up CONFIGURATION” on page 218.
   - You should have a SRD file updated with information about backed up objects in the chosen successful backup session. See “Updating and Editing the System Recovery Data (SRD)” on page 521.
   - In the case of a Cell Manager recovery, you need a successful IDB backup of the Cell Manager. Refer to “Preparing for IDB Recovery” on page 466 for more information on how to perform a IDB backup.
   - In case of Microsoft Cluster Server, consistent backup includes (in the same backup session):
     - ✓ all nodes
     - ✓ administrative virtual server (defined by the administrator)
     - ✓ if Data Protector is configured as a cluster aware application, Cell Manager virtual server and IDB.

     See “Restoring the Microsoft Cluster Server Specifics” on page 572 for details.
   - The disk with the boot partition requires free disk space that is needed for the Data Protector disaster recovery installation (15 MB) and active DR OS installation. Additionally, you also need as much free disk space, as required for the restore of the original system.
5. Copy the contents of `<Data_Protector_home>\Depot\DRSetup` or `\i386\tools\DRSetup` (located on Data Protector installation

medium) for 32 bit Windows Client or Cell Manager on three floppy disks (**drsetup diskettes**) or *<Data_Protector_home>*\Depot\DRSetup64 or \i386\tools\DRSetup64 (Data Protector installation medium) for 64 bit Windows systems on four floppy disks. In case of a disaster, save the updated SRD file of the crashed client to the first floppy disk (disk1). Only one set of drsetup diskettes is required per site for all Windows systems, but you must always copy an updated SRD file of the crashed client on the first floppy disk. If multiple SRD files are found, Data Protector will ask you to select the appropriate version.

6. In order to re-create disk partitions to their initial state prior to the disaster, record the following information for each partition (it will be needed during the recovery process):

   - partitions length and order
   - drive letters assigned to the partitions
   - partitions filesystem type

   This information is stored in the SRD file. The `-type` option in the diskinfo section of the SRD file shows the partition filesystem type for a particular partition:

Table 12-1 How to Determine the Filesystem Type from the SRD File

Type number	Filesystem
1	Fat12
4 and 6	Fat32
5 and 15	Extended partition
7	NTFS
11 and 12	Fat32
18	EISA
66	LDM partition

The table on the next page is an example of the preparation for the disaster recovery. Note that data in the table belongs to a specific system and cannot be used on any other system. Refer to “Windows Manual Disaster Recovery Preparation Template” on page A-49 for an empty template which can be used when preparing for the Assisted Manual Disaster Recovery.

Table 12-2 Example of the AMDR Preparation Template

client properties	computer name	ANDES
	hostname	andes.company.com
drivers		hpn.sys, hpncin.dll
Windows Service Pack		Windows SP3
TCP/IP properties	IP address	3.55.61.61
	default gateway	10.17.250.250
	subnet mask	255.255.0.0
	DNS order	11.17.3.108, 11.17.100.100
medium label / barcode number		“andes - disaster recovery” / [000577]
partition information and order	1st disk label	
	1st partition length	31 MB
	1st drive letter	
	1st filesystem	EISA
	2nd disk label	BOOT
	2nd partition length	1419 MB
	2nd drive letter	C:
	2nd filesystem	NTFS/HPFS
	3rd disk label	
	3rd partition length	
	3rd drive letter	
	3rd filesystem	

Recovery

Follow the procedure below to recover a Windows system using Assisted Manual Disaster Recovery. If you are performing advanced recovery tasks (such as disaster recovery of a Cell Manager or IIS), see also “Advanced Recovery Tasks” on page 572.

1. Install the Windows system from the CD-ROM and install additional drivers if needed. The Windows operating system has to be installed on the same partition as prior to the disaster. Do not install the Internet Information Server (IIS) during the installation of the system. Refer to “Restoring Internet Information Server (IIS) Specifics” on page 578 for more details.

IMPORTANT

If Windows has been installed using the Windows unattended setup, use the same script now to install Windows to ensure that the `<$SystemRoot$>` and `\Documents and Settings` folders are installed to the same position.

2. When the Windows Partition Setup screen appears, proceed as follows:

   - If an vendor-specific partition (e.g. EISA Utility Partition) existed on the system before the crash, create (if it does not exist due to the crash) and format a “dummy” FAT partition using the EUP information gathered from the SRD file. The EUP will be later on recovered to the space occupied by the “dummy” partition. Create and format a boot partition immediately after the “dummy” partition. To do this, you need the data as described in “Preparation” on page 527.
   - If an EUP did not exist on the system before the crash, create (if the boot partition does not exist due to the crash) and format the boot partition as it existed on the disk before the crash. To do this, you need the data as described in “Preparation” on page 527.

   Install Windows into its original location, i.e. the same drive letter and directory as in the original system before the disaster. This information is stored in the SRD file.

NOTE During the installation, do not add the system to the previous location where the Windows domain resided, but add the system to a workgroup instead.

3. Install TCP/IP protocol. If DHCP was not used before the disaster, configure the TCP/IP protocol as prior to the disaster by providing the following information: hostname of the crashed client, its IP address, default gateway, subnet mask and DNS server. Make sure that the field labeled `Primary DNS suffix of this computer` contains your domain name

NOTE By default, Windows install the Dynamic Host Configuration Protocol (DHCP) during the Windows setup.

4. Create a new temporary disaster recovery account in the Windows Administrators group and add it to the Data Protector Admin group on the Cell Manager. See “Adding or Deleting a User” on page 137.

   The account must not have existed on the system before the disaster. The temporary *Windows* account will be removed at a later time during this procedure.

5. Log off and log in to the system using the newly created account.
6. If the information in the SRD file is not up to date (for example, because you changed the backup device after the disaster) and you are performing an offline recovery, edit the SRD file before continuing with this procedure. See “Recovery Using an Edited SRD File” on page 580.
7. Run the `drstart.exe` command from the *<Data_Protector_home>*`\Depot\drsetup\Disk1` (Windows Cell Manager) or `\i386\tools\drsetup\Disk1` (Data Protector installation medium) directories.
   If you have prepared the drsetup diskettes (see “Preparation” on page 527), you can also execute the `drstart.exe` command from the first diskette.
8. `Drstart.exe` first scans the current working directory, floppy and CD drives for the location of disaster recovery setup files (`Dr1.cab`

and omnicab.ini). If the required files are found, the drstart utility installs the disaster recovery files in the *<%SystemRoot%>\system32\OB2DR* directory. Otherwise enter their path in the DR Installation Source text box or browse for the files.

9. If the recovery.srd file is saved in the same directory as dr1.cab and omnicab.ini files, drstart.exe copies recovery.srd file to the <%SystemRoot%>\system32\OB2DR\bin directory and the omnidr utility is started automatically. Otherwise, enter the location of SRD file (recovery.srd) in the SRD Path field or browse for the file. Click Next.

   If multiple SRD files are found on the floppy disk, Data Protector will ask you to select an appropriate version of the SRD file.

   After omnidr successfully finishes, all critical objects required for a proper boot of the system are restored.

10. Remove the temporary *Data Protector* user account (added in step 4) from the Data Protector Admin group on the Cell Manager, unless it existed on the Cell Manager before the disaster recovery.

11. Reboot the computer, log on and verify that the restored applications are running.

12. Additional steps are required if you are recovering a Cell Manager or performing advanced recovery tasks (such as restoring MSCS or IIS, editing the DRecoveryKB.cfg and SRD files). See “Restoring the Data Protector Cell Manager Specifics” on page 569 and “Advanced Recovery Tasks” on page 572 for more information.

13. Use Data Protector to restore user and application data.

The temporary DR OS will be deleted after the first login except in the following cases:

- You have interrupted the Disaster Recovery Wizard during the 10 seconds pause after it has found the DR installation and SRD file on the backup medium, and have selected the Use Debugs option.
- You have manually started the omnidr command with the no_reset or debug options.
- Disaster recovery fails.

Disk Delivery Disaster Recovery of a Windows Client

To perform the Disk Delivery Disaster Recovery, use a working Data Protector client (Data Protector disaster recovery host) to create the new disk while connected to this client. The administrator has to ensure before the disaster that enough data is collected to correctly format and partition the disk. However, Data Protector automatically stores the relevant information as part of the configuration backup.

The recovered partitions are:

- the boot partition
- the system partition
- the partitions containing Data Protector

Any remaining partitions can be recovered by using the standard Data Protector recovery procedure.

For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

TIP This method is specially useful with hot swap hard disk drives, because you can disconnect a hard disk drive from a system and connect a new one while the power is still on and the system is operating.

Requirements

- The partitions have to be the same size or larger than the partitions on the failed disk. This way the information stored on the failed disk can be restored to the new one. Also, the type of filesystem format has to match (FAT, NTFS).
- The system on which the disk is created and the system in which the disk is used have to use the same sector mapping/addressing (SCSI BIOS enabled/disabled; EIDE: both systems have to use the same addressing mode: LBA, ECHS, CHS).

Limitations

- Disk Delivery Disaster Recovery is not supported for Microsoft Cluster Server.
- RAID is not supported. This includes software RAIDs (fault-tolerant volumes and dynamic disks).
- Internet Information Server (IIS) Database, Terminal Services Database and Certificate Server Database are not restored automatically during Phase 2. They can be restored on the target system using the standard Data Protector restore procedure.

Preparation

Complete a few steps in order to prepare for disaster recovery. See also “Preparing for a Disaster Recovery” on page 519, for the general preparation procedure for all disaster recovery methods before completing the steps listed in this section.

IMPORTANT Prepare for disaster recovery *before* a disaster occurs.

In order to recover from a disaster quickly, efficiently and effectively, you need the following:

- The last valid known full backup of the client that you want to recover.
- A new hard disk to replace your crashed disk.
- A Data Protector hosting system, which has to be of the same operating system as the crashed client and must have the same hardware I/O path required to connect the new disk.

In order to re-create disk partitions to their initial state prior to the crash, record the following information for each partition (it will be needed during the recovery process):

- partitions length and order
- drive letters assigned to the partitions
- partitions filesystem type

You can refer to Table 12-2 on page 531 as an example of the preparation for the Disk Delivery disaster recovery. Refer to “Windows Manual Disaster Recovery Preparation Template” on page A-49 for an empty template which can be used when preparing for the Disaster Recovery.

Recovery

This section provides the procedure for recovering your Windows client using the Disk Delivery method. See also “Advanced Recovery Tasks” on page 572

With the Disk Delivery method on Windows, use a Data Protector disaster recovery host (DR host) to restore the last valid known full backup of your crashed disk to a new hard disk connected to the client. Then replace your crashed disk on the faulty system with this new hard disk.

Disk Delivery Disaster Recovery Procedure

The actual Disk Delivery Disaster Recovery procedure consists of the following steps:

1. Connect the new disk to a DR host.
2. Reboot the DR host to recognize the new disk.
3. Use Data Protector GUI on disaster recovery host and switch to the `Restore` context and click the `Tasks` tab. Select the `Disaster Recovery` item in the Scoping Pane, select the client from the drop down list and check the `Disaster recovery with disk delivery` in the Results Area.
4. For each of the critical objects, select an object version that will be restored and click `Next`.
5. If partitioning has not already been done, partition the new disk using the Disk Administrator. Use the partition information you have gathered as part of the preparation for Disk Delivery disaster recovery.
6. When partitioning the system, assign partitions in the same order as prior to the time that the full backup was performed. This simplifies drive letter reassignment after the restore and prevents a possibility of failure at system restart because of an inappropriate path to the system partition in the `boot.ini` file.

IMPORTANT

Assign drive letters for Windows mountpoints. In this case you must have enough unassigned drive letter available in order to be able to assign a drive letter for each mount point.

7. Perform all necessary drive letter mappings by right clicking on the original drive letter. This is necessary because drive letters on hosting and original system can be different.
8. Press `Finish`.
9. Remove the new disk from the DR host, and then connect it to the target system.
10. Power on the target system.
11. Use the standard Data Protector restore procedure to restore user and application data. This completes the recovery of the client.

Disk Delivery can also be a valuable method in case one of disks in a multi boot system has crashed, and the user can still boot at least one configuration.

NOTE

Data Protector does not restore volume-compression flag after recovery. All files, that were compressed at backup time, will be restored as compressed but you will have to manually set volume compression if you want any new files created to be compressed as well.

Enhanced Automated Disaster Recovery of a Windows System

Enhanced Automated Disaster Recovery (EADR) is a fully automated Data Protector recovery method for Windows clients and Cell Manager, where user intervention is reduced to minimum. For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

The EADR procedure for Windows platforms collects all relevant environment data automatically at backup time. During configuration backup, data required for temporary DR OS setup and configuration is packed in a single large **DR OS image file** and stored on the backup tape (and optionally on Cell Manager) for each backed up client in the cell.

In addition to this image file, a Phase 1 startup information (stored in the **P1S** file), required for correct formatting and partitioning of the disk is stored on the Cell Manager. When a disaster occurs, EADR Wizard is used to restore the DR OS image from the backup medium (if it has not been saved on the Cell Manager during the full backup) and convert it to a **disaster recovery CD ISO image**. CD ISO image can then be burned on a CD using any burning tool and used to boot the target system.

Data Protector then automatically installs and configures DR OS, formats and partitions the disks and finally recovers the original system with Data Protector as it was at the time of backup.

IMPORTANT Perform a new backup and prepare a new DR CD after each hardware, software or configuration change. This also applies to any network configuration changes, such as change of IP address or DNS server.

The recovered volumes are:

- the boot partition
- the system partition
- the partitions containing Data Protector

Any remaining partitions can be recovered using the standard Data Protector recovery procedure.

The following sections explain the limitations, preparation, and recovery that pertains to EADR of the Windows clients. See also “Advanced Recovery Tasks” on page 572.

Before selecting this method of disaster recovery, consider the following requirements and limitations:

Requirements

- The Data Protector Automatic Disaster Recovery component must be installed on clients for which you want to enable recovery using this method and on the system, where the DR CD ISO image will be prepared. See *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- The hardware configuration of the target system must be the same as of the original system. This includes SCSI BIOS settings (sector remapping).
- Replacement disks have to be attached to the same host bus adapter on the same bus.
- Boot partition has to be larger than 100 MB or disaster recovery will fail.
- An additional 200 MB of free disk space is required on the boot partition at backup time. If this disk space is not available, the disaster recovery fails. If you had applied the Compress Drive on the original partition, you must have 400 MB free.
- All drivers required for boot must be installed under `<%SystemRoot%>` folder.
- Network must be available when you boot the system in Safe Mode with Networking or in Directory Services Restore Mode (Domain Controller only), but you must do the backup of the system after it was booted with normal boot process.
- The system’s BIOS must support bootable CD extensions as defined in the El-Torito standard and read/write access to hard disk drive using LBA addressing via INT13h function XXh. The BIOS options can either be checked in the user’s manuals of the system or by inspecting the system setup before the boot.

- When backing up the client, the default 64 kB block size should be used to write to the device if you plan to perform an offline restore. This is the only default block size available on Windows when performing disaster recovery. To verify that the default 64 kB block size is set, choose `Advanced` in the `Properties` box, as shown in Figure 12-1 on page 541.

Figure 12-1 Verifying the Default Block Size

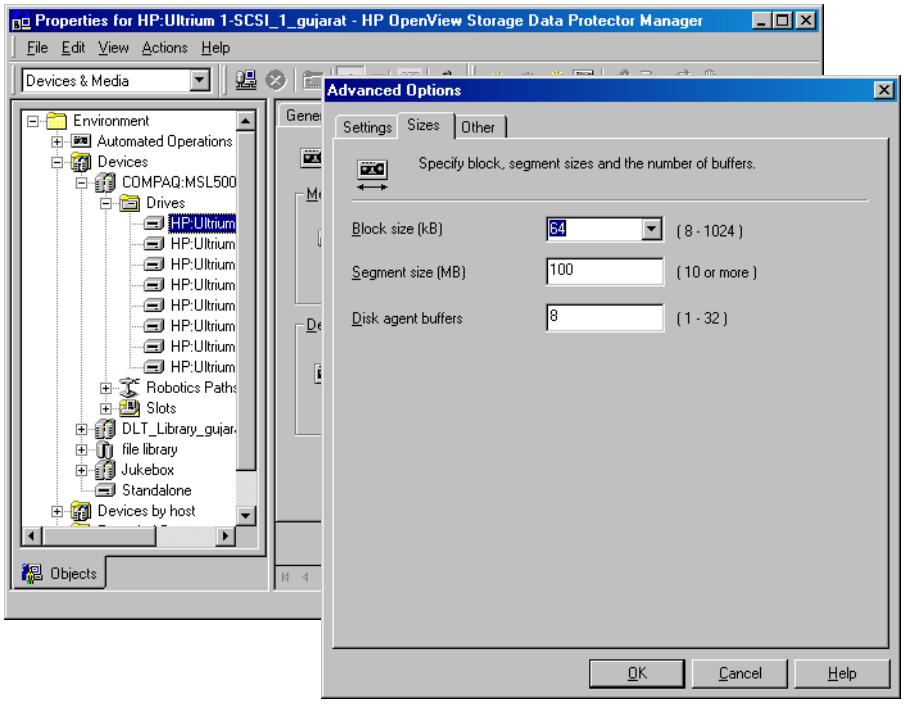

Limitations

Disk and Partition Configuration

- Dynamic disks are not supported (including mirror set upgraded from Windows NT).
- New disk has to be the same size or bigger than the crashed disk. If it is larger than the original disk, the difference will remain unallocated.

- Only vendor specific partitions of type 0x12 (including EISA) and 0xFE are supported for Enhanced Automated Disaster Recovery.

Other

- Multiboot systems that do not use Microsoft's boot loader are not supported.
- Internet Information Server (IIS), Terminal Services Database and Certificate Server Database are not restored automatically during Phase 2. They can be restored on the target system using the standard Data Protector restore procedure.

Preparation

See also "Preparing for a Disaster Recovery" on page 519, for the general preparation procedure for all disaster recovery methods before completing the steps listed in this section. See also "Advanced Recovery Tasks" on page 572.

IMPORTANT Prepare for disaster recovery *before* a disaster occurs.

Prerequisites

- Perform a full client backup (including the CONFIGURATION). See "Backing Up Filesystems (Logical Disk Drives)" on page 213 and "Backing Up CONFIGURATION" on page 218.

Microsoft Cluster Server

- In case of Microsoft Cluster Server, consistent backup includes (in the same backup session):
  - — all nodes
  - — administrative virtual server (defined by the administrator)
  - — if Data Protector is configured as a cluster aware application, Cell Manager virtual server and IDB.

  See "Restoring the Microsoft Cluster Server Specifics" on page 572 for details.

After you performed the backup, merge the P1S files for all nodes in the MSCS, so that P1S file of each node contains information on the shared cluster volumes configuration. Refer to "Merging P1S Files of all nodes for EADR" on page 576 for instructions.

DR Image File

Data required for temporary DR OS installation and configuration (**DR image**) is packed in a single large file and stored on the backup medium and optionally on the Cell Manager during a full client backup. If you want to save the full disaster recovery image file to the Cell Manager for all clients in the backup specification, perform the following steps:

1. In the Context List, select `Backup`.
2. In the `Scoping pane`, expand the `Backup Specifications` and then Filesystem.
3. Select the backup specification you will use for a full client backup (create it if you have not created it already).
4. In the Results Area, click `Options`.
5. Under `Filesystem Options` click `Advanced`.
6. Click the `WinFS Options` and check the `Copy full DR image to disk` check box.

Figure 12-2 WinFS Options Tab

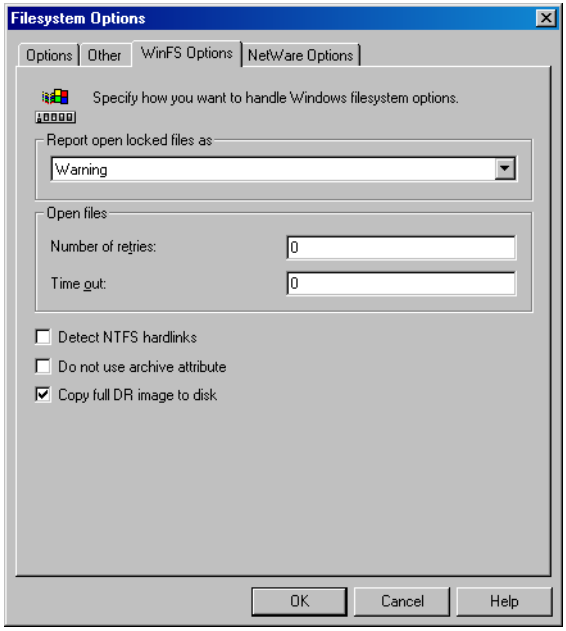

To copy the DR image files only for particular clients in the backup specification, perform the following steps:

1. In the Context List, select `Backup`.
2. In the `Scoping pane`, expand the `Backup Specifications` and then `Filesystem`.
3. Select the backup specification you will use for a full client backup. If you have not created it yet, do so using instructions in "Creating a Backup Specification" on page 199.
4. In the `Results Area`, click `Backup Object Summary`.
5. Select the client for which you would like to store the DR image file onto the Cell Manager and click `Properties`.
6. Click the `WinFS Options` and select `Copy full DR image to disk`.

Saving the full DR image to the Cell Manager is useful if you plan to burn the disaster recovery CD on the Cell Manager, because it is much faster to obtain the DR image from the hard disk than to restore it from a backup medium. The DR image file is by default saved into *<Data_Protector_home>*`\Config\server\dr\p1s` (Windows Cell Manager) or into `/etc/opt/omni/server/dr/p1s` (UNIX Cell Manager) with the name *<client name>*`.img`. To change the default location, specify a new global variable `EADRImagePath =` *<valid_path>* (for example, `EADRImagePath = /home/images` or `EADRImagePath = c:\temp`) in the global options file. Refer to "Global Options File" on page 613 for information.

TIP

If you do not have enough free disk space in the destination directory, you can create a link to another volume on UNIX or create a mount point on Windows.

DRecoveryKB.cfg File

The purpose of this file is to provide a flexible method to enable Data Protector to include drivers (and other needed files) in the DR OS to cover systems with specific boot relevant hardware or application configurations. The default `DRecoveryKB.cfg` file already contains all files necessary for industry standard hardware configurations.

Create and execute a test plan using the default version of the `DRecoveryKB.cfg` file. If the DR OS does not boot normally or cannot access network, then you may need to modify the file. Refer to "Editing the DRecoveryKB.cfg File" on page 579.

Phase 1 Startup File (P1S)

In addition to the DR image file, a **Phase 1 Startup file** (**P1S**) is created during full backup. It is saved on backup medium and on the Cell Manager into *<Data_Protector_home>*`\Config\server\dr\p1s` directory (Windows Cell Manager) or into `/etc/opt/omni/server/dr/p1s` directory (UNIX Cell Manager) with the filename equal to the hostname (for example, `computer.company.com`). It is a Unicode UTF-8 encoded file that contains information on how to format and partition all disks installed in the system, whereas the updated SRD file contains only system information and data about backup objects and corresponding media.

After a disaster occurs, you can use the EADR Wizard to merge DR image, SRD and P1S files with disaster recovery installation into a **disaster recovery CD ISO image**, which can be burned on a CD using any CD burning tool that supports the ISO9660 format. This **disaster recovery CD** can then be used to perform automated disaster recovery.

IMPORTANT

Disaster recovery CD has to be prepared in advance for the Cell Manager.

Additional steps are required if you are preparing disaster recovery CD of a Microsoft Cluster node. See "Restoring the Microsoft Cluster Server Specifics" on page 572.

IMPORTANT

It is recommended to restrict access to backup media, DR images, SRD files and disaster recovery CDs due to security reasons.

Preparing DR CD ISO Image

To prepare a DR CD ISO image, perform the following steps:

1. In the Context List, select `Restore`.
2. Click the `Tasks` navigation tab and select `Disaster Recovery` in the Scoping Pane.

3. From the drop down list in the Results Area, select the client you would like to recover.
4. Click `Enhanced Automated Disaster Recovery` and then `Next`.
5. For each critical object select an appropriate object version and click `Next`.
6. If you have saved the DR image file on the Cell Manager, specify its location, otherwise click `Restore from backup medium`. Click `Next`.
7. Select the destination directory where you want to place the ISO CD image (`recovery.iso`) and click `Finish` to create the ISO CD image.

CAUTION If you place a new ISO CD image to a location where a `recovery.iso` is already located, the old ISO CD image will be overwritten by the new one without a warning.

8. Burn the disaster recovery ISO CD image on a CD using any CD burning tool that supports the ISO9660 format.

IMPORTANT Perform a new backup and prepare a new DR CD after each hardware, software or configuration change. This also applies to any network configuration changes, such as change of IP address or DNS server.

Recovery

You need the following to successfully perform a disaster recovery on the crashed system:

- A new hard disk to replace your crashed disk.
- A successful full client backup of the client that you want to recover.
- The Data Protector disaster recovery CD.

EADR Procedure

The following is a step-by-step procedure for performing EADR of a Windows system:

1. Unless you are performing an offline disaster recovery, add the client's local Administrator account to the Data Protector Admin user group on the Cell Manager. See "Adding or Deleting a User" on page 137.

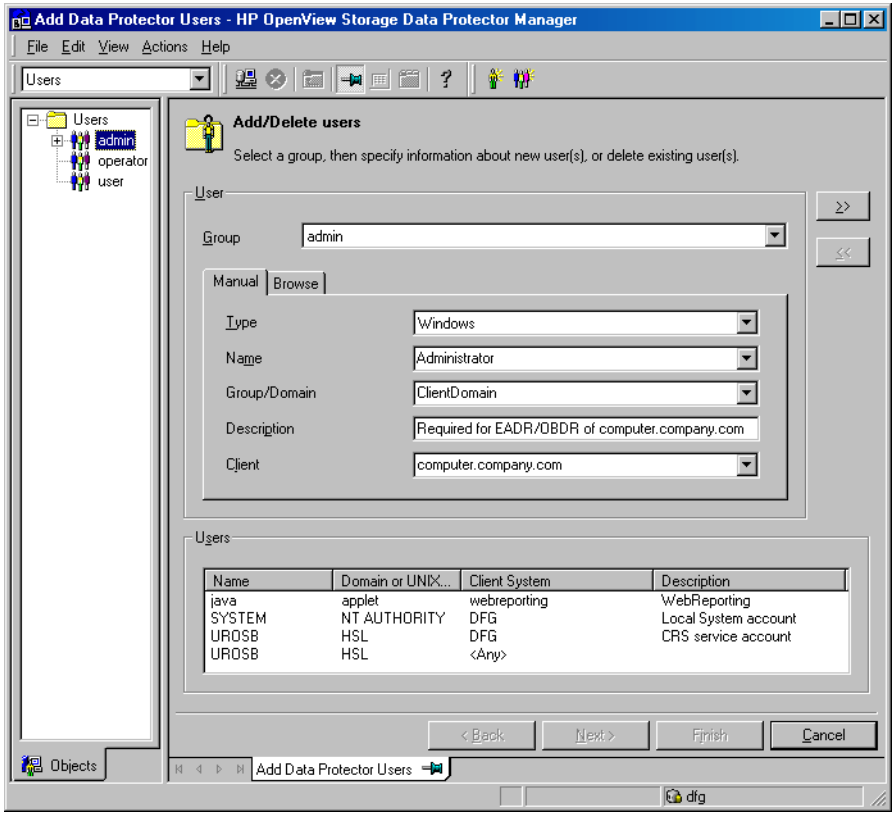

2. Boot from the disaster recovery CD of the original system.
3. Press **F12** when the following message is displayed: `To start recovery of the machine <HOSTNAME> press F12.`

4. Select the scope of recovery and press **Enter**. There are 4 different scopes of recovery:

   - `Reboot`: Disaster recovery is not performed and the computer is rebooted.
   - `Default Recovery`: Critical volumes are recovered. All other disks are not partitioned and formatted and are ready for Phase 3.
   - `Minimal Recovery`: Only system and boot disks are recovered (available for EADR and OBDR only).
   - `Full with Shared Volumes`: Available for MSCS only. This option should be used if all nodes in the MSCS have crashed and you are performing Enhanced Automated Disaster Recovery of the first node. It will recover all volumes in the Restore Set including cluster shared volumes that were locked by the backed-up node at backup time.

     If at least one node is up and the MSCS service is running, than shared volumes will not be restored because the node keeps them locked. In this case, you should use `Default Recovery`.

5. After you have selected the scope of the recovery, Data Protector starts setting up the DR OS directly to the hard disk. You can monitor the progress and, when the DR OS is set up, the system reboots.
6. Wait for 10 seconds when prompted `To start recovery of the machine <HOSTNAME> press F12`, to boot from the hard disk and not from the CD.
7. If you are performing an offline recovery and the information in the SRD file is not up to date (for example, because you changed the backup device after the disaster), edit the SRD file, before you can continue with this procedure. Refer to “Recovery Using an Edited SRD File” on page 580.
8. Data Protector will then reestablish the previous storage structure within the selected scope of recovery and restore all critical volumes. The temporary DR OS will be deleted after the first login, except in the following cases:

   - `Minimal Recovery` is selected.
   - You have interrupted the Disaster Recovery Wizard during the 10 seconds pause after it has found the DR installation and SRD file on the backup medium, and have selected the `Use Debugs` option.

- You have manually started the `omnidr` command with the `no_reset` or `debug` options.
- Disaster recovery fails.

9. Remove the client's local Administrator account created in step 1 from the Data Protector Admin user group on the Cell Manager, unless it existed on the Cell Manager before the disaster recovery.
10. Additional steps are required if you are recovering a Cell Manager or performing advanced recovery tasks (such as restoring MSCS or IIS, editing the `DRecoveryKB.cfg` and SRD files). See "Restoring the Data Protector Cell Manager Specifics" on page 569 and "Advanced Recovery Tasks" on page 572 for more information.
11. Restore user and application data using the standard Data Protector restore procedure.

NOTE Data Protector does not restore the volume-compression flag after recovery. All files that were compressed at backup time will be restored as compressed, but you will have to manually set the volume compression if you want any newly created files to be compressed as well.

One Button Disaster Recovery of a Windows System

One Button Disaster Recovery (OBDR) is a fully automated Data Protector recovery method for Windows clients and Cell Manager, where user intervention is reduced to minimum. For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

OBDR collects all relevant environment data automatically at backup time. During backup, data required for temporary DR OS setup and configuration is packed in a single large OBDR image file and stored on the backup tape. When a disaster occurs, OBDR device (backup device, capable of emulating CD-ROM) is used to boot the target system directly from the tape which contains the OBDR image file with disaster recovery information.

Data Protector then installs and configures the disaster recovery operating system (DR OS), formats and partitions the disks and finally restores the original operating system with Data Protector as it was at the time of backup.

IMPORTANT Perform a new backup after each hardware, software or configuration change. This also applies to any network configuration changes, such as change of IP address or DNS server.

The recovered volumes are:

- the boot partition
- the system partition
- the partitions containing Data Protector

Any remaining partitions can be recovered using the standard Data Protector recovery procedure.

The following sections explain the requirements, limitations, preparation and recovery pertaining to One Button Disaster Recovery on Windows systems. See also “Advanced Recovery Tasks” on page 572.

Requirements

- Data Protector Automatic Disaster Recovery and User Interface components must be installed on the systems for which you want to enable recovery using this method. See *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- It is essential to have an OBDR capable computer configuration: the system's BIOS must support bootable CD extensions as defined in the El-Torito standard and read/write access to hard disk drive using LBA addressing via INT13h function XXh. The OBDR device must conform to the same standard when emulating the CD-ROM. The BIOS options can either be checked in the user's manuals of the system or by inspecting the system setup before the boot.

  For more information about supported systems, devices and media, please refer to the HP StorageWorks Tape Hardware Compatibility Table on the World Wide Web:

  http://www.openview.hp.com/products/datapro/spec_0001.html. Also see the *HP OpenView Storage Data Protector Software Release Notes*.
- The hardware configuration of the target system must be the same as of the original system. This includes SCSI BIOS settings (sector remapping).
- Replacement disks have to be attached to the same host bus adapter on the same bus.
- An additional 200 MB of free disk space is required on the boot partition at backup time. If this disk space is not available, the disaster recovery fails. If you had applied the Compress Drive on the original partition, you must have 400 MB free.
- All drivers, required for boot must be installed under the `<%SystemRoot%>` folder.
- Network must be available when you boot the system in Safe Mode with Networking or in Directory Services Restore Mode (Domain Controller only), but you must do the backup of the system after it was booted with normal boot process.
- A media pool with a `Non-appendable` media usage policy and `Loose` media allocation policy has to be created for the OBDR capable device. Only the media from such pool can be used for disaster recovery.

- When backing up the client, the default 64 kB block size should be used to write to the device if you plan to perform an offline restore. This is the only default block size available on Windows when performing disaster recovery. To verify that the default 64 kB block size is set, choose `Advanced` in the `Properties` box, as shown in Figure 12-3 on page 552.

Figure 12-3 Verifying the Default Block Size

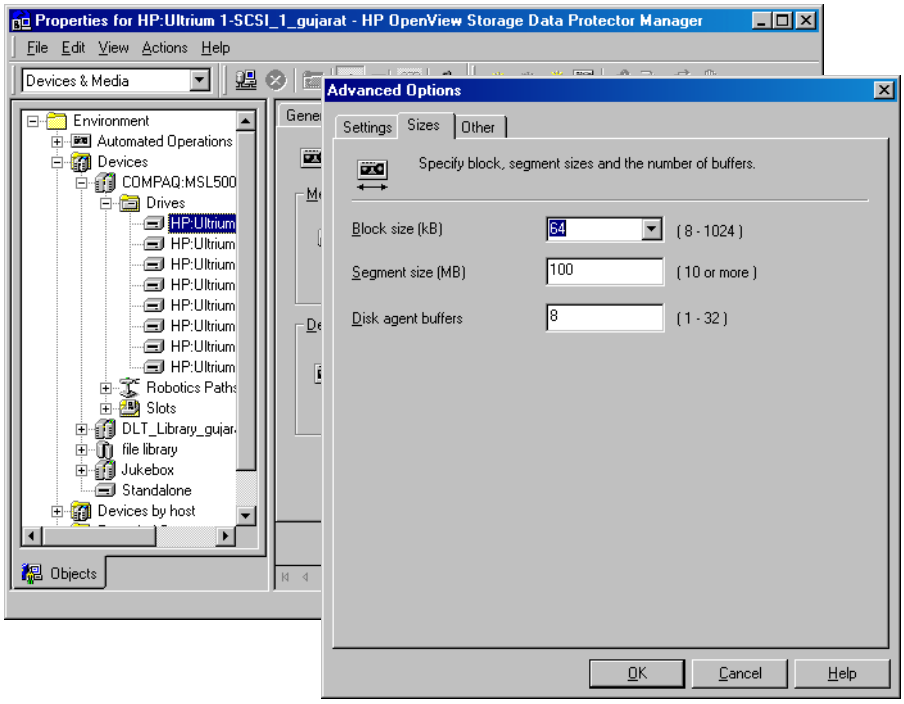

Limitations

General

- Multiboot systems that do not use Microsoft's boot loader are not supported.
- Internet Information Server (IIS) Database, Terminal Services Database and Certificate Server Database are not restored automatically during Phase 2. They can be restored on the target system using the standard Data Protector restore procedure.

- One Button Disaster Recovery backup session can only be performed for one selected client or Cell Manager on the same OBDR device at a time. This has to be done on a single, locally attached OBDR capable device.

Disk and Partition Configuration

- Dynamic disks are not supported (including mirror set upgraded from Windows NT).
- New disk has to be the same size or bigger than the crashed disk. If it is larger than the original disk, the difference will remain unallocated.
- Only vendor specific partitions of type 0x12 (including EISA) and 0xFE are supported for OBDR.

Preparation

See also "Preparing for a Disaster Recovery" on page 519, for the general preparation procedure for all disaster recovery methods before completing the steps listed in this section. See also "Advanced Recovery Tasks" on page 572.

IMPORTANT Prepare for disaster recovery *before* a disaster occurs.

Create a media pool for DDS or LTO media with `Non-appendable` media usage policy (to ensure that this will be the only backup on tape) and `Loose` media allocation policy (because the tape is formatted during OBDR backup). In addition, select this media pool as a default media pool for the OBDR device. Refer to "Creating a Media Pool" on page 149 for more information. Only media from such pool can be used for OBDR.

Microsoft Cluster Server

In case of Microsoft Cluster Server, consistent backup includes (in the same backup session):

- all nodes
- administrative virtual server (defined by the administrator)
- if Data Protector is configured as a cluster aware application, Cell Manager virtual server and IDB.

See "Restoring the Microsoft Cluster Server Specifics" on page 572 for details.

To enable an automatic restore of all shared disk volumes on the MSCS using the OBDR method, move all volumes temporarily to the node for which you are preparing the OBDR boot tape so that shared disk volumes are not locked by another node during the OBDR backup. It is namely impossible to collect enough information for configuring the disk during Phase 1 for shared disk volumes that are locked by another node during the backup.

OBDR Backup

Use the following steps to perform OBDR backup locally on the system, for which you want to enable recovery using OBDR:

1. In the Context List, select `Backup`.
2. Click `Tasks` navigation tab and check `One Button Disaster Recovery Wizard` in the `Scoping Pane`.
3. From the drop-down list in the `Results Area`, select the client for which you would like to perform OBDR backup and click `Next`.
4. All critical objects are already selected (including the IDB in case of the Cell Manager OBDR backup) and can not be deselected. Manually select any other partitions you want to keep, because during the recovery procedure, Data Protector deletes all partitions from your system. Click `Next`.
5. Select the locally attached OBDR device you are going to use for backup and click `Next`.
6. Select backup options. See “Using Backup Options” on page 269 for details.
7. Click `Next` to proceed to the `Scheduler` page, which can be used to schedule the backup. See “Scheduling Unattended Backups” on page 250 for more information.
8. Click `Next` to display the `Backup Object Summary` page, in which you can review the backup options.

NOTE

In the `Summary page`, you cannot change a previously selected backup device or the order in which the backup specifications follow one another (move up and move down functionalities are not available). Only OBDR non-essential backup objects can be deleted as well as general object properties can be viewed.

However, a backup object's description can be changed.

9. In the final page of the Backup wizard, you can save the backup specification, start the interactive backup, or preview the backup.

   It is recommended to save the backup specification so that you can schedule or modify it later.

Modifying an OBDR Backup Specification

Once a backup specification is saved, you can edit it. Right-click the backup specification and select `Properties`. You are offered to treat the modified backup specification as a standard Data Protector backup specification or as an OBDR backup specification. Save it as an OBDR backup specification to keep it in the original One Button Disaster Recovery format. If saved as a standard backup specification, it is not usable for OBDR purposes.

10. Click `Start Backup` to run the backup interactively. The `Start Backup` dialog box appears. Click `OK` to start the backup.

A bootable image file of the system, containing all information required for installation and configuration of temporary DR OS, will be written at the beginning of the tape to make it bootable.

IMPORTANT

Perform a new backup and prepare a bootable backup medium after each hardware, software or configuration change. This also applies to any network configuration changes, such as change of IP address or DNS server.

DRecoveryKB.cfg File

The purpose of this file is to provide a flexible method to enable Data Protector to include drivers (and other needed files) in the DR OS to cover systems with specific boot relevant hardware or application configurations. The default `DRecoveryKB.cfg` file already contains all files necessary for industry standard hardware configurations.

Create and execute a test plan using the default version of the `DRecoveryKB.cfg` file. If the DR OS does not boot normally or cannot access network, then you may need to modify the file. Refer to “Editing the DRecoveryKB.cfg File” on page 579.

CAUTION It is recommended to restrict access to backup media due to security reasons.

Recovery

You need the following to successfully perform a disaster recovery on the crashed system:

- A new hard disk to replace your crashed disk (if needed).
- A bootable backup medium with all critical objects of the client that you want to recover.
- An OBDR device connected locally to the target system.

OBDR Procedure

The following is a step-by-step procedure for performing a One Button Disaster Recovery of a Windows system:

1. Unless you are performing an offline disaster recovery, add the client's local Administrator account to the Data Protector Admin user group on the Cell Manager. See "Adding or Deleting a User" on page 137.

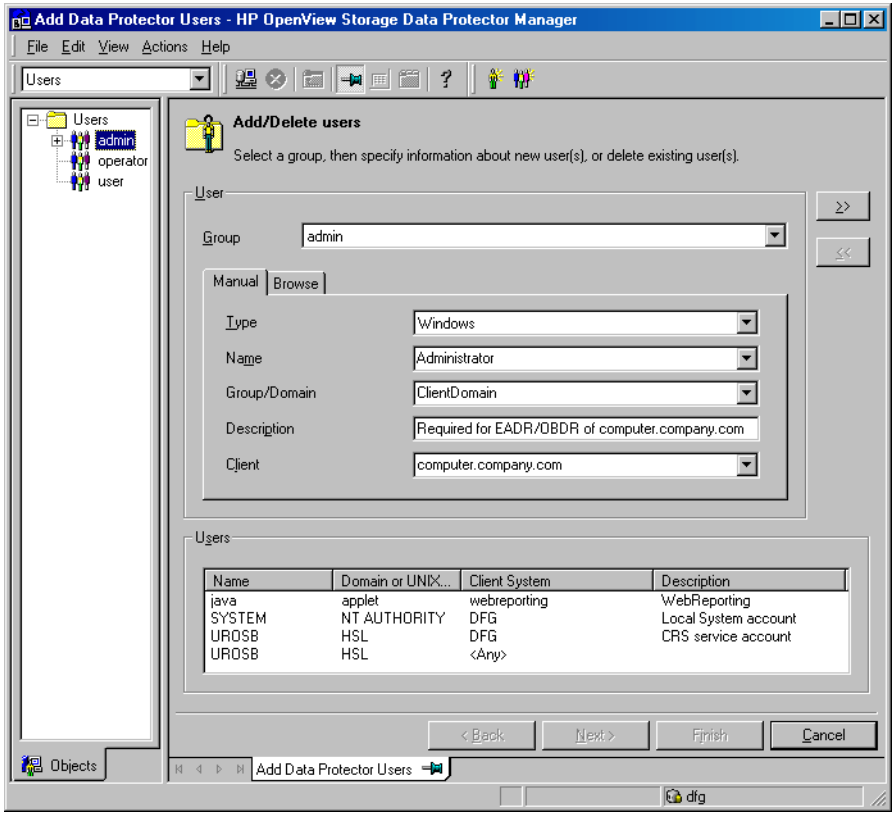

2. Insert the tape containing the image file and your backed up data into an OBDR device.
3. Shut down the target system and power off the tape device.
4. Power on the target system and while it is being initialized, press the eject button to the tape device and power it on. For details see the device documentation.

5. In the screen that appears, select the scope of recovery and press **Enter**. There are 4 different scopes of recovery:

   - `Reboot`: Disaster recovery is not performed and the computer is rebooted.
   - `Default Recovery`: Critical volumes are recovered. All other disks are not partitioned and formatted and remain empty and ready for Phase 3.
   - `Minimal Recovery`: Only system and boot disks are recovered (available for EADR and OBDR only).
   - `Full with Shared Volumes`: Available for MSCS only. This option should be used if all nodes in the MSCS have crashed and you are performing One Button Disaster Recovery of the first node. It will recover all volumes in the Restore Set including cluster shared volumes that were locked by the backed-up node at backup time.

TIP

To enable automatic restore of all shared disk volumes in the MSCS, move all volumes temporarily to the node, for which you are preparing OBDR boot tape. It is namely impossible to collect enough information to configure disks in Phase 1 for shared disk volumes that are locked by another node at backup.

   If at least one node is up and running than shared volumes will not be restored because the node keeps them locked. In this case, you should use `Default Recovery`.

6. After you have selected the scope of recovery, Data Protector starts setting up the DR OS directly to the hard disk. You can monitor the progress and, when the DR OS is set up, the system reboots.
7. If you are performing an offline recovery and the information in the SRD file is not up to date (for example, because you changed the backup device after the disaster), edit the SRD file, before you can continue with this procedure. Refer to “Recovery Using an Edited SRD File” on page 580.

8. Data Protector will then reestablish the previous storage structure and restore all critical volumes. The temporary DR OS will be deleted after the first login, except for the following cases:
   - `Minimal Recovery` is selected.
   - You interrupted the Disaster Recovery Wizard during the 10 seconds pause (after it had found the DR installation and the SRD file on the backup medium), and selected the `Use Debugs` option.
   - You manually started the `omnidr` command with the `no_reset` or `debug` options.
   - Disaster recovery fails.
9. Remove the client's local Administrator account created in step 1 from the Data Protector Admin user group on the Cell Manager, unless it existed on the Cell Manager before the disaster recovery.
10. Additional steps are required if you are recovering a Cell Manager or performing advanced recovery tasks (such as restoring MSCS or IIS, editing the `DRecoveryKB.cfg` and SRD files). See "Restoring the Data Protector Cell Manager Specifics" on page 569 and "Advanced Recovery Tasks" on page 572 for more information.
11. Restore the user and application data using the standard Data Protector restore procedure.

NOTE Data Protector does not restore the volume-compression flag after recovery. All files that were compressed at backup time will be restored as compressed, but you will have to manually set the volume compression if you want any new files created to be compressed as well.

Automated System Recovery

Automated System Recovery (ASR) is an automated system on Windows systems, which reconfigures a disk to its original state (or resizes the partitions if the new disk is larger than the original disk) in the case of a disaster. This includes disk partitioning and logical volume configuration (file formats, drive letter assignments, volume mountpoints, and volume characteristics). ASR thus enables the Data Protector `drstart.exe` command to install the active DR OS which provides Data Protector disk, network, tape and file system access.

Data Protector then recovers the target system to the original system configuration and finally restores all user data.

For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

IMPORTANT Perform a full client backup after each hardware, software or configuration change and update the ASR diskettes. This also applies to any network configuration changes, such as change of the IP address or DNS server.

IMPORTANT Create the ASR set for the Cell Manager in advance, because you will not be able to obtain the ASR archive file after the disaster. ASR sets for other systems can be created using Cell Manager when a disaster occurs.

The recovered volumes are:

- the boot partition
- the system partition
- the partitions containing Data Protector

Any remaining partitions can be recovered using the standard Data Protector recovery procedure.

The following sections explain the requirements, limitations, preparation, and recovery pertaining to Automated System Recovery on Windows systems. See also “Advanced Recovery Tasks” on page 572.

Requirements

- Data Protector Automatic Disaster Recovery component must be installed on systems for which you want to enable recovery using ASR. See the *HP OpenView Storage Data Protector Installation and Licensing Guide*.
- DHCP server must be configured in your network to enable online ASR.

Hardware Configuration

- The hardware configuration of the target system must be identical to that of the original system, except for hard disk drives, video cards and network interface cards. If you have replaced a network card or a video card, you will have to manually configure it.
- Floppy disk drive must be installed.
- Floppy and CD drives must be connected to IDE or SCSI controllers. External devices such as USB or PCMCIA devices are not supported.

Hard Disk Drives

- The target system must have the same number of physical disks with critical volumes as the original system.
- Replacement disks must be attached to the same host bus adapter on the same bus.
- The storage capacity of each replacement disk on the target system must be bigger than or equal to the capacity of the corresponding disk on the original system. In addition, disk geometry of the replacement disk must be the same as on the replaced disk.
- All disks on the target system must have 512 bytes-per-sector.
- All disks used in ASR must be accessible to the system (hardware RAID must be configured, SCSI disks must be correctly terminated, etc.)
- When backing up the client, the default 64 kB block size should be used to write to the device if you plan to perform an offline restore. This is the only default block size available on Windows when

performing disaster recovery. To verify that the default 64 kB block size is set, choose `Advanced` in the `Properties` box, as shown in Figure 12-4:

Figure 12-4 Verifying the Default Block Size

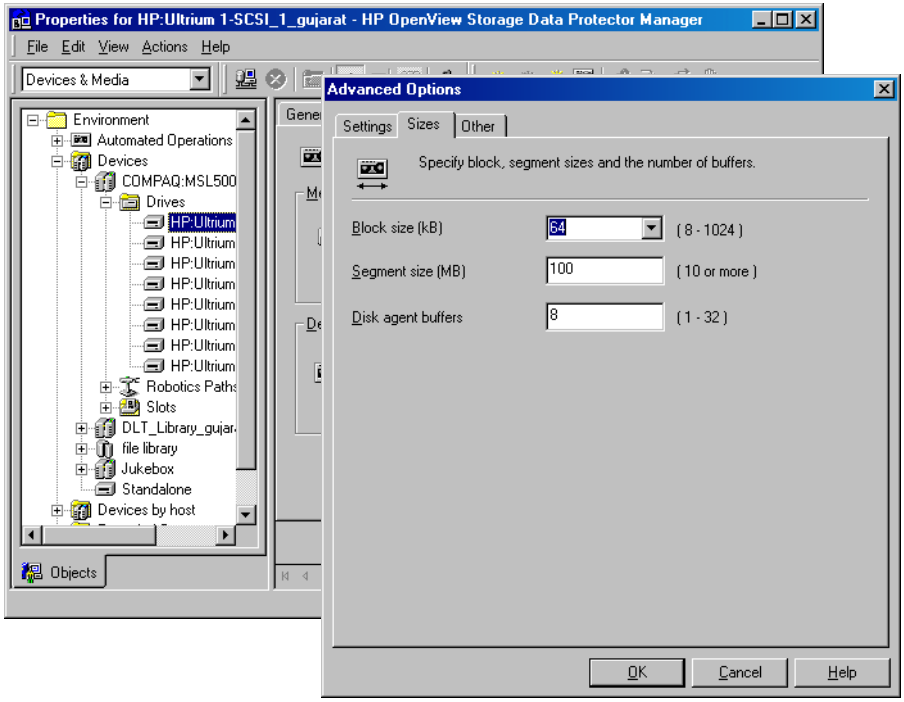

Limitations

- Windows XP Home Edition does not support ASR.
- Multiboot systems that do not use Microsoft's boot loader are not supported.
- Internet Information Server (IIS) Database, Terminal Services Database, and Certificate Server Database are not restored automatically during Phase 2. They can be restored on the target system using the standard Data Protector restore procedure.
- Data stored on vendor specific partitions is not automatically restored

during ASR. The partitions will be recreated during the ASR but you will have to restore the data manually using the vendor specific procedure for restoring data. However, you can restore data on EISA utility partition using the standard Data Protector restore procedure.

- Only those local backup devices are supported, that can be installed by Windows during OS installation (no additional drivers are required).

Preparation

See also "Preparing for a Disaster Recovery" on page 519, for the general preparation procedure for all disaster recovery methods before completing the steps listed in this section. See also "Advanced Recovery Tasks" on page 572 in order to prepare for disaster recovery.

IMPORTANT Prepare for disaster recovery *before* a disaster occurs.

Prerequisite

- A full client backup (including the configuration) is a prerequisite for successful ASR. See "Backing Up Filesystems (Logical Disk Drives)" on page 213 and "Backing Up CONFIGURATION" on page 218.

  In case of Microsoft Cluster Server, consistent backup includes (in the same backup session):

  — all nodes

  — administrative virtual server (defined by the administrator)

  — if Data Protector is configured as a cluster aware application, Cell Manager virtual server and IDB.

  See "Restoring the Microsoft Cluster Server Specifics" on page 572 for details.

After you performed the full client backup prepare an ASR set. An ASR set is a collection of files stored on three or four diskettes, required for proper reconfiguration of the replacement disk (disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and the user data that was backed up during the full client backup. These files are stored as an ASR archive file on the Cell Manager (in `<Data_Protector_home>\Config\server\dr\asr` on Windows or in `/etc/opt/omni/server/dr/asr/` on UNIX) as well as

on the backup medium. ASR archive file is extracted to three diskettes for 32-bit Windows system or four diskettes for 64-bit Windows system after a disaster occurs. You need these diskettes to perform ASR.

NOTE Create the ASR set for the Cell Manager in advance, because you will not be able to obtain the ASR archive file after the disaster.

Creation of ASR Set

Perform the following steps to create an ASR set:

1. Perform a full client backup.
2. Insert a diskette in the floppy drive.
3. In the `HP OpenView Storage Data Protector Manager`, switch to the `Restore` context.
4. Click the `Tasks` navigation tab and select `Disaster Recovery` in the Scoping Pane.
5. From the drop down list in the Results Area, select the client for which you would like to create an ASR set.
6. Click `Create Automated System Recovery set` and then click `Next`.

Figure 12-5 Creating ASR Set

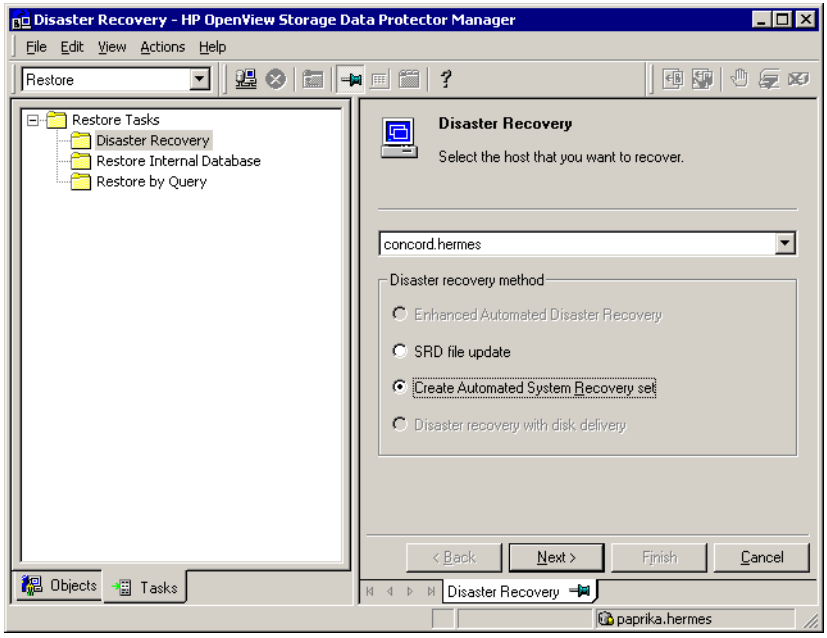

Data Protector will obtain the ASR archive file from the Cell Manager. If it is not saved on the Cell Manager, the Disaster Recovery wizard will offer you to recover it from the backup medium.

7. For each critical object, select the appropriate object version and click `Next`.
8. ASR archive file created during a full client backup is downloaded from the Cell Manager. Select the destination location where you want your ASR archive file extracted and select the `Copy DR installation` check box to copy DR installation files to the same location. The recommended destination is your floppy drive because you will need these files stored on diskettes (ASR set) to perform ASR.

   Data Protector will create three diskettes for a 32 bit Windows system and four diskettes for a 64 bit Windows system. ASR set for the Cell Manager has to be prepared in advance, while you can prepare ASR diskettes for other systems using the Cell Manager when a disaster occurs.

Once the ASR set is created, you have to update only the first diskette (which contains ASR information) after each hardware, software or configuration change. This also applies to any network configuration changes, such as a change of the IP address or DNS server. In order to update the first diskette from the ASR set, repeat the whole procedure, but you do not have to select the `Copy DR installation` check box. This option copies the DR installation files (to a selected destination), which do not need to be updated.

IMPORTANT It is recommended to restrict access to ASR diskettes due to security reasons.

Local Devices If you are using a locally attached device for ASR, test if it is supported. To do so, perform the following steps:

1. Run `devbra -dev` from the command prompt (from *<Data_Protector_home>*\bin).
2. Rename the `scsitab` file (located in *<Data_Protector_home>*) and run `devbra -dev` from the command prompt again.
3. Compare the both outputs of the `devbra -dev` command. If they are identical, ASR using this device is possible, otherwise copy the `scsitab` file to the first ASR diskette. You have to copy the `scsitab` file only the first time you are preparing the ASR set. You do not have to copy it when you are only updating the ASR set. Refer to the “Support of New Devices” on page 48 for more information.
4. Rename the `scsitab` file back to the original name.

Recovery

To successfully perform a disaster recovery of the crashed system, you need the following:

- A new hard disk to replace your crashed disk.
- A successful full client backup of the client that you want to recover.
- Updated ASR set.
- Windows installation medium.

ASR Procedure

The following is a step-by-step procedure for performing ASR:

1. Boot from the Windows installation medium.
2. Press **F2** during the start of the OS setup to enter the ASR mode.
3. If the information in the SRD file on the ASR diskette is not up to date (for example, because you changed the backup device after the disaster) and you are performing an offline recovery, edit the SRD file before continuing with this procedure. See “Recovery Using an Edited SRD File” on page 580.
4. Provide the first (updated) diskette from the ASR set.
5. After reboot, Disaster Recovery Wizard pops-up and requires input for the `DR installation source` and `SRD Path`. DR installation and SRD file are both located on the first diskette of the ASR set (`a:\`).
6. Unless you are performing an offline disaster recovery, add the client’s local system account to the Data Protector Admin user group on the Cell Manager. See “Adding or Deleting a User” on page 137.

Enter the same information as in Figure 12-6 on page 568.

Figure 12-6 User Name for ASR

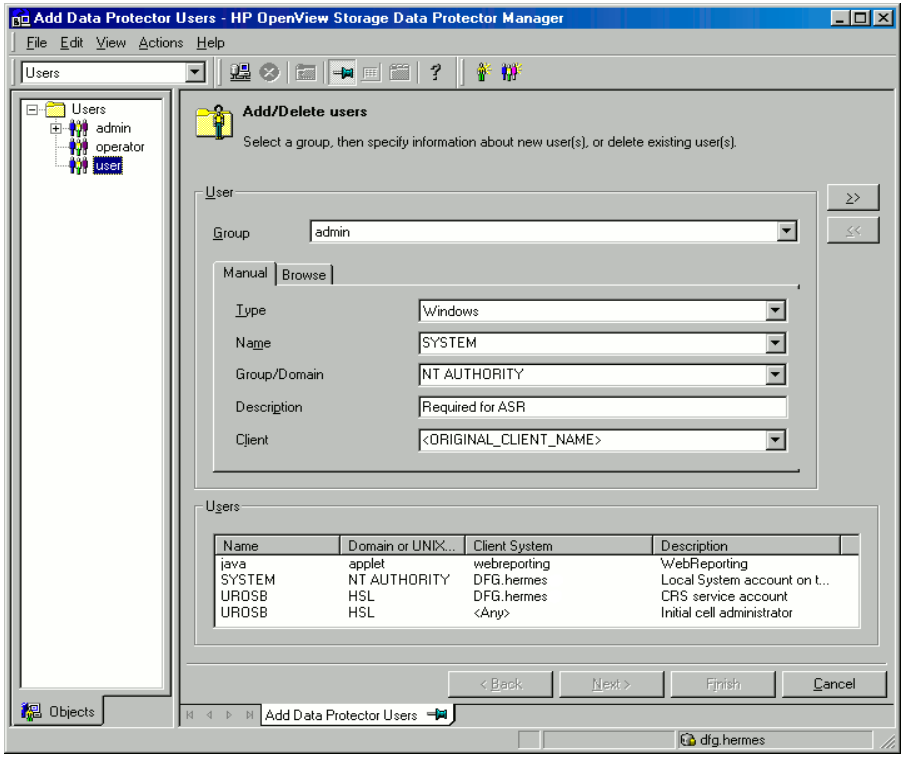

7. Change diskette(s) when prompted.

   Original storage structure will be automatically reestablished and all critical data automatically restored based on the information in the ASR set.

8. Reboot the system when prompted and remove the Windows installation medium and ASR diskette.

9. Remove the client's local system account (created in step 6) from the Data Protector Admin user group on the Cell Manager, unless it existed on the Cell Manager before the disaster recovery.

10. Restore user and application data using the standard Data Protector restore procedure.

Restoring the Data Protector Cell Manager Specifics

This section explains additional steps for particular methods that should be performed when restoring Windows Cell Manager.

Making IDB consistent (all methods)

The procedure described in this section should only be used after you have performed the general disaster recovery procedure.

To make the IDB consistent, import the medium with the last backup so that the information about the backed up objects is imported to the database. In order to do so, perform the following steps:

1. Using the Data Protector GUI, recycle the medium or media with the backup of the partitions that remain to be restored for enabling the medium or media to be imported in the IDB. Refer to “Recycling Media” on page 110 for more information on how to do this. Sometimes it is not possible to recycle a medium since Data Protector keeps it locked. In such a case stop Data Protector processes and delete the `\tmp` directory by running the following commands:
   ```
   <Data_Protector_home>\bin\omnisv -stop
   del <Data_Protector_home>\tmp\*.*
   <Data_Protector_home>\bin\omnisv -start
   ```
2. Using the Data Protector GUI, export the medium or media with the backup of the partitions that remain to be restored. Refer to “Exporting Media from Data Protector” on page 112 for more information on how to do this.
3. Using the Data Protector GUI, import the medium or media with the backup of the partitions that remain to be restored. Refer to “Importing Media” on page 100 for more information on how to do this.

Enhanced Automated Disaster Recovery Specifics

Two additional steps are required in Phase 0 if you are recovering Windows Cell Manager using Enhanced Automated Disaster Recovery:

- Disaster recovery CD for the Cell Manager should be prepared in advance.

IMPORTANT Perform a new backup and prepare a new DR CD after each hardware, software or configuration change. This also applies to any network changes, such as change of IP address or DNS server.

- In addition to the Cell Manager, you should save the updated SRD file of the Cell Manager on several secure locations as a part of the disaster recovery preparation policy, because the SRD file is the only file in Data Protector where information about objects and media is stored, when IDB is not available. If the SRD file is saved only on the Cell Manager, it is not accessible if the Cell Manager fails. See “Preparation” on page 527.

IMPORTANT It is recommended to restrict access to backup media, DR images, SRD files and disaster recovery CDs.

One Button Disaster Recovery Specifics

Since the IDB is not available if the Cell Manager has crashed, you have to know the location of OBDR bootable medium.

IMPORTANT Perform a new OBDR backup and prepare a new bootable medium after each hardware, software or configuration change. This also applies to any network changes, such as change of IP address or DNS server.

IMPORTANT It is recommended to restrict access to backup media.

Automated System Recovery Specifics

An additional step is required in Phase 0 if you are recovering Windows Cell Manager using Automated System Recovery (ASR):

- ASR diskette for the Cell Manager should be prepared in advance.

IMPORTANT Perform a new backup and update the ASR diskette after each hardware, software or configuration change. This also applies to any network changes, such as change of IP address or DNS server.

IMPORTANT It is recommended to restrict access to backup media and ASR diskettes.

Advanced Recovery Tasks

This section provides explanation of the steps you will need to take if you want to perform advanced recovery tasks such as restoring Microsoft Cluster Server and Internet Information Server.

Restoring the Microsoft Cluster Server Specifics

This section provides explanation of the steps you will need to take if you want to perform disaster recovery of a Microsoft Cluster Server (MSCS). For concepts and general information please refer to the clustering section in the *HP OpenView Storage Data Protector Concepts Guide* and "Cluster Integrations with Data Protector" on page 733 in the *HP OpenView Storage Data Protector Administrator's Guide*.

Select the disaster recovery method that is appropriate for your cluster and include it in your disaster recovery plan. Consider the limitations and requirements of each disaster recovery method before making your decision. Perform tests from the test plan.

Possible Scenarios

There are two possible scenarios for disaster recovery of a MSCS:

- at least one of the nodes is up and running
- all nodes in the cluster have experienced a disaster

IMPORTANT

MSCS can be recovered using any disaster recovery method except for Disk Delivery Disaster Recovery. All specifics, limitations and requirements pertaining a particular disaster recovery method you are going to use also apply for the disaster recovery of a MSCS. For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

All prerequisites for disaster recovery (i.e. consistent and up-to-date backup, updated SRD file, all faulty hardware replaced...) must be met to recover MSCS.

Consistent backup for MSCS should include (in the same backup session):

- all nodes
- administrative virtual server (defined by the administrator)
- if Data Protector is configured as a cluster aware application, Cell Manager virtual server and IDB.

Disaster Recovery of a Secondary Node

This is the basic scenario for disaster recovery of a MSCS. The following must be true in addition to other prerequisites for disaster recovery:

- at least one of the cluster nodes is functioning properly
- the cluster service is running on that node
- all physical disk resources must be online (i.e. owned by the cluster)
- all normal cluster functionality is available (the cluster administration group is online)
- the Cell Manager is online

In this case, the disaster recovery of a cluster node is the same as the disaster recovery of a Data Protector client. You should follow the instructions for the specific disaster recovery method that you will use to restore the secondary node.

NOTE Only local disks are restored, because all shared disks are online and owned by the working node(s) during recovery and locked.

After the secondary node has been recovered, it will join the cluster after boot.

You can restore the MSCS database after all nodes have been recovered and have joined the cluster to ensure its coherency. The MSCS database is part of the `CONFIGURATION` on Windows. See “Restoring the Windows CONFIGURATION” on page 350.

Disaster Recovery of the Primary Node

In this case all nodes in the MSCS are unavailable and the cluster service is not running.

The following must be true in addition to other prerequisites for disaster recovery:

- the primary node must have write access to the quorum disk (the quorum disk should not be locked)
- the primary node must have write access to all IDB volumes, when recovering the Cell Manager
- all other nodes must be shut down until all physical disk resources are online

In this case, restore the primary node with the quorum disk first. The IDB has to be restored as well if the Cell Manager has been installed in the cluster. Optionally you can restore the MSCS database. After the primary node has been restored, you can restore all remaining nodes.

NOTE

The MSCS service uses a hard disk signature written into the MBR of every hard disk to identify physical disks. If the shared cluster disks have been replaced, this means that the disk signatures were changed during Phase 1 of disaster recovery. As a consequence, the Cluster Service will not recognize the replaced disks as valid cluster resources, and cluster groups depending on those resources will fail. See “Restoring Hard Disk Signatures On Windows” on page 577 for more information.

Perform the following steps to restore the primary node:

1. Perform disaster recovery of the primary node (including the quorum disk).
   - Assisted Manual Disaster Recovery: All user and application data on the quorum disk will be restored automatically by the `drstart` command. (`-full_clus` option)
   - EADR and OBDR: When you are asked to select the scope of recovery, select `Full with Shared Volumes` to restore quorum disk.
   - Automated System Recovery: All user and application data on the quorum disk will be automatically restored.

TIP To enable automatic restore of all shared disk volumes in the MSCS using OBDR method, move all volumes temporarily to the node for which you are preparing OBDR boot tape. It is namely impossible to collect enough information to configure the disk in Phase 1 for shared disk volumes that are locked by another node.

2. Reboot the computer.
3. Restore the cluster database. MSCS database is part of the `CONFIGURATION` on Windows. See “Restoring the Windows CONFIGURATION” on page 350.

NOTE The MSCS service must be running in order to be able to restore the MSCS database. Therefore it can not be restored automatically during Phase 2 of disaster recovery. However, the cluster database can be restored at the end of Phase 2 using the standard Data Protector restore procedure.

4. Make the IDB consistent if you are recovering a Cell Manager. See “Making IDB consistent (all methods)” on page 569.
5. The quorum and IBD volumes are restored. All other volumes are left intact and are claimed by the recovered primary node if they are not corrupted.

   If they are corrupted you have to:

   a. disable the cluster service and cluster disk driver (the steps required to do so are described in MSDN Q176970)
   b. reboot the system
   c. reestablish the previous storage structure
   d. enable the cluster disk driver and cluster service
   e. reboot the system
   f. restore user and application data
6. Restore the remaining nodes. See “Disaster Recovery of a Secondary Node” on page 573.

Merging P1S Files of all nodes for EADR

Another step is required for EADR after backup has been performed. It is impossible to collect enough information to configure the disk in Phase 1 for shared disk volumes that are locked by another node during backup. This information is necessary to enable the restore of all shared cluster volumes. To include information on shared cluster volumes in the P1S files for all nodes in the cluster, do one of the following:

- After a full client backup has been performed, merge the information on shared cluster volumes in the P1S files for all nodes in the cluster, so that the P1S file of each node contains information on the shared cluster volumes configuration.
- Move all shared cluster volumes temporarily to the node which you are going to back up. This way all required information about all shared cluster volumes can be collected, but only that node can be the primary node.

Merge

To merge the P1S files of all nodes, execute the `mmerge.cmd` command from the *<Data_Protector_home>*`\bin\drim\bin`:

`mmerge p1sA_path ... p1sX_path`

Where `p1sA` is the full path of the first node's P1S file and `p1sX` is the full path of the P1S file of the last node in the MSCS. Merged P1S files will be saved in the same directory as the source P1S files with the `.merged` appended to their filename (for example, `computer.company.com.merged`). Move the original files to another location and then rename the merged P1S files back to the original name (delete the `.merged` extension).

UNIX Cell Manager

The `mmerge.cmd` command works only on Windows systems with Data Protector Automatic Disaster Recovery component installed. If you are using a UNIX Cell Manager, copy the P1S files to a Windows client which has Automatic Disaster Recovery component installed and merge the files. Rename the merged P1S files back to the original name and copy them back to the Cell Manager.

Example

Example for merging P1S files for MSCS with 2 nodes: `mmerge` *<Data_Protector_home>*`\Config\server\dr\p1s\node1.company.com` *<Data_Protector_home>*`\Config\server\dr\p1s\node2.company.com`. Enclose the path in quotes on Windows if the path contains a space character. The merged files will be `node1.company.com.merged` and

node2.company.com.merged. Rename the files back to their original names (you will have to rename the source P1S files first): node1.company.com and node2.company.com.

Restoring Hard Disk Signatures On Windows

The MSCS service uses a hard disk signature written into the MBR of every hard disk to identify physical disks. If the shared cluster disks have been replaced, this means that the disk signatures were changed during Phase 1 of disaster recovery. As a consequence, the Cluster Service will not recognize the replaced disks as valid cluster resources, and cluster groups depending on those resources will fail. This applies only to the restore of the active node, since shared cluster resources are operational as long as at least one of the nodes is up and running and claims ownership of the resources. This problem does not apply to EADR and OBDR critical disks because the original disk signatures of all EADR/OBDR critical disks are automatically recovered. In case you have replaced any other disks, you will have to restore their hard disk signatures as well.

The most critical shared disk is the cluster quorum resource. If it has been replaced, than the original disk signature must be restored, or the cluster service will not start.

During Phase 2, the MSCS Database is restored into the \TEMP\ClusterDatabase directory on the system volume. After the system is rebooted, the cluster service will not be running, because the quorum resource will not be identified due to the changed hard disk signature in Phase 1. This can be resolved by running the clubar utility (located in the *<Data_Protector_home>*\bin\utilns), which restores the original hard disk signature. After clubar successfully finishes, the cluster service is automatically started.

Example

At the command prompt type clubar r c:\temp\ClusterDatabase force q: to restore a MSCS Database from c:\temp\ClusterDatabase.

For more information on clubar usage and syntax, see the clubar.txt file located in the *<Data_Protector_home>*\bin\utilns.

If the Data Protector shared disk on the Cell Manager is different from the quorum disk, it has to be restored as well. To restore the signature of the Data Protector shared disk and any other application disk, you should use the dumpcfg.exe utility included in the Windows 2000 Resource Kit. For details on using dumpcfg.exe, run dumpcfg /? or see

the Windows 2000 Resource Kit documentation. For more information on the problems with hard disk signatures on Windows, see MSDN article Q280425.

You can obtain the original hard disk signatures from the SRD files. The signature is a number following the `volume` keyword in the SRD file.

Example

```
-volume 5666415943 -number 0 -letter C -offslow 32256
-offshigh 0 -lenlow 320430592 -lenhigh 2 -fttype 4 -ftgroup
0 -ftmember 0
```

```
-volume 3927615943 -number 0 -letter Q -offslow 320495104
-offshigh 2 -lenlow 1339236864 -lenhigh 0 -fttype 4 -ftgroup
0 -ftmember 0
```

The number following the `-volume` keyword is the signature of the hard disk. In this case the SRD file stores information about a local hard disk (with drive letters C) and quorum disk (with drive letter Q). The signature of the quorum disk is stored only in the SRD file of the active node (at backup time), because it keeps the quorum disk locked and thus prevents other nodes from accessing the quorum disk. It is therefore recommended to always back up the whole cluster, because you need the SRD files of all nodes in the cluster, since only all SRD files together include enough information to configure the disk in Phase 1 for shared disk volumes. Note that a hard disk signature stored in the SRD file is represented as a decimal number, whereas `dumpcfg` requires hexadecimal values.

Restoring Internet Information Server (IIS) Specifics

Internet Information Server (IIS) is not supported for disaster recovery. To perform Assisted Manual Disaster Recovery of an IIS, follow these steps (in addition to the steps required for Assisted Manual disaster recovery):

1. Do not install the IIS during clean installation of the system.
2. Stop or uninstall the `IIS Admin Service`, if it is running.
3. Run the `drstart` command.
4. The IIS Database is restored as a plain file (with the filename `DisasterRecovery`) into the default IIS location (`%SystemRoot%\system32\inetsrv`).

5. After the successful boot, restore the IIS Database using the standard Data Protector restore procedure or IIS Backup/Restore snap-in. Note that this may take quite some time.

Troubleshooting

1. If any of the IIS dependant services (for example, SMTP, NNTP) do not start automatically, try to start them manually.
2. If this fails, stop the `IIS Admin Service` and restore the `%SystemRoot%\system32\inetsrv\MetaBase.bin` file, using the `overwrite` option.

NOTE

`%SystemRoot%\system32\inetsrv` is the default location of `IIS Service`. If you have installed the service into other location, use this location as a destination for restore of `MetaBase.bin` file.

3. Start the `IIS Admin Service` and all dependant services.

Editing the DRecoveryKB.cfg File

Some drivers have their functionality split into several separate files which are all required for the driver to function properly. Sometimes, it is impossible for Data Protector to identify all driver files during the creation of DR image file, if they are not listed in the `DRecoveryKB.cfg` file on a case-by-case basis. In this case, they will not be included in the disaster recovery operating system and as a consequence, some driver or service will not be operational after the boot of the DR OS.

`The DRecoveryKB.cfg` file is located in the *<Data_Protector_home>*`\bin\drim\bin` directory and stores information on the location of driver files, located under the *%SystemRoot%* directory. When you execute the test plan, make sure that all required services are running and that all drivers are operational after the boot of the OS.

If you want to back up these drivers, add information about dependant files to the `DRecoveryKB.cfg` file in the appropriate format as described in the instructions at the beginning of the `DRecoveryKB.cfg` file.

The easiest way to edit the file is to copy and paste an existing line and just replace it with the relevant information. Note that the path separator is "/" (forward slash). White space is ignored except inside quoted-pathname so the depend entry can therefore span several lines. You can also add comment lines that start with a "#" (pound) sign and extend to the end of line.

After you finished editing the file, save it to the original location. Then perform another full client backup as described in "Preparation" on page 542, to include the added files in the DR image.

Due to the numerous configurations of system hardware and applications, it makes it impossible to provide an "out of the box" solution for all possible configurations. Therefore you can modify this file to include drivers or other files at your own risk.

Any modification to this file are at your own risk and as such not supported by Hewlett-Packard.

WARNING **It is required to create and execute a test plan to be sure the recovery will work after you have edited the `DRecoveryKB.cfg` file.**

Recovery Using an Edited SRD File

Editing the SRD File

Information about backup devices or media stored in the SRD file may be out of date at the time you are performing disaster recovery. This is not a problem if you are performing an online recovery, because the required information is stored in the IDB on the Cell Manager. But if you are performing an offline recovery, the information stored in the IDB is not accessible.

For example, a disaster stroke not only the Cell Manager, but also a backup device connected to it. If you replace the backup device with a different backup device after the disaster, the information on backup devices stored in the updated SRD file (`recovery.srd`) will be wrong and the recovery will fail. In this case, edit the updated SRD file before performing Phase 2 of disaster recovery to update the wrong information and thus enable a successful recovery.

To edit the SRD file, open it in a text editor and update the information that has changed.

TIP You can display the device configuration information using the `devbra -dev` command.

For example, if the client name of the computer you are trying to recover has changed, replace the value of the `-host` option. You can also edit the information about the:

- Cell Manager client name (`-cm`).
- Media Agent client (`-mahost`).
- Logical device or drive (library) name (`-dev`).
- Device type (`-devtype`).

  Refer to the “Initial Configuration of Library Device” on page 72 for possible `-devtype` option values.
- Device SCSI address (`-devaddr`).
- Device policy (`-devpolicy`).

  Policy can be defined as `1` (Standalone), `3` (Stacker), `5` (Jukebox), `6` (external control), `8` (Grau DAS exchanger library), `9` (STK Silo medium library) or `10` (SCSI-II Library).
- Robotics SCSI address (`-devioctl`).
- Library slot (`-physloc`)
- Logical library name (`-storname`)

After you have edited the file, save it in Unicode format to the original location.

Example for Changing a MA Client

You performed a disaster recovery backup using a backup device connected to the client `old_mahost.company.com`. At the time of disaster recovery, the same backup device is connected to the client `new_mahost.company.com` with the same SCSI address. To perform a disaster recovery, replace the `-mahost old_mahost.company.com` string in the (updated) SRD file with `-mahost new_mahost.company.com`, before performing the Phase 2 of disaster recovery.

If the backup device has a different SCSI address on the new MA client, modify the value of the `-devaddr` option in the updated SRD file accordingly.

Example for Changing a Backup Device and MA Client

To perform disaster recovery using another device than the one which was used for the backup (MA client is the same), modify the following option values in the updated SRD file: `-dev`, `-devaddr`, `-devtype`, `-devpolicy`, and `-devioctl`. If you are using a library device for restore, modify also the values of the following options in the SRD file: `-physloc`, and `-stornam`e.

For example, you performed backup for disaster recovery purposes using an HP StorageWorks Ultrium standalone device with the device name `Ultrium_dagnja`, connected to the MA host `dagnja` (Windows). However, for the disaster recovery you would like to use an HP StorageWorks Ultrium robotics library with the logical library name `Autoldr_kerala` with drive `Ultrium_kerala` connected to the MA client `kerala` (Linux).

First, run the `devbra -dev` command on `kerala` to display the list of configured devices and their configuration information. You will need this information to replace the following option values in the updated SRD file:

```
-dev "Ultrium_dagnja" -devaddr Tape4:1:0:1C -devtype 13
-devpolicy 1 -mahost dagnja.company.com
```

with something like:

```
-dev "Ultrium_kerala" -devaddr /dev/nst0 -devtype 13
-devpolicy 10 -devioctl /dev/sg1 -physloc "     2 -1"
-stornam "AutoLdr_kerala" -mahost kerala.company.com.
```

The procedure on using the edited SRD file for disaster recovery is different for each disaster recovery method. Specific details are explained in the sections pertaining to disaster recovery methods.

IMPORTANT You should restrict access to the SRD files due to security reasons.

AMDR/ASR

Perform the following before proceeding with the normal AMDR/ASR recovery procedure:

1. Open the `recovery.srd` file (located on the first `drsetup` / ASR diskette) in a text editor and make the necessary changes.
2. Save the file to its original location in Unicode format.

EADR/OBDR

Perform the following additional steps before proceeding with the normal EADR/OBDR recovery procedure:

1. When the `Disaster Recovery Wizard` appears, press any key to stop the wizard during the countdown, select the `Install only` option and click `Finish`. This option will install only the temporary operating system to the target system and thus finish Phase 1 of disaster recovery. Phase 2 of disaster recovery will not start automatically if the `Install only` option is selected.

Figure 12-7 Install only option in the Disaster Recovery Wizard

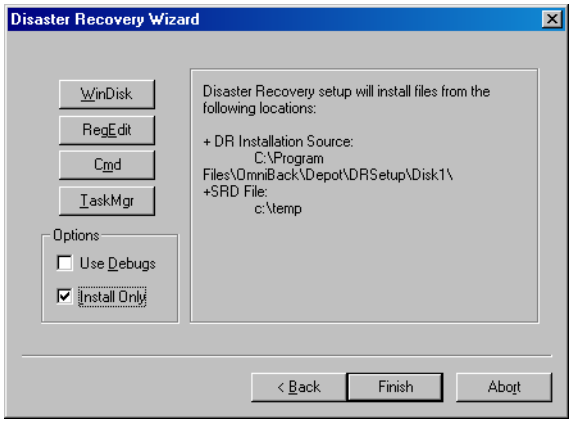

2. Run `Windows Task Manager` (press **alt+ctrl+del** and select `Task Manager`).
3. Click `File` and then `New task (Run...)`. Type `notepad c:\DRSYS\System32\OB2DR\bin\recovery.srd` and press **Enter**. The SRD file will be opened in the Notepad.
4. Edit the SRD file. For details on how to edit it, refer to "Updating and Editing the System Recovery Data (SRD)" on page 521.
5. After you have edited and saved the SRD file, run the following command from `c:\DRSYS\System32\OB2DR\bin`:

   ```
   omnidr -drimini c:\$DRIM$.OB2\OBRecovery.ini
   ```

6. Proceed with the next step in the normal EADR/OBDR recovery procedure.

Manual Disaster Recovery of an HP-UX Client

This chapter explains the procedure that should be used to recover a HP-UX client from a disaster.

The procedure is based on the Ignite-UX product; an application primary developed for HP-UX system installation and configuration tasks, which offers (in addition to a powerful interface for the system administration) preparation and recovery of the system from a disaster.

While Ignite-UX is focused on the disaster recovery of the target client, Data Protector has to be used to restore the user and application data in order to complete the *Phase 3* of disaster recovery.

This chapter cannot cover the full functionality of Ignite-UX. For detailed information please refer to the “Ignite-UX Administration Guide”.

Concept

Ignite-UX offers 2 different approaches to prepare a system for and recover a system from a disaster:

- Using custom installation medium (**Golden Image**)
- Using system recovery tools (**make_tape_recovery**, **make_net_recovery**)

While the usage of Golden Image is most suitable for IT environments with a large number of basically identical hardware configurations and OS releases, the usage of the system recovery tools supports the creation of recovery archives, which are customized for your individual systems.

Both methods allow the creation of bootable installation media like DDS-Tapes or CD’s. Using these media, the system administrator is able to perform a local disaster recovery directly from the system console of the failed client.

In addition, both methods can also be used to run a network based recovery of the client by assigning the failed client a suitable Golden Image or the previously created “recovery archive”. In this case, the client boots directly from the Ignite Server and runs the installation from the assigned depot, which has to be located on a NFS share on your network.

Use Ignite-UX GUI where it is supported.

Using Custom Installation Medium

Overview

Large IT environments often consist of a large number of systems that are based on identical hardware and software. Installation of OS, applications and required patches can be significantly reduced if a complete snapshot of the installed system is used to install other systems. Ignite-UX includes a feature, which allows you to modify parameters like networking or filesystem settings and add software like Data Protector to the image (with Ignite-UX command `make_config`) before you assign such a Golden Image to another system. This feature can thus be used to recover a system from a disaster.

Creating a "Golden Image"

Steps to Create a Golden Image

The following steps explain how to create a Golden Image of a client system on a target system, which will share the image via NFS to your network. In this example, Data Protector client is already installed on the client system and will be included in the "Golden Image" without additional configuration steps.

1. Copy the `/opt/ignite/data/scripts/make_sys_image` file from your Ignite-UX Server into a temporary directory on the client system.
2. Run the following command on the client node to create a compressed image of the client on another system: `make_sys_image -d <directory of the archive> -n <name of the archive>.gz -s <IP address of the target system>`

   This command will create a gzipped file depot in the specified directory on the system defined with the `-d` and `-s` options. Make sure that your HP-UX client has granted a passwordless access to the target system (an entry in the `.rhosts` file with the name of the client system on the target system) otherwise the command will fail.
3. Add the target directory to the `/etc/exports` directory on the target system and export the directory on the target server (`exportfs -av`).

4. On the Configuring Ignite-UX server, copy the archive template file `core.cfg` to `archive_<name>.cfg`:
   `cp /opt/ignite/data/examples/core.cfg /var/opt/ignite/data/<OS_Release>/archive_<name>.cfg`

Example

```
cp /opt/ignite/data/examples/core.cfg
/var/opt/ignite/data/Rel_B.11.11/archive_HPUX11_11_DP50_CL.cfg
```

5. Check and change the following parameters in the copied configuration file:
   - In the `sw_source` section:

```
load_order = 0
source_format = archive
source_type="NET"
# change_media=FALSE
post_load_script = "/opt/ignite/data/scripts/os_arch_post_l"
post_config_script =
"/opt/ignite/data/scripts/os_arch_post_c"
nfs_source = "<IP Target System>:<Full Path>"
```

   - In the matching `OS archive` section:

```
archive_path = "<archive_name>.gz
```

6. Determine the "impacts" entries by running the command `archive_impact` on your image file and copy the output in the same "OS archive" section of your configuration file:
   `/opt/ignite/lbin/archive_impact -t -g <archive_name>.gz`

Example

```
/opt/ignite/lbin/archive_impact -t -g
/image/archive_HPUX11_11_DP50_CL.gz


impacts = "/" 506Kb
impacts = "/.root" 32Kb
impacts = "/dev" 12Kb
impacts = "/etc" 26275Kb
impacts = "/opt" 827022Kb
```

```
impacts = "/sbin" 35124Kb
impacts = "/stand" 1116Kb
impacts = "/tcadm" 1Kb
impacts = "/usr" 729579Kb
impacts = "/var" 254639Kb
```

7. To make Ignite-UX aware of the new created depot, add an `cfg` entry to the `/var/opt/ignite/INDEX` file with the following layout:

```
cfg "<This_configuration_name>" {
description "<Description of this configuration>"
"/opt/ignite/data/<OS>/config"
"/var/opt/ignite/data/<OS>/ archive_<name>.cfg
}
```

Example

```
cfg "HPUX11_11_DP50_Client" {
description "HPUX 11.i OS incl Patches and DP50 Client"
"/opt/ignite/data/Rel_B.11.11/config"

"/var/opt/ignite/data/Rel_B.11.11/archive_HPUX11_11_DP50_CL.cfg
"
}
```

8. Make sure that one or more IP addresses reserved for booting clients are configured in the `/etc/opt/ignite/instl_boottab` file. The number of IP addresses is equal to the number of parallel booting clients.

After the above described procedure is completed, you have a Golden Image of an HP-UX client (with a specific hardware and software configuration), which can be used to recover any client of a similar layout.

Repeat these steps to create a Golden Image for all systems with different hardware and software configuration.

NOTE Ignite-UX enables you to create a bootable tape or CD based on the created Golden Image. Please refer to the Ignite-UX Administration Guide for more information.Ignite-UX enables you to create a bootable tape or CD based on the created Golden Image. Please refer to the Ignite-UX Administration Guide for more information.

Recovery

Recovery Using a Golden Image

To recover an HP-UX client by applying the Golden Image, which is located on a NFS share on your network, perform the following steps:

- On the Client System

  1. Replace the faulty hardware.
  2. Boot the HP-UX client from the Ignite-UX server: `boot lan.<IP-address Ignite-UX server>install`.
  3. Select `Install HP-UX` when the `Welcome to Ignite-UX` screen appears.
  4. Choose `Remote graphical interface running on the Ignite-UX server` from the UI Option screen.
  5. Respond to the Network configuration dialog.
  6. The system is now prepared for a remote Ignite-UX Server controlled installation.

- On the Ignite-UX Server

  7. Right click the client icon in the Ignite-UX GUI and select `Install Client - New Install`.
  8. Select the Golden Image you want to install, check the settings (network, filesystem, time zone,...) and click the `Go!` button.
  9. You can check the installation progress by right clicking the client icon and choosing `Client Status`...
  10. After the installation has finished, restore additional user and application data using the standard Data Protector restore procedure.

Using System Recovery Tools

Overview

The usage of the system recovery tools, bundled with the Ignite-UX, enables you a fast and easy recovery from a disk failure. The recovery archive of system recovery tools includes only essential HP-UX directories. However, it is possible to include other files and directories (for example, additional volume groups or the Data Protector files and directories) in the archive to speed up the recovery process.

`make_tape_recovery` creates a bootable recovery (installation) tape, customized for your system and enables you unattended disaster recovery by connecting the backup device directly to the target system and booting the target system from the bootable recovery tape. The backup device has to be locally connected to the client during the creation of the archive and recovery of the client.

`make_net_recovery` allows you to create a recovery archive over the network onto the Ignite-UX server or any other specified system. The target system can be recovered across subnets after booting ether from a bootable tape created by the Ignite-UX `make_boot_tape` command or the system boots directly from the Ignite-UX server. Booting directly from the Ignite-UX server can be automated with the Ignite-UX `bootsys` command or interactively specified on the boot console.

Creating Recovery Archives

The easiest way to create a recovery archive of an HP-UX client is to use the Ignite-UX GUI on the Ignite-UX server. All GUI commands can also be executed from the command line. Refer to the “Ignite-UX Administration Guide” for more information.

Prerequisites

Before you are able to prepare your system for disaster, the Ignite-UX fileset has to be installed on the client in order to enable the Ignite-UX server to communicate with the client.

Make sure that the revisions of the Ignite-UX fileset on the Ignite-UX server and on the client are the same. The simplest way to keep everything consistent is to install Ignite-UX from a depot build on the Ignite-UX server. This depot can be constructed by running the following command on the Ignite-UX server:

```
pkg_rec_depot -f
```

This creates an Ignite-UX depot under `/var/opt/ignite/depots/recovery_cmds`, which can be specified as a source directory by `swinstall` on the client for the Ignite-UX software installation.

After you have installed Ignite-UX on the client node, you can use the GUI on the Ignite-UX server to create recovery archives using `make_net_recovery` or `make_tape_recovery`.

Creating an Archive Using make_tape_recovery

Perform the following steps to create an archive using `make_tape_recovery`:

1. Make sure that a backup device is connected to the HP-UX client.
2. Start the Ignite-UX GUI by executing the following command: `/opt/ignite/bin/ignite &`.
3. Right click the client icon and select `Create Tape Recovery Archive`.
4. Select a tape device, if more than one device is connected to the HP-UX client.
5. Select the volume groups you want to include into the archive.
6. The tape creation process will now begin. Check the status and log file on the Ignite-UX server by right clicking the client icon and selecting `Client Status`.

NOTE

Ignite-UX recommends the usage of 90m DDS1 backup tapes to ensure that the tape will work with all DDS with any DDS drive.

Creating an Archive Using make_net_recovery

The procedure for creating a recovery archive using `make_net_recovery` is almost the same as using `make_tape_recovery`. The advantage is that there is no need for a locally attached backup device, as the recovery archive is stored on the Ignite-UX server by default.

1. Start the Ignite-UX GUI by executing the following command: `/opt/ignite/bin/ignite &`
2. Right click the client icon and select `Create Network Recovery Archive`.

3. Select the destination system and directory. Make sure that there is enough space to store the compressed archive.
4. Select the volume groups you want to include into the archive.
5. The archive creation process will now begin. Check the status and log file on the Ignite-UX server by right clicking the icon and selecting `Client Status`.

NOTE

Ignite-UX allows you to create bootable archive tape out of the compressed archive file. See the chapter *Create a Bootable Archive Tape via the Network* in the *Ignite-UX Administration Guide*.

Recovery

Recovery From the Backup Tape

To recover a system from a disaster using the bootable tape created by `make_tape_recovery` follow the steps below:

1. Replace the faulty hardware.
2. Make sure that the tape device is locally connected to the crashed HP-UX client and insert the medium with the archive you want to restore.
3. Boot from the prepared recovery tape. To do so, type in `SEARCH` at the boot admin menu to get a list of all available boot devices. Determine which one is the tape drive and type in the boot command: `boot <hardware path>` or `boot P<number>`.
4. The recovery process starts automatically.
5. After the recovery has completed successfully, restore additional user and application data using the standard Data Protector restore procedure.

Recovery From the Network

To recover an HP-UX client from a disaster via the network, follow the instructions on how to perform recovery with a Golden Image. Make sure you have selected the desired archive for the installation.

- On the Client
  1. Replace the faulty hardware.
  2. Boot the HP-UX client from the Ignite-UX server: `boot`

lan.<*IP-address Ignite-UX server*> install

3. Select `Install HP-UX` from the `Welcome to Ignite-UX` screen.
4. Choose `Remote graphical interface running on the Ignite-UX server` on the UI Option screen.
5. Respond to the Network configuration dialog.
6. The system is now prepared for a remote installation controlled from the Ignite-UX Server.

- On the Ignite-UX Server

7. Right click the client icon within the Ignite-UX GUI and select `Install Client - New Install`.
8. Under `Configurations:` select the `Recovery Archive` you want to install, check the settings (network, filesystem, time zone,...) and click the `Go!` button.
9. You can check the installation progress by right clicking the client icon and choosing `Client Status...`
10. After the recovery has completed successfully, restore additional user and application data using the standard Data Protector restore procedure.

Disk Delivery Disaster Recovery of a UNIX Client

To perform a Disk Delivery Disaster Recovery of a UNIX client, connect a bootable disk that contains a minimal OS installation and Data Protector Disk Agent to the crashed system. The administrator has to ensure (before the disaster) that enough data has been collected to correctly format and partition the disk.

For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

Limitations

- This description does not cover the recovery of a cluster environment. Depending on the configuration of the cluster environment, additional steps and modification to the environment are necessary.
- RAID is not supported.
- Auxiliary disk should be prepared on a system of the same hardware class as the target system.

Preparation

Preparation for this disaster recovery method should be performed on several levels: gathering the information for your backup specification, preparing the disk, preparing your backup specification (pre-exec), and executing the backup. All of these preparatory steps are necessary before executing disaster recovery of the client.

Gathering Information

This section provides a list of items that need to be executed for each target system at backup time, in order to perform successful disaster recovery. If the information is collected as part of a pre-exec command, it is important to document the location of these files in the Disaster Recovery plan so that the information can be found once disaster strikes. Also version administration (there is a collection of the “auxiliary information” per backup) has to be considered.

- If the system that will be backed up has application processes active at low run levels, establish a state of *minimal activity* (modified *init 1 run level*) and enter the single user mode to prevent errors after recovery (see "Consistent and Relevant Backup" on page 520). Consult your operating system documentation for details.

HP-UX Example

1. Move some kill links from `/sbin/rc1.d` to `/sbin/rc0.d` and complement the changes for the boot-up section. The kill links include the basic services that would otherwise be suspended by moving to run level 1, and they are needed for the backup. For an example, see "Disaster Recovery: Move Kill Links on HP-UX 11.x" on page A-26.
2. Ensure that `rpcd` is configured on the system (configure the variable `RPCD=1` within the file `/etc/rc.config.d/dce`).

   This prepares the system so that it enters the state of minimal activity. The state can be characterized as follows:

   — `Init-1 (FS_mounted, hostname_set, date_set, syncer_running)`

   — Network must be running

   — The following processes should also be running: `inetd`, `rpcd`, `swagentd`

Solaris Example

1. Move the `rpc` kill link from `/etc/rc1.d` to `/etc/rc0.d` and complement the change for the boot-up section. The kill links include the basic services that would otherwise be suspended by moving to run level 1, and they are needed for the backup.
2. Ensure that `rpcbind` is configured on the system.

   This prepares the system so that it enters the state of minimal activity. The state can be characterized as follows:

   — `Init 1`

   — Network must be running

   — The following processes should also be running: `inetd`, `rpcbind`.

Tru64

1. If the system is powered down, boot up the system and enter the System Reference Manual (SRM) console (the firmware console).

2. Execute the following command from the SRM console to enter the single user mode:
   - `boot -fl s` to boot using already generated vmunix file
   - `boot -fi genvmunix -fl s` to boot into the single user mode with the generic kernel.
3. If the system is already powered up and running, change from the current run level to single-user mode by executing the following command: `init s`

AIX

- No action is required, because the `alt_disk_install` command, used to prepare the auxiliary disk, ensures consistent disk image without entering the state of minimal system activity.

Creating an Auxiliary Disk

- If you want to work with the auxiliary boot disk, you have to prepare it. Only one bootable auxiliary disk is required per site and platform. This disk has to contain the operating system and network configuration, and has to be bootable.

Preparing the Backup Specification

- Provide a Pre-exec script that performs the following:
  - Collects all the necessary information about the environment and puts it in an available location in case of a disaster recovery. It is suggested to put it onto a different system which can be accessed easily. The information should cover:
    - ✓ Physical and logical storage structure of the storage
    - ✓ Current logical volume structure (for example, on HP-UX, using `vgcfgbackup` and `vgdisplay -v`)
    - ✓ ServiceGuard configuration data, disk-mirroring, striping
    - ✓ Filesystems and mountpoints overview (`for example, on HP-UX, using bdf` or copy of `/etc/fstab`)
    - ✓ System paging space information, for example, on HP-UX, using the output of the `swapinfo` command
    - ✓ I/O-structure overview (`for example, on HP-UX, using ioscan -fun` and `ioscan -fkn`)
    - ✓ Client network settings

  - — An emergency copy of the data can also be put into the backup itself. If done so, the information has to then be extracted prior to the actual recovery.
  - — Consider logging out all users from the system.
  - — Shut down all applications, unless the application data gets backed up separately, for example, using online database backup.
  - — You may want to restrict network access to the system, so that no one can log on to the system while the backup is running (for example, on HP-UX, overwrite `inetd.sec` and use `inetd -c`).
  - — If needed, enter the state of minimal system activity (for example, on HP-UX, use `sbin/init 1`; wait 60; check if run_level 1 is reached). Note that this is a modified “init 1" state.
- Provide a post-exec script that elevates the system to the standard run-level, restarts applications, and so on.
- Setup a backup specification for the client on the Data Protector Cell Manager. It should include all the discs (with disc discovery) and include the pre- and post-exec scripts.

Testing the Procedure

- Execute this backup procedure and repeat it on a regular basis, or at least at every major system configuration change, especially any change in the logical volume structure (for example, using LVM on HP-UX).

Recovery

This section describes how to restore a system to the state when the backup was done. You will need the following to successfully perform a Disk Delivery Disaster Recovery:

- A new hard disk to replace your crashed disk.
- An auxiliary disk containing the relevant operating system and the Data Protector agents.
- A successful full backup of the client that you want to recover.

The following steps need to be performed:

1. Replace the faulty disk with a new disk of comparable size.

2. Attach the auxiliary disk (which contains the relevant operating system and the Data Protector client) to the system and make it the boot device.
3. Boot from the auxiliary operating system.
4. Reconstruct the logical volume structure if applicable (for example, using LVM on HP-UX). Use the saved data for the non-root volume groups (for example, with `vgcfgrestore` or SAM on HP-UX).
5. Additionally, the root volume group to be restored has to be created on the repaired disk (for example, using `vgimport` on HP-UX). It will not look like a root volume group during the restore process. This is because the OS from the auxiliary disk will be running. For more information on `vgimport`, see its man page.
6. Make the new disk bootable.
7. Reconstruct any other storage structures like mirror, striping, service guard, and so on from the data saved on a secondary storage device during backup.
8. Create the filesystems and mount them as required by the data from the backup; use similar but not the original mountpoint names (like `/etc_restore` for `/etc`, and so on).
9. Remove any files in the mountpoints to be restored, they must be clean.
10. Start the Data Protector GUI and open a connection to the Cell Manager. Import the system with the auxiliary disk into the cell.
11. Select the version from which you want to restore. First list all the required media for the restore and make sure they are available. Restore all the required mountpoints including the (future) root-volume to the system, using the option `Restore As <new_mountpoint>`. The root-volume from the backup is restored to the root-volume on the repaired disk. Nothing is restored to the currently-running auxiliary operating system on the auxiliary disk.
12. Shut down the system that was just restored.
13. Disconnect the auxiliary disk from the system.
14. Reboot the system from the new (or repaired) disk.

NOTE Instead of using an auxiliary disk, the new disk can also be temporarily connected to a client that has to have a Disk Agent installed. After being restored, it can be connected to the faulty system and booted.

Manual Disaster Recovery of a UNIX Cell Manager

Manual Disaster Recovery is a basic method, that involves recovering the system by reinstalling it in the same way as it was initially installed. In addition, Data Protector is used to then restore all files, including the operating system.

Limitation

For details on supported operating systems, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

This description does not cover the recovery of a cluster environment. Depending on the configuration of the cluster environment, additional steps and modification to the environment are necessary.

Preparation

Perform the same preparatory steps without the steps pertaining to the auxiliary disk, as for Disk Delivery Disaster Recovery of an HP-UX or Solaris client. See “Preparation” on page 594 for reference. In addition to completing those steps, you also have to complete the following:

1. The IDB has to be backed up regularly, ideally in a separate backup specification, scheduled after the backup of the Cell Manager.
2. The IDB and configuration backup must run to a specific device located on the Cell Manager system, to make the administrator aware that the medium in the device contains the most recent version of the IDB.

Recovery

Use the following method to recover your UNIX Cell Manager.

Prerequisites

You will need the following to successfully perform a disaster recovery:

- Media containing the last valid known backup of the root partition of the Cell Manager and IDB.

- A device connected to the Cell Manager.

Procedure

The following steps need to be performed to recover a Cell Manager:

1. Replace the crashed disk.
2. Boot your system from the installation media of your operating system.
3. Reinstall the operating system. Refer to your system administrator's manual for instructions. During the installation, using the data gathered during the preparation phase (pre-exec script), re-create and configure the physical and logical storage structure of the storage, logical volume structure, filesystem and mountpoints, network settings and other.
4. Reinstall the Data Protector on the Cell Manager.
5. Restore the latest backup of your database and `/etc/opt/omni` to a temporary directory. This simplifies the restore of all other files from media. Note that you cannot restore the database directly. See Chapter 6, "Restoring Data," for instructions. This includes stopping all Data Protector processes with the `/opt/omni/sbin/omnisv -stop` command. This ensures that no files will be in use.
6. Remove the `/etc/opt/omni` directory and replace it with the `/etc/opt/omni` directory from the temporary area. This re-creates the previous configuration.
7. Start Data Protector processes with the `/opt/omni/sbin/omnisv -start` command.
8. Start the Data Protector user interface and restore all the files used from your backup.
9. Reboot the system.

Your Cell Manager should now be successfully recovered.

Troubleshooting Disaster Recovery on Windows

This section provides explanation of the steps you will need to take if you happen to encounter problems with disaster recovery procedures on Windows systems.

General Troubleshooting

Autodr.log

`Autodr.log` is a log file located in the *<Data_Protector_home>*`\tmp` directory and contains messages relevant to the automatic disaster recovery methods (EADR, ODBR, ASR). You should inspect it if an error has occurred. `Autodr.log` logs many different messages, mostly for development and support purposes. Only some of them are relevant to you and indicate that an error has occurred. These error messages are usually logged at the end of the log file with a traceback appended.

There are four types/levels of messages in the `autodr.log` (note that they do not correspond to the same report levels for messages that are reported at the end of a backup session in the Data Protector GUI):

- `Critical error`: the error is so serious that the backup of the object can not continue and will be aborted.
- `Error`: the error may be critical, but it depends on different factors.

  For example, `autodr.log` reports an error that some driver has not been included in the disaster recovery operating system. The missing driver may be the reason for the recovered system not to be operational after the recovery or only for some non-critical service not to be running after the boot of the operating system. It depends on which driver has not been backed up.
- `Warning` and `Info`: These are not error messages and usually do not mean that anything is wrong.

Some of the most common messages stated in the `autodr.log` file are:

- `filename 'not a pe file'`: This means that a certain file is not Portable-Executable. This warning does not indicate that anything is wrong.

- `unsupported location`: Data Protector notices that a certain file that is required by a service or a driver that will be included in the disaster recovery operating system (DR OS), is not located under the *`%SystemRoot%`* directory.

  Such drivers are often used by the antivirus and remote control software (for example pcAnywhere). This message is important, because it can mean that the service/driver that requires the missing file, will not be operational after the boot. It depends on which service or driver was affected, if the disaster recovery will fail or succeed. A possible solution for this problem is copying the missing file into the *`%SystemRoot%`* directory and changing its path in the Windows Registry. Note that incorrect editing of the Windows Registry may severely damage your system.

Problems Logging on to the System After Disaster Recovery Finishes

Problem

You may receive the following error message after the system is recovered:

```
The system cannot log you on to this domain, because the
system's computer account in its primary domain is
missing or the password on that account is incorrect.
```

This type of message is usually caused by one of the following reasons:

- After collecting all information for successful disaster recovery (including full backup), you reinstalled Windows and (re)inserted into the offending domain.
- After collecting all information for successful disaster recovery (including full backup), you removed your system from the offending domain and later (re)inserted it into the same or some other domain.

Action

In cases like this, Windows generates new system security information, which is incompatible with information that is restored during disaster recovery. The solution is the following:

1. Log on to the system locally with an Administrator account.

2. In the `Control Panel`, click `Network` and, using the `Identification` tab, remove the system from its current domain to a temporary workgroup (for example, TEMP). After this is done, reinsert the system into the domain from which it was previously removed. You need a domain administrator's password.
3. After the computer is again in the proper domain, click `OK` in the `Network` window. Windows will force you to reboot the system.
4. To update this new state with disaster recovery, you should perform all necessary procedures (collecting system data, backup) once more, as described in the "Preparing for a Disaster Recovery" section.

Disaster Recovery from a Copy

Problem

You cannot perform a disaster recovery from a media copy or an object copy.

Data Protector by default uses the original media set to perform a disaster recovery. Thus, copy object versions are not displayed in the Disaster Recovery Wizard of the Data Protector GUI.

Action

To perform a disaster recovery from a media copy or an object copy, if your original media set is not available or is damaged, proceed as follows:

- Object copy: Export all media in the original media set from the IDB and then regenerate the SRD file. Data Protector then offers you the first available copy of the original media set in the Disaster Recovery Wizard.

  Refer to "Exporting Media from Data Protector" on page 171 and "Updating and Editing the System Recovery Data (SRD)" on page 521.

- Media copy: In the SRD file, replace the media IDs of the original media with the media IDs of the media copies. Data Protector then offers you the first available copy of the original media set in the Disaster Recovery Wizard.

  Refer to "Updating and Editing the System Recovery Data (SRD)" on page 521.

Troubleshooting Assisted Manual Disaster Recovery

Problem

Drstart reports: "Can not copy <filename>"

This error is reported because the `drstart` utility can not copy the specified file. One of the reasons may be that the file is locked by the system. For example, if `drstart` cannot copy `omniinet.exe`, it might be because the `Inet` service is already running. This is not a normal scenario and should not happen after a clean install.

Action

A dialog box will appear asking you whether you would like to proceed with copying the rest of the files. If you click `Yes`, `drstart` will skip the locked file and continue copying other files. This will solve the problem if the file is locked by the system, as the process required for the disaster recovery is already running and therefore the file does not need to be copied.

You can also close the `drstart` utility by clicking the `Abort` button.

Troubleshooting Disk Delivery Disaster Recovery

Problem

"Cannot Find Physical Location of Drives Selected for Disk Delivery"

When using the Disk Delivery method for disaster recovery, it is possible that you will receive the following error: "Cannot find physical location of drives selected for disk delivery." Objects will be restored when creating a partition on the new disk if you select a drive letter that has not been used before. The better solution would be:

Action

Disaster recovery checks disk information before restoring objects. An internal function reads the Registry value Information, which is created by the Disk Administrator. If the Disk Administrator is started several times, the `Information` value becomes corrupted (format is changed during update) - the parsers fail in such cases. If you delete the `HKEY_LOCAL_MACHINE\SYSTEM\DISK Information` key and restart the Disk Administrator, the function will succeed.

Problem **"No Operating System Found"**

Action After performing disaster recovery, if the final boot of a Windows system fails with "`No Operating System Found`", check the `boot.ini` file for information about where the partition information is located. See Step 4 in the section "Recovery" on page 537 for additional information.

Problem **Disk Delivery Disaster Recovery of a Media Agent Client**

If you are performing a Disk Delivery disaster recovery, Data Protector first tries to connect to the original client where the backup device was attached (the Media Agent client) in order to use the same device for restore. However, when you are performing Disk Delivery disaster recovery of the Media Agent client where the backup has been made, Data Protector will not be able to connect to it and will proceed with offline restore and search for a local device for the restore. If there is no local device attached, Data Protector will issue a notification that there is no local device attached and will abort the disaster recovery.

Action There are two methods to avoid this:

- Move the media to another pool. This way you assign the media to the new device. Then proceed with Disk Delivery disaster recovery.
- The third method involves preparation prior to the disaster. If you have two Media Agent clients in the cell, you can back up of the first Media Agent client to another and vice versa before the disaster to avoid problems when performing Disk Delivery disaster recovery of a Media Agent client.

Troubleshooting Enhanced Automated Disaster Recovery and One Button Disaster Recovery

Problem **Automatic DR information could not be collected**

When using EADR or OBDR, it is possible that you will receive the following error: "`Automatic DR information could not be collected. Aborting the collecting of system recovery data`"

Action

- Check if all storage devices are configured correctly. If Device Manager reports a device as "`Unknown Device`", install the proper device drivers before you can perform EADR/OBDR. A similar entry would appear in `autodr.log` (located in

<Data_Protector_home>\tmp) if improperly configured storage devices are attached to your system:
DRIM_WIN_ERROR 13 SetupDiGetDeviceRegistryProperty

- There must be enough registry space available. It is recommended to set the maximum registry size to at least twice that of the current registry size. If there is not enough registry space available, a similar entry would appear in the autodr.log:
  ERROR registry 'Exception while saving registry'
  ...
  WindowsError: [Errno 1450] Insufficient system resources exist to complete the requested service.

If the problem persists, uninstall the Data Protector Automatic Disaster Recovery component (so that at least Manual Disaster Recovery and Disk Delivery Disaster Recovery will work) and contact technical support.

Problem **Some Non-critical Errors Were Detected**

When using EADR or OBDR, it is possible that you will receive the following error: "Some non-critical errors were detected during the collecting of Automatic DR data. Please review the Automatic DR log file."

A non-critical error detected during the execution of the Automatic Disaster Recovery module, means that such backup can most likely still be used for disaster-recovery purposes. Possible reasons for non-critical errors are stored in autodr.log (located in *<Data_Protector_home>*\tmp):

Action

- Services or drivers outside of the <%SystemRoot%> folder (for example, virus scanners). Autodr.log would contain a similar error message:

  ERROR safeboot 'unsupported location' 'intercheck support 06' 2 u'\\??\\D:\\Program Files\\Sophos SWEEP for NT\\icntst06.sys'.

  You can ignore this error message, as it does not affect the success of disaster recovery.

Problem **Blank Screen During Recovery**

Certain system configurations have been encountered where the video display does not work, if Windows is started in safe mode. This error is not related to Data Protector and can occur even if only Windows is installed.

Action If the screen is blank during disaster recovery, this does not mean that the recovery has failed. You can monitor the progress of disaster recovery on the Cell Manager or use `ping` and `telnet 5555` (or appropriate) commands from another client to see if the target system responds. Other indicators that the recovery is still in progress are that the device is working and that hard disk lights are blinking.

If the target system responds to `ping` and `telnet 5555` commands, but hard disk lights are not blinking and the device is not active, it is possible that the auto logon failed. Press **Enter** to log on using the administrator's account with a blank password.

The display on the restored system will then work just as it did at backup time.

Problem **Network is Not Available During Restore**

Action Ensure that the problem is not with switch, cables, etc. Another possibility is also that the DNS server (as configured at backup time) is offline during the restore. Since the configuration of the DR OS is the same as at backup time, the network will not be available. In this case perform offline restore and change the DNS settings after recovery. You can also edit the registry (`HKey_Local_Machine\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters`) before Phase 2 is started. In this case reboot before Phase 2 for the changes to take effect. After Phase 2 finishes, you can correct the settings before Phase 3 can be started.

CAUTION Editing the registry incorrectly can result in failed disaster recovery.

Problem **Auto Logon Does Not Work**

Action Sometimes auto logon does not work and you have to manually log on using an administrator's account with a blank password.

Problem **Computer Stops Responding**

Action Check if the CD/tape is readable. Do not reuse CD-RWs/tapes too many times.

Problem **Cannot Create a CD ISO Image for EADR of MSCS**

Action The quorum disk has to be backed up in order to be able to create an CD ISO Image.

Troubleshooting Automated System Recovery

Network Problems During ASR

Problem Network problems can be the cause of different problems during ASR.

For example, the target system has two network adapters installed and one of them had been disabled when the disaster recovery backup was performed. During ASR, all devices are enabled by default. If both network adapters are active on the target system during ASR, the network may not be configured properly, resulting in problems connecting to the Cell Manager and Media Agent client. In this case, Data Protector will switch to offline or local recovery, display a connection error or ASR will fail.

Action To resolve the error, follow the normal ASR recovery procedure and press **F8** when the following text is displayed in the Disaster Recovery wizard:

```
Press F8 in the next 5 seconds to skip network
configuration...
```

This will revert from Data Protector ASR network configuration to the standard Microsoft ASR network configuration.

13 Customizing the Data Protector Environment

In This Chapter

This chapter describes how you can customize Data Protector to better suit your needs. The chapter consists of the following sections:

IMPORTANT

For specific information on Data Protector limitations and recommendations, see the *HP OpenView Storage Data Protector Software Release Notes*. For details about adding security to your Data Protector cell, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Global Options File

Global options affect the entire Data Protector cell, and cover various aspects of Data Protector, such as time-outs and limits. All global options are described in the global options file, which you can edit in order to customize Data Protector. It is located in the `/etc/opt/omni/server/options` directory on the UNIX Cell Manager and in the `<Data_Protector_home>\Config\server\options` directory on Windows Cell Manager. The file is named `global`.

How to Use Global Options

Each option has a hash mark, or pound sign (#), which comments out the option and provides an explanation of the option in the text following the hash mark. For options not described in this guide, refer to the file itself.

To use a global option, uncomment the line that has the option name and set an appropriate value. To uncomment a line, simply remove the '#' mark.

NOTE

Most users should be able to operate Data Protector without changing the global options.

Most Often Used Variables

The following list includes the most often used global variables. See the Global Options file for a complete description.

- `MediaView`: Changes the fields and their order in the Media Management context.
- `MaxBSessions`: Changes the default limit of five concurrent backups.
- `InitOnLoosePolicy`: Allows Data Protector to automatically initialize blank or unknown tapes under a `loose` media policy.
- `MaxMAperSM`: Increases the default limit of concurrent devices per backup session. (Maximum device concurrency is 32.)
- `DCDirAllocation`: Determines the algorithm used for selecting into which `dcbf` directory a new detail catalog binary file goes. Three algorithms are available: fill in sequence (default), balance size, and

balance number.

- `DailyMaintenanceTime`: Determines the time after which the daily maintenance tasks can begin, using the twenty-four hour clock notation. By default, this time is set to 12:00 (Noon). For a list of daily maintenance tasks, refer to “Data Protector Checking and Maintenance Mechanism” on page 725.
- `DailyCheckTime`: Determines the time after which daily check can begin, using the twenty-four hour clock notation. By default, this time is set to 12:30 P.M. If you do not wish to perform a daily check, you can disable it. For a list of daily check tasks, refer to “Data Protector Checking and Maintenance Mechanism” on page 725.

Using Omnirc Options

The omnirc options (variables) are most useful for troubleshooting or overriding other settings affecting the behavior of the Data Protector client only. However, even advanced users should not use them unless their operating environment demands it. The Disk Agents and Media Agents use the values of these options.

These options are found in the following locations:

Locations

- /opt/omni/.omnirc on HP-UX and Solaris clients
- /usr/omni/.omnirc on other UNIX clients
- *<Data_Protector_home>*\omnirc on Windows clients
- sys:\usr\omni\omnirc on Novell NetWare clients

How to Use Omnirc Options?

Installation provides a template for the omnirc file (.omnirc.TMPL or omnirc.TMPL, depending on the platform). This file is not active. To create an active omnirc file, copy the template file to omnirc (or .omnirc) and edit it. To use a specific option, uncomment the line (remove the '#' character) and edit the value if necessary.

- When creating the omnirc file (either by copying or by using an editor), verify its permissions. On UNIX, permissions will be set according to your umask settings and may be such that some processes may be unable to read the file.

  Set the permissions to 644 manually.

- When changing the omnirc file, you have to restart the Data Protector services/daemons on the Data Protector client where you modified the omnirc file. This is mandatory for the crs daemon on UNIX and recommended for Data Protector CRS and Inet services on Windows. Specifically on Windows, restarting is not required when adding or changing entries, only when removing entries (or renaming the file).

Setting Omnirc Options During Disaster Recovery on Windows

If you need to set an omnirc option during the disaster recovery on Windows (except for Disk Delivery Disaster Recovery), perform the following steps:

1. When the `Disaster Recovery Wizard` appears, press any key to stop the wizard during the countdown.

Figure 13-1 Disaster Recovery Wizard

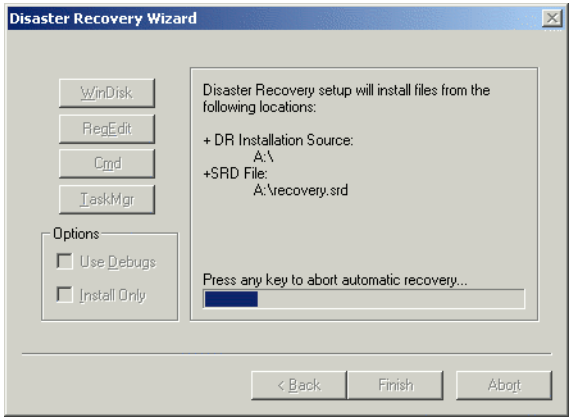

2. Click `cmd` to start the command prompt.
3. Run the following command:

   ```
   echo <Variable> > %systemroot%\system32\OB2DR\omnirc
   ```

   where *<Variable>* is the omnirc option exactly as it should be written in the `omnirc` file.

   For example:

   ```
   echo OB2RECONNECT_RETRY=1000 >
   %systemroot%\system32\OB2DR\omnirc
   ```

   This command creates an `omnirc` file in the disaster recovery operating system with the `OB2RECONNECT_RETRY` variable set to `1000` seconds.
4. Close the command prompt and click `Next` in the `Disaster Recovery Wizard` to proceed with disaster recovery.

NOTE

When using special characters in variable names in the `omnirc` file, take into account operating system specific limitations regarding supported characters for setting environment variables. For example, on UNIX systems, variables cannot contain any of the following characters: `Space Tab / : * " < > |`.

Most Often Used Variables

The most often used omnirc variables include:

- **OB2BLKPADDING_n:** This is a set of variables that can be used to specify the number of empty blocks written to the media at the initialization time.
- **OB2DEVSLEEP:** Changes the sleep time between each retry while loading a device.
- **OB2ENCODE:** Allows a user to always turn on data encoding, regardless how the backup options are set in the backup specification.
- **OB2OEXECOFF:** Allows a user to restrict or disable any object `pre-` and `post-exec` scripts defined in backup specifications for a specific client.
- **OB2INCRDIFFTIME** and **OB2CHECKCHANGETIME:**

  The **OB2CHECKCHANGETIME** and the **OB2INCRDIFFTIME** variables are relevant only for UNIX clients. During an incremental backup, the Disk Agent checks the "last modification" and "last inode change" time attributes to detect which files changed since the last backup. However, the "last inode change" attribute can only be respected when the Data Protector `Do not preserve access time attributes` backup option is enabled and the Data Protector `Lock files during backup` backup option is disabled. This variable allows more control over when to use the "last inode change" time for incremental backups.

  The **OB2INCRDIFFTIME** variable specifies an "incremental latency" period (in minutes) that is enforced when checking the "last inode change" time with incremental backups. The referential time received from the Session Manager (time of the previous backup) is first incremented by the specified period and then compared to the "last inode change" to qualify for backup. This variable takes effect only when the **OB2CHECKCHANGETIME** variable is set to 2.

- **OB2RECONNECT_ACK:** Defines how long Data Protector should wait for the message of acknowledgment (default 1200 seconds). In other words, if the agent does not get an acknowledgment in `OB2RECONNECT_ACK` seconds, it will assume that the socket connection is no longer valid.
- **OB2RECONNECT_RETRY:** Defines how long Data Protector should wait before trying to reconnect after a connection failure either between a

  — Disk Agent and Media Agent (during backup) or

  — Backup Session Manager and Disk Agents or Media Agents

  The default value is 600 seconds. In other words, the LAN/WAN line between the Backup Session Manager and DA/MA or between Disk Agent and Media Agent during backup cannot be down more than `OB2RECONNECT_RETRY` seconds.
- **OB2REXECOFF:** Allows a user to disable any remote session `pre-` and `post-exec` scripts for a specific client.
- **OB2SHMEM_IPCGLOBAL:** This option should be set to 1 on HP-UX clients that have both the Disk Agent and a Media Agent installed, in case the following error occurs during the backup:

  ```
  Cannot allocate/attach shared memory (IPC Cannot Allocate
  Shared Memory Segment)

  System error: [13] Permission denied) => aborting
  ```
- **OB2VXDIRECT:** Enables direct (without cache) reading for Advanced VxFS filesystems, as well as improving performance.
- **OB2PORTRANGE:** This option limits the range of port numbers that Data Protector uses when allocating listen ports dynamically. This option is typically set to enable the administration of a cell through a firewall. Note that the firewall needs to be configured separately and that the specified range does not affect the `Inet` listen port.

Example

```
OB2PORTRANGE=40000-40199
```

This sets the port range to ports from 40000 to 40199.

- **OB2PORTRANGESPEC:** This option allows you to specify a range of port numbers for every binary. This mechanism gives you more control over the ranges and helps to keep their sizes smaller. Note that the firewall needs to be configured separately and that the specified range does not affect the `Inet` listen port.

  For configuration examples, refer to “Firewall Support” on page 627.

Selecting the Language for the Data Protector GUI

Data Protector is localized in multiple languages. In order to correctly display international characters for localized language catalogs (Data Protector GUI and messages), a few prerequisites have to be met.

The language you have selected for your Data Protector GUI also influences the display of international characters in session messages and file names in the Data Protector GUI.

The available choice of localized language catalogs depends on which languages were selected during the installation of Data Protector. If the localized language catalogs are not installed, you can install additional Data Protector language catalogs on your existing GUI client. Refer to the online Help index keyword “components of Data Protector, adding”.

To determine which languages are installed for Data Protector, perform the following steps:“

1. In the Context List, click `Clients`.
2. In the Scoping Pane, expand `Clients` to display the list of the installed clients.
3. Click the GUI client and check the installed components in the Results Area.

The English localized catalog is always installed and is the default language for Data Protector GUI.

Prerequisites

The following prerequisites must be met:

- The desired language support must be installed on the Data Protector GUI client.
- On Windows systems, the localized Data Protector GUI must be run on the appropriate localized Windows operating system. For example, Japanese language support for Data Protector GUI works correctly only on Japanese localized Windows.
- On UNIX systems, the desktop environment must be started in the locale, which is used to encode your file names. For example, in the SJIS environment, start the desktop in the SJIS locale.

If the above prerequisites are met, the Data Protector GUI will be started in the language corresponding to the locale set on your operating system. For example, on Windows, if the French locale is set in the `Regional options` in the Windows Control Panel, the French language catalog will be used for the Data Protector GUI if available.

Settings for the File Name Encoding in GUI

In the Data Protector GUI, a specific setup and configuration is required to ensure the correct display of international characters in file names and session messages. For background information, refer to *HP OpenView Storage Data Protector Concepts Guide*.

Prerequisites

The following prerequisites apply for the correct display of international characters on the GUI system:

- In case of upgrade to Data Protector A.05.50, the file names in the IDB have been converted. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.
- The appropriate locales (UNIX) and fonts for the selected character set are installed on the Data Protector GUI system. For example, to see Japanese characters in the Windows GUI running on an European system, install Japanese fonts.
- On UNIX systems with Data Protector GUI, set the appropriate locale *before* starting the GUI. Refer to “Correct Display of International Characters in the Data Protector GUI on UNIX” on page 623.

To correctly display international characters in GUI running on Windows and UNIX systems, select the same character encoding that was in effect on the system, where the files were created.

Limitations

The following limitations apply even though the correct character encoding is selected:

- Some characters may be displayed incorrectly even though the correct encoding is selected from the GUI. The reason for this are some minor differences between implementations of code pages on Windows and UNIX operating systems. Some characters could therefore not be mapped correctly, if Data Protector GUI is run on a different platform as the client being configured. In the worst case, only a few characters could be displayed incorrectly, but this will not affect your backups or restore.
- Unfortunately, names of GUI items in the Data Protector GUI (such as backup devices and backup specifications) on UNIX, created under a certain locale, may appear corrupted when viewed in a different

locale. Regardless of the corrupted display of the GUI items names, such GUI objects are still usable. For example, you configured a backup device and named it using non-ASCII characters. In this case, its name may appear corrupted, if GUI is run in a locale that uses only ASCII. Although its name appears corrupted in the GUI, you can still perform backups and restores using this device.

Correct Display of International Characters in the Data Protector GUI on UNIX

On UNIX Data Protector GUI systems, the appropriate locale (considering your cell configuration) must be set prior to starting Data Protector GUI, to enable encoding switching in the GUI and thus proper display of international characters.

Setting the locale influences the display and encoding of files on your system, as well as which of the installed language catalogs is used for the Data Protector GUI. The specific file name related settings are used, for example, in the file system browser when selecting files for backup, in the IDB browser when selecting files to restore, in session messages, and so on.

There are two principal alternatives for the selection of encoding:

- Use a UTF-8 based locale. This allows you to dynamically switch the encoding for file names without restarting the Data Protector GUI.
- Use any locale of your choice. This will show you file names correctly using any known locale.

It is recommended to set the locale as follows:

- If you have uniform locale settings on all clients in your cell, then use the same locale settings also for the Data Protector GUI on UNIX.
- In Data Protector GUI on UNIX systems in heterogeneous environments (different operating systems with different local settings in one cell), the locale must be set to a locale that uses UTF-8 encoding, *prior to* starting Data Protector GUI.
- In heterogeneous environments, if the appropriate encoding is not available in the Data Protector GUI, set the desired locale and restart the GUI.

To list all locales, that are installed on your system, run the following command: `locale -a`.

The locales are installed during the installation of the operating system, but it is possible to install them also after the operating system installation is finished. For detailed steps refer to the operating system documentation.

To set the locale, execute the following command:

```
export LANG=<locale>
```

You can list the locales and just copy and paste the name of the locale that uses UTF-8 encoding. To see if the locale has been set successfully, open a new terminal and check if the appropriate characters are available.

Example

The following is an example of setting the locale to a locale that uses UTF-8 encoding on UNIX and then running the GUI in that locale:

```
export LANG=C.utf8
xomni&
```

This enables you to switch the character encoding in Data Protector GUI on UNIX.

Changing the Default Character Encodings in the Data Protector GUI

The default character encoding in your Data Protector GUI is set according to:

- On UNIX systems: the locale in which your desktop environment is started.
- On Windows systems: the locale set in the `Regional options` in the Windows Control Panel.

After the IDB conversion (after upgrade) is finished, all files can be backed up and restored regardless of the character encoding used. To correctly display international characters in file names and session messages, the proper encoding has to be selected in the Data Protector GUI. Some character encodings are available in the Data Protector GUI by default but you can also replace the default character encodings with any other, if the appropriate locales/code pages are installed on your system.

Windows GUI

To replace the default character encoding with some other character encoding in the Windows GUI, perform the following steps:

1. In the Windows `Control Panel`, click `Regional Options`.
2. Click the `Advanced` button. In the code page conversion table, browse for the character encoding you would like to add to the Data Protector GUI and remember its code.

   If the code page for your character encoding is not installed, install it before continuing with the next step.
3. Open the Windows `Registry Editor` and browse for the following key:
   `\HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\OmnibackII\Gui\Core`.
4. The default character encodings are listed under the respective registry strings with the name `REG_KEY_CP<number>`. Double click the registry string to replace. For example, double click the `REG_KEY_CP5` registry string to replace Shift-JIS with some other character encoding.
5. Type in the character encoding code in the `Value Data` field. The code has to match the code in the Windows code page conversion table (see step 2). For example, type `949` to add the ANSI/OEM Korean character encoding to the Data Protector GUI.
6. Double click the corresponding `REG_KEY_NAME_CP<number>` registry string and edit the `Value Data` field. Type in the name of the character encoding as you would like it to appear in the Data Protector GUI in the encodings list.

   For example, if you modified the `REG_KEY_CP5` registry string in step 4, edit the `REG_KEY_NAME_CP5` registry string to change the name of the character encoding.
7. Restart the Data Protector GUI for the changes to take effect.

UNIX GUI

To replace the default character encoding with some other character encoding in the UNIX GUI, perform the following steps:

1. Open the terminal and run the following program:
   `/opt/omni/bin/xregedit`

2. Browse for the following key:
   `\HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\OmnibackII\Gui\Core`.

   This key is created when you enable the encoding switching in the Data Protector GUI for the first time, that is when you set a locale to the locale that uses the UTF-8 encoding and start the Data Protector GUI.

3. The default character encodings are listed under the respective registry strings with the name `REG_KEY_LC<number>`. Double click the registry string you would like to replace. For example, double click the `REG_KEY_LC5` registry string to replace Shift-JIS with some other character encoding.

4. Type in the character encoding code in the `Value Data` field. The code has to match the code in the `/usr/lib/nls/iconv/config.iconv` file. For example, type `cp864` to add the Arabic character encoding to the Data Protector GUI.

   Note that the appropriate locale must be installed on the system.

5. Double click the corresponding `REG_KEY_NAME_CP<number>` registry string and edit the `Value Data` field. Type in the name of the character encoding as you would like it to appear in the Data Protector GUI in the encodings list.

   For example, if you modified the `REG_KEY_LC5` registry string in step 4, edit the `REG_KEY_NAME_CP5` registry string to change the name of the character encoding.

6. Restart the Data Protector GUI for the changes to take effect.

Firewall Support

This section describes how to configure Data Protector in an environment where the Data Protector processes communicate across a firewall.

Communication in Data Protector

Data Protector processes communicate using TCP/IP connections. Every Data Protector system accepts connections on port 5555 by default. In addition, some processes dynamically allocate ports on which they accept connections from other Data Protector processes.

To enable Data Protector processes to communicate across a firewall, Data Protector allows you to limit the range of port numbers from which dynamically allocated ports are selected. Port ranges are defined on a per system basis. It is possible to define a port range for all Data Protector processes on a specific system, as well as to define a port range for a specific Data Protector agent only.

Configuration Mechanism

The port allocation behavior can be configured through two `omnirc` variables: `OB2PORTRANGE` and `OB2PORTRANGESPEC`. By default, both variables are not set and ports are assigned dynamically by the operating system.

Limiting the Range of Port Numbers

For All Data Protector Processes

You can limit the port range for all Data Protector processes on a system by using the `OB2PORTRANGE` variable in the `omnirc` file:

```
OB2PORTRANGE=<start_port>-<end_port>
```

Data Protector processes use dynamically allocated ports and select ports from this range. The port range is allocated by taking the first available port, starting with port "start_port". If there is no available port within the specified range, the port allocation fails and the requested operation is not done. Refer to Table 13-1 on page 629 for information on port consumption.

NOTE

The OB2PORTRANGE variable only applies to dynamically allocated ports. It does not affect the usage of the default Data Protector port number 5555.

Defining a port range for the Data Protector processes limits the port usage of Data Protector. It does not prevent other applications from allocating ports from this range as well.

For a Specific Data Protector Agent

In many cases it is not required that all Data Protector agents communicate across a firewall. For example, one specific agent can be outside a firewall, while all other components are inside of it. In such environments it is useful to limit the range of port numbers only for the specific agent. This allows you to define a much smaller port range and so reduce the need of open ports through the firewall.

You can limit the port range on a system on which a specific agent runs by using the OB2PORTRANGESPEC variable in the omnirc file:

```
OB2PORTRANGESPEC=<AGENT>:<start_port>-<end_port>;...
```

All agent processes check the OB2PORTRANGESPEC for range restrictions. If there is a range defined for an agent process, all dynamically allocated ports select from this specified range. The port range is allocated by taking the first available port, starting with port "start_port". If there is no available port within the specified range, the port allocation fails and the requested operation is not done. See “Examples of Configuring Data Protector in Firewall Environments” on page 634 for information on how to calculate the required range of port numbers.

The table below lists all possible Data Protector agent identifiers that can be used in the `OB2PORTRANGESPEC` variable. Note that agent processes that do not dynamically allocate listen ports are not listed in the following table.

Table 13-1 Agent Identifiers

Data Protector Component	Agent Identifier	Description	Port Consumption
Cell Manager	BSM	Backup Session Manager	1 port per concurrently running BSM
	RSM	Restore Session Manager	1 port per concurrently running RSM
	DBSM	Database Session Manager	1 port per concurrently running DBSM
	xSM	Wildcard matching all Session Managers	1[a] + 1 port per concurrently running Session Manager
	MMD	Media Management Daemon	1 port
	CRS	Cell Request Server Service	1 port
Media Agent	BMA-NET	Backup Media Agent[b]	1 port per concurrently running Media Agent
	RMA-NET	Restore Media Agent[b]	1 port per concurrently running Media Agent
	xMA-NET	Wildcard matching all Media Agents[b]	1 port per concurrently running Media Agent

a. This additional port is required during database operations such as filename purges or database upgrades.

b. BMA and RMA fork two processes, the main process and a NetIO process. The listen port is allocated by the BMA-NET / RMA-NET process.

NOTE

The OB2PORTRANGESPEC variable only applies to dynamically allocated ports. It does not affect the usage of the default Data Protector port number 5555.

Defining a port range for a specific Data Protector agent process limits the port usage of this agent. It does not prevent other processes (applications or other Data Protector agents) from allocating ports from this range as well.

Using Both Variables Together

If both variables OB2PORTRANGESPEC and OB2PORTRANGE are set, OB2PORTRANGESPEC overrides the settings of OB2PORTRANGE.

For example, the setting

```
OB2PORTRANGESPEC=BMA-NET:18000-18009
```

```
OB2PORTRANGE=22000-22499
```

limits the port range used by a Media Agent to port numbers 18000-18009, while all other Data Protector processes use port numbers from the range 22000-22499.

By using both variables it is possible to force a specific agent to use only a dedicated port range (OB2PORTRANGESPEC) and, at the same time, prevent other Data Protector processes from selecting port numbers from this range.

Port Usage in Data Protector

The following section provides two tables that describe the port requirements of the different Data Protector components. Table 13-2 breaks down the different Data Protector components and shows to which other components they can connect. It also defines the destination specification for the firewall rules. Table 13-3 gives the same list of components but shows from which other components they can accept connections. It also determines the source port of the firewall rule.

The following table provides a list of all Data Protector components. The first two columns list the process identifiers and their listen ports. The last two columns list all applicable connecting processes.

Table 13-2

Listening Component		**Connecting Component**	
Process	**Port**	**Process**	**Source Port**
Cell Manager			
Inet	5555	Application Agent	N/A[a]
		GUI/CLI	N/A[a]
CRS	Dynamic	Application Agent	N/A[a]
		GUI/CLI	N/A[a]
MMD	Dynamic	xSM	N/A[a]
		CLI (from CM)	N/A[a]
xSM	Dynamic	GUI/CLI	N/A[a]
		xMA[b]	N/A[a]
		xDA[b]	N/A[a]
		Application Agent	N/A[a]
Disk Agent			
Inet	5555	xSM	N/A[a]
xDA	Does not accept connections		
Media Agent			
Inet	5555	xSM	N/A[a]
xMA	Does not accept connections		
xMA-NET	Dynamic	xDA	N/A[a]
		Application Agent	N/A[a]

Table 13-2

Listening Component		**Connecting Component**	
Process	**Port**	**Process**	**Source Port**
Application Host			
`Inet`	5555	xSM	N/A[a]
Application Agent	Does not accept connections		

a. The source port of a connection is always assigned by the operating system and cannot be limited to a specific range.
b. Only for backup sessions with the reconnect feature enabled. The Disk Agent and a Media Agent communicate with the Cell Manager using the existing TCP connection. The connection in this column is only established after the original connection is broken.

When writing the firewall configuration rules, the process in the first column must be able to accept new TCP connections (SYN bit set) on the ports defined in the second column, from the process listed in the third column.

In addition, the process listed in the first column must be able to reply to the process in the third column on the existing TCP connection (SYN bit not set).

For example, the `Inet` process on a Media Agent system must be able to accept new TCP connections from the Cell Manager on port 5555. A Media Agent must be able to reply to the Cell Manager using the existing TCP connection. It is not required that a Media Agent is capable of opening a TCP connection.

The following table provides a list of all Data Protector components. The first two columns list all applicable connecting processes, while the last two columns list the process identifiers and their listen ports. Processes that do not initiate connections are not listed (for example, `Inet`).

Table 13-3

Connecting Component		Listening Component	
Process	**Port**	**Process**	**Port**
Cell Manager			
xSM	N/A[a]	xMA[b]	5555
	N/A[a]	xDA[b]	5555
	N/A[a]	Application Agent[b]	5555
	N/A[a]	MMD[c]	Dynamic
User Interface			
GUI/CLI	N/A[a]	`Inet` on CM	5555
	N/A[a]	CRS	Dynamic
	N/A[a]	BSM	Dynamic
	N/A[a]	RSM	Dynamic
	N/A[a]	MSM	Dynamic
	N/A[a]	DBSM	Dynamic
CLI (Cell Manager only)	N/A[a]	MMD	Dynamic
Disk Agent			
xDA	N/A[a]	xMA-NET	Dynamic
	N/A[a]	xSM[d]	Dynamic
Media Agent			
xMA	N/A	xSM[d]	Dynamic
	N/A[a]	UMA[b, e]	5555
Application Agents			

Table 13-3

Connecting Component		Listening Component	
Process	**Port**	**Process**	**Port**
Application Agent	N/A[a]	`Inet` on CM	5555
	N/A[a]	CRS	Dynamic
	N/A[a]	RSM	Dynamic
	N/A[a]	BSM	Dynamic
	N/A[a]	xMA-NET	Dynamic

a. The source port of a connection is always assigned by the operating system and cannot be limited to a specific range.
b. To be more precise, it is the `Inet` process that accepts the connection on port 5555 and then starts the requested agent process. The agent process inherits the connection.
c. This applies only to the MMD on the system running the CMMDB in a Manager-of-Managers (MoM) environment.
d. Only for backup sessions with the reconnect feature enabled.
e. Connections to the Utility Media Agent (UMA) are only required when sharing a library across several systems.

Examples of Configuring Data Protector in Firewall Environments

The following section provides examples on how to configure Data Protector in four different firewall environments.

Example 1: Disk Agent and Media Agent Installed Outside, Other Components Installed Inside a Firewall

You can configure your backup environment so that the Cell Manager and GUI are in the intranet and some Disk Agents and Media Agents are in the Demilitarized Zone (DMZ):

Figure 13-2 Configuration Diagram

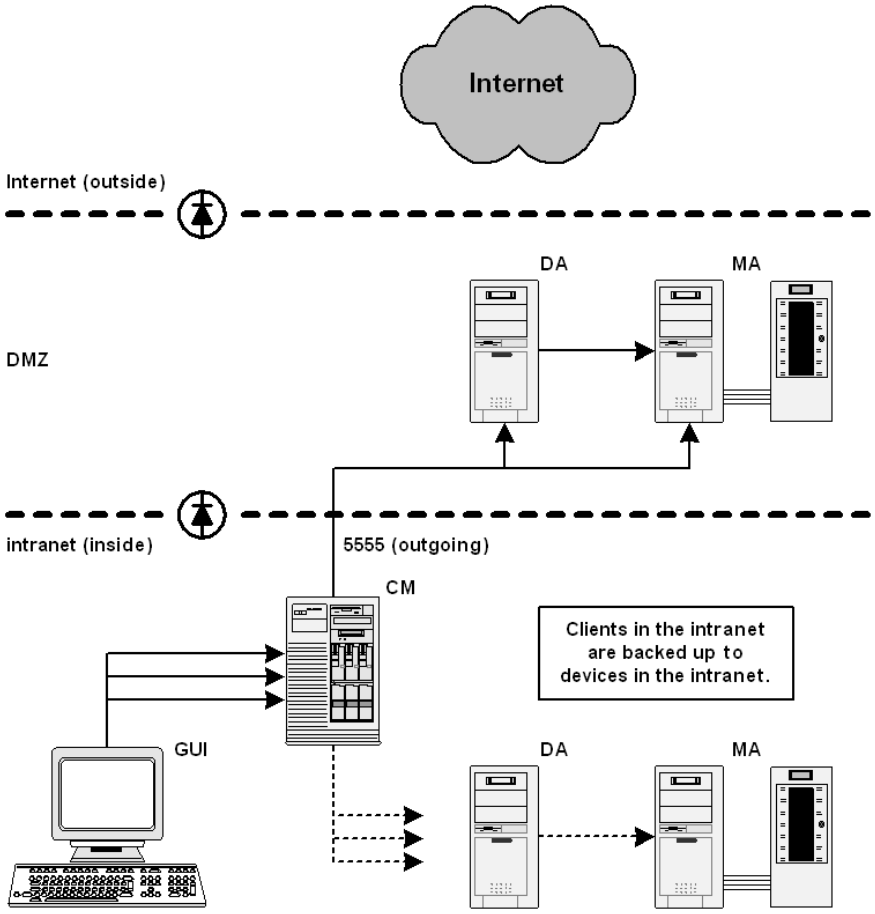

The following two items define the port range settings for this configuration:

1. In order to determine which processes need to communicate across the firewall, see Table 13-2 for the Disk Agent and a Media Agent. It shows that the Disk Agent and a Media Agent need to accept connections from the Session Manager on port 5555. This leads to the following rules for the firewall:

   ✓ Allow connections from the CM system to port 5555 on the DA system

   ✓ Allow connections from the CM system to port 5555 on the MA system

   This table also shows that a Media Agent needs to accept connections from the Disk Agent. However, since these two agents do not communicate through the firewall, you do not need to define a firewall rule for them.

2. See also Table 13-3 for the Disk Agent and a Media Agent.

   This table also shows that both agents may connect to the Session Manager and that a Media Agent may need to connect to a utility Media Agent (UMA). However, this only occurs when shared tape libraries are used or the `Reconnect broken connections` option is enabled. See “Backup Specification Options” on page 280 for information on this option.

Port Range Settings

Since all connections that need to go through the firewall connect to the fixed port number 5555, you do not need to define `OB2PORTRANGE` or `OB2PORTRANGESPEC` variables in this environment.

Limitations

- Remote installation of clients across the firewall is not supported. You need to install clients locally in the DMZ.
- This cell can back up clients in the DMZ, as well as clients in the intranet. However, each group of clients must be backed up to devices configured on clients that are on the same side of the firewall.

IMPORTANT If your firewall does not restrict connections from the intranet to the DMZ, it is possible to back up clients in the intranet to devices configured on clients in the DMZ. However, this is not recommended, as the data backed up in this way becomes more vulnerable.

- If a device in the DMZ has robotics configured on a separate client, this client must also be in the DMZ.
- This setup does not allow the backup of databases or applications using Application Agents on the clients in the DMZ. For details on Application Agents in the DMZ, refer to “Example 4: Application Agent and Media Agent Installed Outside, Other Components Installed Inside a Firewall” on page 644.

Example 2: Disk Agent Installed Outside, Other Components Installed Inside a Firewall

You can configure your backup environment so that the Cell Manager, a Media Agent, and GUI are in the intranet and some Disk Agents are in the DMZ:

Figure 13-3 Configuration Diagram

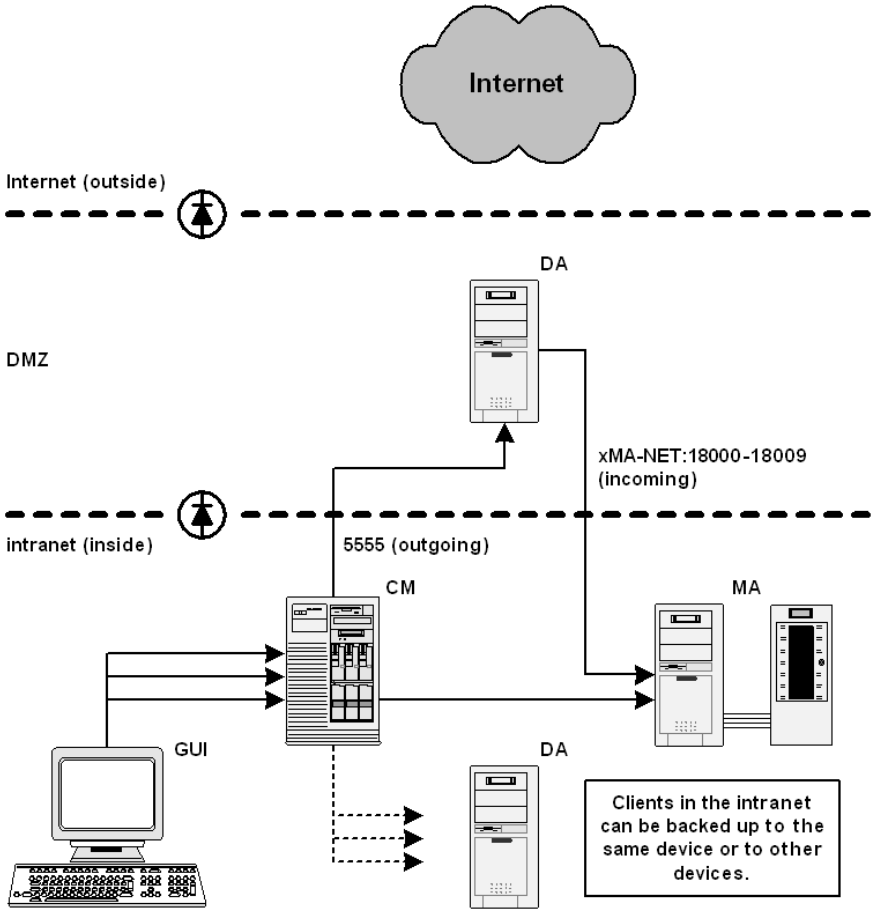

The following three items define the port range settings for this configuration:

1. In order to determine which processes need to communicate across the firewall, see Table 13-2 (Disk Agent column). It shows that the

Disk Agent needs to accept connections from the Session Manager on port 5555. This leads to the following rule for the firewall:

✓ Allow connections from the CM system to port 5555 on the DA system

2. See also Table 13-3 for the Disk Agent. It shows that the Disk Agent connects to a dynamically allocated port on a Media Agent. Since you do not want to open the firewall for communication between the Disk and a Media Agent in general, you need to limit the range of ports from which a Media Agent can allocate a listen port.

   See Table 13-1 for the port consumption of a Media Agent. A Media Agent requires only one port per running Media Agent. For example, if you have four tape devices connected, you may have four Media Agents running in parallel. This means that you need at least four ports available. However, since other processes may allocate ports from this range as well, you should specify a range of about ten ports on the MA system:

   ```
   OB2PORTRANGESPEC=xMA-NET:18000-18009
   ```

   This leads to the following firewall rule for the communication with a Media Agent:

   ✓ Allow connections from the DA system to port 18000-18009 on the MA system

NOTE This rule allows connections from the DMZ to the intranet, which is a potential security risk.

3. Table 13-3 also shows that the Disk Agent needs to connect to the Session Manager (BSM/RSM) when the `Reconnect broken connections` option is enabled. You can specify a required port range on the CM system analogous to the previous item.

   ```
   OB2PORTRANGESPEC=xSM:20100-20199
   ```

NOTE All Session Managers allocate ports from this range, not only the one communicating through the firewall.

Limitations

- Remote installation of clients across the firewall is not supported. You need to install clients locally in the DMZ.
- This setup does not allow the backup of databases or applications using Application Agents on the clients in the DMZ. For details on Application Agents in the DMZ, refer to “Example 4: Application Agent and Media Agent Installed Outside, Other Components Installed Inside a Firewall” on page 644.

Example 3: GUI Installed Inside, Other Components Installed Outside a Firewall

You can configure your backup environment so that the entire cell is in the DMZ and only the Graphical User Interface is in the intranet:

Figure 13-4 Configuration Diagram

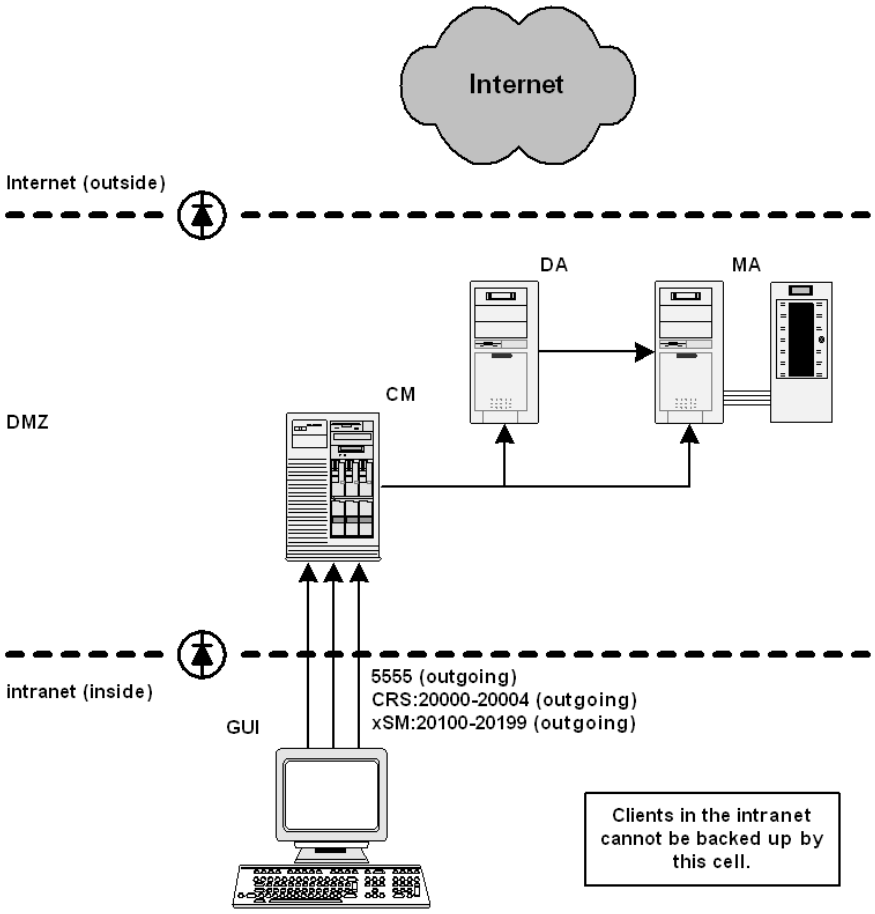

The following three items define the port range settings for this configuration:

1. Table 13-2 and Table 13-3 show that the GUI does not accept any connections. However, it needs to connect to the following processes

on the Cell Manager:

Table 13-4

Process	Port
`Inet`	5555
CRS	Dynamic
BSM	Dynamic
RSM	Dynamic
MSM	Dynamic
DBSM	Dynamic

This leads to the following firewall rule for the connection to the `Inet` listen port:

✓ Allow connections from the GUI system to port 5555 on the CM system

2. Table 13-1 shows that CRS requires only one port. However, since other processes may allocate ports from this range as well, you should specify a range of about five ports on the CM system. The port range could be defined as follows:

   ```
   OB2PORTRANGESPEC=CRS:20000-20004
   ```

   The resulting firewall rule for the connection to the CRS process is:

   ✓ Allow connections from the GUI system to ports 20000-20004 on the CM system

3. For the Session Manager, the situation is much more complex. Every Session Manager requires only one port. However, the number of Session Managers (BSM, RSM, MSM, DBSM) heavily depends on the backup environment. The minimum requirement can be estimated with the following formula:

$$NoOfPorts = NoOfConcurrentSessions + NoOfConnectingGUIs$$

Port Range Settings on the Cell Manager

For example, if there are 25 backup and five restore sessions running and two GUIs opened, you need to have at least 32 ports available. However, since other processes may allocate ports from this range as well, you should specify a range of about 100 ports on the CM system. The port range could be defined as follows:

```
OB2PORTRANGESPEC=xSM:20100-20199
```

or:

```
OB2PORTRANGESPEC=BSM:20100-20139;RSM:20140-20149;DBSM:20150-20199
```

Limitations

For this configuration almost all Data Protector functionality is available, including remote installation and online backup of databases and applications.

- This cell cannot be a part of a MoM environment if centralized media management or centralized licensing is used and the MoM cell is inside.
- All backup clients must be in the DMZ. The GUI client cannot be backed up by a Media Agent from the DMZ. The GUI can also be run from a client that is a member of another cell located in the intranet, provided that both cells use the same `Inet` listen port.

Example 4: Application Agent and Media Agent Installed Outside, Other Components Installed Inside a Firewall

You can configure your backup environment so that the Cell Manager and GUI are in the intranet and some Application Agents (SAP R/3, Oracle ...) and Media Agents are in the DMZ:

Figure 13-5 Configuration Diagram

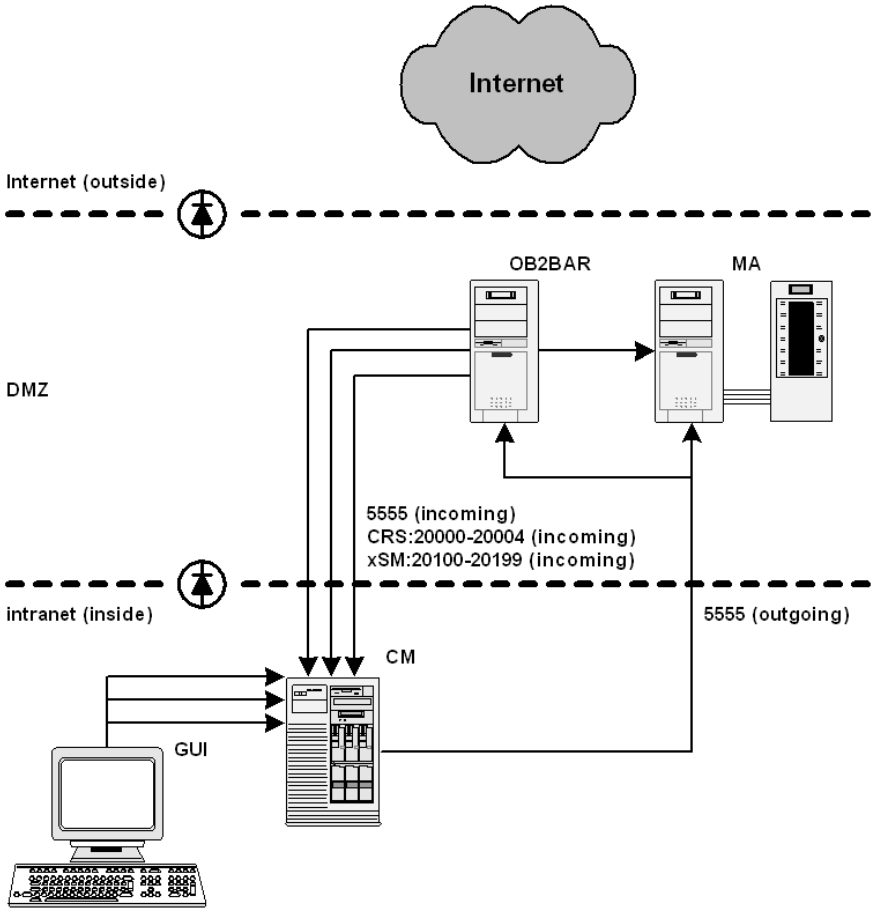

The following three items define the port range settings for this configuration:

1. Table 13-2 shows that Application Agents connects to the following processes on the Cell Manager:

Table 13-5

Process	Port
`Inet`	5555
CRS	Dynamic
RSM	Dynamic
BSM	Dynamic
DBSM	Dynamic
xMA-NET	Dynamic

   Here, the application Agent connects to a Media Agent. However, this connection does not go through the firewall and so you do not need to specify a port range.

   This leads to the following firewall rule for the connection to the `Inet` listen port.

   ✓ Allow connections from the Application Agent system to port 5555 on the CM system

NOTE This rule allows connections from the DMZ to the intranet, which is a potential security risk.

2. Table 13-1 shows that CRS requires only one port. However, since other processes may allocate ports from this range as well, you should specify a range of about five ports on the CM system. The port range could be defined as follows:

   ```
   OB2PORTRANGESPEC=CRS:20000-20004
   ```

The resulting firewall rule for the connection to the CRS process is:

✓ Allow connections from the Application Agent system to ports 20000-20004 on the CM system

3. For the Backup and Restore Session Manager, the situation is more complex. Every backup and restore session is started by one Session Manager, and every Session Manager requires one port. Additionally, an Application Agent may need to start some DBSMs. For Microsoft Exchange, Microsoft SQL, and Lotus Notes/Domino Server integrations, one DBSM will be started. For Oracle and SAP R/3 integrations, "concurrency + 1" DBSMs will be started. The port range for the Session Managers needs to be added to the `OB2PORTRANGESPEC` variable on the CM system:

Port Range Setting on the Cell Manager

```
OB2PORTRANGESPEC=CRS:20000-20004;xSM:20100-20199
```

Therefore, the firewall rule for the connections to the Session Managers is the following:

✓ Allow connections from the Application Agent system to ports 20100-20199 on the CM system

Limitations

- Remote installation of clients across the firewall is not supported. You need to install clients locally in the DMZ.
- This cell can back up clients in the DMZ, as well as clients in the intranet. However, each group of clients must be backed up to devices configured on clients that are on the same side of the firewall.

IMPORTANT

If your firewall does not restrict connections from the intranet to the DMZ, it is possible to back up clients in the intranet to devices configured on clients in the DMZ. However, this is not recommended, as the data backed up in this way becomes more vulnerable.

- If a device in the DMZ has robotics configured on a separate client, this client must also be in the DMZ

14 Troubleshooting

In This Chapter

If you have problems with Data Protector, use the suggestions in this chapter to get back on track, including information on:

For an overview and hints on the performance aspects of the Data Protector, refer to "Performance Considerations" on page A-8.

Backup devices (such as tape drives) are subject to specific Data Protector licenses. See the *HP OpenView Storage Data Protector Software Release Notes* for details.

Before Calling Your Support Representative

In order to speed up the process of solving your problem, you should prepare before reporting a problem to HP Customer Support Service. See the suggestions below for preliminary steps you can take.

Ensure that:

- You are not running into known limitations that cannot currently be overcome. For specific information on Data Protector limitations and recommendations, as well as on known Data Protector and non-Data Protector problems, see the *HP OpenView Storage Data Protector Software Release Notes*.
- Your problems are not related to third-party hardware and software. If they are, contact the third-party vendor for support.
- You have the latest Data Protector patches installed. Patches can be obtained from the HP OpenView Web site: http://support.openview.hp.com/patches/patch_index.jsp. The list of OS patches is available in the *HP OpenView Storage Data Protector Software Release Notes*.
- For integrated backups, the backup is not failing because the application is down.
- The debug logs or redo logs filesystem has not overflown.
- The application data filesystem has not overflown.

Collect the following data about the problem you encountered:

- A description of your problem, including the session output (or equivalent output, depending on the type of problem).
- Output from the `omnidlc` command for the Cell Manager and for all clients involved. Refer to “The Omnidlc Command” on page 661 for more information on `omnidlc` command.

Data Protector Log Files

If you encounter problems using the Data Protector application, you can use information in the log files to determine your problem.

Location of Data Protector Log Files

The Data Protector log files are located in the following directories:

- On Windows systems: *<Data_Protector_home>*\log
- On HP-UX and Solaris systems: /var/opt/omni/log and /var/opt/omni/server/log
- On other UNIX systems: /usr/omni/log
- On Novell NetWare systems: SYS:\USR\OMNI\LOG

Format of Data Protector Log Files

Most Data Protector log file entries are of the following format:

<time_stamp> <process:PID:Thread_ID> <source_file and branch> <Data Protector_version> <log_entry_message>

For example:

```
09/06/00 16:20:04 XOMNI.11561.0 ["/src/lib/ipc/ipc.c
/main/r31_split/10":3414] A.04.10 b325[ipc_receiveDataEx]
buffer 102400 bytes to small to receive data 796226418 bytes
=> ignored
```

Log Files and Their Contents

The table below describes the information found in Data Protector log files:

Table 14-1 Data Protector Log Files

Log File	Description
debug.log	Unexpected conditions are logged to this file. While some can be meaningful to you, it will be used mainly by the support organization.
Ob2EventLog.txt	Data Protector events that occurred during Data Protector operation and all Data Protector notifications are logged into this file. The Event Log represents a centralized Data Protector event depository.
inet.log	Requests made to the `Data Protector Inet` service are logged to this file. It can be useful to check the recent activity of Data Protector on clients.
IS_install.log	This file contains a trace of the remote installation and is located on the Installation Server.
media.log	Each time a medium is used for backup, initialized, or imported, a new entry is made to this log. The `media.log` can be used in IDB recovery to find the tape with the database backup and to find out which media were used after the last backup of the database.
omnisv.log	Contains information on when Data Protector services were stopped and started.
purge.log	Contains traces of the background purge of the IDB.

Table 14-1 Data Protector Log Files

Log File	Description

RDS.log	Contains IDB logs. The file resides on the Cell Manager:
	On Windows:
	`<Data_Protector_home>\db40\datafiles\catalog`
	On UNIX:
	`/var/opt/omni/server/db40/datafiles/catalog`

sanconf.log	Contains session reports generated by the `sanconf` command.
sm.log	Contains errors that occurred during backup and restore sessions, such as errors in parsing the backup specifications.
upgrade.log	This log is created during the upgrade and contains UCP (upgrade core part) and UDP (upgrade detail part) messages.
OB2_Upgrade.log (UNIX only)	This log is created during the upgrade and contains traces of the upgrade process.
sap.log, oracle8.log, informix.log, sybase.log, db2.log	Application specific logs contain traces of the integration calls between the application and Data Protector. The files are located on application servers and can be used for troubleshooting integrations.

Debugging

You should collect debugs only when the support organization requires them to resolve a technical issue. When Data Protector runs in debug mode, it creates debug information that consumes a large amount of disk space. Consult the support organization about the detail level that should be applied and environmental conditions for running Data Protector in the debug mode.

Limiting the Maximum Size of Debugs

Circular Debugging

Data Protector can run in a special debugging mode called circular debugging. In this mode, debugging messages are added until the size of the debug file reaches a preset size (*n*). The counter is then reset and the oldest debugging messages are overwritten. This limits the trace file size, but does not affect the latest records.

When to Use Circular Debugging

Using this mode is recommended only if the problem occurs near the end of the session or if Data Protector aborts or finishes soon after the problem has occurred.

Estimating the Required Disk Space

With circular debugging turned on, an estimate of the maximum required disk space is as follows:

- On Media Agent client(s): `2*`*`n`* `[kB]` for each running MA in a backup or restore.
- On Disk Agent client(s): `2*`*`n`* `[kB]` for each mount point in a backup or restore.
- On the Cell Manager client: `2*`*`n`* `[kB]`.
- On a integration client: `2*`*`n`* `[kB] * parallelism`.
- For Inet and CRS debugging, the upper limit cannot be reliably determined, because separate debug traces are produced for various actions.

Ways of Debugging

You can start Data Protector in the debug mode in different ways and use it to generate debug traces. For more details about debugging options refer to the section “Debug Syntax” on page 656.

IMPORTANT When Data Protector runs in the debug mode, debug information is generated for every action. For example, if you start a backup specification in the debug mode, Disk Agents deliver output on each client backed up in this backup specification.

Debugging Using the Data Protector GUI

To set the options for debugging using the Data Protector GUI, in the `File` menu, click `Preferences`, and then click the `Debug` tab. Specify the debug options and restart the GUI. The GUI will be restarted in the debug mode.

Debugging Using the Trace Configuration File

Another way to set debugging options is to edit the trace configuration file (`/etc/opt/omni/server/options/trace` on UNIX and *<Data_Protector_home>*`\Config\server\Options\trace` on Windows).

Debugging Using the OB2OPTS Variable

Debugging parameters for Data Protector integrations can be set using the `OB2OPTS` environment variable. For information about the `OB2OPTS` variable contact your Support Representative.

Debugging Scheduled Sessions

To debug scheduled sessions, edit the schedule file (`/etc/opt/omni/server/schedules` or `/etc/opt/omni/server/barschedules` on UNIX and *<Data_Protector_home>*`\Config\server\Schedules` or *<Data_Protector_home>*`\Config\server\BarSchedules` on Windows). Debugging parameters must be added in the first line of the file.

NOTE Before you edit the file, make a copy of it, as the changes have to be reverted when debugging is no longer desired.

Example of a Modified Schedule

```
-debug 1-99 sch.txt
-full
-only 2002
-day 14 -month Dec
-at 22:00
```

Debug Syntax

Almost all Data Protector commands can be started with an additional `-debug` parameter that has the following syntax:

`-debug 1-99[,C:<n>][,T:<s>][,U] <XYZ> [<host>]`

where:

`1-99` is the debug range. The range should always be specified as 1-99 unless instructed otherwise. Optional parameters (size of debug files, timestamp, and the Unicode flag) are to be specified as a part of the ranges parameter and separated by commas as follows:

- `C:<n>` limits the size of debug files to *n* kilobytes. The minimum value is 4 (4 kB) and the default value is 1024 (1 MB).
- `T:<s>` is the timestamp resolution, where the default value is 1, 1000 means the resolution is one millisecond and the value 0 means timestamps are turned off. The timestamp resolution and size limit for circular debugging are supplied as a part of the ranges parameter.
- `U` is the Unicode flag. If it is specified, the debug files (on Windows) are written in the Unicode format.

`<XYZ>` is the debug postfix, for example DBG_01.txt

`<host>` is the list of hostnames where debugging is turned on.

NOTE On some platforms (Novell NetWare, MPE), millisecond resolution is not available.

The list of hostnames limits the systems where debugging is turned on during the execution of the Data Protector command. If there are multiple systems on the list, they should be delimited by spaces. The entire list must be within quotation marks, for example: `"host1.company.com host2.company.com"`.

Trace File Name

The debug postfix option is used for creating the trace files in the following directory:

- On UNIX systems: `/tmp`
- On Windows systems: *<Data_Protector_home>*`\tmp`
- On Novell NetWare systems: `SYS:\USR\OMNI\TMP`

The files are named

`OB2DBG_`*<did>*`__`*<Program>*`_`*<Host>*`_`*<pid>*`_`*<XYZ>*

where:

<did> (debugging ID) is the process ID of the first process that accepts the debugging parameters. This ID is used as an ID for the debugging session. All further processes will use this ID.

<Program> is the code name of the Data Protector program writing the trace.

<Host> is the name where the trace file is created.

<pid> is the process ID.

<XYZ> is the postfix as specified in the `-debug` parameter.

Once the backup or restore session ID(*<sid>*) is determined, it will be added to the file name:

`OB2DBG_`*<did>*`_`*<sid>*`_`*<Program>*`_`*<Host>*`_`*<pid>*`_`*<XYZ>*

Processes that add the *<sid>* are BMA/RMA, xBDA/xRDA, and other processes started by the session, but not by the BSM/RSM itself.

NOTE The session ID is intended to help you identify sets of debug files. Other debug files may belong to the same session and you may have to provide them as well.

trace.log A `trace.log` file is generated on the Cell Manager, containing information where (on which hosts) debug files are generated and which debugging prefixes are being used. Note that this file does not contain a complete list of all generated files.

OB2DBGDIR The default location of trace files can be changed on a per system basis with the `omnirc` variable `OB2DBGDIR`. For more details about `omnirc` variables, refer to “Using Omnirc Options” on page 615.

INET Debug on UNIX

To debug `Inet` on UNIX systems, edit the `/etc/inetd.conf` file and change the following line:

1. `omni stream tcp nowait root /opt/omni/lbin/inet inet -log /var/opt/omni/log/inet.log`

   to

   `omni stream tcp nowait root /opt/omni/lbin/inet inet -log /var/opt/omni/log/inet.log -debug 1-140 SSF`

2. After the file has been changed and saved, run the `/etc/inetd -c` command to apply the changes.

NOTE If you enable `Inet` debugs, all integrations will generate trace log files.

INET Debug on Windows

To debug `Data Protector Inet` on Windows systems, launch the Windows Service Control Manager and restart the `Data Protector Inet` service with the following startup parameters:

`-debug 1-140 <`*`POSTFIX`*`>`

NOTE If you enable `Inet` debugs, all integrations will generate trace log files.

CRS Debug on UNIX

To debug CRS on UNIX systems:

1. Stop the CRS by entering the `/opt/omni/lbin/crs -shutdown` command.
2. Restart the CRS with the debug option by entering the `/opt/omni/lbin/crs -debug 1-140 <POSTFIX>` command.

CRS Debug on Windows

To debug Data Protector CRS on Windows systems, launch the Windows Service Control Manager and restart the `Data Protector` CRS service with the following startup parameters:

`-debug 1-140 <POSTFIX> <Cell_Manager_name>`

NOTE Use the `-debug` option carefully because execution traces can become quite large.

CRS Debug in the Microsoft Cluster Environment

In the Data Protector shared directory, edit the `<Data_Protector_home>\Config\server\options\Trace` file. Add the following lines:

```
ranges=1-99,110-500

postfix=DBG

select=obpkg.rc.aus.hp.com
```

From the Cluster Administrator GUI, take the CRS service resource (OBVS_MCRS) offline.

CAUTION Do not stop the CRS from Windows Service Control Manager as it will cause the Data Protector package to failover.

CRS Debug in the MC/Service Guard Environment

To debug Data Protector CRS in MC/Service Guard environment, follow the procedure:

1. Open the file `/etc/opt/omni/server/options/trace` and uncomment and set the required debugging options. Close and save the file.
2. Start debug collection by entering the `/opt/omni/lbin/crs -redebug` command.

To stop the debug collection, set all debugging options in the `/etc/opt/omni/server/options/trace` file to an empty string, save the file, and then issue the `/opt/omni/lbin/crs -redebug` command.

Collecting Data to be Sent to HP Customer Support Service

Since Data Protector operates in large network environments, the data needed by the HP support service might sometimes be difficult to gather. Data Protector provides a tool for collecting and packing log, debug and getinfo files to be sent to the HP support service.

After Data Protector debugging has been enabled (as described in the sections “Debugging” on page 654 or “Example of Collecting Data to be Sent to HP Customer Support Service” on page 668), the Data Protector `omnidlc` command can be used to compress the data needed by the HP support service. The command transfers the data from selected clients or Cell Managers to the Cell Manager or MoM, where it is then automatically packed.

Using the `omnidlc` command, it is possible to selectively collect the data, for example, it is possible to collect only log files from a certain client, or only debug files that were created during a particular Data Protector session, and similar.

For more information on `omnidlc` command, refer to “The Omnidlc Command” on page 661.

NOTE

The `omnidlc` command cannot be used to collect the Data Protector installation execution traces. For more information on how to create and collect the Data Protector installation execution traces, refer to the “Creating Installation Execution Traces” section in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

The Omnidlc Command

Overview

The `omnidlc` command collects Data Protector debug, log and getinfo files from the Data Protector cell (by default, from every client) or from a MoM environment (by default, from every Cell Manager). To collect data from Cell Managers in a MoM environment, the command must be run from the MoM. To collect data from clients in a MoM environment, the command must be run from their Cell Managers.

The collected data can then be:

- compressed and/or packed on the Cell Manager or on the MoM
- deleted on clients
- the disk space required for the collected data can be displayed

NOTE Only one of the above three actions is possible to be performed at one time.

The `omnidlc` command also provides a means of unpacking and uncompressing the packed or compressed files, and a means of limiting the scope of the collected data.

Limitations

- The *debug* files that are not created in Data Protector default debug directory, but in one or several other directories must be copied to the `/tmp` (UNIX systems) or to the `<Data_Protector_home>\tmp` (Windows systems) directory in order to be collected using the `omnidlc` command. The command collects only the *debug* files from the Data Protector default debug directory, which is `/tmp` on UNIX systems or `<Data_Protector_home>\tmp` on Windows systems.
- To add files other than the collected files to the package, copy the files to one of the following directories before running the `omnidlc` command: `dlc/<client>/getinfo`, `dlc/<client>/log`, or `dlc/<client>/tmp`. You cannot add directories, but only files. If the files are not copied to one of the specified directories, the package cannot be unpacked during the unpack phase.
- The command can only be used on Cell Managers and MoMs.
- It cannot be used to collect the Data Protector installation execution traces.
- The Data Protector GUI debug files for systems other than Cell Manager can only be gathered using the `-hosts` option. The description of the `-hosts` option is provided further on in this section.
- To collect debug files in a cluster, the command must be run using the `-hosts` option; the cluster nodes hostnames must be specified as the argument for the option. In a cluster, if the `-hosts` option is not specified, the data is collected from the active node. The description of the `-hosts` option is provided further on in this section.

- When a MA agent launches an UMA process on some other client, UMA debugs are not gathered. In this case, the `-hosts` or `-allhosts` option should be used. It is not recommended to use the `-allhosts` option in large environments.

Syntax

The following is the syntax of the `omnidlc` command:

```
omnidlc {-session <sessionID> | -did <debugID> | -postfix
string | -no_filter} [-allhosts | -hosts <list>] [-pack
<filename> | -depot [<directory>] | -space | -delete_dbg]
[-no_getinfo] [-no_logs] [-no_debugs] [-no_compress]
[-verbose]
```

```
omnidlc -localpack [<filename>]
```

```
omnidlc -unpack [<filename>]
```

```
omnidlc -uncompress <filename>
```

Limiting the Scope of Collected Data

To limit the scope of collected data the following `omnidlc` command options are used:

```
{-session <sessionID> | -did <debugID> | -postfix string |
-no_filter} [-allhosts | -hosts <list>] [-no_getinfo]
[-no_logs] [-no_debugs]
```

Limit the scope of the collected data by combining the following features:

- The data is collected from all systems in the cell (using the `-allhost` option). This is the default behavior.
- The data is collected only from the selected systems (using the `-hosts <list>` option). The *<list>* argument are the hostnames of the clients (or in a MoM environment, Cell Managers), separated by spaces.
- The getinfo, log files or debug log files can be excluded from the collected data (using the `-no_getinfo`, `-no_logs`, or `-no_debugs` options).
- The debug files are collected only from a specific session (using the `-session <sessionID>` option).
- The debug files matching a specific debugID are collected (using the `-did <debugID>` option).
- The debug files matching a specific postfix are collected (using the `-postfix <string>` option).

NOTE

When using the `-did` or the `-session` option, the debug, log and getinfo files from clients other than Cell Manager can sometimes not be gathered. In such a case, use the `-hosts` or the `-allhosts` option additionally. Note that the `-allhosts` option can be time consuming.

In order not to limit the scope of collected data, use the `-no_filter` option. If this option is used, you must specify either the `-allhosts` or the `-host <list>` option.

Compression on the Clients

The collected data is then, by default, compressed on the clients involved. It is also possible to disable the compression on the clients using the `-no_compress` option. The compressed files are added the `.gz` extension. The compressed files can be uncompressed using the `-uncompress <filename>` option.

Operations on Collected Data

To select the type of operation performed on the collected compressed or uncompressed data from the clients, the following `omnidlc` command options are used:

```
[-pack <filename> | -depot [<directory>] | -space |
-delete_dbg]
```

The collected compressed or uncompressed data from the clients is subject to one of the following operations:

- The data is sent over the network to the Cell Manager (default behavior), where it can be:
  - Packed and saved in the current directory as the `dlc.pck` file, if the `-pack <filename>` option is not used. If the `-pack <filename>` option is used, the data is packed and saved either in the current directory (if the `<filename>` argument is specified as a file name) or in the specified directory under the specified file name (if the `<filename>` argument is specified as a full path name).

    The packed file includes a generated directory structure that includes the hostnames, paths and the (compressed) collected files of the clients involved.

    The packed files can be unpacked using the `-unpack [<filename>]` option.

— Left unpacked and saved to the specified directory (using the `-depot [<directory>]` option). If the `<directory>` is not specified, the files are saved in the `<Data_Protector_home>\tmp\dlc` directory on Windows Cell Managers, or in the `/tmp/dlc` on UNIX Cell Managers. If the `<directory>` is specified, the collected files are saved to the `dlc` directory of the specified directory.

The directory for the packed or unpacked files is generated as follows:

```
./dlc/client_1/tmp/debug_files
./dlc/client_1/log/log_files
./dlc/client_1/getinfo/get_info.txt
./dlc/client_2/tmp/debug_files
./dlc/client_2/log/log_files
./dlc/client_2/getinfo/get_info.txt
...
```

- Deleted on the clients (using the `-delete_dbg` option). Note that only the debug log files are deleted; the get info and log files are not deleted.
- The amount of disk space required on the Cell Manager to gather the data is displayed (using the `-space` option)

Segmentation of Data

If the data is sent over the network to the Cell Manager and if the file to be sent is larger than 2GB, the file is split in 2GB chunks before it is compressed (it can be left uncompressed) and sent to the Cell Manager. Every chunk retains the file name and is added the first extension ranging from `s001` to `s999`. The second extension (`.gz`) is not added if the files are not compressed. Additionally, on the Cell Manager side, if the size of all collected compressed or uncompressed files exceeds 2GB, the collected files are packed in 2GB sized (original size) packages and added an extension ranging from `s001` to `s999`.

Additional Operations

- If the collected data was sent to the Cell Manager, the collected compressed or uncompressed, *unpacked* files (if the `-depot [<directory>]` option was used) can then be packed using the `-localpack [<filename>]` option. This option packs the directory structure of the current directory (must be the directory containing

the `dlc` directory generated by the `-depot` option) as the *<filename>* argument. If the *<filename>* argument is not specified, the `dlc.pck` file is created in the current directory.

This option is equivalent to the `-pack` *<filename>* option, but is to be used only if the data is collected using the `-depot` [*<directory>*] option.

- To unpack collected packed data, use the `-unpack` [*<filename>*] option. If the *<filename>* argument is not specified, the `dlc.pck` file from the current directory is unpacked. The data is always unpacked to the `dlc` directory in the current directory. Use this option when the collected (compressed or uncompressed) data was packed on the Cell Manager ether using the `-pack` *<filename>* option or the `-localpack` [*<filename>*] option.
- To uncompress an unpacked compressed single file, use the `-uncompress` *<filename>* option. Use this option after the packed data is unpacked using the `-unpack` [*<filename>*] option.
- To enable verbose output, specify the `-verbose` option. By default, verbose output is disabled.

Examples

1. To collect and compress all debug, log and getinfo files from the cell, and pack them in the “dlc.pck” file in the current directory on Cell Manager, using the verbose output, execute the following command:

   ```
   omnidlc -no_filter -allhosts -verbose
   ```

2. To collect only the log and debug files (without the getinfo files) from the clients “client1.company.com” and “client2.company.com” to the directory “c:\depot” on the Cell Manager, without compressing and packing the files, execute the following command:

   ```
   omnidlc -no_filter -hosts client1.company.com
   client2.company.com -depot c:\depot -no_getinfo
   -no_compress
   ```

3. To collect the log, debug, and getinfo files from the client “client1.company.com”, compress and pack them to the “c:\pack\pack.pck” file on the Cell Manager, execute the following command:

   ```
   omnidlc -hosts client1.company.com -pack c:\pack\pack.pck
   ```

4. To delete all debug files for the session with the ID “2003/08/27-9”, execute the following command:

   ```
   omnidlc -session 2003/08/27-9 -delete_dbg
   ```

5. To display disk space needed on the Cell Manager for the uncompressed debug files with the debugID “2351” from the client “client.company.com”, execute the following command:

   ```
   omnidlc -did 2351 -hosts client.company.com -space
   -no_getinfo -no_logs -no_compress
   ```

6. To pack the directory structure in the current directory (must be the directory containing the `dlc` directory generated by the `-depot` option) to the “dlc.pck” file in the same directory, execute the following command:

   ```
   omnidlc -localpack
   ```

7. To unpack the “dlc.pck” file to the “dlc” directory of the current directory, execute the following command:

   ```
   omnidlc -unpack
   ```

Example of Collecting Data to be Sent to HP Customer Support Service

Follow the procedure below to collect debug, log, and getinfo files for problems occurring during backup sessions involving one client and the Cell Manager:

1. Reduce the error environment as much as possible:
   - Create a backup specification that contains just one or a few files or directories.
   - Include only one failing client in the debug run.
2. Create an `info` text file that contains the following:
   - Hardware identification of the Cell Manager, Media Agent, and Disk Agent clients. For example, `HP-9000 T-600 Series; Vectra XA.`
   - The SCSI controller's name, for example, `onboard_type/Adaptec xxx/...` for Windows Media Agent clients.
   - Topology information obtained from the `omnicellinfo -cell` command output.
   - The output of the `devbra -dev` command if you have issues with backup devices.
3. Discuss the technical issue with the support organization and request the following information:
   - Debug level (For example, "1-99." This is a command option needed later.).
   - Debug scope (client only, Cell Manager only, every system)
4. Exit all user interfaces and stop all other backup activities in the cell.
5. In case you need to collect the CRS or Inet debugs as well, you need to restart the CRS and Inet services on the Cell Manager in the debug mode as described in "Debugging" on page 654.

6. On the Cell Manager run the following command to start the GUI in debug mode:

   - On Windows systems: `manager -debug 1-140 error_run.txt`
   - On UNIX systems: `xomni -debug 1-140 error_run.txt`

   You can define the postfix of the trace file names created by substituting the `error_run` text with your preference.

7. Reproduce the problem using Data Protector.
8. Exit all user interfaces to quit the debug mode.

   If you collected CRS and Inet debugs as well, you have to and restart the Data Protector services on the Cell Manager without the debug option as described in “Debugging” on page 654.

9. On the Cell Manager, run the following command:

   ```
   omnidlc -postfix error_run.txt
   ```

   The command compresses the log, getinfo, and debug files with the `error_run.txt` postfix on the client, while sending them over the network to the Cell Manager, where they are packed and saved in the `dlc.pck` file in the current directory. Refer to “The Omnidlc Command” on page 661 for more information.

10. Email the packed files (`dlc.pck`) to the support organization.
11. Delete the created debug file (with the `error_run.txt` postfix) on the client by running the following command on the Cell Manager:

    ```
    omnidlc -postfix error_run.txt -delete_dbg
    ```

Browsing Troubleshooting Messages

Data Protector provides an interactive online troubleshooting utility, where you can get a detailed explanations of your error messages, including suggestions for correcting problems.

When you receive an error message from Data Protector, the error number is presented as a clickable link. To see detailed information about the error, click the link. The error message dialog appears providing extensive information about the error. Click `Details` to see a detailed description of the error message and the actions you can take to avoid or solve the problem.

Error message dialog consists of the following:

- `Error Message`: Exact message as it appears.
- `Description:` Detailed description of the error message.
- `Action:` Possible actions to take to solve or avoid the problem.

Figure 14-1 Sample Error Message Dialog

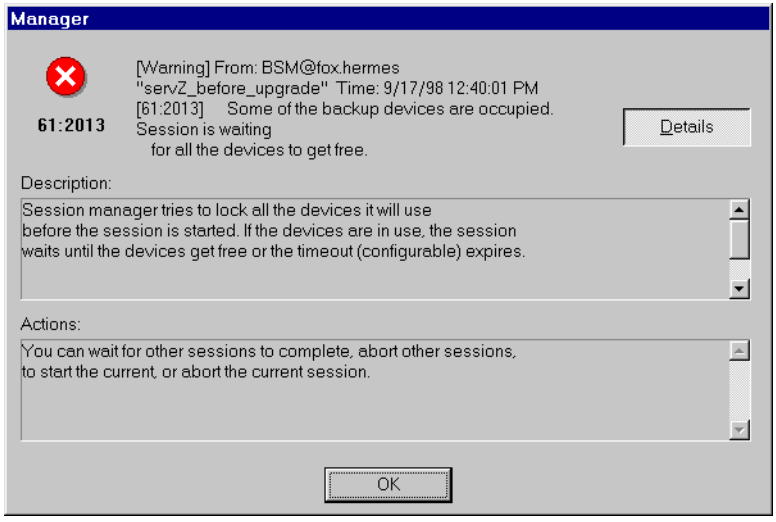

When You Cannot Access Online Troubleshooting

If the user interface cannot be started, you can access the troubleshooting file. This is a text file containing all Data Protector error messages each of which includes the following information:

- MESSAGE: The error message as it appears in Data Protector.
- DESCRIPTION: A detailed or extended information about the error.
- ACTION: Actions you can take to solve or avoid the problem.

The troubleshooting file is only available in the directory where the Cell Manager is installed. It can be found in the following locations:

- On UNIX: `/opt/omni/gui/help/C/Trouble.txt`
- On Windows: *<Data_Protector_home>*`\help\enu\Trouble.txt`

An example of an error message is shown below:

```
MESSAGE:
[12:5]  Internal error in (\"p\":num) => process aborted
This is an unexpected condition and is likely due to a
combination of circumstances involving both this product
and the operating system.

Report this error to your post-sales Data Protector
Support Representative.

DESCRIPTION:

An internal error occurred. The process was not able to
recover and aborted ungracefully immediately after
reporting this condition.

ACTION:

Before contacting your post-sales Data Protector Support
Representative, please gather as much information as
possible:

* Write down product version and build number.
* Make a note of the circumstances that cause this error.
* Save session output to a file (e.g. session.txt).
```

```
* Collect all log files (*.log) in
<Data_Protector_home>/log directories
on all hosts involved in the situation when this error
occurred
(i.e. host running VBDA, host running BMA and host
running BSM).
```

Description of Common Problems

If you have problems with Data Protector, find the problem area listed below that most closely matches the problem you are having:

- Networking and Communication, on page 674
- Service Startup, on page 680
- Device Usage, on page 685
- Starting Backup and Restore Sessions, on page 693
- User Interface Startup, on page 709

Certain functionality of Data Protector is subject to particular license requirements. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information on licensing.

Troubleshooting Networking and Communication

The section addresses the following networking and communication problems:

Hostname Resolution Problems

Hostname resolution is a very common problem in a Data Protector environment. It means that host A is unable to communicate with host B.

The table below shows Data Protector components and how they should communicate within the Data Protector environment. Communication among hosts means that host A in the table should resolve host B by its fully qualified domain name (FQDN). Resolving a host means that host A can interpret the FQDN and determine its IP address.

Table 14-2 Data Protector Components Name Resolution

Host A	Host B
Disk Agent Client Host	Media Agent Client Host
Disk Agent Client Host	Cell Manager Host
Disk Agent Client Host	MoM Server Host
Media Agent Client Host	Disk Agent Client Host
Media Agent Client Host	Cell Manager Host
Media Agent Client Host	MoM Server Host
Cell Manager Host	Media Agent Client Host
Cell Manager Host	Disk Agent Client Host

Table 14-2 Data Protector Components Name Resolution

Host A	Host B
Cell Manager Host	MoM Server Host
MoM Server Host	Disk Agent Client Host
MoM Server Host	Media Agent Client Host
MoM Server Host	Cell Manager Host

DNS Resolution Problem

Test DNS resolution among hosts using the `omnicheck` command. Refer to the "Verifying DNS Connections within Data Protector Cell" section in the *HP OpenView Storage Data Protector Installation and Licensing Guide* and to `omnicheck` man page for more information on how to use the command.

Enter the following command:

```
omnicheck -dns
```

This will check all DNS connections needed for normal Data Protector operating.

Problem

If the response to the `omnicheck` command is:

`<client_1> connects to <client_2>, but connected system presents itself as <client_3>`

The message may occur when the `hosts` file on *`client_1`* is not correctly configured or the hostname of the *`client_2`* does not match its DNS name.

If the response to the `omnicheck` command is:

`<client_1> failed to connect to <client_2>`

The message may occur when the `hosts` file on *`client_1`* is not correctly configured or *`client_2`* is unreachable (for example, disconnected).

Action

Consult your network administrator. Depending on how your environment is configured to perform name resolution, you may need to resolve this problem either in your DNS configuration or by editing the `hosts` file located in the following directories:

- On Windows: `<%SystemRoot%>\System32\drivers\etc`
- On UNIX: `/etc`

Problem

The response to the `omnicheck` command is:

`<client_1>` `cannot connect to` *`<client_2>`*

This means that the packet has been sent, but not received because of the timeout.

Action

Check for and resolve any network problems on the remote host.

Checking the TCP/IP setup

An important aspect of the TCP/IP configuration process is the setup of a hostname resolution mechanism. Each system in the network must be able to resolve the address of the Cell Manager and all machines with Media Agents and physical media devices. The Cell Manager must be able to resolve the names of all systems in the cell.

Action

Once you have the TCP/IP protocol installed, you can use the `ping` and `ipconfig` utilities to verify the TCP/IP configuration. For detailed steps, refer to the online Help index keyword “checking, TCP/IP setup”.

HOSTS file resolution problem

Action

If you encounter resolution problems when using the `Hosts` file, do the following:

- On Windows: edit the `LMHosts` file in the `<%SystemRoot%>\System32\drivers\etc` directory.
- On UNIX: edit the `/etc/hosts` file.

Client Fails with "Connection Reset by Peer"

On Windows, default configuration parameters of the TCP/IP protocol may cause connections to break. This can be due to a high network or computer usage, unreliable network, and connections between different operating systems.

The connection breaks and the system displays the error: `[10054] Connection reset by peer.`

Action

You can configure the TCP/IP protocol to use 8 instead of the default 5 retransmissions. It is better not to use higher values because each increment doubles the timeout. The setting applies to all network connections, not only to connections used by Data Protector.

On Windows, apply the change to the Cell Manager first.

If you run the UNIX Cell Manager and the problem persists, apply the change to any problematic Windows clients.

1. Add a new `DWORD` parameter `TcpMaxDataRetransmissions` and set its value to `0x00000008(8)` under the following registry key:

   ```
   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters
   ```

   ```
   MaxDataRetries:(DWORD):8
   ```

CAUTION

Making a mistake in editing the registry can cause your system to become unstable and unusable.

2. Restart the system after making this change.

Client Fails with "The Client Is not a Member of any Cell"

When performing a Data Protector operation on a client, the Cell Manager information is not found on the client. The operation fails with the following error:

```
The Client is not a member of any cell.
```

Action

- If the client is listed in the `Clients` context of the Data Protector GUI, proceed as follows:
  1. In the `Clients` context, expand `Clients`, right-click the client, and select `Delete`.
  2. Click `No`.
  3. Right-click `Clients` and select `Import Client`.
  4. Enter the name of the client and click `Finish`.
- If the client is not listed in the `Clients` context of the Data Protector GUI, proceed as follows:
  1. In the `Clients` context, right-click `Clients`, and select `Import Client`.
  2. Enter the name of the client and click `Finish`.

Excessive Logging to inet.log File

Problem

If the clients are not secured and the Cell Manager is configured in the MC/ServiceGuard environment or has multiple names or IP numbers, `inet.log` file may contain many entries of the following type:

```
A request 0 came from host name.company.com which is not a
cell manager of this client.
```

This happens because the client, which is not secured, recognizes only the primary hostname of the Cell Manager. Requests from any other host are allowed, but logged to `inet.log` file.

Action

Secure the client and the Cell Manager nodes. Requests from the hosts listed in the `allow_hosts` file will not be logged to `inet.log`. Requests from other hosts are denied.

If this workaround is for any reason not possible in your environment, you can secure the clients and specify `*` as an IP address range for the systems you want to allow access. This means that your clients will accept requests from all systems (any IP address) and will practically not be secured, but you will still resolve the excessive logging issue.

IMPORTANT

All possible hostnames for the Cell Manager nodes should be listed in the `allow_hosts` file on each client. This enables access to the client also in case of a failover.

If you accidentally lock-out a client, you can manually edit the `allow_hosts` file on this client.

Troubleshooting Data Protector Services and Daemons

The Data Protector services and daemons run on the Cell Manager. Run the `omnisv -status` command to check whether services are running.

If the Data Protector services seem to be stopped or have not been installed on the target Data Protector client, ensure first that you don't have a name resolution problem. Refer to "Troubleshooting Networking and Communication" on page 674 for more information.

You can run into the following problems with Data Protector services and daemons:

Problems Starting Data Protector Services on Windows

You do not have permission to start the services

The following error message displays:

`Could not start the <Service_Name> on <System_Name>.`

`Access is denied.`

Action

The system administrator should grant you permission to start, stop, and modify services on the system that you administer. Run the `services.msc` (located in the `<%SystemRoot%>\system32` directory) as an administrator by right-clicking the file while holding down the **Shift** button and selecting `Run as` from the pop-up menu. Then provide administrator's user name and password.

Changed service account properties

If the service account does not have permission to start the service or if the service account properties (the password, for example) have been changed, you get the following error message:

```
The Data Protector Inet service failed to start due to the
following error:

The service did not start due to a logon failure.
```

Action

1. Modify the service parameters: in the Windows Control Panel, go to `Administrative Tools, Services`.
2. If this does not solve the problem, contact your system administrator to set up the account with appropriate permissions. The account should be a member of the `Admin` group and should have the `Log on as a service` user right set.

A specific service has not been found

The location of the service is registered in the `ImagePath` key. If the executable does not exist in the location specified under this key, the following error message is displayed:

```
Could not start the <Service_Name> on <System_Name>. The
system can not find the file specified!
```

Action

1. On the Cell Manager, copy the *<Data_Protector_home>*`\db40` and *<Data_Protector_home>*`\Config` directories to a safe location before uninstalling Data Protector.
2. Copy the *<Data_Protector_home>*`\db40` and *<Data_Protector_home>*`\Config` directories back in place.
3. Uninstall the current Data Protector installation either on the client or on the Cell Manager, and then reinstall the software.

This guarantees a clean installation with all the binaries in place.

MMD fails upon starting the CRS service

If the Data Protector CRS service fails to start and `mmd.exe` invokes a *Dr.Watson* diagnosis, this points to a corruption in the database log files.

Action

1. Delete the `mmd.ctx` file on the *<Data_Protector_home>*`\tmp` directory and the problems should be resolved.

2. Restart the services using the `omnisv -start` command.

RDS does not work on the Windows TSE Cell Manager

Use TCP transport instead of local transport by modifying the *<Data_Protector_home>*\db40\datafiles\catalog\velocis.ini file:

Under `TCP Configuration`, set `Enabled` to `yes`.

Problems Starting Data Protector Daemons on UNIX

The following daemons run on the UNIX Cell Manager:

- Data Protector CRS daemon: `/opt/omni/lbin/crs`
- IDB daemon: `/opt/omni/lbin/rds`
- Data Protector Media Management daemon: `/opt/omni/lbin/mmd`

The `Data Protector Inet` service (`/opt/omni/lbin/inet`) is started by the system inet daemon when an application tries to connect to the Data Protector port, which is by default port number 5555.

Normally, these daemons are started automatically during the system's start-up.

To manually stop, start, and get the status of Data Protector daemons, log on to the Cell Manager as root.

Stopping Daemons

To stop the Data Protector daemons, enter the following command in the `/opt/omni/sbin` directory:

```
omnisv -stop
```

Starting Daemons

To start the Data Protector daemons, enter the following command in the `/opt/omni/sbin` directory:

```
omnisv -start
```

Checking the Status of the Daemons

To check the running status of the Data Protector daemons, enter the following command in the `/opt/omni/sbin` directory:

```
omnisv -status
```

There are several possible reasons why the Data Protector daemon has failed to start:

Raima Velocis server daemon could not be started

```
/opt/omni/sbin/omnisv -start
```

```
Could not start Raima Velocis server daemon.
```

Action

See `/var/opt/omni/server/db40/datafiles/catalog/RDS.log` for details.

Check that you have all IDB files in the `/var/opt/omni/server/db40` directory. Compare the list of files in the `/opt/omni/newconfig/var/opt/omni/server/db40` to the list of files in the `/var/opt/omni/server/db40` directory. Ensure that these directories are mounted.

Raima Velocis server daemon is apparently not running

If any of the Data Protector commands terminate with following message:

```
[12:1166] Velocis daemon error - the daemon is probably not running
```

Action

Check if the database server is really not running using following command: `/opt/omni/sbin/omnisv -status`

- If the database server is not running, start it by running: `/opt/omni/sbin/omnisv -start`
- If the database server is running, then it is likely either that the `/var/opt/omni/server/db40` directory does not exist or some of the files are missing. This can happen if someone has accidentally removed the directory or some of the IDB files. Recover the IDB. Refer to “Recovering the IDB” on page 494.

Data Protector Cell Manager daemon could not be started

```
/opt/omni/sbin/omnisv -start
```

```
Could not start the Cell Manager daemon.
```

Action

See `/var/opt/omni/tmp/omni_start.log` for details.

Ensure that the following configuration files exist:

- /etc/opt/omni/server/options/global
- /etc/opt/omni/server/options/users/UserList
- /etc/opt/omni/server/options/ClassSpec

Data Protector Processes

Table 14-3 shows which processes run and where they run while Data Protector is idle, or doing a backup, a restore or a media management session.

Table 14-3 Which Processes Run Where, and When

	Idle	Backup	Restore	Media Management
Windows Cell Manager	rds.exe, crs.exe, omniinet.exe, bsm.exe	rds.exe, mmd.exe, omniinet.exe, mmd.exe	rds.exe, omniinet.exe, mmd.exe, crs.exe, rsm.exe	rds.exe, omniinet.exe, mmd.exe, crs.exe, msm.exe
UNIX Cell Manager	rds, mmd,crs	rds, mmd, crs, bsm	rds, mmd, crs, rsm	rds, mmd, crs, msm
Windows Disk Agent Client	omniinet.exe	omniinet.exe, vbda.exe	omniinet.exe, vrda.exe	omniinet.exe
UNIX Disk Agent Client		vbda	vrda	
Windows Media Agent Client	omniinet.exe	omniinet.exe, bma.exe	omniinet.exe, rma.exe	omniinet.exe, mma.exe
UNIX Media Agent Client		bma	rma	mma

Troubleshooting Devices and Media

This section describes solutions to the following problems that can arise while using backup devices:

Problems involving device SCSI addresses are explained in detail in Appendix B of the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Cannot Access Exchanger Control Device on Windows

Data Protector uses the SCSI mini-port driver to control backup drives and libraries. Data Protector may fail to manage devices if other device drivers are loaded on the same system. The error message `Cannot access exchanger control device` appears when device operations such as media formatting or scanning are started.

Action

Run the `<Data_Protector_home>\bin\devbra -dev` command on the system where the devices are located, to list all physical devices that are configured on the system. If any of the SCSI addresses have the `CLAIMED` status value, they are used by another device driver.

Disable the Windows robotic driver. For detailed steps, refer to the online Help index keyword “robotics drivers”.

Device Open Problem

The error message `Cannot open device (not owner)` appears when trying to use a DDS device.

Action

Check whether you are using a medium that is incompatible with the Media Recognition System. Media used with DDS drives must comply with the Media Recognition System.

Using Unsupported SCSI HBAs/FC HBAs on Windows

System fails due to usage of unsupported SCSI HBAs/FC HBAs with backup devices.

Typically, the problem occurs when the SCSI device was accessed by more than one Media Agent at the same time or when the length of the transferred data defined by the device’s block size was larger than the length supported by the SCSI HBA/FC HBA.

Action

You can change the `Block size` in the `Advanced Backup Options` for the backup specification.

For information on supported SCSI HBAs/FC HBAs, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

For detailed steps, refer to the online Help index keyword “setting advanced options for devices and media”.

Automatic Recovery Upon Library Reconfiguration Failure

Configuration errors are reported during modification of an existing library configuration using the `sanconf` command, after the device list file has been altered. The library configuration remains only partially created.

You can recover the previous library configuration if you reuse the file with list of hosts in your SAN environment and scan the hosts with `sanconf` again.

Action

Proceed as follows:

1. Scan the hosts in the cell by running the following command:

   ```
   sanconf -list_devices mySAN.txt -hostsfile hosts.txt
   ```

2. Configure your library using the saved configuration file. Run the following command:

   ```
   sanconf -configure mySAN.txt -library
   <LibrarySerialNumber> <LibraryName>
   [<RoboticControlHostName>] [<DeviceTypeNumber>]
   -hostsfile hosts.txt
   ```

   The previous successful library configuration is automatically recovered.

Later, if you add, remove or modify the library and configuration with the `sanconf` command fails for any reason, you can repeat the above procedure to restore the successful configuration.

Medium Quality Statistics

This functionality is used to detect any problems with media while they're still in their early stages. Before each medium is ejected from a drive, Data Protector queries the `SCSI log sense` command for medium read and write statistical information. The information is written to the `media.log` file.

The medium quality statistics feature is disabled by default. To enable it, set the following global variable: `Ob2TapeStatistics=1` in the Global Options file.

The path of the Global Options file is:

- on UNIX: `/etc/opt/omni/server/options`
- on Windows: *<Data_Protector_home>*`\Config\server\Options`

If you receive media related errors during write operations, or if the medium is marked as poor, you can check the `media.log` file for media errors statistics. You can do this also when receiving media related errors during read operations.

`Media.log` file contains the following error statistics:

Error statistics	Explanation
`errsubdel=`*n*	errors corrected with substantial delays
`errposdel=`*n*	errors corrected with possible delays
`total=`*n*	total number of re-writes
`toterrcorr=`*n*	total number of errors corrected and recovered while writing
`totcorralgproc=`*n*	total number of times correction algorithm processed
`totb=`*n*	total bytes processed (write)
`totuncorrerr=`*n*	total number of uncorrected errors (write)

where *n* stands for number of errors.

If a parameter has the value `-1`, it means that the device does not support this statistic parameter. If all parameters have the value `-1`, it can either mean that during processing the tape quality statistics an error occurred or the device does not support medium quality statistics at all.

Although the tape statistical results are reported in bytes for `total bytes processed`, this is not true for all devices. LTO and DDS devices report data sets and groups, respectively, and not bytes.

Examples

Here are a few examples from the `media.log` file:

- Log sense write report for DLT/SDLT devices - total bytes processed.

  ```
  Media ID from tape= 0fa003bd:3e00dbb4:2310:0001; Medium Label=
  DLT10; Logical drive= dlt1; Errors corrected no delay= 0; Errors
  corrected delay= 0; Total= 13639; Total errors corrected= 13639;
  Total correction algorithm processed= 0; Total bytes processed=
  46774780560; Total uncorrected errors= 0
  ```

  46774780560 bytes of native data after compression were processed (a full DLT8000 tape).

- Log sense write report for LTO devices - total data sets processed.

  ```
  Media ID from tape=0fa003bd:3e0057e6:05b7:0001; Medium Label=
  ULT2; Logical drive=ultrium1; Errors corrected no delay= 0;
  ```

```
Errors corrected delay= 0; Total= 0;Total errors corrected= 0;
Total correction algorithm processed= 0; Total bytes processed=
47246; Total uncorrected errors= 0
```

One data set is 404352 bytes. To calculate the amount of total bytes processed, use the following formula:

```
47246 data sets * 404352 bytes = 19104014592 bytes after
compression a full tape).
```

- Log sense write report for DDS devices - total groups processed.

```
Media ID from tape= 0fa0049f:3df881e9:41f3:0001; Medium Label=
Default DDS_5; Logical drive= DDS; Errors corrected no delay=
-1; Errors corrected delay= -1; Total= -1; Total errors
corrected= 0; Total correction algorithm processed= 154; Total
bytes processed= 2244; Total uncorrected errors= 0
```

DDS1/2: One group is 126632 bytes
DDS3/4: One group is 384296 bytes

To calculate the amount of total bytes processed, use the following formula:

```
2244 groups * 126632 bytes = 284162208 bytes after compression
(a 359 MB backup on DDS2).
```

359 MB of data was backed up, resulting in 271 MB of native data on tape.

Medium Header Sanity Check

Data Protector performs a medium header sanity check before a medium is ejected from a drive to validate the medium header.

The medium header sanity check is enabled by default. The global variable can be set by uncommenting the following line in the Global Options file: `Ob2HeaderCheck=1`.

Problem

In case the medium header sanity check detects any header consistency errors on the medium, an error message is displayed and all the objects on the medium are marked as failed.

If the medium header is corrupt, all objects on the affected medium are marked as failed and the medium state is marked as poor.

Action

Export the medium from the IDB and restart the failed session using a different medium.

Cannot Use Devices After Upgrading to Data Protector A.05.50

Problem

After upgrading to Data Protector A.05.50, you cannot use devices that were configured as different device types in previous releases. For example, you cannot use 9940 devices that were configured as 9840 devices, 3592 devices that were configured as 3590 devices, or SuperDLT devices that were configured as DLT devices. The following error occurs:

```
[Critical] From: BMA@ukulele.company.com "SDLT"  Time:
2/22/2003 5:12:34 PM
[90:43]  /dev/rmt/1m
Invalid physical device type => aborting
```

Action

Manually reconfigure these devices using the `mchange` command, located on the Cell Manager in the following directories:

- On HP-UX: `/opt/omni/sbin/utilns/HPUX`
- On Solaris: `/opt/omni/sbin/utilns/SOL`
- On Windows: *<Data_Protector_home>*`\bin\utilns\NT`

Command Syntax

`mchange -pool` *`PoolName`* `-newtype` *`NewMediaClass`*

where:

`PoolName` is the name of the media pool with devices that are currently configured and should be reconfigured (for example, Default DLT. Default T3590, or Default T9840).

`NewMediaClass` is the new media type of the devices, for example, T9940 for 9940 devices, T3592 for 3590 devices, and SuperDLT for SuperDLT device.

Example

```
mchange -pool "Default DLT" -newtype "SuperDLT"
```

The command changes media types for all media, drives and libraries that use the defined media pool. After you have executed this command for each device you wanted to change, move the media associated with the reconfigured devices from the current media pool to the media pool corresponding to these media.

For example, move the media associated with the reconfigured 9940 devices to the Default T9940 media pool, and the media associated with the reconfigured SuperDLT devices to the Default SuperDLT media pool. For related procedures, refer to the online Help.

Problems with Device Serial Number

Problem

Data Protector reports the following error when performing any operation involving the problematic backup device (such as backup or restore, format, scan, and so forth) or robotics:

`Device <DeviceName> could not be opened (Serial number has changed).`

The error is reported when the device path points to a device with a different serial number as the number stored in the IDB. This can happen in the following cases:

- you misconfigured the device (for example, using the `omniupload` command, or if you configured an incorrect device file)
- you replaced the physical device without updating the corresponding logical device (reloading the new serial number)
- misconfigured path in a multipath device

Action

Perform the following steps:

1. In the Data Protector GUI, switch to Devices & Media context.
2. In the Scoping Pane, expand `Devices`, right click the problematic device, and click `Properties`.
3. In the Results Area, click the `Control` tab and enable the `Automatically discover changed SCSI address` option.
4. Click the `Reload` button to update the device serial number in the IDB.

Cannot Find the Device File for the XCopy Engine on an External FC Bridge

Problem

You cannot locate the device file for the external FC bridge when configuring the XCopy engine.

Action

1. Using the FC bridge administration utility, check if the `Active Fabric` setting on the FC bridge is turned `ON`. If it is not, turn it `ON`.
2. On the backup system run the following command: `ioscan -fkn`

   In the output, the name of the external FC bridge and the device file for the external FC bridge are listed. You should see something similar to the following in the output:

   ```
   ctl 24 0/2/0/0.2.24.25A.05.10.0.5 sctl CLAIMED DEVICE HP
   A4688A

   /dev/rscsi/c19t0d5
   ```

Cannot Find the Device File for the XCopy Engine on an Internal FC Bridge

Problem

You cannot locate the device file for the internal FC bridge when configuring the XCopy engine.

Action

1. Using the backup device `Interface Manager` telnet utility, check that you have installed the License Key to enable direct backup. If you have not installed it, install it.
2. On the backup system run the following command: `ioscan -fkn`

   In the output, the name of the internal FC bridge and the device file for the internal FC bridge are listed. You should see something similar to the following in the output:

   ```
   ctl 5 0/8/0/0.1.16.255.0.0.2 sctl CLAIMED DEVICE HP
   C7200FC Interface

   /dev/rscsi/c18t0d7
   ```

Other Common Problems

Other common problems are hardware-related.

Action

Check the SCSI communication between the system and the device, such as adapters or SCSI cables and their length. Try running an OS-provided command, such as `tar`, to verify that the system and the device are communicating.

Troubleshooting Backup and Restore Sessions

You may run into the following problems while running or starting backup and restore sessions:

File Names or Session Messages Are Not Displayed Correctly in GUI

When using the Data Protector GUI, some file names or session messages containing non-ASCII characters are displayed incorrectly. This happens when a wrong character encoding is used to display file names and session messages in the GUI.

Action

To view these objects correctly, specify the appropriate encoding in the Data Protector GUI by selecting `Encoding` from the `View` menu, then selecting the appropriate coded character set.

To enable encoding switching in GUI on UNIX, set the locale to a locale that uses UTF-8 character encoding prior to starting GUI.

Refer to online Help index keyword "`internationalization`" for internationalization limitations tables.

Full Backups Are Performed Instead of Incrementals

There are several reasons, outlined below, that Data Protector might run a full backup despite the fact that you specified an incremental backup.

No previous full backup

Before performing an incremental backup of an object, Data Protector requires a full backup. Data Protector uses a full backup as a base for comparison to determine which files have changed and consequently need to be included in the incremental backup. If a protected full backup is not available for this comparison, a full backup is performed.

Action

Set the protection for the full backup.

The description has changed

A backup object is defined by the client, mount point, and description. If any of these three values changes, Data Protector considers it as a new backup object, and performs a full backup instead of an incremental.

Action

Use the same description for full and incremental backups.

Trees have changed

The protected full backup already exists but with different trees as its incremental backup. There are two possible reasons for this:

- You have changed the trees in the backup specification of the protected full backup
- You have created multiple backup specifications with the same backup object but different trees specified for the backup object.

Action

If you have multiple backup specifications with the same backup object, change the (automatically generated) universal description of the backup object. Data Protector will consider them as new objects and a full backup will be run. After the full backup is performed, incremental backups will be possible for all consecutive backups.

The backup owner is different

If your backups are configured to run as private, the person starting the backup is the owner of the data. For example, if USER_1 performs a full backup and USER_2 tries to start an incremental backup, the incremental backup will be executed as a full backup. This is because the data for USER_1 is private and cannot be used as a base for the USER_2's incremental backup.

The same problem occurs if USER_1 performs a full backup, then USER_2 performs an object copy session, and the original is exported or overwritten. USER_1 cannot perform an incremental backup because the full backup (the copy) now belongs to USER_2.

Action

Configure backup session `Ownership` in the `Advanced Backup Options` for the backup specification. The backup owner should be a user from the `Admin` user group. This will make all backups owned by this user, regardless of who actually starts the backup session.

Unexpected Mount Request for a Standalone Device

There are several situations, described below, that may cause Data Protector to issue a mount request for a standalone device while media are available in the backup device.

The media in the device are in a media pool that has the Non Appendable policy

Even though there is still available space on the media, the media will not be used because of the `Non Appendable` policy of the pool.

Action

Modify the media pool policy to `Appendable` to enable the appending of backups to the media until the media are full.

The media in the device are not formatted and the media pool to be used has a Strict policy

If your pool uses a `Strict` media allocation policy, media that are not formatted will not be used for backup. If no formatted media are available, Data Protector issues a mount request.

Action

If you would like Data Protector to automatically format unformatted media, set the media pool policy to `Loose` and change global variable `InitOnLoosePolicy` to 1.

The media in the device are not formatted and the media pool to be used has a Loose policy

If your pool uses a `Loose` media allocation policy, media are not automatically formatted.

Action

If you would like Data Protector to automatically format unformatted media, you need to change global variable `InitOnLoosePolicy` to 1.

The media in the device are formatted but are different from those in the preallocation list

The media in the device are formatted but are different from those in the preallocation list of the backup specification, and the pool specified has a `Strict` policy

If you use a preallocation list of media in combination with the `Strict` media policy, the exact media specified in the preallocation list need to be available in the device when a backup is started. If the exact media are not available, a mount request is issued.

Action

To use media available in the device in combination with the preallocation list, modify the media pool allocation policy to `Loose`.

Unexpected Mount Request for a Library Device

There are several situations, described below, that may cause Data Protector to issue a mount request for a library device while media are available in the library.

The media in the library are not formatted and the media pool with the media used for backup has a Strict policy

If your pool uses a `Strict` media allocation policy, unformatted media are not used for backup. If no formatted media are available in the library, Data Protector issues a mount request.

Action

If you would like Data Protector to automatically format unformatted media that are available in the library, set the media pool policy to `Loose`. This can be modified in the media pool `Properties`.

The media in the library are formatted but are different from those in the preallocation list

The media in the library are formatted but are different from those in the preallocation list of the backup specification, and the media pool specified has a `Strict` policy.

If you are using a preallocation list of media in combination with the `Strict` policy and the exact media specified in the preallocation list are not available in the device when backup is started, a mount request is issued.

Action

The exact media specified in the preallocation list need to be available in the device when the backup is started.

To use other media, if available in the device, in combination with the preallocation list, modify the media pool allocation policy to `Loose`.

To use any available media in the device without the preallocation list, remove the preallocation list from the backup specification. Do this by changing backup device options for the backup specification.

Unexpected Mounted Filesystems Detected

When restoring a disk image, you may get a message that the disk image being restored is a mounted file system and will not be restored:

```
Object is a mounted filesystem => not restored.
```

This happens when an application on the disk image leaves some patterns on the disk image. The patterns confuse the system call that verifies whether the eventually mounted filesystem on the disk image is mounted or not, so the system call reports that there is a mounted filesystem on the disk image.

Action

1. Before you start a restore erase the disk image on the Data Protector client with the disk image being restored by entering the following commands:

   ```
   prealloc null_file 65536
   ```

   ```
   dd if=null_file of=<device_file>
   ```

   where *<device_file>* is a device file for the disk image being restored.

2. Start the restore.

Data Protector Fails to Start a Scheduled Session

The scheduled sessions no longer run

The scheduled sessions no longer run since the Data Protector system account, which is supposed to start scheduled sessions, is not in the `Admin` user group on the Cell Manager.

This account is added to the Data Protector `Admin` group on the Cell Manager at installation time. If this is modified and the permission for this account is removed, or if the service account changes, the scheduled sessions no longer run.

Action

Add the Data Protector account to the `Admin` user group on the Cell Manager.

The session fails and Data Protector issues the session status No licenses available.

A backup session is started only after Data Protector has checked the available licenses. Otherwise, the session fails and Data Protector issues the session status `No licenses available`.

Action

Obtain information on available licenses by running the `omnicc -check_licenses -detail` command. Refer to the `omnicc` man page for more information.

Request new licenses and apply them to the Data Protector system. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for licensing details.

Data Protector Backup sessions are not started at all (UNIX-specific)

Action

Run the `crontab -l` command to check whether the `omnitrig` program is included in the `crontab` file. If the following line does not appear, the `omnitrig` entry was automatically added by Data Protector:

```
0,15,30,45 * * * * /opt/omni/sbin/omnitrig
```

Stop and start the Data Protector daemons by running the `omnisv -stop` and the `omnisv -start` commands in the `/opt/omni/sbin` directory.

Data Protector Fails to Start an Interactive Session

Every time a backup is started, permission to start a backup session is required and checked for the user who is currently running Data Protector. If the user does not have sufficient permission, the session cannot be started.

Action

Check and change the user rights for the particular user. Refer to Chapter 4, “Configuring Users and User Groups,” on page 127.

Poor Backup Performance on Novell NetWare Server

Backup performance on a Novell NetWare Server may be poor. Backup does not run continuously, but intermittently. This is a well-known problem caused by the system `TCPIP.NLM`.

Action

Set the following parameters:

- NW5.1/NW6.0: `SET TCP DELAYED ACKNOWLEDGEMENT = OFF`
- NW5.0: `SET TCP DELAYED ACK = OFF`

This increases backup performance without any secondary effects.

Data Protector Fails to Start Parallel Restore Media Agent on Novell NetWare Clients

Data Protector UNIX session manager sometimes fails to start restore Media Agents in parallel on Novell NetWare clients with an error message like, for example, `Could not connect to inet` or `Connection reset by peer`. It is possible that some parallel restore sessions are completed without errors, while other restore sessions are not even started.

Action

A workaround for this problem is to set the `SmMaxAgentStartupRetries` global variable in the Data Protector global options file (located in `/etc/opt/omni/server/options/global`) to `2` or more (max. 50). This variable specifies the maximum number of retries for the session manager to restart the failed agent before it fails. Refer to "Global Options File" on page 613 for more information about the Data Protector global options file.

Novell NetWare Cluster Shared Volumes not Backed up During Full Server Backup

Shared cluster volumes on Novell NetWare cluster are not backed up during full server backups. A possible cause for this problem is improper handling of SMS clustered resources, causing clustered volumes to be skipped because the TSA module was loaded with the cluster support enabled.

Action

Run the `TSAFS /NoCluster` command on the active node to disable the cluster support.

Backup Protection Expiration

When scheduling backups, you have set the same protection period for full and incremental backups, which means that incremental backups are protected for the same duration as the relevant full backup. The consequence of this is that your data will actually only be protected until the full backup expires. You cannot restore incremental backups that have been based on expired full backups.

Action

Configure the protection for your full backups so that they are protected for longer than your incremental backups.

The time difference between the protection for the full backup and the incremental backup should be the amount of time between the full backup and the last incremental backup before the next full backup. For example, if you run incremental backups Monday through Friday and full backups on Saturday, you should set the protection of the full backup to at least 6 days more than for the incremental backups. This will keep your full backup protected and available until your last incremental backup expires.

Troubleshooting Application Database Restores

A poorly-configured DNS environment could cause problems with database applications. If you try to restore a database and it fails with the message `Cannot connect to target database` or `Cannot create restore set`, the problem is as follows:

When backing up the database on a system, the agent that starts on the system logs the system's name to the database as *<system.company.com>*. The Restore Session Manager wants to restore to the *<system_name.company.com>*, but it cannot because it does not know this system as *<system_name.company.com>*, but only as *<system_name>*. The system name cannot be expanded to the full name because the DNS is improperly configured. This situation can also be the other way around, where DNS is configured on the Cell Manager and not on the Application Client.

Action

Set up the TCP/IP protocol and configure DNS properly. Refer to Appendix B in the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information.

Problems with non-ASCII Characters in File Names

In mixed platform environments, there are some limitations regarding handling of file names containing non-ASCII characters in Data Protector GUI, if IDB has not yet been converted to a new internal character encoding. For more information, refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

Action

Convert the IDB to the new internal character encoding and then upgrade Disk Agents on your clients.

Action

If you do not perform the conversion of the IDB, the workaround for trees that cannot be selected for backup or restore is to select a tree above the desired tree, assuming that this parent tree is one that can be successfully specified (for example, its name consists of ASCII characters only).

For backups, this means that more data will be backed up. Usually, this is not an issue since typically entire disks or at least major trees are backed up (for example `/home` or `\My Documents`).

For restores, you can choose to restore the parent tree to a new location, using the `Restore as` or `Restore to new location` option, to prevent any damage by restoring more than just the desired file or directory.

For restores, when in doubt, you are encouraged to restore one tree/file per restore session. A message `"Nothing restored"` will give you a clear indication that the tree was not restored. There are other possible reasons, for example if default file conflict handling is used (`Keep most recent`) this message may also indicate that the files are already on the disk and were not overwritten. `Restore Into` option on the other hand would restore the files into the specified path. When only a few files are restored you may also use the `List restored data` option.

Refer to online Help index keyword "`internationalization`" for internationalization limitations tables.

File Library Device Disk Full

When using a file library device, you may receive a mount request with the following message:

`There is no disk space available for file library "`*`File Library Device`*`". Please add some new disk space to this library.`

Action

You need to create more space on the disk where the file library is located. It is possible to create disk space in any of the following ways:

- Create space on the disk where the files are being backed up.
- Add more disks to the system where the file library device resides.

Files Are Restored With a Wrong File Name After IDB Conversion

IDB conversion has already been performed, but files are restored with a wrong file name.

Action

If IDB data for a specific Data Protector client has already been converted, upgrade the Disk Agent on that client prior the restore, if file names include non-ASCII characters. If you do not upgrade the Disk Agent, the file may be restored with a wrong file name. You can see the status of IDB conversion in the Data Protector GUI in the `Monitor` Context List.

Intermittent “Connection Refused” Error Messages

The backup session aborts with a critical error message:

```
Cannot connect to Media Agent on system
computer.company.com, port 40005 (IPC Cannot Connect
System error: [10061] Connection refused)
```

This problem can occur if a Media Agent is running on a non-server edition of Windows and the Disk Agent concurrency is set to more than `5`. Due to the TCP/IP implementation on non-server editions of Windows, the operating system can accept only 5 incoming connections simultaneously.

Action

Set the Disk Agent concurrency to `5` or less.

It is recommended to use server editions of Windows for systems involved in intensive backup operations, for example, for the Cell Manager, Media Agent clients, Application Agent clients, file servers, and so forth.

Backup or Restore on a TruCluster Server Is Aborted with a Critical Error

Problem

Backup or restore session aborts with the following error message:

```
Internal error in ("ma/xma/bma.c) => process aborted

This is an unexpected condition and is likely due to a corrupted
media or combination of circumstances involving both this product
and the operating system.
```

This error can occur when:

- The backup device used for backup is configured on a cluster virtual server.
- A filesystem being backed up resides on a cluster virtual server.

Action

Set the following `omnirc` variables on a TruCluster Server:

- `OB2BMANET=1`
- `OB2RMANET=1`
- `OB2RDANET=1`
- `OB2BDANET=1`

For information on the `omnirc` file, see “Using Omnirc Options” on page 615.

Restore Problems if the Cell Manager Is Configured in a Cluster

Problem

A backup with a cluster-aware Data Protector Cell Manager with the `Restart backup of all objects` backup option enabled was performed. A failover occurred during the backup and the backup session was restarted on another cluster node and successfully finished. When trying to restore from the last backup, the following error is reported although the session finished successfully:

```
You have selected a version that was not successfully
completed. If you restore from such a backup, some or all the
files may not be restored correctly.
```

If system times on Cell Manager cluster nodes are not synchronized, it is possible that the failed backup has a newer timestamp than the restarted backup. When selecting data for restore, the last backup version is selected by default, resulting in a restore from the failed backup.

Action

To restore from the last successful backup, select the correct backup version for restore. To prevent such errors, it is recommended to configure a time server on your network. This will ensure automatic synchronization of system times on your Cell Manager cluster nodes.

Restore Fails After Upgrading the MoM Manager

Problem

The following error messages may be displayed: `Unknown internal error`, `Started session manager got bad options`, or `Cannot get information from backup host`.

After upgrading the MoM Manager/CMMDB Server to Data Protector A.05.50, you cannot perform a filesystem or integration *restore* of an older Data Protector client using the Data Protector A.05.50 MoM GUI.

Action

Use either the old MoM GUI for restore or upgrade the clients to Data Protector A.05.50.

Troubleshooting Object Copy Sessions

You may run into the following problems while running object copy sessions:

Fewer Objects Are Copied Than Expected

With post-backup or scheduled object copy, the number of objects that match the selected filters is higher than the number of objects that are actually copied.

The following message is displayed:

```
Too many objects match specified filters.
```

Action

- Tighten the criteria for object version selection.
- Increase the maximum number of objects copied in a session by increasing the `CopyAutomatedMaxObjects` variable value in the global options file.

Not All Objects in the Selected Library Are Copied

With post-backup or scheduled object copy, some objects that reside on media in the selected library are not copied. This happens if an object does not have a complete media set in the selected library.

Action

Insert the missing media into the selected library, or select the library that has a complete media set for these objects.

Mount Request for Additional Media Is Issued

In an interactive object copy session from the Media starting point, you selected a specific medium. A mount request for additional media is issued. This happens if an object residing on the medium spans to another medium.

Action

Insert the required medium into the device and confirm the mount request.

Troubleshooting Data Protector Installation

If you run into problems while installing the Data Protector software, check the system's log files on UNIX and setup log files on Windows to determine the problem:

System	Log File
UNIX (local installation)	`/var/adm/sw/swinstall.log` `/var/adm/sw/swagent.log`
UNIX (remote installation)	`/var/opt/omni/log/IS_install.log`
Windows (local installation)	*<System_disk>*`:\`*<Temp>*`\OB2_Setup_ui_`*<Date>*`_`*<Time>*`.txt`
Windows (remote installation)	*<System_disk>*`:\`*<Temp>*`\OB2_Setup_exe_`*<Date>*`_`*<Time>*`.txt`

If the setup log files were not created, run the installation with the `-debug` option.

Problems with Remote Installation of Windows Clients

When using Data Protector remote installation to update Windows clients, you get the following error:

```
Error starting setup process, err=[1326] Logon failure:
unknown user name or bad password.
```

The problem is that the `Data Protector Inet` service on the remote computer is running under a user account that does not have access to the OmniBack II share on the Installation Server computer. This is most probably a local user.

Action

Change the user name for the `Data Protector Inet` service that can access the OmniBack II share.

Name Resolution Problems when Installing the Windows Cell Manager

During the installation of the Data Protector Cell Manager on Windows, Data Protector detects and warns you if the DNS or the LMHOSTS file is not set up as required. In addition, Data Protector notifies you if the TCP/IP protocol is not installed on your system.

Name resolution fails when using DNS or LMHOSTS

If the name resolution fails, the "`error expanding hostname`" message is displayed and the installation is aborted.

- If you encounter resolution problems when using DNS, you get a warning message about your current DNS configuration.
- If you encounter resolution problems when using LMHOSTS file, you get a warning message to check your LMHOSTS file configuration.
- If you have not configured either DNS or LMHOSTS, you get a warning message to enable the DNS or the LMHOSTS resolution in the TCP/IP properties dialog.

Action Check your DNS or LMHOSTS file configuration or activate it. Refer to "Hostname Resolution Problems" on page 674.

The TCP/IP protocol is not installed and configured on your system

If the TCP/IP protocol is not installed and configured on your system, the installation is aborted.

Data Protector uses the TCP/IP protocol for network communications; it must be installed and configured on every client in the cell.

Action Check the TCP/IP setup. For detailed steps, refer to the online Help index keyword "checking, TCP/IP setup".

Troubleshooting User Interface

This section describes a solution to the following problem that can arise while using the Data Protector graphical user interface (GUI):

Corrupted Names of GUI Objects in the Data Protector GUI on UNIX

Problem

Names of GUI objects in the Data Protector GUI (such as backup devices and backup specifications) on UNIX appear corrupted.

If these GUI objects were created under a certain locale, they may appear corrupted when viewed in a different locale. Regardless of the corrupted display of the GUI object names, such GUI objects are still usable.

For example, you configured a backup device and named it using non-ASCII characters. In this case, it's name may appear corrupted, if GUI is run in a locale that uses only ASCII. Although it's name appears corrupted in the GUI, you can still perform backups and restores using this device.

Action

You can either recreate these objects in a locale that uses UTF-8 encoding or still use the old locale on the system, where the Data Protector GUI is running (but then you will not be able to switch encodings in GUI and thus use the internationalization features of Data Protector).

Troubleshooting User Interface Startup

Data Protector user interface start-up problems are usually the result of services not running, services not being installed, or problems with network communication.

Inet Is Not Responding on the Cell Manager

The following message appears:

```
Cannot access the system (inet is not responding). The Cell
Manager host is not reachable, is not up and running, or has
no Data Protector software installed and configured on it.
```

Action

If communication between the systems is not the problem, check the installation using `telnet`.

It is possible that some components were not or were improperly installed. Review the steps in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

If the installation is correct, run the `omnisv -status` command to check whether the services on the Cell Manager are running properly.

No Permissions to Access the Cell Manager

The following message appears:

```
Your Data Protector administrator set your user rights so
that you do not have access to any Data Protector
functionality.

Contact your Data Protector administrator for details.
```

Action

Contact the Data Protector administrator to add you as a user and give you appropriate user rights in the cell. Refer to Chapter 4, “Configuring Users and User Groups,” on page 127.

Connection to a Remote System Refused on Windows or Novell NetWare

The response of the `telnet <hostname> 5555` command is `Connection refused`.

Action

If the `Data Protector Inet` service is not running on the remote system, run the `omnisv -start` command to start it.

If Data Protector is not installed on the remote system, install Data Protector on the remote system.

Troubleshooting the IDB

This section provides troubleshooting for the following problems using the IDB:

File Names Are Not Logged into the IDB During Backup

When performing backups using Data Protector, file names are not logged into the IDB.

Action

Check if you have selected the `Log none` option for backup.

A possible cause for not logging the file names in the IDB on Windows Cell Manager is also that file name conversion in the IDB was running while the client was being backed up. In this case, the backup was performed using the `No log` option and hence no data was written to the IDB for this client in this session. You can check the session messages of that specific backup session for warnings.

Problems While Running the User Interface

IDB is corrupted

Any of the following messages can be displayed:

- `Database is corrupted.`
- `Interprocess communication problem.`
- `Cannot open Database/File.`
- `Error - Details Unknown.`

Action Recover the IDB. For more information, refer to “Recovering the IDB” on page 494.

The IDB Session Manager is not running on the Cell Manager

If the IDB Session Manager is not running on the Cell Manager when Data Protector tries to access or use the IDB, the `Interprocess communication problem error` message is displayed.

- On Windows Cell Manager, the Data Protector process `dbsm.exe` is not displayed among the processes in the Windows Task Manager.
- On UNIX Cell Manager, the `/opt/omni/lbin/dbsm` is not displayed when listing the Data Protector processes using the `ps -ef | grep omni` command.

Action Close and restart the Data Protector GUI.

Libraries (Executables) Missing

On Windows Cell Managers, the following library files should exist in the *<Data_Protector_home>*`\bin` directory:

- `libob2ecmn.dll, libob2eadm.dll, libob2ecdb.dll, libob2emmdb.dll, _eadm32.dll, _erdm32.dll`

On UNIX Cell Managers, the following library files should exist in the `/opt/omni/lib` directory:

- `libob2ecmn.sl, libob2eadm.sl, libob2ecdb.sl, libob2emmdb.sl, _eadm.sl, _erdm.sl`

The RDS service/process cannot be started

If one or several shared library files are missing, the `omnisv -status` command informs you that the RDS service/process is down, while all other services/processes are running.

Action

Reinstall Data Protector and reboot your Cell Manager. This will reinstall the shared libraries and restart the RDS service/process.

Data Files (Directories) Missing

The following IDB data files (directories) should exist in *<Data_Protector_home>*`\db40` (on the Windows Cell Manager) and in the in `/var/opt/omni/server/db40` (on the UNIX Cell Manager):

- `datafiles\catalog`
- `datafiles\cdb`
- `datafiles\mmdb`
- `dcbf`
- `logfiles\rlog`
- `logfiles\syslog`
- `meta`
- `msg`

One or several IDB data files or directories are missing

If one or several IDB data files or directories are missing, the following errors are displayed when Data Protector tries to access or use the IDB:

- `Database network communication error`
- `Cannot open database/file`

Action

Reinstall Data Protector and reboot your Cell Manager. This will reinstall the IDB data files and directories.

Temporary Directory Missing

The following temporary directories should exist on the Cell Manager:

- On Windows: <*Data_Protector_home*>\tmp
- On UNIX: /var/opt/omni/tmp

The Data Protector GUI cannot connect to the Cell Manager

When Data Protector GUI tries to connect to the Cell Manager, the following error message is displayed if Data Protector temporary directory is missing:

```
Cannot access the Cell Manager system. (inet is not
responding) The Cell Manager host is not reachable or is not
up and running or has no Data Protector software installed
and configured on it.
```

Action

1. Close the Data Protector GUI.
2. Run the `omnisv -stop` command on the Cell Manager to stop the Data Protector services/processes:
   - On Windows: <*Data_Protector_home*>\bin\omnisv -stop
   - On UNIX: /opt/omni/sbin/omnisv -stop
3. On the Cell Manager, manually create the temporary directory:
   - On Windows: <*Data_Protector_home*>\tmp
   - On UNIX: /var/opt/omni/tmp
4. Run the `omnisv -start` command to start the services/processes.
   - On Windows: <*Data_Protector_home*>\bin\omnisv -start
   - On UNIX: /opt/omni/sbin/omnisv -start
5. Restart the Data Protector GUI.

Problems During Backup and Import

The BSM or RSM is terminated during the IDB backup or import session

If the BSM or RSM get terminated during the IDB backup or import session, the following error message is displayed:

```
IPC Read Error System Error: [10054] Connection reset by
peer
```

In the Internal Database context, the session status of the IDB backup or import session is still marked as In progress but the session is actually not running.

Action

1. Close the Data Protector GUI.
2. Run the omnidbutil -clear command to set the status of all sessions that are actually not running but are marked as In Progress or Failed, to Failed.
3. Run the omnidbutil -show_locked_devs command to see if any devices and media are locked by Data Protector.
4. If there are, run the omnidbutil -free_locked_devs to unlock them.
5. Restart the Data Protector GUI.

The MMD is terminated during the IDB backup or import session

If the media management daemon MMD is terminated during the IDB backup or import session, the following two error messages are displayed:

- Lost connection to MMD
- IPC Read Error System Error: [10054] Connection reset by peer

Use the following methods to check whether the MMD services/processes are running:

- The omnisv -status command informs you that the MMD service/process is down.
- On UNIX, the Data Protector MMD (/opt/omni/lbin/mmd) is not displayed when listing the Data Protector processes using the ps -ef | grep omni command.

  On Windows, the Data Protector MMD process (mmd.exe) is not listed among processes in the Windows Task Manager.

Action

1. Close the Data Protector GUI.
2. Run the omnisv -stop command to stop the Data Protector services/processes.
3. Run the omnisv -start command to start the Data Protector

services/processes.

4. Run the `omnisv -status` command to check if all the services/processes are running.

The DC binary files are corrupted or missing

If the DC binary files are corrupted or missing, the error message `Open of Detail Catalog Binary File failed` is displayed when browsing backed up objects in the `Restore` context.

- The `omnidbcheck -bf` command reports that one or several DC binary files are missing or of incorrect size, or the `omnidbcheck -dc` command reports that one or several DC binary files are corrupted.
- The `debug.log` file on the Cell Manager, located in the *<Data_Protector_home>*`\log\debug.log` (Windows systems) or in the `/var/opt/omni/log/debug.log` (UNIX systems) contains one or several entries on Data Protector not being able to open a DC binary file.

Action

Recreate DC binary files by importing catalog from media. For more information refer to “Handling Minor Database Corruption in the DCBF Part” on page 498.

Performance Problems

The number of IDB objects and IDB objects' sizes are too large

When browsing object versions and single files for restore, it can take a long time before the information is read from the IDB and displayed.

Action

Set the time interval, which will be used when browsing object versions for restore. You can change this time interval in the `Restore` context when searching for the specific object version you want to restore.

Set the *default* time interval used when browsing object versions for restore.

1. In the Data Protector GUI, click the `File` menu and then click `Preferences`.
2. Click the `Restore` tab and in the `Search interval` drop-down list, select the search interval. Select `Interval` if you want to set an absolute search interval, or `None` if you want all object versions to be

listed.

3. Click OK to apply the change.

The IDB Is Running Out of Space

A part of the IDB is running out of space. The IDB Space Low or IDB Tablespace Space Low notification is issued.

Action

Extend the IDB size. For more information, refer to "Extending the IDB Size" on page 485.

MMDB and CDB Are Not Synchronized

The MMDB and CDB may not be synchronized when the following is true:

- The MMDB and CDB contain information from different periods in time. This may be the result of importing the CDB and the MMDB (the omnidbutil -readdb command) from files generated in separate export (the omnidbutil -writedb command) sessions.
- In a MoM environment, when the local CDB and CMMDB are not synchronized. This may be the result of the CMMDB restore.

Data Protector reports when an object in the IDB has no medium assigned or when the data protection for a medium is not correctly set.

Action

In a one-cell environment:

- Run the omnidbutil -cdbsync *<Cell_Server_Hostname>* command in the /opt/omni/sbin (UNIX Cell Manager) or in the *<Data_Protector_home>*\bin (Windows Cell Manager) directory to synchronize the MMDB and CDB.

In a MoM environment:

- Run the omnidbutil -cdbsync *<Cell_Server_Hostname>* command in the /opt/omni/sbin (UNIX Cell Manager) or in the *<Data_Protector_home>*\bin (Windows Cell Manager) directory with the CMMDB installed (MoM). Run this command for every Cell Manager in the MoM environment by specifying its hostname as the argument.

IDB Purge Performance Problems

Problem

The file versions purge in the IDB is extremely slow.

Action

Check if the following message is logged for the current purge session in the *<Data_Protector_home>*\log\server\purge.log file:

```
Multiple passes needed. This will decrease the performance
of the purge session. To improve performance increase the
amount of memory a purge session is allowed to use.
```

If the log file contains this message, abort the session and increase the value of the PurgeBufferSize option in the global options file. Refer to "Global Options File" on page 613 for information on how to edit the global options file. Then restart the purge session.

IDB Fails Due to Memory Allocation Problems on HP-UX

Problem

The RDS service fails on HP-UX during IDB maintenance or query operations because of memory allocation problems.

Action

Perform the following steps:

1. Set the following environmental variable in the omnirc file on the Cell Manager: _M_ARENA_OPTS=1:32. Refer to "Using Omnirc Options" on page 615.
2. Restart the Data Protector services. Refer to the omnisv man page for details.

Troubleshooting Reporting and Notifications

If you use Outlook XP or Outlook 98/2000 with the latest security patch installed, you following problem appears: when you add a report to a report group specifying email as a send method, and then try to start a report group, GUI hangs. The same happens if you configure a notification and select the email send method. The cause of the problem is that Outlook requires user interaction before sending an email notification. This feature cannot be disabled since it is a part of the Outlook security policy. To solve this problem, start a report from the CLI:

```
omnirpt -report licensing -email <email_address>
```

When a warning asking whether you allow sending email on your behalf appears, click `Yes` to receive a notification.

For more information on how to customize security settings, refer to *HP OpenView Storage Data Protector Software Release Notes*.

Troubleshooting Data Protector Online Help

Data Protector online Help consists of two parts: Help Topics and the Help Navigator. Help Navigator is context-sensitive help, explaining screens and options in the Data Protector GUI, while Help Topics provide conceptual information, procedure instructions, and examples.

The Help system you use depends on the platform (Windows or UNIX) on which you are running Data Protector. You use HTML Help on Windows systems and WebHelp on UNIX systems.

Troubleshooting Online Help on Windows

When accessing online Help on Windows systems, you can run into the following Help Navigator display problem:

The Help Navigator contents do not change in parallel with the Data Protector windows.

Action

1. If you use `Microsoft HTML Help` mode (default option), ensure that the button shown below is enabled.

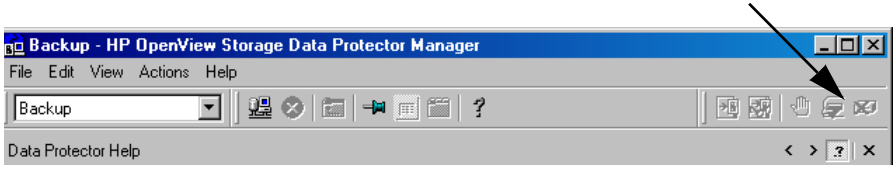

2. If you use `Default HTML Browser` mode (an external HTML browser for displaying the help files) go to `File` menu, click `Preferences` and enable the `Check the box to enable the context-sensitive help navigator` option. Then restart the Help Navigator.

Troubleshooting Online Help on UNIX

Online Help Start and Display Problems

Data Protector supports Mozilla and Netscape Navigator for online Help viewing. If your browser (HTML viewer) is not properly set, you can run into online Help start and display problems. You need to set the browser as follows:

Action

1. In the `File` menu, click `Preferences`, and then `Settings` to open the `HTML Viewer Settings` window.
2. In the `Location of executable script or binary file` text box, enter the location of your browser (for example, `/opt/mozilla/bin` for Mozilla or `/opt/netscape` for Netscape Navigator).
3. In the `Command to start viewer` text box, enter the command that will start the browser. For Netscape Navigator, enter `netscape $HTML$` or `mozilla &HTML$` for Mozilla.

WebHelp on Mozilla

Search functionality does not work properly in the Data Protector WebHelp on Mozilla with the default Mozilla security settings. To enable search, it is recommended to create a new Mozilla profile and use it only for viewing Data Protector WebHelp.

Perform the following steps:

1. Run the `Mozilla Profile Manager` by executing the following command: `/opt/mozilla/mozilla -profilemanager`.
2. Create a new profile named `Data Protector WebHelp` and start Mozilla using this profile.
3. In the `Edit` menu, select `Preferences` and then expand `Privacy & Security`.
4. Click `SSL` and deselect the `Sending form data from an unencrypted page to an unencrypted page` warning option. Click `OK`.

The changed security option will be saved only in the newly created profile without changing other user profiles. This profile will enable you to search in the Data Protector WebHelp without compromising your system security, provided that you use it only for displaying Data Protector WebHelp.

Troubleshooting ADIC/GRAU DAS and STK ACS Libraries Installation and Configuration

The section addresses the following ADIC/GRAU DAS and STK ACS Libraries installation and configuration issues:

ADIC/GRAU DAS Library Installation Failed.

Problem

Your ADIC/GRAU DAS library installation failed.

Action

Perform the following steps:

1. Install a Media Agent on the system controlling the GRAU robotics (PC/robot).
2. Install a Media Agent on the PCs where a drive is connected (PC/drive).
3. Copy `aci.dll + winrpc.dll + ezrpcw32.dll` to `winnt\system32` and *<Data_Protector_home>*`\bin` directory.
4. Create `aci` directory on PC/robot.
5. Copy `dasadmin.exe` to this directory.
6. Copy `portmapper` and `portinst` to `aci` directory.
7. Start `portinst` to install `portmapper` (only on PC/robot).
8. Install `mmd` patch on the CM.
9. The PC needs to be rebooted; then open Windows `Control Panel`. Go to `Administrative Tools, Services` and check if `portmapper` and both rpc services are running.
10. Go to the OS/2 PC within the GRAU library, edit the `/das/etc/config` file:

`cd /das/etc/`

execute: `"e config"`

Within this config file you need to add a client called OMNIBACK containing the IP address of the PC/robot.

Not Possible to See Drives

Problem

You cannot see any drives.

Action

Execute the following commands from PC/robot:

1. `dasadmin listd`
2. `dasadmin` all DLT7000 UP *<AMUCLIENT>*
3. `dasadmin mount` *<VOLSER>* (then you need to push the `UNLOAD` button on the drive)
4. `dasadmin dismount` *<VOLSER>* or `dasadmin dismount -d` *<DRIVENAME>*)

Where:

- *<AMUCLIENT>* = `OMNIBACK`
- *<VOLSER>* for example = `001565`
- *<DRIVENAME>* is for example = `DLT7001`
- `all` stands for allocate

If you are not successful with these commands (communication to DAS Server (OS/2), try to execute these commands on the OS/2 PC. You can find the `dasadmin` command in `/das/bin/` directory.

If you execute theses commands from the OS/2 PC, use *<AMUCLIENT>* = `AMUCLIENT`.

1. Login to the AMU client. The common login are the following:

   `user: Administrator pwd: administrator`

   `user: Supervisor pwd: supervisor`

2. It may be necessary to set the media type:

   set `ACI_MEDIA_TYPE` set `ACI_MEDIA_TYPE=DECDLT`

3. To reboot the library, proceed as follows:

   Shutdown OS/2 and then switch off robotics.

   Restart OS/2 and when OS/2 is ready, the AMU log will display that the robotics is not ready. Then, switch on robotics.

GRAU CAPs not Configured Properly

Problem GRAU CAPs are not configured properly.

Action You can only move media from the CAP to a slot and then to a drive, using the device's robotics. You have to use `import` and `export` commands. For example:

```
import CAP: I01
import CAP range: I01-I03
export CAP: E01
export CAP range: E01-E03
```

The Library Operations Fail

Problem The libraries cannot be used (operations fail).

Action The following syntax is used when you use the Data Protector `uma` utility to manage the GRAU and STK library drives:

`uma -pol <POLNUMBER> -ioctl <LIBRARYNAME> -type <MEDIATYPE>`

Where:

- pol `8` for GRAU
- pol `9` for STK

For example: `uma -pol 8 -ioctl grauamu`

The default media type is DLT.

Check Whether Data Protector Functions Properly

The following sections provide an overview of the Data Protector Checking and Maintenance Mechanism and an overview of things to be checked in order to determine whether Data Protector is properly configured in your backup environment.

Data Protector Checking and Maintenance Mechanism

Data Protector provides its own checking and maintenance mechanism, which is performing the following checking and maintenance tasks on a daily basis:

Maintenance Tasks

- Deletes obsolete DC binary files, sessions, and related messages every day at 12:00 (Noon) by default.
- Finds any free (unprotected) media in pools with the `Use free pool` and `Move free media to free pool` options set and deallocates the found free media to a free pool by issuing the following command every day at 12:00 (Noon) by default:

  ```
  omnidbutil -free_pool_update
  ```

  For more information on the `omnidbutil` command, refer to the `omnidbutil` man page. For more information on the above mentioned options, refer to Chapter 5, “Managing Media,” on page 143.

Checks

Every day at 12:30 P.M. by default, starts checks for the following Data Protector notifications:

- IDB Space Low
- IDB Tablespace Space Low
- Not Enough Free Media
- Health Check Failed
- User Check Failed
- Unexpected Events

- License Will Expire
- IDB Purge Needed

Licensing Check

Every day at 12.30 P.M. by default, Data Protector checks the Data Protector licenses and in case of missing or expired licenses reports them to the Data Protector Event Log.

For more information on Data Protector notifications, refer to “Data Protector Notifications” on page 414. Any notification that is triggered is, by default, sent to the Data Protector Event Log. For more information on the Data Protector Event Log, refer to “Data Protector Event Log” on page 430.

The default schedule values for maintenance tasks and checks can be changed by changing the `DailyMaintenanceTime` and the `DailyCheckTime` options in the Data Protector global options file. Refer to “Global Options File” on page 613 for more information on global options.

The User Check Failed Notification

The `User Check Failed` notification automates the task of checking whether your backup environment is functioning normally. Note that the definition of “normal” depends on your backup environment (backup policy, network configuration, hardware used, etc.). For an overview of items to be checked in an “average” backup environment, refer to “Overview of Items to Be Checked” on page 727. For more information on Data Protector notifications, refer to “Data Protector Notifications” on page 414.

The `User Check Failed` notification executes the command or script entered as an input parameter to this notification and triggers the notification if the return value of the executed command or of any of the executed commands in the script is other than zero. The command/script should be created in the `/opt/omni/lbin` (on UNIX systems) or `<Data_Protector_home>\bin` (on Windows systems) directory of the application system. The `User Check Failed` notification can be configured to be sent using various send methods (e-mail, broadcast message, SNMP traps, log file, etc.) when it is triggered. It can also be configured to start a Report Group when it is triggered.

Thus, scripts containing checks specified in accordance with your backup environment can be developed and configured in a `User Check Failed` notification. Data Protector, using its maintenance and checking mechanism, then prompts you whenever something goes wrong in your backup environment.

All *configured* `User Check Failed` notifications are by default scheduled to be started every day at 00:00 (Midnight) and are, if triggered, sent to Data Protector Event Log.

Overview of Items to Be Checked

In order to ensure that Data Protector is functioning properly and to identify potential problems before they arise, it is recommended that you perform regular checks as described in the following sections.

Using the `User Check Failed` notification, it is possible to automate these checks by developing scripts including these checks. Some of the checks (for example the `omnihealthcheck` and `omnitrig -run_checks` commands) are already automated by the means of Data Protector checking and maintenance mechanism.

Check the Data Protector Cell Manager

1. Run the `omnihealthcheck` command to check the following:

   - whether the Data Protector services (`rds`, `crs`, `mmd`, `omnitrig`, and `OmniInet`) are active
   - whether the Data Protector Media Management Database is consistent
   - whether at least one backup of the IDB exists

   The exit code of the command is 0 (OK) only if all three checks completed successfully (exit code for every check was 0). Exit values other than 0 indicate that one or more of the checks failed.

   For more information on exit codes, refer to `omnihealthcheck` man page.

2. Run the `omnidbcheck -core` command to check the core parts of the IDB.

   The exit code of the command is 0 (OK) only if the check completed successfully. Exit values other than 0 indicate that the check failed.

For more information on exit codes, refer to `omnihealthcheck` man page.

3. Check the critical parts of IDB using the `omnidbcheck -critical` command. For more information on the `omnidbcheck` command, refer to `omnidbcheck` man page.

   The exit code of the command is 0 (OK) only if the check completed successfully. Exit values other than 0 indicate that the check failed. For more information on exit codes, refer to `omnidbcheck` man page.

Check whether backups are configured properly

1. Run the backup preview for crucial backup specifications. Refer to Chapter 6, “Backup,” on page 195 for more information on previewing backups. Successfully completed previews prove that:
   - All clients in the backup specification are accessible from the Cell Manager.
   - All files are accessible.
   - The amount of data to be backed up is determined.
   - All backup devices are configured properly.
2. Run the `omnirpt -report dl_sched` command to check whether the backup specifications are scheduled in compliance with your backup policy. For more information on `omnirpt` command, refer to `omnirpt` man page. The command will list all backup specifications and their schedule.

Verify the Data Protector installation

Verify the installation using the Data Protector GUI, `Clients` context, to check whether the Data Protector software components are up and running on the Cell Managers or the clients. Refer to the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information on how to verify the Data Protector installation.

Inspect the Data Protector log files

Inspect the following Data Protector log files and identify possible problems:

- `event.log`

- `debug.log`
- `purge.log`

For more information on Data Protector log files, refer to “Data Protector Log Files” on page 651.

Run the Notifications Checks

Any Data Protector notification that is triggered is sent to Data Protector Event Log by default. You can also run the `omnitrig -run_checks` command to start checks for the following notifications:

- ✓ IDB Space Low
- ✓ IDB Tablespace Space Low
- ✓ Not Enough Free Media
- ✓ Health Check Failed
- ✓ User Check Failed
- ✓ Unexpected Events
- ✓ License Warning
- ✓ License Will Expire
- ✓ IDB Purge Needed

For more information on Data Protector notifications, refer to “Data Protector Notifications” on page 414. For more information on Data Protector Event Log, refer to “Data Protector Event Log” on page 430.

Check Other System Resources

Inspect the following operating system log files and identify possible problems:

- On UNIX systems: `/var/adm/syslog/syslog.log`
- On Windows systems: inspect the `Windows Event Viewer` and its `Security`, `System` and `Application` logs.

Check whether IDB System Configuration Backups are Being Made Regularly

Check the Data Protector recovery file, `obrindex.dat`, to make sure that the IDB and configuration files, needed for successful recovery of a system, are created regularly. For more information on `obrindex.dat` file, refer to “Preparing for IDB Recovery” on page 466.

15 Integrations with Other Applications

In This Chapter

This chapter gives detailed information on how to integrate the following applications with Data Protector:

For information on integrations with other applications, such as Microsoft SQL, Oracle, and many more, refer to the *HP OpenView Storage Data Protector Integration Guide*. For a list of supported integrations, see the Data Protector documentation overview in the preface of this manual.

NOTE Some functionality is subject to specific Data Protector licenses. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for details.

Cluster Integrations with Data Protector

See the *HP OpenView Storage Data Protector Software Release Notes* for details on the supported cluster software on specific operating systems, level of cluster support and for supported configurations.

See the *HP OpenView Storage Data Protector Concepts Guide* for more information about cluster support and cluster concepts.

See the *HP OpenView Storage Data Protector Integration Guide* for details on Data Protector integrated database applications in a cluster.

Cluster Concepts and Terminology

What Is a Cluster?

A **cluster** is a group of two or more independent computers that appear on the network as a single system. This group of computers is managed as a single system and is designed to:

- Ensure that mission-critical applications and resources are as highly available as possible
- Tolerate component failures
- Support either the addition or subtraction of components

Figure 15-1 shows a typical cluster containing the following components:

Figure 15-1 A Typical Cluster

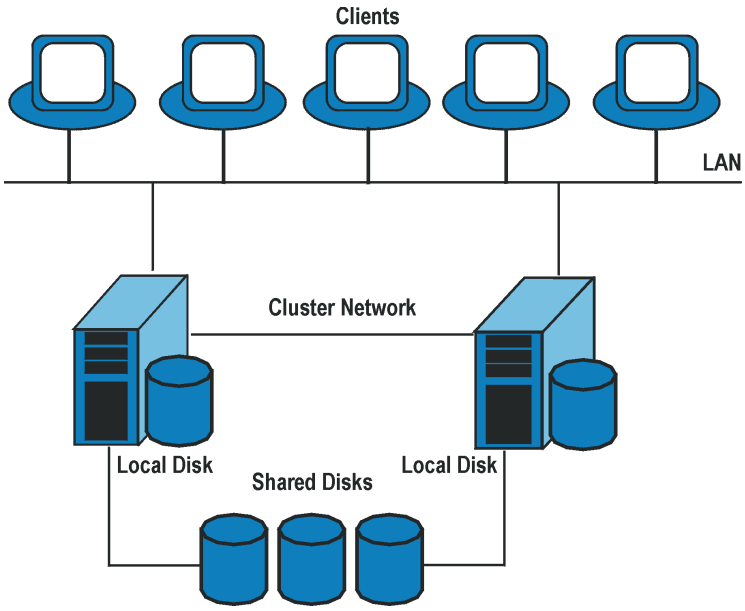

- Cluster nodes (two or more)
- Local disks
- Shared disks (shared between nodes)

Cluster Nodes

Cluster nodes are computers that compose a cluster. They are physically connected to one or more shared disks.

Shared Disks

The **shared disks volumes** (MSCS) or **shared volume groups** (MC/SG) or **shared pools** (Novell NetWare Cluster) contain mission-critical application data as well as specific cluster data needed to run the cluster. In MSCS and Novell NetWare clusters, a shared disk/pool is exclusively active on only one cluster node at a time. In MC/SG clusters, the other node can activate the disk in the read only mode.

Cluster Network

Cluster network is a private network that connects all cluster nodes. It transfers the internal cluster data called **heartbeat of the cluster**. The heartbeat is a data packet with a time stamp that is distributed among all cluster nodes. Each cluster node compares this packet and determines which cluster node is still operational so that appropriate ownership of the **package** (MC/SG, Veritas Cluster) or **group** (MSCS) can be determined.

What is a Package or Group?

A package (MC/SG, Veritas Cluster) or a group (MSCS) is a collection of resources that are needed to run a specific **cluster-aware** application. Each cluster-aware application declares its own critical resources. The following resources must be defined in each group or package:

- Shared disk volumes (MSCS)
- Shared volume groups (MC/SG, Veritas Cluster)
- Network IP names
- Network IP addresses
- Cluster-aware application services

What Is a Virtual Server?

Disk volumes and volume groups represent shared physical disks. A network IP name and a network IP address are resources that define a **virtual server** of a cluster-aware application. Its IP name and address are cached by the cluster software and mapped to the cluster node on which the specific package or group is currently running. Since the group or package can switch from one node to another, the virtual server can reside on different machines in different time frames.

What Is a Failover?

Each package or group has its own preferred node on which it normally runs. Such a node is called a primary node. A package or group can be moved to another cluster node (one of the secondary nodes). The process of transferring a package or group from the primary cluster node to the secondary is called **failover** or switchover. The secondary node accepts the package or group in case of failure of the primary node. A failover can occur for many different reasons:

- Software failures on the primary node
- Hardware failures on the primary node
- The administrator intentionally transfers the ownership because of maintenance on the primary node

NOTE In MSCS environment, Cluster Service components (for example, Database Manager) maintain a coherent image of the central cluster database, which stores information regarding changes in the status of a node, resource, or group. Cluster database must be stored on the cluster's shared disk volume.

Cluster-Aware Databases and Applications

Data Protector integrates with cluster-aware applications that have already been installed on the cluster as virtual servers, by using the application's virtual server configuration.

To back up the cluster-aware application, use its virtual server name when configuring the backup specification.

Microsoft Cluster Server Integration

As a part of its high-availability functionality and support, Data Protector provides an integration with the Microsoft Cluster Server (MSCS). See the *HP OpenView Storage Data Protector Software Release Notes* for details on the supported cluster software on specific operating systems, level of cluster support and for supported configurations.

NOTE

This section provides specific information for integration of Data Protector and Microsoft Cluster Server.

It is assumed that you are familiar with clustering concepts and concepts related to the Microsoft Cluster Server.

Refer to the following manuals for more information:

- Microsoft Cluster Server online documentation.
- *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information on how to install Data Protector.
- *HP OpenView Storage Data Protector Software Release Notes* for last minute information on the current Data Protector release.

Licensing and Microsoft Cluster Server

When you purchase a license for the Data Protector Cell Manager, note that the license will be bound to the virtual server and will work regardless of which physical node inside a Microsoft Cluster Server runs the Data Protector Cell Manager.

The integration is provided on two levels, Cell Manager or client:

- The Data Protector Cell Manager can be installed on the Microsoft Cluster Server, thus providing higher availability of the Data Protector Cell Manager.
- Data Protector cluster client supports a filesystem backup in a cluster environment and backup of the cluster-aware applications.

Cell Manager on Microsoft Cluster Server

The Data Protector Cell Manager can be installed on the 32-bit Microsoft Cluster Server. This enables an automatic migration of the Data Protector services from one cluster node to another in case of failover.

Installation

See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information on how to install Data Protector cluster Cell Manager.

After setup finishes, the Data Protector cluster cell has the following systems automatically added:

- All cluster nodes
- All cluster virtual servers

Clients on Microsoft Cluster Server

Data Protector can back up a full cluster (local and shared disks) and applications running in a cluster environment.

Installation

To back up a cluster-aware application the Data Protector client software must be installed locally on all the cluster nodes. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information on how to install a cluster-aware client.

Configuration

After the installation, virtual server hostname of the client must be imported to the Data Protector cell. See the Figure 15-2 on page 739 and the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.

NOTE

If you want an application backup to be cluster-aware, that is, access it through its virtual server, also this application integration module has to be installed on each application preferred owners (nodes). Only this way the Data Protector integration agents can start on cluster nodes where the application currently resides.

Figure 15-2 **Importing Cluster Virtual Server Hostnames to a Cell on Microsoft Cluster Server**

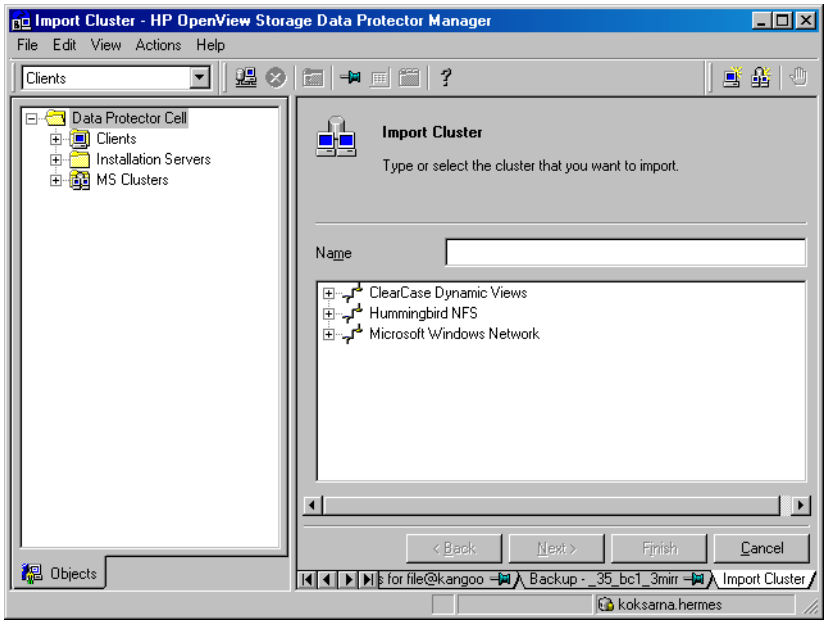

Backing Up Data in a Cluster (MSCS)

When backing up data that reside on cluster node disks, you need to distinguish between:

- Local cluster node disks
- Shared cluster node disks

In the Data Protector GUI, you can see only local disks listed for each cluster node. On the other hand, you can see cluster virtual server items that contain only shared disks for the group in which they are defined. This prevents creation of a backup specification for backing up shared disks. Such backup would fail in case the shared disks are not available on a specific cluster node.

To distinguish between local cluster node disks and shared cluster node disks, Data Protector queries the MSCS database for a list of physical cluster disk resources. All cluster disks presented as proprietary cluster disk resources (e.g. NetRAID 4 disk type) are treated as local cluster node disks.

However, when creating a backup specification, you can see three or more systems that can be backed up:

- Primary node (selected when backing up local disks)
- Secondary node(s) (selected when backing up local disks)
- Virtual server(s) (selected when backing up shared disks)

Backing Up Local Disks

To back up cluster local disks, proceed as follows:

1. Install and configure the Data Protector Disk Agent and cluster component on each cluster node that has the local disks you want to back up.
2. Configure a backup specification for specific cluster node and select which of its local disks you want to back up.

Backing Up Shared Disks

To back up cluster shared disks, proceed as follows:

1. Install (locally) the Data Protector cluster client software on each cluster node. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.
2. Import virtual server hostname (Microsoft Cluster Server) to the Data Protector cell. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.
3. Configure a backup specification for the virtual server and select the shared disks you want to back up.

Managing Cluster-Aware Backups

In the Data Protector cluster Cell Manager, the backup session is cluster-aware. You can set options that define backup behavior if a failover of Data Protector or other cluster-aware applications occurs.

Failover of Data Protector

If a failover of the cluster-aware Data Protector occurs during backup, all running and pending backup sessions fail. In the Data Protector GUI and in the backup specification, you can set one of the options that define automatic backup session restart at failover of Data Protector. See Figure 15-3 on page 742.

Automating Restart of Failed Sessions To modify a backup specification, either filesystem or integration, so that the running backup sessions are automatically restarted at failover of the Cell Manager, perform the following steps:

1. In the `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context, expand the `Backup Specifications` item, and select the backup specification that you would like to modify.
2. In the Results Area, click `Options`.
3. Under the `Backup Specification Options`, click `Advanced`.
4. In the `Backup Options` window, click `Clustering` and select one of the `Automatic session restart` options.

Figure 15-3 Advanced Backup Specification Options-Clustering

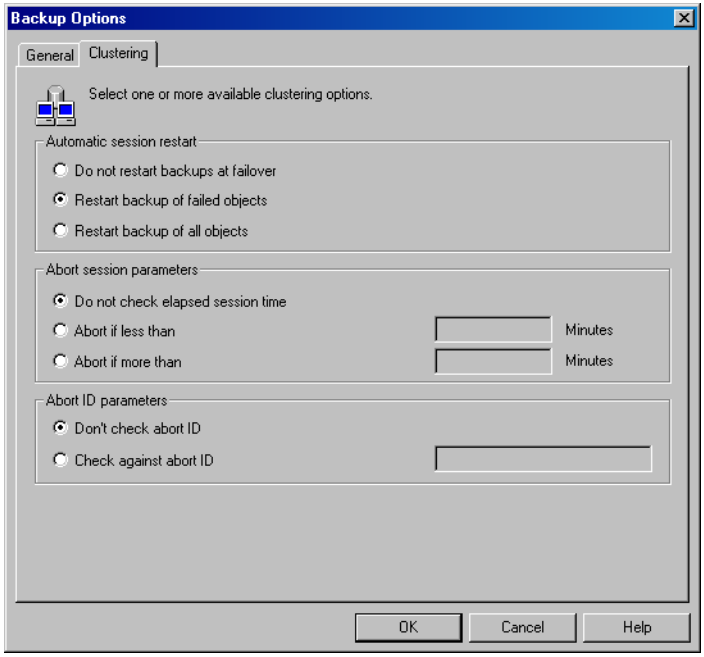

Do not Restart Backups At Failover

When the `Do not restart backups at failover` option is selected, sessions that failed are not restarted. This is the default option.

Restart Backup of Failed Objects

The `Restart backup of failed objects` option is only valid for a filesystem backup specification and specifies that completed objects within the filesystem backup specification will not be restarted. Only objects that failed (running or pending at the moment of the failover) will be restarted. This can minimize the backup time in case failover occurs after some backup objects have been completed.

Restart Backups of All Objects

The `Restart backups of all objects` option is valid for both filesystem and integration backup specifications. When this option is selected, the entire session will be restarted after failover, including the objects that have been completed.

Failover of Application Other Than Data Protector

As the Data Protector cluster Cell Manager is a storage application within a cluster environment, it has to be aware of other applications that might be running within the cluster. If they are running on a node other than Data Protector and if some application fails over to the node where Data Protector is running, this will result in a high load on this node. The node that previously managed only backup operations has now to handle critical application requests as well. Data Protector allows you to define what should happen in such a situation so that the critical application data is protected and the load is balanced again. You can:

- Abort all running backup sessions
- Abort specific running backup sessions
- Inhibit the Data Protector cluster Cell Manager for a specific time frame

Aborting All Running Sessions If the backup is less important than the application, Data Protector can automatically abort all running sessions to balance the load after failover of the application.

To define this option use the `omniclus` command. This command is used as part of a script that is run when a failover of the application occurs. You need to create this script in advance and define it as a new resource type in the application group.

To create the script that will abort all running sessions at failover of the application other than Data Protector, perform the following steps:

1. In the *<Data_Protector_home>*\bin directory create a batch file with the following command line:

   ```
   omniclus.exe -clus <Data Protector_virtual_server>
   -session * -abortsess
   ```

NOTE The * wild card represents all sessions. It can be replaced with the name of a specific backup specification in order to abort only this specific backup session.

2. Open the Windows Cluster Administrator and add a new resource to the application group. For `Resource type` select `Generic Application`. For `Possible owners` select the node on which this

script will be run. This is the node where Data Protector is running. In the Generic Application Parameters window, enter the path name of batch file (for example, c:\program_files\omniback\bin\clus.bat) and directory of the omniclus command. This command resides in the <Data_Protector_home>\bin directory.

Examples

To abort all running sessions on the server obsv.company.com use the following command line:

```
omniclus.exe -clus obvs.company.com -session * -abortsess
```

To abort only session from a backup specification backup_1 on the server obsv.company.com use the following command line:

```
omniclus.exe -clus obvs.company.com -session backup_1
-abortsess
```

Aborting Running Sessions Based on a Logical ID If a specific running backup session is more important than the application, Data Protector can continue this session. To balance the load after a failover, you can abort all backup sessions except an important one using its abort ID. You define this option by using the Data Protector GUI and scripting.

Proceed as follows:

Data Protector GUI

1. In the Data Protector GUI, modify the backup specification with the following steps:
   a. In the HP OpenView Storage Data Protector Manager, switch to the Backup context, expand the Backup Specifications item, and select the backup specification that you would not like to be aborted at failover of the application.
   b. In the Results Area, click Options.
   c. Under Backup Specification Options, click Advanced.
   d. In the Backup Options window, click Clustering. Select Check against abort ID and enter a backup specification ID that will represent this specification and will be used in the command line.

Command Line

2. In the batch file, modify the omniclus command as follows:

   ```
   omniclus.exe -clus <Data_Protector_virtual_server>
   -session <backup_specification> -abortsess -abortid
   <logical_operator_ID>
   ```

Example

In the Data Protector GUI you have configured a backup specification with abort ID = 10. Use the following command line to abort all backup sessions except one with abort ID = 10 on the server `obsv.company.com`:

```
omniclus.exe -clus obvs.company.com -session * -abortsess
-abortid != 10
```

Aborting Sessions Based on Elapsed Session Time To balance the load after a failover you can abort backup sessions based on how long they have already been running. If a specific running backup session is just ending, Data Protector can continue the session. If the backup session has just started and if it is not important, Data Protector can abort the session. You define this option by using the Data Protector GUI and scripting.

Proceed as follows:

Data Protector GUI

1. In the Data Protector GUI, modify the backup specification with the following steps:
   a. In the `HP OpenView Storage Data Protector Manager`, switch to the `Backup` context, expand the `Backup Specifications` item, and select the backup specification that you would like to be aborted based on elapsed session time.
   b. In the Results Area, click `Options`.
   c. Under `Backup Specification Options`, click `Advanced`.
   d. In the `Backup Options` window, click `Clustering`. Select `Abort if less` than or `Abort if more than` and enter the minutes that will represent this specification. It will be aborted if the specified condition is fulfilled when a failover occurs.

Command Line

2. In the batch file, modify the `omniclus` command as follows:

   ```
   omniclus.exe -clus <Data_Protector_virtual_server>
   -session * -abortsess
   ```

NOTE When the command is run, the elapsed time for each backup specification is checked and the session is aborted if the specified conditions are met. For example, in the Data Protector GUI specify that the backup specification is aborted if it has been running for less than 30 minutes. When the failover occurs and when the `omniclus` command is started, the session is aborted if it has been running for less than 30 minutes, otherwise it continues.

Temporarily Disabling Backup Sessions To balance the load after a failover, you can also disable the Cell Manager for some time. All running session are continuing but you cannot start new backups until the Cell Manager is enabled again. You define this only by using scripting.

Command Line In the batch file, modify the `omniclus` command as follows:

```
omniclus.exe -clus <Data_Protector_virtual_server> -inhibit
minutes
```

Examples To disable new backups on the server `obvs.company.com` for 20 minutes, use the following command line:

```
omniclus.exe -clus obvs.company.com -inhibit 20
```

To disable new backups until the Cell Manager is enabled again, use the following command line:

```
omniclus.exe -clus obvs.company.com -inhibit *
```

To enable backups again, run the following command line in CLI:

```
<Data_Protector_home>\bin\omniclus -clus obvs.company.com
-inhibit 0
```

MC/ServiceGuard Integration

As part of its high-availability support, Data Protector provides a full integration of the Data Protector Cell Manager with MC/ServiceGuard on HP-UX systems. For details on supported operating system versions, supported configurations, and level of cluster support, refer to the *HP OpenView Storage Data Protector Software Release Notes*.

NOTE

This section provides specific information for integration of Data Protector and MC/ServiceGuard.

It is assumed that you are familiar with clustering concepts and concepts related to MC/ServiceGuard.

Refer to the following manuals for more information:

- *Managing MC/ServiceGuard* for more information on MC/ServiceGuard.
- *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information on how to install Data Protector.
- *HP OpenView Storage Data Protector Software Release Notes* for last minute information on the current Data Protector release.

Licensing and MC/ServiceGuard

When you purchase a license for the Data Protector Cell Manager, note that the license will be bound to the virtual server and will work regardless of which physical node inside an MC/ServiceGuard cluster runs the Data Protector Cell Manager, so long as the package is running on one of the nodes.

Cell Manager on MC/ServiceGuard

Prerequisites

- In an MC/ServiceGuard cluster environment, a Data Protector Cell Manager should have its own package. Before installing Data Protector Cell Manager on MC/ServiceGuard, you need to get the following information from your network administrator:
  - Virtual server name (the hostname specified in the cluster package)

— Package IP or virtual IP-address

In addition, you will also need to create a volume group on a shared disk.

- Ensure that the cluster nodes and the package IP (virtual IP) are on the same subnet.
- If you have DNS in your environment, ensure that all the cluster nodes and the package IP are registered with the DNS server.

Installation

Install all hosts in the cluster using the standard procedure for installing the Cell Manager on UNIX as described in the *HP OpenView Storage Data Protector Installation and Licensing Guide*.

IMPORTANT

If you need to add additional software components on cluster nodes using the GUI, make sure that the node to which you add the components is active.

Configuration

Prerequisites for Configuration

Before you start configuring Data Protector with MC/ServiceGuard, check the following:

- The cluster should be installed and running.
- Decide which systems are going to be the Primary Cell Manager and the Secondary Cell Manager(s).
- Systems chosen to be the Primary Cell Manager and the Secondary Cell Manager(s) must have MC/ServiceGuard installed, with recommended patches, and must be configured as members of the same cluster. For instructions on MC/ServiceGuard installation and configuration, refer to the *Managing MC/ServiceGuard* manual.
- Data Protector Cell Manager, with recommended patches, and all other Data Protector software components for the integrations you want to have in the cluster must be installed on the Primary node and each of the Secondary nodes. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.

Configuring the Primary and Secondary Cell Managers

The following sections explain how to configure the Primary and Secondary Cell Managers.

NOTE

The following sections provide step-by-step examples to configure the Primary and Secondary Cell Managers. Directory and file names, numbers, and other variables will differ from the following examples according to your environment.

Configuring the Primary Cell Manager

When configuring the Primary Cell Manager, you should first create a volume group. If you are using ob2 disk as a cluster lock disk, you should already have created a volume group for it. If you are not, follow the steps:

1. Create a volume group on a shared disk accessible to both Cell Managers (for example, `/dev/vg_ob2cm`), with the following steps:

   a. Create a directory for a new volume group:

   ```
   mkdir /dev/vg_ob2cm
   ```

NOTE

The shared volume group will contain the IDB and configuration files. Keep this in mind when considering the size of the shared disk.

   b. List all existing volume groups on the system to look for the next available minor number:

   ```
   ll /dev/*/group
   ```

   c. Create a group file for the volume group:

   ```
   mknod /dev/vg_ob2cm/group c 64 0x010000
   ```

   d. Prepare the disk(s) to be used within the volume group:

   ```
   pvcreate -f /dev/rdsk/c0t1d0
   pvcreate -f /dev/rdsk/c1t2d0
   ```

   e. Create the new volume group:

   ```
   vgcreate /dev/vg_ob2cm /dev/dsk/c0t1d0 /dev/dsk/c1t2d0
   ```

2. Create a logical volume for that group (for example, `/dev/vg_ob2cm/lv_ob2cm`), with the following steps:

   a. Create a new logical volume:

      ```
      lvcreate -L 100 -n lv_ob2cm /dev/vg_ob2cm
      ```

      The number 100 presents the size of the partition in MB. The `etc/opt/omni` and `var/opt/omni` Data Protector directories will be located there.

   b. Create a journaled filesystem on the logical volume:

      ```
      newfs -F vxfs /dev/vg_ob2cm/rlv_ob2cm
      ```

NOTE If you want to mirror the new logical volume, refer to the HP-UX LVM documentation on the configuration steps.

3. Set volume group properties according to the cluster documentation, with the following steps:

   a. Deactivate the volume group from regular mode:

      ```
      vgchange -a n /dev/vg_ob2cm
      ```

   b. Mark the volume group for the cluster use:

      ```
      vgchange -c y /dev/vg_ob2cm
      ```

NOTE If this is a cluster lock disk and you are using a later version of MC/ServiceGuard like 11.09, this is done automatically.

   c. Use the volume group in the exclusive mode:

      ```
      vgchange -a e /dev/vg_ob2cm
      ```

4. Mount the logical volume to a directory (for example, `/omni_shared`), with the following steps:

   a. Create a mount point directory:

      ```
      mkdir /omni_shared
      ```

b. Mount the filesystem to the mount point directory:

```
mount /dev/vg_ob2cm/lv_ob2cm /omni_shared
```

5. Modify the `/etc/opt/omni/server/sg/sg.conf` template file.

IMPORTANT

The `SHARED_DISK_ROOT` variable must contain the name of the mount point directory (for example, `SHARED_DISK_ROOT=/omni_shared`).

The `CS_SERVICE_HOSTNAME` variable must contain the name of the virtual Cell Manager, as it is known to the network. Each package in the cluster requires its own virtual IP address and its virtual server network name (for example, `CS_SERVICE_HOSTNAME=ob2cl.company.com`).

6. Configure the Primary Cell Manager. Make sure not to be positioned in the `/etc/opt/omni` or `/var/opt/omni` directory or their subdirectories when running the script. Make also sure to have no mounted subdirectories in the `/etc/opt/omni` or `/var/opt/omni`. Run:

   ```
   /opt/omni/sbin/install/omniforsg.ksh -primary
   ```

   Note that after running this script, the Data Protector services are stopped and will be restarted later on.

7. Unmount the mount point directory (Data Protector shared directory):

   ```
   umount /omni_shared
   ```

8. Deactivate the volume group you created:

   ```
   vgchange -a n /dev/vg_ob2cm
   ```

9. Export the volume group you created on the Primary Cell Manager with the following steps:

   a. From system1 (Primary Cell Manager) export the LVM configuration information with map file `/tmp/lvm_map`:

   ```
   vgexport -p -m /tmp/lvm_map /dev/vg_ob2cm
   ```

   b. Transfer the map file over to system2 (Secondary Cell Manager):

   ```
   rcp /tmp/lvm_map second_system:/tmp/lvm_map
   ```

Configuring the Secondary Cell Manager

To configure the secondary Cell Manager on system2, proceed as follows:

1. On system2 set up the volume group to be imported, with the following steps:

   a. Create a directory for the volume group to be imported:

   ```
   mkdir /dev/vg_ob2cm
   ```

   b. List all existing volume groups on the system to look for the next available minor number:

   ```
   ll /dev/*/group
   ```

   c. Create a group file for the volume group:

   ```
   mknod /dev/vg_ob2cm/group c 64 0x010000
   ```

   d. Import the volume group with map file `/tmp/lvm_map`:

   ```
   vgimport -m /tmp/lvm_map -v /dev/vg_ob2cm
   /dev/dsk/c0t1d0 /dev/dsk/c1t2d0
   ```

2. Set volume group properties according to the cluster documentation, with the following steps:

   a. Mark the volume group for the cluster use:

   ```
   vgchange -c y /dev/vg_ob2cm
   ```

NOTE If this is a cluster lock disk and you are using a later version of MC/ServiceGuard like 11.09, this is done automatically.

   b. Use the volume group in the exclusive mode:

   ```
   vgchange -a e /dev/vg_ob2cm
   ```

3. Mount the logical volume to the mount point directory, with the following steps:

   a. Create the same mount point directory as you have created on the Primary Cell Manager (`/omni_shared`):

   ```
   mkdir /omni_shared
   ```

   b. Mount the filesystem to the mount point directory:

   ```
   mount /dev/vg_ob2cm/lv_ob2cm /omni_shared
   ```

4. Configure the Secondary Cell Manager:

   ```
   /opt/omni/sbin/install/omniforsg.ksh -secondary
   /omni_shared
   ```

5. Unmount the mount point directory (Data Protector shared directory):

   ```
   umount /omni_shared
   ```

6. Deactivate the volume group you imported:

   ```
   vgchange -a n /dev/vg_ob2cm
   ```

Configuring the Cell Manager Package

NOTE

The following section provides step-by-step examples to configure the Data Protector package. Directory and file names, numbers, and other variables will differ from the following examples according to your environment. The cluster configuration file name `cluster.conf` and the Data Protector package name `ob2cl` is used also as an example. You should follow the names given to you by your network or domain administrator.

Note that the Data Protector daemons are not running anymore on either cluster node.

Prerequisites

- The Data Protector Cell Manager should be installed and configured on both cluster nodes as explained in the previous section.
- Before configuring the Data Protector cluster package, you should have a cluster configuration file created and edited.

Configuring Data Protector Package

On the Primary Cell Manager node proceed as follows:

1. Check the cluster configuration file for errors:

   ```
   cmcheckconf -C /etc/cmcluster/cluster.conf
   ```

   If there are errors, fix them.

   If there are no errors, enable the configuration:

   ```
   cmapplyconf -C /etc/cmcluster/cluster.conf
   ```

2. Start the cluster:

   ```
   cmruncl
   ```

3. Create the directory in the `/etc/cmcluster` directory that will hold the Data Protector package:

   ```
   mkdir /etc/cmcluster/ob2cl
   ```

4. Change to the `/etc/cmcluster/ob2cl` directory:

   ```
   cd /etc/cmcluster/ob2cl
   ```

5. Create a package configuration file in the Data Protector package directory:

   ```
   cmmakepkg -p /etc/cmcluster/ob2cl/ob2cl.conf
   ```

6. Create a package control file in the Data Protector package directory:

   ```
   cmmakepkg -s /etc/cmcluster/ob2cl/ob2cl.cntl
   ```

7. Modify the Data Protector package configuration file (for example, `/etc/cmcluster/ob2cl/ob2cl.conf`). Refer to the example of this file in "Example of the Package Configuration File" on page A-28.

   In this file, modify the following fields:

Modifying the Configuration File

- `PACKAGE_NAME`

  Enter the Data Protector cluster package name. For example:

  ```
  PACKAGE_NAME ob2cl
  ```

- `NODE_NAME`

  Enter the names of the nodes. First enter the name of the primary (original) node, then the name(s) of the secondary node(s). For example:

  ```
  NODE_NAME partizan
  ```

  ```
  NODE_NAME lyon
  ```

- `RUN_SCRIPT`, `RUN_SCRIPT_TIMEOUT`, `HALT_SCRIPT`, `HALT_SCRIPT_TIMEOUT`

  Enter the name of the Data Protector package control file (script) and adjust the time-out for the execution of the script. By default, there is no time-out. For example:

  ```
  RUN_SCRIPT /etc/cmcluster/ob2cl/ob2cl.cntl
  ```

```
RUN_SCRIPT_TIMEOUT NO_TIMEOUT
HALT_SCRIPT /etc/cmcluster/ob2cl/ob2cl.cntl
HALT_SCRIPT_TIMEOUT NO_TIMEOUT
```

- SERVICE_NAME, SERVICE_FAIL_FAST_ENABLED , SERVICE_HALT_TIMEOUT

  Enter the service information. For the service name, you can enter any name but note that you will use the same name in the control file afterwards. For example:

```
SERVICE_NAME omni_sv
SEVICE_FAIL_FAST_ENABLED NO
SERVICE_HALT_TIMEOUT 300
```

- SUBNET

  Enter the subnet of the cluster. For example:

```
SUBNET 10.17.0.0
```

8. Modify the Data Protector package control file (for example, /etc/cmcluster/ob2cl/ob2cl.cntl). Refer to the example of this file in “Example of the Package Control File” on page A-38.

   In this file, modify the following fields:

Modifying the Control File

- VG [n]

  Specify the volume group used by this package. For example:

```
VG [0] = /dev/vg_ob2cm
```

- LV [n], FS [n], FS_MOUNT_OPT [n]

  Specify the logical volume and filesystem mount information:

```
LV [0] = /dev/vg_ob2cm/lv_ob2cm
FS [0] = /omni_shared
FS_MOUNT_OPT[0]=" "
```

- IP, SUBNET

  Specify the IP and the subnet information used by this package. For example:

```
IP [0] = 10.17.3.230
SUBNET [0] = 10.17.0.0
```

- SERVICE_NAME, SERVICE_CMD, SERVICE_RESTART

  Specify the service name, command, and restart parameters.

IMPORTANT

The service name must be the same as that used in the configuration file. The service command (the SERVICE_CMD variable) must be the one used in the example below.

For example:

```
SERVICE_NAME [0] = omni_sv
SERVICE_CMD [0] = "/etc/opt/omni/server/sg/csfailover.ksh
start"
SERVICE_RESTART [0] = "-r 2"
```

To make sure that the Cell Manager package is restarted at failover, set the SERVICE_RESTART parameter to -R (to restart the service for an infinitive number of times; this is not recommended) or to "-r *<number of restarts>*" (to restart the service for defined number of times).

9. Check and propagate the Data Protector cluster package files, with the following steps:

   a. Copy the package control file to other nodes within the cluster:

   ```
   remsh system2 "mkdir /etc/cmcluster/ob2cl"
   rcp /etc/cmcluster/ob2cl/ob2cl.cntl
   system2:/etc/cmcluster/ob2cl/ob2cl.cntl
   ```

   b. Enable the Data Protector shared disk as a cluster volume group (created before) on all cluster nodes:

   ```
   vgchange -c y /dev/vg_ob2cm
   ```

   c. Check the Data Protector package:

   ```
   cmcheckconf -P /etc/cmcluster/ob2cl/ob2cl.conf
   ```

   d. If the check was successful, add the Data Protector package:

   ```
   cmapplyconf -P /etc/cmcluster/ob2cl/ob2cl.conf
   ```

   e. Start the package:

   ```
   cmrunpkg ob2cl
   ```

The cluster should be configured and the Data Protector Cell Manager package should be up and running.

f. Import the virtual server (the hostname specified in the cluster package) manually (for example, by using the `omnicc` command):

```
omnicc -import_host <virtual_hostname> -virtual
```

g. If the Data Protector Installation Server was also installed on the MC/ServiceGuard (default), you have to import this Installation Server (for example, by using the `omnicc` command):

```
omnicc -import_is <virtual_hostname>
```

h. In order to run the Data Protector graphical user interface on the secondary node, you have to open the Data Protector graphical user interface and add the root user of the secondary node to the `admin` user group. Refer to “Adding or Deleting a User” on page 137.

Clients on MC/ServiceGuard

Data Protector can back up a full cluster (local and shared disks) and applications running in a cluster environment.

Installation

To back up a cluster-aware application, the Data Protector client must be installed locally on all the cluster nodes. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for information on how to install a cluster-aware client.

Configuration

You need to import the application virtual server (the hostname specified in the application cluster package) to the cell.

If the Cell Manager and the application are in the same cluster, you need to move the Cell Manager package to the application node before importing the virtual server. Proceed as follows:

1. Stop the Cell Manager package (for example `ob2cl`):

```
cmhaltpkg ob2cl
```

2. Run the Cell Manager package on the application node:

```
cmrunpkg -n <node_name> ob2cl
```

NOTE

When using the Data Protector GUI, import each virtual server as a client. See Figure 15-4 on page 758 and the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.

Figure 15-4 Importing an Application Cluster Package to a Cell on MC/ServiceGuard

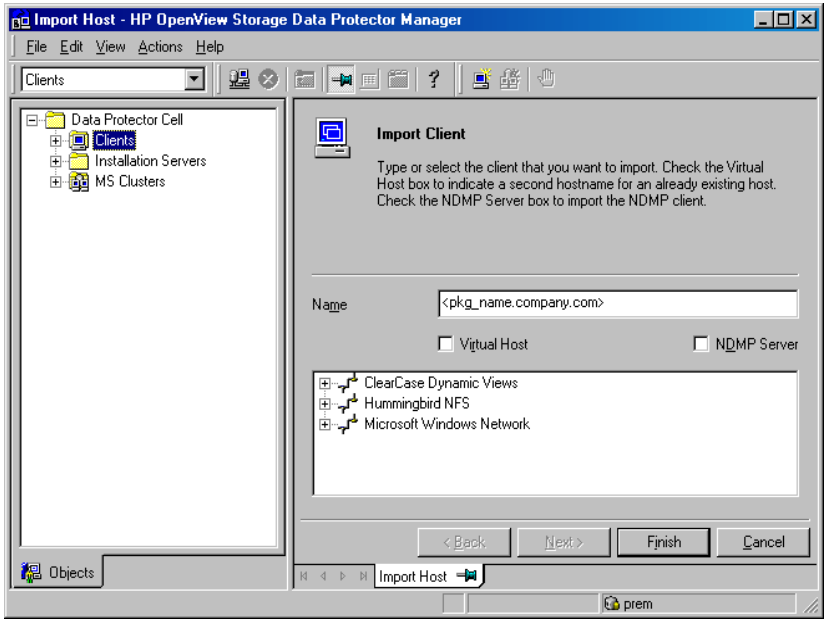

Backing Up Data in a Cluster (MC/SG)

This section provides an overview of how to back up specific data in a cluster environment. For additional information on backing up data in a cluster, see “Backing Up Data in a Cluster (MSCS)” on page 739.

NOTE When backing up a virtual server, the object ownership will acquire the ownership of the stationary host on which the cluster package is running. Therefore, when a failover occurs, the same object backup is showing a different ownership. To avoid this, set the ownership in the backup specification to the virtual server.

Backing Up Local Disks

To back up cluster local disks, proceed as follows:

1. Install and configure the Data Protector Disk Agent component on each cluster node that has the local disk(s) you want to back up.
2. Configure a backup specification for specific cluster node using the physical node name and select which of its local disks you want to back up.

Backing Up Shared Disks

To back up cluster shared disks, proceed as follows:

1. Install (locally) and configure the Data Protector cluster client software on all the cluster nodes. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for instructions.
2. Import the virtual server (the hostname specified in the cluster package) to the Data Protector cell.
3. Configure a backup specification and select the virtual server. Define the shared disks you want to back up.

Veritas Cluster Integration

Clients on Veritas Cluster

Data Protector can only be used to back up local or shared disks in a Veritas Cluster environment.

Cluster aware operation is not supported for Data Protector with Veritas Clusters.

Installation

Data Protector has to be installed locally on each client, and each client has to be imported to the cell. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for step-by-step instructions.

To configure Veritas Cluster with Data Protector, you need the Data Protector user interface.

Refer to the following for more information:

- Veritas Cluster documentation.
- *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information on how to install Data Protector.
- *HP OpenView Storage Data Protector Software Release Notes* for last minute information on the current Data Protector release.

Configuration

To be able to back up local disks on cluster nodes, the individual nodes have to be imported into the Data Protector Cell Manager.

Backing Up Local Disks

Disks local to the systems in the cluster are visible when you browse a system where a disk is locally connected.

To back up local disks:

1. Install the Data Protector Disk Agent on each system with the local disk you want to back up.
2. Configure a backup of the local system in the cluster and define the local disks you want to back up.

Backing Up Shared Disks

A shared disk can only be backed up as a local disk, as described above. It can however be backed up from any of the cluster nodes between which it is shared.

For example, to back up a disk shared between two nodes:

1. Install the Data Protector Disk Agent on each system that shares the disk.
2. Define a backup specification for the disk as a “local disk” on each system.
3. If you want to safeguard the backup of the shared disk further, you could create a post-exec within each backup specification that checks for errors and starts a backup on the other system, if the first fails.

Novell NetWare Cluster Integration

Clients on Novell NetWare Cluster

Data Protector can only be used to back up local disks or cluster shared pools in a Novell NetWare Cluster environment.

Cluster aware operation is not supported for Data Protector with Novell NetWare Clusters. In case of failover, backup or restore sessions have to be restarted manually.

Installation

Data Protector has to be installed locally on each client, and each client has to be imported to the cell. See the *HP OpenView Storage Data Protector Installation and Licensing Guide* for step-by-step instructions.

To configure Novell NetWare Cluster with Data Protector, you need the Data Protector user interface.

Refer to the following for more information:

- Novell NetWare Cluster documentation.
- *HP OpenView Storage Data Protector Installation and Licensing Guide* for more information on how to install Data Protector.
- *HP OpenView Storage Data Protector Software Release Notes* for last minute information on the current Data Protector release.

NOTE

It is not possible to add a device on Novell NetWare virtual server.

Configuration

To be able to back up local disks on cluster nodes, the individual nodes have to be imported into the Data Protector cell. To be able to back up cluster shared pools, virtual server has to be imported into the cell as well.

Backing Up Local Disks

Disks local to the systems in the cluster are visible when you browse a system where a disk is locally connected.

To back up local disks:

1. Install the Data Protector Disk Agent on each system with the local disk you want to back up.
2. Configure a backup of the local system in the cluster and define the local disks you want to back up.

Backing Up Shared Cluster Pools

A cluster shared pool can only be backed up via the virtual server. When the virtual server is selected for backup, only cluster shared pools are displayed as available pools for backup.

For example, to back up a pool shared between two nodes:

1. Install the Data Protector Disk Agent on each system that shares the pools.
2. Import the cluster virtual server into the cell.
3. Create a backup specification that includes all pools on the virtual server and start the backup.

Data Source Integration (DSI)

What Is DSI?

The Data Source Integration (DSI) allows you to use the HP OpenView Performance Agent to log data, define alarms, and access metrics from sources of data other than the metrics logged by the HP OpenView Performance Agent `scopeux` collector. Data Protector provides a sample script and configuration file that show you how to use the Data Protector reporting command-line interface with Data Source Integration to log data about the Data Protector environment, and backup and restore sessions.

What Can You Measure?

Some examples of what can be measured using the DSI integration are:

- Database size
- Media usage
- Media status
- Number of systems
- Amount of data per system
- Full and incremental backup figures.

Overview of Configuration

In order to use DSI, you have to:

- Identify what data you want to log
- Write a script to query data from Data Protector
- Set up a class specification file
- Compile the class specification file
- Start the logging process.

Data Protector provides a sample Korn shell (`ksh`) script and class specification file that, by default, log two metrics: the number of clients in cell and the size of IDB size. The script and class specification file can be easily modified for collecting other information from Data Protector. The scripts are supported on UNIX systems.

Configuring the Integration

To configure the Data Protector DSI integration, follow these steps:

1. Write a script to collect data.

   First select which data you want to log. Data Protector provides a reporting command `omnirpt` located in the `/opt/omni/bin/` directory. This command can be used to gather various information about the Data Protector environment. See the `omnirpt` man page for more information on the command. Secondly, write a script that in an infinite loop queries for the selected data and writes it to standard output.

2. Create the class specification file.

   The class specification file defines what data you want to log and how you want it to be logged. Data Protector provides a sample class specification file `obdsi.spec` in the `/etc/opt/omni/server/dsi` directory. Refer to the DSI manual for the complete syntax of the class specification file.

3. Compile the class specification file.

   Use the `sdlcomp` command from the `/opt/perf/bin` directory to compile the class specification file. In order to compile the Data Protector sample class specification file, use:

   ```
   sdlcomp obdsi.spec OmniBack.log OmniBack
   ```

4. Configure `perflbd.rc`

   Before you start modifying `perflbd.rc` file, you have to stop the mwa services. You do this using the following command:

   ```
   /opt/perf/bin/mwa stop
   ```

   Now you can edit the file `/var/opt/perf/perflbd.rc`. If you are configuring Data Protector sample metrics, add the following line to the file. Note that this has to be added as a single line:

   ```
   DATASOURCE=OMNIBACKII
   LOGFILE=/etc/opt/omni/server/dsi/OmniBack.log
   ```

5. Start the logging process.

   Start the script that collects your data and pipe its output using `dsilog` command. In case of Data Protector sample metrics, use the following command (in one line):

   ```
   obdsi.ksh | /opt/perf/bin/dsilog OmniBack.log OMNIBACKII
   ```

Application Response Measurement (ARM) Integration

What Is the ARM Integration?

Data Protector supports the emerging standard for measuring the response time of transactions in distributed environments, the Application Response Measurement (ARM) interface. Data provided by Data Protector can be used in ARM-compliant system management and monitoring tools, such as HP OpenView Performance Agent. Such tools can log this information for trend analysis, reporting, or alert-based notifications. The collected data can be viewed and analyzed by HP OpenView Performance Manager or some other tool.

How to Install the ARM Integration

For the installation, all you need is the ARM 2.0 compatible RPM agent and the ARM 2.0-compliant library installed on the Cell Manager. It does not matter whether you install them before or after the Data Protector installation.

With a UNIX Cell Manager, you need to replace the dummy library `/opt/omni/lib/arm/libarm.sl` (HP-UX) or `/opt/omni/lib/arm/libarm.so` (Solaris) with the appropriate ARM library that actually logs transactions, or create a link to it. It is recommended to create a link. For example, in case of HP OpenView Performance Agent on an HP-UX 11.x Cell Manager, you need to link the above mentioned file to the `/opt/perf/lib/libarm.sl` file. Note that the `/opt/perf/lib/libarm.sl` file links to `libarm.1`.

Windows Cell Managers require no additional steps for setting up the ARM Integration.

What Can Be Measured?

The following information can be measured with the ARM integration:

- Overall session duration
- Disk Agent read times
- Disk Agent network write times
- Media Agent network read times
- Media Agent data write times
- Session Manager write to database time
- Database purge duration

The following table shows the supported ARM transactions:

Table 15-1 ARM Transactions

Transaction Name	Additional Information	Transaction Description
BS-*<Backup_specification>*	Time	Duration of a backup session
RS-*<Session_ID>*	Time	Duration of a restore session
BO-*<Object_name>*	Time	Duration of a backup of a specific object
DP	Number of purged records and IDB size (MB)	Duration of the IDB purge
DC	IDB size (MB)	Duration of the IDB check

ManageX Integration

What Is the ManageX Integration?

ManageX integration is supported on those Windows systems where ManageX is running. It allows the operator using ManageX to check Data Protector operation and backup status.

What Is Supported?

The integration supports the following:

- Sends the Data Protector messages with the severity levels you choose to the ManageX console.
- Checks if all Data Protector services are running and sends a messages to the ManageX console if one of the services stops.

Configuring the Integration

To configure the ManageX integration with Data Protector perform the following steps:

1. Enable Data Protector message forwarding on the Cell Manager:
   a. In the global file set `EventLogMessages=1`. Refer to “Global Options File” on page 613 for more information.
   b. Stop and restart the Data Protector services.
2. To set the Data Protector severity levels you want to receive in the ManageX console, delete or add them in the *<Data_Protector_home>*`\config\server\managex\filter` file. By default, all severity levels (`normal | warning | minor | major | critical`) are listed in this file.
3. Distribute the policies from the ManageX to Data Protector Cell Manager using the ManageX console. The Data Protector policies are in the folder Backup Applications.

Access Points for System and Management Applications

This section provides information on Data Protector access points for System and Management applications.

Introduction

The Data Protector HP OpenView Integrations allow you to administer, monitor and measure the performance of Data Protector processes using System and Application Management applications such as:

- HP OpenView Vantage Point Operations
- HP OpenView DSI
- HP OpenView ManageX

As a generic interface for these applications, Data Protector provides the following access points:

- SNMP traps
- User Interfaces (Data Protector GUI and CLI, Web reporting interface)
- Data Protector log files
- Windows Application Log

Depending on the application integrated with Data Protector, any or only some of the access points can be utilized. Data Protector already provides a set of predefined reports and actions that can be performed using the applications. They are described in the Chapter 15, “Integrations with Other Applications,” on page 731.

Data Protector Access Points

SNMP Traps

SNMP traps allow a System and Application Management application to receive and process an SNMP trap message when a Data Protector event occurs or when an SNMP trap is sent as a result of Data Protector

checking and maintenance mechanism. For more information on Data Protector checking and maintenance mechanism, refer to “Data Protector Checking and Maintenance Mechanism” on page 725.

On HP-UX and Solaris, there are two Data Protector files residing on the Cell Manager, that specify the behavior of Data Protector SNMP traps:

- `/etc/opt/omni/server/snmp/OVdest`

  This file contains the names of the systems to receive the Data Protector SNMP traps. It has the following format:

  `trap-dest: <hostname1>`

  `trap-dest: <hostname2>`

  `...`

- `/etc/opt/omni/server/snmp/OVfilter`

  This file contains the severity level of the Data Protector SNMP trap messages that are to be filtered out (will not be sent by Data Protector). It has the following format:

  <message_level>

  <message_level>

  `...`

  Where *<message_level>* can be any of the following: (`normal | warning | minor | major | critical`).

On Windows systems, the destination is set in the Windows SNMP service configuration.

NOTE

On Windows systems, you need to configure the SNMP service first. For information on how to configure the Windows SNMP service, refer to “SNMP Send Method” on page 423.

The SNMP traps sent by Data Protector contain the following information:

- **Enterprise Event ID**

  Each event is marked with an Enterprise Event ID (EID) used to designate the type of entity that has sent the event. The EID for the events, sent by the OpenView entity, is ".1.3.6.1.4.1.11.2.17.1".

- **Generic Event ID**

  Each event is also marked with a Generic Event ID (GID). For standard SNMP traps, the GID tells `ovtrapd` which standard SNMP trap was generated. For other types of events, the GID is **6**, meaning that the sending entity has used a Specific Event ID to further qualify the event. Data Protector uses GID 6 only.

- **Specific Event ID**

  Events with GID=6 are also marked with a Specific Event ID (SID). The use of SIDs allows enterprises to define their own custom set of event definitions. (**59047936**, used by Data Protector, is the number for the Application Alert traps which is a subtype of the existing SNMP-Traps for the HP OpenView traps.)

- **Variables**

The Table 15-2 on page 771 shows the format of SNMP traps sent by Data Protector together with exemplary values.

Table 15-2 Data Protector SNMP Traps Format

MIB ID	Meaning	Exemplary Value
1.3.6.1.4.1.11.2.17.1.1.0	Application type	1
1.3.6.1.4.1.11.2.17.1.2.0	Hostname of the Cell Manager	machine.company.com
1.3.6.1.4.1.11.2.17.1.3.0	Trap message type	Either NOTIFICATION or nothing
1.3.6.1.4.1.11.2.17.1.4.0	Application name	HP Data Protector
1.3.6.1.4.1.11.2.17.1.5.0	Severity of the message	critical

Table 15-2 Data Protector SNMP Traps Format

MIB ID	Meaning	Exemplary Value
1.3.6.1.4.1.11.2.17.1.6.0	The message	Error on device “DLT_1” occurred
1.3.6.1.4.1.11.2.17.2.7.0	Parameter list	Mount request for device name=DLT_1

Command-Line Interface, Graphical User Interface and Web Reporting Interface

The Data Protector CLI provides comparable functionality as it is provided in Data Protector GUI. Using the Data Protector CLI you can:

- Start the Data Protector GUI and sub-GUIs. For a list of the Data Protector sub-GUIs, refer to “Graphical User Interface” on page 6.
- Configure and start Data Protector actions such as backup, restore and IDB purge. For a list of possible Data Protector actions, refer to “Data Protector Commands” on page A-7.
- Configure and start Data Protector reports using the Data Protector `omnirpt` CLI command. For more information about reporting, refer to “Data Protector Reporting” on page 388.
- Start the Java user interface to configure and start Data Protector reports. For more information about web reporting, refer to “Configuring Reports and Notifications on the Web” on page 426.

You can use Data Protector commands for scripts that provide the input data to System and Application Management application.

Data Protector Log Files

Some System and Application Management applications, such as HP OpenView Vantage Point Operations, allow you to specify when and which log files should be monitored for a specific log entry. If the specified entry is detected in the file, an action can be specified. In VPO this is called *Log file encapsulation*.

You can configure such a System and Application Management application to monitor Data Protector log files for specific log entries (Data Protector events) and define an action that is to be executed in case a particular Data Protector event is detected.

For more information on Data Protector log files refer to “Data Protector Log Files” on page 651. Note that there is no log files formatting specification provided. For Data Protector log files exemplary entries, refer to “Data Protector Log Files Example Entries” on page A-44.

Windows Application Log

Some System and Application Management applications, such as ManageX, monitor the Windows Application Log.

To enable automatic forwarding of all Data Protector messages and messages about the Data Protector services (if they are stopped) to Windows Application Log, set the `EventLogMessages` variable in the Data Protector global options file to `1`. For more information on Data Protector global options file refer to “Global Options File” on page 613.

Examples

Verifying Data Protector Processes

Data Protector provides a means of checking if its required processes are running by the means of the `omnisv -status` CLI command.

The `omnisv -status` command provides you with the status of the required Data Protector processes (when the command is started).

omnisv -status

To get the status of required Data Protector processes enter the following command:

```
omnisv -status
```

Data Protector Health Check Failed Notification

The User Health Check notification is triggered and sent only if any of the required processes are not running or if the IDB is not operational. The Health Check Failed notification by default checks these conditions every day at 12:00 (Noon) and is (if the conditions are met), by default sent to Data Protector Event Log. You can change the scheduled time by changing the `DailyMaintenanceTime` variable, using the twenty-four hour clock notation, in the Data Protector global options file. For more information on Data Protector global options file refer to “Global Options File” on page 613. You can also redirect the notification to be sent, for example as an SNMP trap.

To check every day at the scheduled time if the required Data Protector processes are running and if the IDB is operational, and to be notified by an SNMP trap if any of the processes are not running or if database is not operational, configure the Health Check Failed notification as described in the “Data Protector Notifications” on page 414.

To check the conditions of the Health Check Failed notification interactively, enter the following command:

```
omnihealthcheck
```

Refer to the `omnihealthcheck` man page for more information on `omnihealthcheck` command.

Getting the Results of the Last Night’s Backup

You can get the report on the results of the last night’s backups using the Data Protector reporting functionality. For more information on Data Protector reporting functionality refer to “Data Protector Reporting” on page 388 and to the `omnirpt` man page - more than 30 different reports, each having many different options, can be run.

To get the HTML report on the last night’s backup in the file `report.html` enter the following command:

```
omnirpt -report list_sessions -timeframe 24 24 -html -log
report.html
```

A Further Information

In This Appendix

This chapter gives information on the following topics:

Backing Up and Restoring UNIX Specifics

This section explains how to backup specific UNIX formats, including VxFS, Enterprise Filesystems, and Context Dependent Filesystems.

VxFS Snapshot

What Is VxFS?

VxFS allows you to back up a filesystem while it is being used by some other application. This is called an online backup and is done by creating a snapshot of a filesystem and backing up this snapshot.

You create a snapshot of a filesystem when you mount the VxFS filesystem to a temporary directory. At this point you also specify the filesystem you want to snap.

A **snapshot** is a copy of the filesystem at a specific moment in time you mount the VxFS filesystem to a temporary directory.

You can perform normal backups without using the VxFS snapshot feature by simply configuring a backup as for any other filesystem. In this case you cannot back up files that are in use.

You configure a backup of this temporary directory, which is actually a mountpoint to the snapshot of the filesystem as it was at the moment of the mount.

When the backup is finished, you unmount the snapshot filesystem so that it can be used for other purposes.

How to Configure VxFS Backup?

If you want to use the VxFS online backup functionality, you must configure the backup as follows:

1. You have to have an empty or unused partition created on your system that can be used by VxFS for a snapshot. See your system administrator's manual for instructions.

   The recommended size for the snapshot filesystem is up to 15% of the snapped filesystem, if the filesystem is used heavily use during the backup. Normally, the size should be around 5%.

If the amount of data modified on the snapped filesystem is higher than the space available, Data Protector produces `Cannot stat` error messages for all the remaining files to be backed up. You must unmount the snapshot filesystem and repeat the backup procedure.

2. Create a temporary directory to which you will mount the snapshot filesystem.
3. Create shell scripts to mount and unmount the snapshot filesystem to the temporary directory. See "Pre- and Post -exec Script Templates" in the next section for templates of these scripts.
4. Configure a backup of the temporary directory. The mount script must be specified as the Pre- exec command, and the unmount script as the Post-exec command.

Pre- and Post- exec Script Templates

Here are example templates that can be configured as Data Protector Pre- exec and Post- exec commands to mount or unmount the VxFS filesystem.

Example A-1 Pre- exec Script Template

```
# SnapMount.sh
#
# Mounting snapshot filesystem (Pre-exec script)
#
# A script requires 3 parameters:
# 1. a block special file of the snapped FS
# or
# a mount point directory of the snapped FS
# 2. a block special file of the snapshot FS
# 3. a mount point of the snapshot FS
#
# NOTE:
#
# In case of multiple Disk Agents reading from the same
snapshot
# FS,
# the Pre-exec script should contain a kind of
synchronization
# mechanism for following reasons:
#
# 1) an attempt to mount an already mounted snapshot FS,
```

```
# snapping the same FS will cause the Pre-exec script to
fail and
# a DA to abort
#
# 2) an attempt to mount an already mounted snapshot FS,
# snapping some other FS will cause a warning to be
generated,
# script to fail and a DA to abort
#
# 3) a synchronization with the Post-exec script should
be also
# provided because the snapshot FS must not be unmounted
while
# there is other DA reading from the FS.
#


SNAPPED_FS=$1
SNAPSHOT_FS=$2
MOUNT_POINT=$3


mount -F vxfs -e -o snapof=$SNAPPED_FS $SNAPSHOT_FS
$MOUNT_POINT


#
# end SnapMount.sh
#
```

The template below can be used to unmount a VxFS system.

Example A-2 Post- exec Script Template

```
# SnapUnmount.sh
#
# Unmounting snapshot filesystem (Post-exec shell
script)
#
# Script requires 1 parameter:
# - a mount point directory of the snapshot FS
# or
```

```
# - a block special file of the snapshot FS
#
# NOTE
# In case of multiple Disk Agents reading from the same
snapshot
# FS, a kind of synchronization mechanism has to be added
for
# the following reasons:
#
# 1) Post-exec script should unmount snapshot FS only if
there
# is no other DA reading from the snapshot FS
#
# Success/failure of the DA can be checked by examining
# the BDACC environment variable
#

MOUNT_POINT=$1

umount -v $MOUNT_POINT


#
# end SnapUnmount.sh
#
```

Data Protector Commands

For a complete list of supported Data Protector commands, refer to the *HP OpenView Storage Data Protector Command Line Interface Reference* (`CLIReference.pdf`) or the `omniintro` man page on UNIX.

The *HP OpenView Storage Data Protector Command Line Interface Reference* is located in the `<Data_Protector_home>\docs\MAN` directory on Windows or in the `/opt/omni/doc/C/` directory on UNIX.

The documents are available, if you installed the `User Interface` component on Windows or the `OB2-DOCS` component on UNIX.

On UNIX, use `man <command_name>` for more details about the command.

Performance Considerations

This section gives an overview of the most common backup performance factors. It is not meant to discuss performance. Due to the high number of variables and permutations, it is not possible to give distinct recommendations that fit all user requirements and affordable investment levels. Further discussions can be found in the *HP OpenView Storage Data Protector Concepts Guide*.

The Infrastructure

The infrastructure has a high impact on backup and restore performance. The most important factors are the parallelism of data paths and the use of high speed equipment.

Network Versus Local Backups and Restores

Sending data over the network introduces additional overhead, as the network becomes a component to performance consideration. Data Protector handles the datastream differently for the following cases:

Network Datastream Disk to Memory to Network to Memory to Device

Local Datastream Disk to Memory to Device

In order to maximize the performance, it is recommended to use local backup configurations for high volume datastreams.

Devices

The device type and model impacts the performance because of the sustained speed at which a device can write data to a tape (or read data from it). For example:

- DDS/DAT devices typically have a sustained speed of 510 KB/s to 3 MB/s, without compression, depending on the model.
- DLT devices typically have a sustained speed of 1.5 MB/s to 6 MB/s, without compression, depending on the model.
- LTO devices typically have a sustained speed of 10 MB/s to 20MB/s, without compression, depending on the model.

The speed also varies if a device-compression gets used. The achievable compression ratio depends on the nature of the data being backed up. For most cases, using high speed devices with device-compression ON does improve performance. This however is true only if the device(s) stream.

Libraries offer additional advantages because of their fast and automated access to a large number of media. At a backup time loading new or reusable media is needed and at a restore time the media which contain the data to be restored need to be accessed quickly.

High Performance Hardware Other Than Devices

The computer systems themselves, that is, reading the disk and writing to the device, directly impact performance. The systems are loaded during backup by reading the disk or handling software (de-)compression.

The disk read data rate and available CPU are important performance criteria for the systems themselves in addition to the I/O performance and network types.

Using Hardware in Parallel

Using several datapaths in parallel is a fundamental and efficient method to improve performance. This includes the network infrastructure. Parallelism helps in the following situations:

- Several systems can be backed up locally, that is, with the disk(s) and the related devices connected on the same system.
- Several systems can be backed up over the network. Here the network traffic routing needs to be such that the datapaths do not overlap, otherwise the performance will be reduced.
- Several objects (disks) can be backed up to one or several (tape) devices.
- Several dedicated network links between certain systems can be used. For example, system_A has 6 objects (disks) to be backed up, and system_B has 3 fast tape devices. Putting 3 network links dedicated to backup between system_A and system_B is a solution.

- Load Balancing: This is where Data Protector dynamically determines which filesystem should be backed up to which device. Normally, it is best to enable this feature. This is especially true when a large number of filesystems in a dynamic environment are being backed up.

Configuring Backups and Restores

Any given infrastructure must be used efficiently in order to maximize performance. Data Protector offers high flexibility in order to adapt to the environment.

Device Streaming

To maximize a device's performance, it must be kept streaming. A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for some more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. In network-focused backup infrastructures, this deserves attention.

Backups can be setup so that the data from several disk agents is sent to one Media Agent, which sends the data to the device.

Block Size

The device hardware processes data it receives using a device type specific block size. Data Protector allows to adjust the size of the block it sends to the device. The default value is 64kB.

Increasing the block size can improve the performance. Changing the block size should be done *before* formatting tapes. For example, a tape written with the default block size cannot be appended to a tape using a different block size.

Software Compression

Software compression is done by the client CPU when reading the data from the disk. This reduces the data which gets send over the network, but it requires significant CPU resources from the client.

NOTE By default, software compression should be disabled. Software compression should only be used for backup of many systems over a slow network where the data can be compressed before sending it over the network. If software compression is used, hardware compression should be disabled since trying to compress data twice actually expands the data.

Hardware Compression

Hardware compression is done by a device, which receives the original data from the Media Agent client and writes it to the tape in compressed mode. Hardware compression increases the speed at which a tape drive can receive data, because less data is written to the tape.

By default, hardware compression should be enabled. On HP-UX and Solaris, hardware compression should be enabled by selecting a hardware compression device file. On Windows, hardware compression can be selected during the device configuration. Using hardware compression or not should be a conscious decision, because media written in compressed mode cannot be read using the device in uncompressed mode and vice-versa.

Limitations HP Ultrium LTO drives use automatic hardware compression which cannot be disabled. Ensure that you do not enable `Software compression` when you configure an HP Ultrium LTO drive.

Full and Incremental Backups

A basic approach to improve performance is to reduce the amount of backed-up data. Take full advantage of time and resources when planning your full and incremental backups. An important consideration is that there is no need to do full backups of all the systems on the same day, unless necessary. See the *HP OpenView Storage Data Protector Concepts Guide* for more information.

Image Backup Versus Filesystem

It used to be more efficient to back up images (raw volumes) instead of backing up filesystems. This can still be true in some cases, such as with heavily-loaded systems or if the disks contain a large number of scattered files. The general recommendation is to use the filesystem backup.

Object Distribution to Media

There are many ways to configure a backup such that the backup data ends up on the media in just as many different configurations. For example:

- One object goes to one medium, or
- Several objects go to several media, each medium contains data from each object

Under certain conditions, one distribution may be advantageous considering the backup performance, however this may not be the optimal restore configuration.

The challenge is to optimize the setup for a backup (since it is done frequently) and at the same time have an acceptable restore media situation.

Miscellaneous Performance Hints

- Patches:

  Ensure you have installed all patches pertaining the performance on the network.

- On the computers that are Media Agent and Disk Agent clients, set the IP as shown below:

  ```
  IP is local "<MA_And_DA_Client_name>" == true
  ```

- LAN Cards:

  If you use a FDDI card, you can move it up on the bus so that it receives a higher priority. Use `ftp` to transfer large files between the MA and DA systems to see how the speed compares to Data Protector performance. The network cards configured in half-duplex decrease the performance.

- Simulating a high-speed device:

  If you suspect that the sustained data flow to the tape device is too low or that the device does not handle it correctly, you can simulate a very fast device on the Media Agent client by doing the following:

  1. Create a standalone file device and a device file `/dev/null` on UNIX and `nul` on Windows.
  2. Create a separate pool and select `loose` policy.
  3. Set `InitOnLoosePolicy=1` and set data protection to `None`. Perform backups to this device and check if the performance discrepancy between backups to the file device and backups to the real device can be explained. You can also run the `vbda` locally and write directly to a file. Run the commands listed below:

     On HP-UX and Solaris:

     ```
     /opt/omni/lbin/vbda -vol /home -trees /home/jdo -
     out /dev/null -profile
     ```

     On Windows:

     ```
     <Data_Protector_home>\bin\vbda -vol /C -trees
     "/Program Files/OmniBack/bin" -out nul -profile
     ```

     On Novell NetWare:

     ```
     load sys:usr\omni\bin\hpvbda.nlm -vol /sys -tree
     /usr/omni -out \tmp\test
     ```

- Device configuration

  Adjust the device block size if necessary.

- CRC option

  CRC option impacts performance due to the CRC calculation, which is performed by the Media Agent client.

- Logging and Report Level

  If an update of the IDB takes too long, disable logging by setting it to `Log None`. The same way you can filter messages by setting the Report level to `Critical`.

- Data Protector Application Clients

If a restore session of the Application clients (Oracle, SAP R/3) takes too long, decrease the SmWaitforNewClient value, which is by default 5 minutes. Set it to a lower value.

Example of Scheduled Eject of Media

You might want to eject all media that were used for backup during the night every morning at 6.00 AM. To schedule such an operation proceed as follows:

Schedule the Report Group

1. In the Data Protector GUI, select `Reporting`.
2. In the Scoping Pane expand `Reporting` and right click `Reports`. Select `Add Report Group`. The `Add Report Group` wizard is displayed.
3. In the wizard, name your report group and click `Next`. The Data Protector Scheduler is displayed.
4. In the Scheduler, select the starting day and click `Add`. In the `Schedule Report Distribution` dialog window, specify the hour, and that the report is to be generated daily. Click `OK` and then `Finish`.

   The Report Group is now scheduled. Now you can add the report to it.

Add the Report to the Report Group and Configure It

1. In the `Add New Report Wizard`, select `Reports on Media and Pools`.
2. Select the `List of Media` type and name the report. Click `Next`.
3. To eject *all* media, regardless of media pool and location leave all fields set to default settings. Click `Next` four times.
4. Select the `Relative time` and specify `8` for `Started within last hours` and `8` for `Duration hours` text boxes respectively. This will cause only the media that were used for backup in the last eight hours from the point of starting a report to be listed in the report. Click `Next`.
5. In the `Format` and `Send` text boxes, select `Tab` and `External`, respectively. In the `Script` text box, provide the name of the script (HP-UX and Solaris systems) or the batch file containing the command that starts the script (Windows systems). The script is

given in the next section. The script (HP-UX and Solaris systems) or the starting batch file (Windows systems) must reside in the /opt/omni/lbin (HP-UX and Solaris systems) or *<Data_Protector_home>*\bin (Windows systems) directory.

On Windows systems, the contents of the batch file containing command for starting the script is:

<perl_home>\perl.exe
"*<Data_Protector_home>*\bin\omnirpt_eject.pl"

6. Click the >> button to add this recipient. Click Finish.

The Report Group is now scheduled and configured.

Copy the Script to the Specified Directory

Copy or create the script with the name omnirpt_eject.pl in the /opt/omni/lbin (HP-UX and Solaris systems) or *<Data_Protector_home>*\bin directory (Windows systems).

```
#!/usr/contrib/bin/perl

#=========================================================
=======================

#   FUNCTION    Library_Eject

#

#   ARGUMENTS   param 1 = Library to eject from

#               param 2 = Slots to eject

#

#   DESCRIPTION Function ejects specified slots from
specified library

#=========================================================
=======================

sub Library_Eject {

    local ($lib,$slots)=@_;

    print "[Normal] Ejecting slot(s) ${slots}from
library \"$lib\"\n";

    print("[Normal] Executing \"${OMNIBIN}omnimm\"
-eject \"$lib\" $slots\n");
```

```
    $report =`"${OMNIBIN}omnimm" -eject \"$lib\"
$slots`;

    #print "\debug>\n$report\n<debug\n";

    if ($report !~/Final report: (\d+) cartridges out of
(\d+) successfully ejected\./) {

        print "[Critical] Eject has
failed!\n\nReport:\n$report\n";

        return (1);

    }

    print "$report\n";

    if ($1 ne $2) {

        print "[Warning] Not all media successfully
ejected!\n";

        return (2);

    }

    print "[Normal] Eject from library \"$lib\"
successfully completed.\n";

    return (0);

}

#=====================================================
========================

#   FUNCTION    Eject

#

#   ARGUMENTS   none

#

#   DESCRIPTION Function for each library in %List call
Library_Eject

#=====================================================
========================

sub Eject {

    local ($lib,$slot,$result);
```

```
    while (($lib, $slot) = each(%List)) {
        $result |=&Library_Eject($lib,$slot);
    }
    if ($result) {
        return (1);
    } else {
        print "[Normal] All operations successfully
completed.\n";
        return (0);
    }
}
#=======================================================
========================
#   FUNCTION    Omnirpt
#
#   ARGUMENTS   none
#
#   DESCRIPTION Function get slots to eject from omnirpt
report
#=======================================================
========================
sub Omnirpt {
    @lines =<STDIN>;
    for ($i=5;$i<@lines;$i++) {
        @line =split(/\t/,$lines[$i]);
        if ($line[2] =~/^\[([\w:\-\s]+):\s+(\w+)\]/) {
           $List{$1}.=$2.' ';  # $1= "Library name", $2=
"Slot ID"
        }
    }
```

```
    if (!keys(%List)) {
        print "[Warning] No tape(s) to eject.\n";
        return (1);
    }
    return (0);

}
#-----------------------------------------------------------
-----------------
#                                   MAIN
#-----------------------------------------------------------
-----------------
if ($ENV{"OS"}=~/Windows/) {  # Windows NT
      $OMNIBIN ='c:\\program files\\omniback\\bin\\';
    } else {
        local($uname)=`uname -a`;
        chop $uname;
        @uname=split(' ', $uname);
        if ($uname[0]) {
            if ($uname [0] eq 'HP-UX') {
                $OMNIBIN ='/opt/omni/bin/';
            } else {
                $OMNIBIN ='/usr/omni/bin';
            }
        } else {
            exit (1);
        }
    }
```

```
    print "[Normal] Starting eject of media that have
been used in the last 24 hours.\n";


    exit (0) if (&Omnirpt());
    exit (1) if (&Eject());
```

Examples of Pre-Exec and Post-Exec Commands for UNIX

The following scripts are some examples of Pre- and Post- exec commands on UNIX.

Session Pre-Exec: Shut Down Application

The script shuts down an Oracle instance.

```
#!/bin/sh
export ORACLE_HOME=$2
export ORACLE_SQLNET_NAME=$1
if [ -f $ORACLE_HOME/bin/svrmgrl ]; then
$ORACLE_HOME/bin/svrmgrl << EOF
connect sys/manager@$ORACLE_SQLNET_NAME as sysdba
shutdown
EOF
echo "Oracle database \"$ORACLE_SID\" shut down."
exit 0
else
echo "Cannot find Oracle SVRMGRL
($ORACLE_HOME/bin/svrmgrl)."
exit 1
fi
```

Disk Image Pre-Exec: Unmount a Disk Before a Raw Volume Backup

```
#!/bin/sh
echo "The disk will be unmounted!"
umount /disk_with_many_files
if [ $? = 0 ]
then
echo "The disk has been successfully unmounted!"
exit 0
```

```
else
echo "Failed to unmount the disk --> ABORTED!"
exit 1
fi
```

Filesystem Pre-Exec: Report Usage of the Filesystem

```
#!/bin/sh

echo
"================================================"
fuser -cu /var/application_mount_point
echo
"================================================"
exit 0
```

Session Post-Exec: Application Startup

This example Post-exec script will start up the Oracle database.

```
#!/bin/sh
export ORACLE_HOME=$2
export ORACLE_SQLNET_NAME=$1
if [ -f $ORACLE_HOME/bin/svrmgrl ]; then
        $ORACLE_HOME/bin/svrmgrl << EOF
connect sys/manager@$ORACLE_SQLNET_NAME as sysdba
startup
EOF
        echo "Oracle database \"$ORACLE_SID\" started."
        exit 0
else
        echo "Cannot find Oracle SVRMGRL
($ORACLE_HOME/bin/svrmgrl)."
        exit 1
```

```
fi
```

Disk Image Post-Exec: Mount a Disk After the Raw Volume Backup

```
#!/bin/sh
if [ $BDACC != 0]
then
echo "Backup could not read the disk!"
echo "Disk will not be automatically mounted!"
fi
echo "The disk will be now mounted!"
mount /dev/vg05/lvol2 /disk_with_many_files
if [ $? = 0 ]
then
echo "Disk successfully mounted!"
exit 0
else
echo "Failed to mount disk!"
exit 1
fi
```

Filesystem Post-Exec: Log Backup for the Record

```
#!/bin/sh
if  [ ! -f /etc/logfile ]
then
/etc/logfile
fi
echo "Backup finished with code $BDACC on " `date` >>
/etc/logfile
# We do not want a backup to be marked failed even if the
previous
action failed.
exit 0
```

Session Post-Exec: Notify User

```
#!/bin/sh
/opt/omni/bin/omnirpt -report single_session -session
$SESSIONID | \
 mailx -s "Report for $SESSIONID" $OWNER
```

Session Post-Exec: Start Another Backup

```
#!/bin/sh
# First check how the current backup finished
if [ $SMEXIT != 0 -o $SMEXIT != 10 ]
then
echo "Backup not successful --> next backup will not be
started!"
exit 0
fi
if  [ $RESTARTED != 0 ]
then
echo "Restarted backup --> next backup will not be
started!"
exit 0
fi
/opt/omni/bin/omnib -datalist BACKUP_NO_2 -no_mon
exit 0
```

Session Post-Exec: Restart Failed Backup

```
#!/bin/sh
# First check how the current backup finished
if [ $SMEXIT != 0 -o $SMEXIT != 10 ]
then
echo "Backup not successful --> backup will not be
restarted!"
exit 0
fi
if  [ $RESTARTED != 0 ]
```

```
then
echo "Restarted backup --> backup will not be
restarted!"
exit 0
fi
/opt/omni/bin/omnib -restart $SESSIONID -no_mon
exit 0
```

Disaster Recovery: Move Kill Links on HP-UX 11.x

Proceed as shown below on the system which you want to back up to move some links:

```
# The system will go from "run-level" 4 to "run-level 1"

# retaining the inetd, networking, swagentd services up.
The state is called "minimum activity" for backup
purposes (need networking).

# IMPORTANT: ensure the links are present in /sbin/rc1.d
before

# moving and they do have this exact name. You have to
rename them for the rc0.d directory. Put them BELOW the
lowest (original "/sbin/rc0.d/Kxx") "K...-link" in rc0.d

# Move K430dce K500inetd K660net K900swagentd into
../rc0.d BELOW the lowest kill link!!!

echo "may need to be modified for this system"

exit 1

#

cd /sbin/rc1.d

mv K430dce ../rc0.d/K109dce

mv K500inetd ../rc0.d/K110inetd

mv K660net ../rc0.d/K116net

mv K900swagentd ../rc0.d/K120swagentd
```

Creating a libaci.o on AIX

OmniBack II and A.04.10

OmniBack II A.04.10 DAS Agent on AIX uses the library object module named `libaci.a` which has to be created from the library archive file of the same name. Proceed as follows to create the object module:

1. Create the file `libaci.exp` containing the list of modules used by the OmniBack II DAS Agent:

```
#! /usr/omni/lib/libaci.a
aci_initialize
aci_qversion
aci_init
d_errno
aci_view
aci_drivestatus
aci_drivestatus2
aci_driveaccess
aci_mount
aci_dismount
aci_qvolsrange
aci_eject_complete
aci_eject
aci_insert
```

2. Create the object module `libaci.o` by executing following command:

```
ld -L/usr/omni/lib -bM:SRE -e_nostart -lc
-bE:<DAS_PATH>/libaci.exp <DAS_PATH>/libaci.a -o libaci.o
```

   <DAS_PATH> is the path to the directory where the library archive file `libaci.a` and the `libaci.exp` files are located.

3. Copy the library object module `libaci.o` to the `/usr/omni/lib` directory and rename it to `libaci.a`.

IMPORTANT

The full path to the library archive file is *<DAS_PATH>*`/libaci.a`, whereas the full path to the object module used by DAS Agent is `/usr/omni/lib/libaci.a`.

Example of the Package Configuration File

This section gives an example of a package configuration file that you need to modify while configuring Data Protector Cell Manager package in an MC/ServiceGuard environment:

```
**********************************************************************
********

# ****** HIGH AVAILABILITY PACKAGE CONFIGURATION FILE
(template) *******

#
**********************************************************************
********

# ******* Note: This file MUST be edited before it can be used.
********

# * For complete details about package parameters and how to
set them, *

# * consult the MC/ServiceGuard or ServiceGuard OPS Edition
manpages   *

# * or manuals.
*

#
**********************************************************************
********


# Enter a name for this package.  This name will be used to
identify the

# package when viewing or manipulating it.  It must be
different from

# the other configured package names.


PACKAGE_NAME ob2cl
```

```
# Enter the failover policy for this package. This policy will
be used

# to select an adoptive node whenever the package needs to be
started.

# The default policy unless otherwise specified is
CONFIGURED_NODE.

# This policy will select nodes in priority order from the list
of

# NODE_NAME entries specified below.

#

# The alternative policy is MIN_PACKAGE_NODE. This policy will
select

# the node, from the list of NODE_NAME entries below, which is

# running the least number of packages at the time this package
needs

# to start.


FAILOVER_POLICY CONFIGURED_NODE



# Enter the failback policy for this package. This policy will
be used

# to determine what action to take when a package is not
running on

# its primary node and its primary node is capable of running
the

# package. The default policy unless otherwise specified is
MANUAL.

# The MANUAL policy means no attempt will be made to move the
package

# back to its primary node when it is running on an adoptive
node.

#

# The alternative policy is AUTOMATIC. This policy will attempt
to
```

```
# move the package back to its primary node whenever the
primary node
# is capable of running the package.


FAILBACK_POLICY MANUAL



# Enter the names of the nodes configured for this package.
Repeat
# this line as necessary for additional adoptive nodes.
# Order IS relevant.  Put the second Adoptive Node AFTER the
first
# one.
# Example : NODE_NAME  original_node
#           NODE_NAME  adoptive_node


NODE_NAME partizan
NODE_NAME lyon



# Enter the complete path for the run and halt scripts.  In
most cases
# the run script and halt script specified here will be the
same script,
# the package control script generated by the cmmakepkg
command.  This
# control script handles the run(ning) and halt(ing) of the
package.
# If the script has not completed by the specified timeout
value,
# it will be terminated.  The default for each script timeout
is
# NO_TIMEOUT.  Adjust the timeouts as necessary to permit full
```

```
# execution of each script.
# Note: The HALT_SCRIPT_TIMEOUT should be greater than the sum
of
# all SERVICE_HALT_TIMEOUT specified for all services.

RUN_SCRIPT /etc/cmcluster/ob2cl/ob2cl.cntl
RUN_SCRIPT_TIMEOUT NO_TIMEOUT
HALT_SCRIPT /etc/cmcluster/ob2cl/ob2cl.cntl
HALT_SCRIPT_TIMEOUT NO_TIMEOUT



# Enter the SERVICE_NAME, the SERVICE_FAIL_FAST_ENABLED and the
# SERVICE_HALT_TIMEOUT values for this package.  Repeat these
# three lines as necessary for additional service names.  All
# service names MUST correspond to the service names used by
# cmrunserv and cmhaltserv commands in the run and halt
scripts.
#
# The value for SERVICE_FAIL_FAST_ENABLED can be either YES or
# NO.  If set to YES, in the event of a service failure, the
# cluster software will halt the node on which the service is
# running.  If SERVICE_FAIL_FAST_ENABLED is not specified, the
# default will be NO.
#
# SERVICE_HALT_TIMEOUT is represented in the number of seconds.
# This timeout is used to determine the length of time (in
# seconds) the cluster software will wait for the service to
# halt before a SIGKILL signal is sent to force the termination
# of the service.  In the event of a service halt, the cluster
# software will first send a SIGTERM signal to terminate the
# service.  If the service does not halt, after waiting for the
```

```
# specified SERVICE_HALT_TIMEOUT, the cluster software will
send

# out the SIGKILL signal to the service to force its
termination.

# This timeout value should be large enough to allow all
cleanup

# processes associated with the service to complete.  If the

# SERVICE_HALT_TIMEOUT is not specified, a zero timeout will be

# assumed, meaning the cluster software will not wait at all

# before sending the SIGKILL signal to halt the service.

#

# Example: SERVICE_NAME                  DB_SERVICE

#          SERVICE_FAIL_FAST_ENABLED     NO

#          SERVICE_HALT_TIMEOUT          300

#

# To configure a service, uncomment the following lines and

# fill in the values for all of the keywords.

#

#SERVICE_NAME                   <service name>

#SERVICE_FAIL_FAST_ENABLED      <YES/NO>

#SERVICE_HALT_TIMEOUT           <number of seconds>


SERVICE_NAME                   omni_sv

SERVICE_FAIL_FAST_ENABLED      NO

SERVICE_HALT_TIMEOUT           300



# Enter the network subnet name that is to be monitored for
this package.

# Repeat this line as necessary for additional subnet names.
If any of
```

```
# the subnets defined goes down, the package will be switched
to another

# node that is configured for this package and has all the
defined subnets

# available.


SUBNET 10.17.0.0



# The keywords RESOURCE_NAME, RESOURCE_POLLING_INTERVAL,

# RESOURCE_START, and RESOURCE_UP_VALUE are used to specify
Package

# Resource Dependencies.  To define a package Resource
Dependency, a

# RESOURCE_NAME line with a fully qualified resource path name,
and

# one or more RESOURCE_UP_VALUE lines are required.  The

# RESOURCE_POLLING_INTERVAL and the RESOURCE_START are
optional.

#

# The RESOURCE_POLLING_INTERVAL indicates how often, in
seconds, the

# resource is to be monitored.  It will be defaulted to 60
seconds if

# RESOURCE_POLLING_INTERVAL is not specified.

#

# The RESOURCE_START option can be set to either AUTOMATIC or
DEFERRED.

# The default setting for RESOURCE_START is AUTOMATIC.  If
AUTOMATIC

# is specified, ServiceGuard will start up resource monitoring
for

# these AUTOMATIC resources automatically when the node starts
up.
```

```
# If DEFERRED is selected, ServiceGuard will not attempt to
start

# resource monitoring for these resources during node start up.
User

# should specify all the DEFERRED resources in the package run
script

# so that these DEFERRED resources will be started up from the
package

# run script during package run time.

#

# RESOURCE_UP_VALUE requires an operator and a value.  This
defines

# the resource 'UP' condition.  The operators are =, !=, >, <,
>=,

# and <=, depending on the type of value.  Values can be string
or

# numeric.  If the type is string, then only = and != are valid

# operators.  If the string contains whitespace, it must be
enclosed

# in quotes.  String values are case sensitive.  For example,

#

# Resource is up when its value is

# --------------------------------

# RESOURCE_UP_VALUE= UP"UP"

# RESOURCE_UP_VALUE!= DOWNAny value except "DOWN"

# RESOURCE_UP_VALUE= "On Course""On Course"

#

# If the type is numeric, then it can specify a threshold, or a
range to

# define a resource up condition.  If it is a threshold, then
any operator

# may be used.  If a range is to be specified, then only > or
>= may be used
```

```
# for the first operator, and only < or <= may be used for the
second operator.

# For example,

# Resource is up when its value is

# --------------------------------

# RESOURCE_UP_VALUE     = 55    (threshold)

# RESOURCE_UP_VALUE     > 5.1greater than 5.1    (threshold)

# RESOURCE_UP_VALUE     > -5 and < 10between -5 and 10
(range)

#

# Note that "and" is required between the lower limit and upper
limit

# when specifying a range.  The upper limit must be greater
than the lower

# limit.  If RESOURCE_UP_VALUE is repeated within a
RESOURCE_NAME block, then

# they are inclusively OR'd together.  Package Resource
Dependencies may be

# defined by repeating the entire RESOURCE_NAME block.

#

# Example : RESOURCE_NAME
/net/interfaces/lan/status/lan0

#     RESOURCE_POLLING_INTERVAL120

#     RESOURCE_STARTAUTOMATIC

#     RESOURCE_UP_VALUE= RUNNING

#     RESOURCE_UP_VALUE= ONLINE

#

#           Means that the value of resource
/net/interfaces/lan/status/lan0

#         will be checked every 120 seconds, and is considered
to

#           be 'up' when its value is "RUNNING" or "ONLINE".

#
```

```
# Uncomment the following lines to specify Package Resource
Dependencies.

#

#RESOURCE_NAME     <Full_path_name>

#RESOURCE_POLLING_INTERVAL   <numeric_seconds>

#RESOURCE_START              <AUTOMATIC/DEFERRED>

#RESOURCE_UP_VALUE     <op> <string_or_numeric> [and <op>
<numeric>]



# The default for PKG_SWITCHING_ENABLED is YES. In the event of
a

# failure, this permits the cluster software to transfer the
package

# to an adoptive node.  Adjust as necessary.


PKG_SWITCHING_ENABLED YES



# The default for NET_SWITCHING_ENABLED is YES.  In the event
of a

# failure, this permits the cluster software to switch LANs
locally

# (transfer to a standby LAN card).  Adjust as necessary.


NET_SWITCHING_ENABLED YES



# The default for NODE_FAIL_FAST_ENABLED is NO.  If set to YES,

# in the event of a failure, the cluster software will halt the
node

# on which the package is running.  Adjust as necessary.
```

```
NODE_FAIL_FAST_ENABLEDNO
```

Example of the Package Control File

This section gives an example of a package control file that you need to modify while configuring Data Protector Cell Manager package in an MC/ServiceGuard environment:

```
************************************************************
********

# *
*

# *      HIGH AVAILABILITY PACKAGE CONTROL SCRIPT (template)
*

# *
*

# *     Note: This file MUST be edited before it can be used.
*

# *
*

#
************************************************************
********


# UNCOMMENT the variables as you set them.


# Set PATH to reference the appropriate directories.

PATH=/usr/bin:/usr/sbin:/etc:/bin


# VOLUME GROUP ACTIVATION:

# Specify the method of activation for volume groups.

# Leave the default ("VGCHANGE="vgchange -a e") if you want
volume

# groups activated in exclusive mode. This assumes the volume
groups have

# been initialized with 'vgchange -c y' at the time of
creation.
```

```
#
# Uncomment the first line (VGCHANGE="vgchange -a e -q n"), and
comment
# out the default, if your disks are mirrored on separate
physical paths,
#
# Uncomment the second line (VGCHANGE="vgchange -a e -q n -s"),
and comment
# out the default, if your disks are mirrored on separate
physical paths,
# and you want the mirror resynchronization to occur in
parallel with
# the package startup.
#
# Uncomment the third line (VGCHANGE="vgchange -a y") if you
wish to
# use non-exclusive activation mode. Single node cluster
configurations
# must use non-exclusive activation.
#
# VGCHANGE="vgchange -a e -q n"
# VGCHANGE="vgchange -a e -q n -s"
#VGCHANGE="vgchange -a y"
VGCHANGE="vgchange -a e"# Default


# VOLUME GROUPS
# Specify which volume groups are used by this package.
Uncomment VG[0]=""
# and fill in the name of your first volume group. You must
begin with
# VG[0], and increment the list in sequence.
#
# For example, if this package uses your volume groups vg01 and
vg02, enter:
```

```
#          VG[0]=vg01

#          VG[1]=vg02

#

# The volume group activation method is defined above. The
filesystems

# associated with these volume groups are specified below.

#

VG[0]=/dev/vg_ob2cm


# FILESYSTEMS

# Specify the filesystems which are used by this package.
Uncomment

# LV[0]=""; FS[0]=""; FS_MOUNT_OPT[0]="" and fill in the name
of your first

# logical volume, filesystem and mount option for the file
system. You must

# begin with LV[0], FS[0] and FS_MOUNT_OPT[0] and increment the
list in

# sequence.

#

# For example, if this package uses the file systems pkg1a and
pkg1b,

# which are mounted on the logical volumes lvol1 and lvol2 with
read and

# write options enter:

#          LV[0]=/dev/vg01/lvol1; FS[0]=/pkg1a;
FS_MOUNT_OPT[0]="-o rw"

#          LV[1]=/dev/vg01/lvol2; FS[1]=/pkg1b;
FS_MOUNT_OPT[1]="-o rw"

#

# The filesystems are defined as triplets of entries specifying
the logical

# volume, the mount point and the mount options for the file
system. Each
```

```
# filesystem will be fsck'd prior to being mounted. The
filesystems will be

# mounted in the order specified during package startup and
will be unmounted

# in reverse order during package shutdown. Ensure that volume
groups

# referenced by the logical volume definitions below are
included in

# volume group definitions above.

#

#LV[0]=""; FS[0]=""; FS_MOUNT_OPT[0]=""


LV[0]=/dev/vg_ob2cm/lv_ob2cm

FS[0]=/omni_shared

FS_MOUNT_OPT[0]=""


# FILESYSTEM UNMOUNT COUNT

# Specify the number of unmount attempts for each filesystem
during package

# shutdown.  The default is set to 1.

FS_UMOUNT_COUNT=2


# IP ADDRESSES

# Specify the IP and Subnet address pairs which are used by
this package.

# Uncomment IP[0]="" and SUBNET[0]="" and fill in the name of
your first

# IP and subnet address. You must begin with IP[0] and
SUBNET[0] and

# increment the list in sequence.

#

# For example, if this package uses an IP of 192.10.25.12 and a
subnet of
```

```
# 192.10.25.0 enter:
#          IP[0]=192.10.25.12
#          SUBNET[0]=192.10.25.0 # (netmask=255.255.255.0)
#
# Hint: Run "netstat -i" to see the available subnets in the
Network field.
#
# IP/Subnet address pairs for each IP address you want to add
to a subnet
# interface card.  Must be set in pairs, even for IP addresses
on the same
# subnet.
#
IP[0]=10.17.3.230
SUBNET[0]=10.17.0.0


# SERVICE NAMES AND COMMANDS.
# Specify the service name, command, and restart parameters
which are
# used by this package. Uncomment SERVICE_NAME[0]="",
SERVICE_CMD[0]="",
# SERVICE_RESTART[0]="" and fill in the name of the first
service, command,
# and restart parameters. You must begin with SERVICE_NAME[0],
SERVICE_CMD[0],
# and SERVICE_RESTART[0] and increment the list in sequence.
#
# For example:
#          SERVICE_NAME[0]=pkg1a
#          SERVICE_CMD[0]="/usr/bin/X11/xclock -display
192.10.25.54:0"
#          SERVICE_RESTART[0]=""  # Will not restart the
service.
```

```
#
#          SERVICE_NAME[1]=pkg1b
#          SERVICE_CMD[1]="/usr/bin/X11/xload -display
192.10.25.54:0"
#          SERVICE_RESTART[1]="-r 2"    # Will restart the
service twice.
#
#          SERVICE_NAME[2]=pkg1c
#          SERVICE_CMD[2]="/usr/sbin/ping"
#         SERVICE_RESTART[2]="-r 1" # Will restart the service
an infinite
#                                      number of times.
#
# Note: No environmental variables will be passed to the
command, this
# includes the PATH variable. Absolute path names are required
for the
# service command definition.  Default shell is /usr/bin/sh.
#
SERVICE_NAME[0]=omni_sv
SERVICE_CMD[0]="/etc/opt/omni/server/sg/csfailover.ksh start"
SERVICE_RESTART[0]="-r 2"
```

Data Protector Log Files Example Entries

This section provides some typical Data Protector messages that are logged to in some Data Protector log files. This section does not intended to provide further in-depth information on troubleshooting. For a complete list of Data Protector log files and for more in formation on them refer to “Data Protector Log Files” on page 651.

IMPORTANT The contents and format of entries to Data Protector log files are subject to change.

debug.log

```
02/11/00 12:22:01  OMNIRPT.23856.0
["/src/lib/cmn/obstr.c /main/r31_split/2":212] A.03.10
b325

  StrFromUserSessionId: "-detail": not in correct format


03/01/00 14:19:28  DBSM.21294.0
["PANSRC/db/RCS/cmn_srv.c,v 1.40":229] A.03.10 b325

  DB[1] internal error [9] cannot exclusively open
database, it is already opened


03/01/00 14:21:14  DBSM.21393.0
["PANSRC/db/RCS/cmn_srv.c,v 1.40":272] A.03.10 b325

  CDB cell server "bmw" different than current host
"bmw.hermes"


03/01/00 14:21:43  OMNIB.21471.0 ["/src/cli/omnibackup.c
/main/23":2585] A.03.10 b325

[Process] CanBackup failed!
```

```
03/02/00 09:36:51  INET.26130.0 ["/src/lib/ipc/ipc.c
/main/r31_split/10":6920] A.03.10 b325

IpcGetPeer: Could not expand ConnectionIP "10.17.6.227"


03/16/00 19:09:42  BSM.13152.0 ["src/db/cdb/cdbwrap.c
/main/84":1538] A.03.10 bPHSS_21234/PHSS_21235

  DB[1] internal error [-2009] The session is
disconnected


05/17/01 12:00:30  OMNIMM.7515.0 ["lib/cmn/obstr.c
/main/17":187] A.04.00.%B3 b335

  StrToUserSessionId: "0": not in correct format


5/14/01 11:08:53 AM  UPGRADE_CFG.357.356
["integ/barutil/upgrade_cfg/upgrade_cfg.c
/main/27":1472] A.04.00.%B3 b335

[UpgradeSQL] Can not read registry value
HKLM\Software\Hewlett-Packard\OpenView\OmniBackII\Agents
\MS-SQL70\saUser

[UpgradeSQL] Warning:  2, The system cannot find the
file specified.


5/14/01 11:08:54 AM  UPGRADE_CFG.369.368
["integ/barutil/upgrade_cfg/upgrade_cfg.c /main/27":154]
A.04.00.%B3 b335

[GetConfig] Can not read configuration from Cell Server
"brainiac.hermes" with integration "Oracle" and instance
"_OB2_GLOBAL"

[GetConfig] Error: 1012, [12:1012] Can not access the
file.

System error: [2] The system cannot find the file
specified.
```

```
5/14/01 12:41:41 PM  OMNIDBUTIL.98.124
["db/vel_cls_spec.c /main/39":103] A.04.00.%B3 b335

  VELOCIS DB ERROR [0] internal error [-2005] Server
unavailable
```

sm.log

```
3/28/00 03:00:01  BSM.23475.0 ["/src/sm/bsm2/brsmutil.c
/main/r31_split/4":630] A.03.50.%B2 b158

Error connecting to database. Code: 1166.


03/27/01 08:17:06  BSM.2709.0 ["sm/bsm2/bsmutil.c
/main/502":3306] A.04.00.%B1 b281

Error opening datalist OMNIBACK-.
```

inet.log

```
5/15/01 12:19:54 AM  INET.119.122 ["inetnt/allow_deny.c
/main/10":524] A.04.00.%B3 b335

A request 3 came from host bmw.hermes which is not a Cell
Manager of this client


[Critical] From: INET@clio.hermes "clio.hermes"  Time:
03/29/01 09:48:29

[70:5]  Cannot execute '/opt/omni/lbin/ob2rman.exe' (No
such file or directory) => aborting
```

media.log

```
02/04/00 06:57:46 0a110210:3861cbbb:742d:0003 "[CBF492]
BMW_DLT_23" [2000/02/04-8] OmniDB


02/04/00 07:02:38 0a110210:3861cbbb:742d:0003 "[CBF492]
BMW_DLT_23" [2000/02/04-9]
```

```
02/04/00 13:38:56 0a110210:389ac85b:3c6e:0001 "[CBF502]
DLT_ARC_8" [INITIALIZATION]


02/29/00 16:04:25 0a110210:38bbdff4:6d85:0026 "NULL_33"
[AUTOINITIALIZATION]


03/02/00 10:03:25 0a110210:385a24bf:410b:0002 "[CW1231]
BMW_DLT_15" [IMPORT]
```

upgrade.log

```
03/15/01 09:15:38

UCP session started.


03/15/01 09:20:55

UCP session finished.

total running time: 317 seconds


03/15/01 10:00:09

UDP session started.


03/15/01 10:02:54

Abort request from CLI/GUI on handle 0. Terminating
session


03/15/01 10:03:06

UDP session started.


03/15/01 10:26:47

Abort request from CLI/GUI on handle 0. Terminating
session
```

```
03/15/01 12:40:43
Database check error! Can not proceed with upgrade.

03/15/01 13:24:15
System error

03/15/01 13:24:15
Session was aborted by child ASM, marked error=1026

03/15/01 15:27:22
OmniBack II 3.x database not found.

03/15/01 16:33:19
[12:10904] Open of detail catalog binary file failed.

03/16/01 08:39:31
Internal error: Invalid Ct function argument specified.

03/20/01 10:56:57
[12:1165] Database network communication error.

03/22/01 14:38:21
[12:10953] Database is in incorrect state. Database must
be empty before critical upgrade can start.
```

Windows Manual Disaster Recovery Preparation Template

The template on the next page can be used to prepare for Windows Assisted Manual Disaster Recovery, as described in the Chapter 12, “Disaster Recovery,” on page 513.

Table A-1

client properties	computer name	
	hostname	
drivers		
Windows Service Pack		
TCP/IP properties	IP address	
	default gateway	
	subnet mask	
	DNS order	
medium label / barcode number		
partition information and order	1st disk label	
	1st partition length	
	1st drive letter	
	1st filesystem	
	2nd disk label	
	2nd partition length	
	2nd drive letter	
	2nd filesystem	
	3rd disk label	
	3rd partition length	
	3rd drive letter	
	3rd filesystem	

Changing Block Size on Windows Media Agent

In order to increase the maximum block size on a Windows Media Agent client, you have to modify its Registry. After modifying the Registry, restart the computer. Drivers read `MaximumSGList` at boot time. The actual formula that a Windows class driver uses to determine the maximum transfer size is:

```
maximum size = ((number of supported scatter/gather
elements - 1)* 4096)
```

For the typical aic78xx case, it renders the following:

```
((17-1) * 4096) = 64k
```

Windows provides a mechanism to support more scatter/gather elements via the Registry. Start the `regedit32` and add a `DWORD` value in the following Registry key:

```
\\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\aic78
xx\Parameters\Device0\MaximumSGList
```

Use the following formula to calculate the value of the `MaximumSGList`:

$$MaximumSGList = (BlockSize/4096) + 1$$

Example

The `MaximumSGList` value for a 256k block size is 65:
`MaximumSGList` = (256k/4k) + 1 = 64 + 1 = 65

If you have, for example, 3 aic78xx based SCSI channels on your system, change the appropriate ...`\Device0`, ...`\Device1` or ...`\Device2` value. If you want to set all adapters at the same time, specify `MaximumSGList` for ...`\Device\`.... Omitting the numerical reference sets the value for all aic78xx adapters.

B Standalone File and File Jukebox Devices

In This Appendix

This appendix provides an introduction to the configuration and use of standalone file and file jukebox devices as backup solutions using Data Protector. The following sections are included:

Overview of the Standalone File and File Jukebox Devices

The standalone file device and file jukebox device are part of a family of devices designed to backup data to disk. This family includes the following devices:

- Standalone file device
- File jukebox device
- File library device

The standalone file device and file jukebox device are simpler than the file library device and are useful for smaller backups. For information about the more sophisticated file library device, see Chapter 3, “Configuring and Using Disk- Based Devices,” on page 103.

The standalone file device is designed to store up to 1 GB of data. This device cannot be re-configured after it has been created.

The file jukebox device can store large amounts of data (up to 1 TB). It is possible to control and change the configuration of the file jukebox device once it has been created.

Both devices save data in the form of files on disk. Each of these files is the equivalent of a slot in a tape device. The standalone file device has only a single slot. The file jukebox device has multiple slots.

Data being backed up to the devices is first converted to tape format before being saved.

The standalone file and file jukebox devices are created and configured using the Data Protector GUI.

Recommended Configuration

Device Location

It is recommended that the device you are creating is located on a disk other than the one on which the IDB is located. This ensures that there will be an adequate amount of disk space available for the database. Putting the device and internal database on separate disks also improves performance.

If you put the device and the Data Protector internal database on the same disk you must ensure that there is enough space for them both. See “Allocating Disk Space for Future Use” on page 464 for details.

Number of Devices on Disk

If there is more than one data stream being read from or written to the same disk, in other words, if more than one Disk Agent or Media Agent are concurrently accessing the same mount point, which is defined as one disk, read/write performance for that disk will drop significantly.

It is therefore recommended to have only one standalone file device or file jukebox device per disk and only one drive per device, and to avoid other applications transferring large amounts of data from/to the disk when a Data Protector backup/restore is in progress. This situation may be different if using disk arrays.

Slot Size

If a large amount of data (over 1 GB) is to be backed up, it is recommended that you use a file jukebox device rather than a standalone file device. A file jukebox device is more flexible and controllable.

In general, the size of the device to be configured is dependent upon the amount of data to be backed up. However, it is normally recommended that you keep the standalone file device/slot size between 100 MB and 50 GB on Windows, or 100 MB to 2 TB on UNIX.

If, for instance, you have 1 TB of data to back up, the following device configuration is possible:

Windows systems 1 File jukebox device with 100 file slots of 10 GB each.

UNIX systems 1 File jukebox device with 250 file slots of 4 GB each.

Data protection is set for each individual file slot in a file jukebox, so it is possible to recycle a single slot by removing its protection. Therefore, having multiple small slots can increase flexibility and enable more efficient data protection and space retention management.

We recommend the following sizing (of course this can be changed according to requirements, such as small backups). Refer to *HP OpenView Storage Data Protector Installation and Licensing Guide* to see how this affects your licensing arrangements.

Table B-1 Recommended slot sizes on Windows and Unix Platforms

Available disk space	Number of slots	Slot size (GB)
1 TB	100	10
5 TB	250	20
10 TB	250	40

Block Size

Standalone file devices and file jukebox devices are effectively fast devices used for what would otherwise be a tape backup. For this reason, the format used when writing data to these devices is a tape writing format. The main parameters you need to consider when using such a format are:

- **Block size**, which is a legacy from tape devices. Tape drives from different vendors tend to use different block sizes.

  For your standalone file device and file jukebox device drives, the block size should normally be kept the same, otherwise standalone file devices/slots written using one drive may not be recognized if you decide to use a different drive. The default block size is 64 KB.

- **Segment size,** this parameter sets the amount of data between file marks when recording backup data. A file mark table, containing the positions for the file marks for each segment is also written in the standalone file device/slot header.

On Windows the maximum recommended slot size is 50 GB, although the file jukebox device has been tested on Windows with slots of up to 600 GB. On Unix the maximum recommended slot size is 2 TB.

For details about how to set block and segment sizes in a file device refer to the online Help index, enter the keywords “advanced options, devices and media”.

Configuring Standalone File or File Jukebox Devices

Prerequisite

Before you configure these devices on a Windows system, disable the compression option. This can be done using `Windows Explorer`. Right-click the directory to which data will be backed up, select `Properties` and deselect `Compress` under `Attributes`. If `Compress` is selected, Data Protector will not be able to write to the device.

IMPORTANT

Do not use the name of an existing device for configuring these devices, because the existing device will be overwritten.

Do not use the same device name for configuring several devices, because every time the device is accessed it will be overwritten.

Configuring the Devices

To create these devices, specify either the standalone file or the file jukebox device type in the `Add Device` wizard. As a device address, specify a pathname for the device, for example, `C:\My_Backup`. Refer to Figure B-1. For details about the configuration procedure refer to the online Help index and enter the keywords “configuring standalone file devices” or “configuring file jukebox devices”.

Figure B-1 Set Device Path

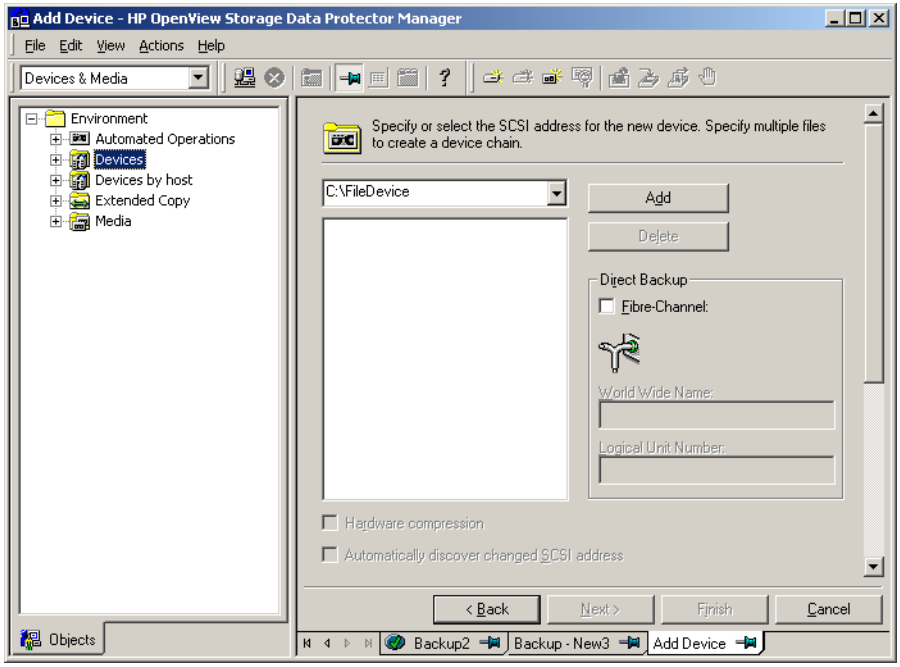

Next Steps

At this point, the device has been specified to Data Protector, but it does not yet actually exist on disk. Before it can be used for backup, you have to format it. For a detailed procedure refer to the online help index keywords “formatting media in a library device” for a jukebox file device and “formatting a medium” for a standalone file device.

Backup and Restore Using Standalone File or File Jukebox Devices

Backup to Standalone File and File Jukebox Devices

In a standard backup specification, select the standalone file device that you want to use as the target device for the backup.

Figure B-2 Backup to a Standalone File or File Jukebox Device

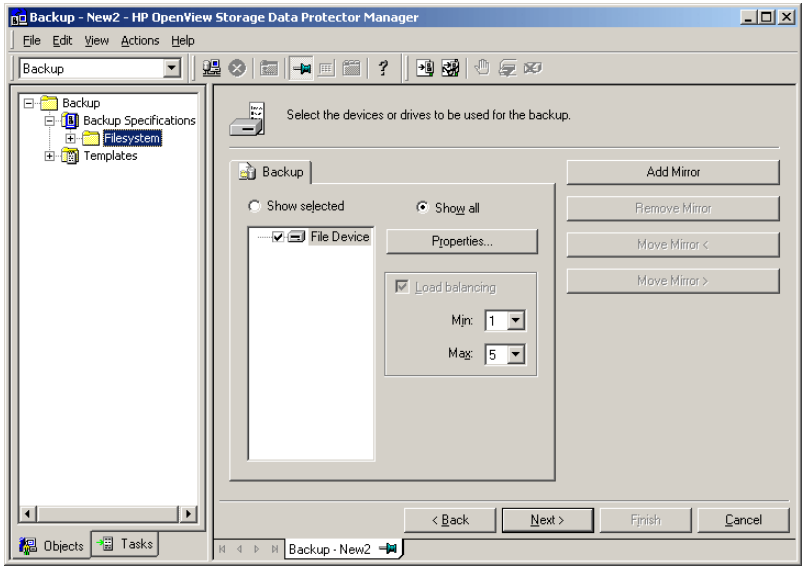

If the device is a file jukebox and is newly created, Data Protector picks the first file slot. If the file jukebox has been used for backup previously, Data Protector picks the same file slot as was used for the previous backup and writes data to that slot. When the slot is full, Data Protector picks the next file slot and continues with the backup.

This process is continued until all slots are full. If there is still backup data to be written at that point, Data Protector checks if any file slots have been recycled. If it finds one, it will use that to continue the backup;

if it does not find one, it will issue a mount request. To continue the backup, you then need to either mark a slot for recycling or create a new one and then confirm the mount request.

Data Protector also issues a mount request if there is not enough free space on the disk to complete a backup but there are still some free slots (which were created by GUI without reserving the space). In this case, you must free some space on the disk and confirm the mount request. Data Protector then continues the backup.

Maintenance of the Standalone File and File Jukebox Devices

If all the device you are using becomes full, you will need to do one of the following before continuing to make backups with it:

- Start moving data to tape, freeing up the file device or one or more file slots.
- Mark file jukebox slots for recycling.
- Add a new slot to the file jukebox.

Recycling File Slots

If you want to re-use file slots without moving the data in them to tape, for instance if you are using a file jukebox as your main backup device and want to restore directly from its slots, you can mark them for recycling. The recycled slots will be re-used for backups and the data within them overwritten. For a detailed procedure describing how slots are recycled refer to the online Help index, enter the keywords “recycling a file jukebox slot.”

IMPORTANT If this method is used, the existing data on the media will be overwritten and lost.

Adding a New File Slot

To create an additional file slot for a file jukebox device, refer to the online Help index for the procedure, enter the keywords “adding a slot”. Before using the slot, you will also need to format it. You can find a detailed description of how to format media in the online Help index by entering the keywords “formatting media in library devices”.

Restore from Standalone File and File Jukebox Devices

At restore time simply select the restore object in the Data Protector GUI and start a normal restore. For detailed information refer to the online Help index, enter the keywords “standard restore procedure.”

This process can be made more convenient by creating a script to run it automatically using the Data Protector CLI.

Glossary

access rights
See **user rights**.

ACSLS *(StorageTek specific term)*
The Automated Cartridge System Library Server (ACSLS) software that manages the Automated Cartridge System (ACS).

Active Directory *(Windows specific term)*
The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access and manage resources regardless of the physical system they reside on.

AML *(EMASS/GRAU specific term)*
Automated Mixed-Media library.

application agent
A component needed on a client to back up or restore online database integrations.
See also **Disk Agent**.

application system *(ZDB specific term)*
A system the application or database runs on. The application or database data is located on source volumes.
See also **backup system** and **source volume**.

archived redo log *(Oracle specific term)*
Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to one (or more) archived log destination(s). This copy is the Archived Redo Log. The presence or absence of an Archived Redo Log is determined by the mode that the database is using:

- ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered from an instance and disk failure. The “hot” backup can be performed only when the database is running in this mode.

- NOARCHIVELOG - The filled online redo log files are not archived.

See also **online redo log.**

archive logging *(Lotus Domino Server specific term)*
Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.

ASR Set
A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk

(disk partitioning and logical volume configuration) and automatic recovery of the original system configuration and user data that was backed up during the full client backup.
These files are stored as an ASR archive file on the Cell Manager (in *<Data_Protector_home>*\Config\Server\dr\asr on a Windows Cell Manager or in /etc/opt/omni/server/dr/asr/ on a UNIX Cell Manager) as well as on the backup medium. The ASR archive file is extracted to three diskettes for 32-bit Windows systems or four diskettes for 64-bit Windows systems after a disaster occurs. You need these diskettes to perform ASR.

autochanger
See **library**

autoloader
See **library**

BACKINT *(SAP R/3 specific term)*
SAP R/3 backup programs can call the Data Protector backint interface program via an open interface, which enables them to communicate with Data Protector software. For backup and restore, SAP R/3 programs issue orders for the Data Protector backint interface.

backup API
The Oracle interface between the Oracle backup/restore utility and the backup/restore media management layer. The interface defines a set of routines to allow the reading and writing of data to the backup media, the creation, searching and removing the backup files.

backup chain
This relates to a situation where full and incremental backups are performed. Based on the level of the incremental backups used (Incr, Incr 1, Incr 2, and so on), simple or rather complex dependencies of incrementals to previous incrementals can exist. The backup chain are all backups, starting from the full backup plus all the dependent incrementals up to the desired point in time.

backup device
A device configured for use with Data Protector, which can write data to and read data from storage media. This can be, for example, a standalone DDS/DAT drive or a library.

backup generation
One backup generation includes one full backup and all incremental backups until the next full backup.

backup ID
An identifier of an integration object that equals the session ID of the backup of this object. The backup ID is preserved when an object is copied, exported, or imported.

backup object
A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database entity or a disk image (rawdisk).

A backup object is defined by:

- Client name: hostname of the Data Protector client where the backup object resides.
- Mount point: the access point in a directory structure (drive on Windows and mount point on UNIX) on the client where the backup object is located.
- Description: uniquely defines backup objects with identical client name and mount point.
- Type: backup object type (for example filesystem or Oracle).

backup owner
Each backup object in the IDB has an owner. The default owner of a backup is the user who starts the backup session.

backup session
A process that creates a copy of data on storage media. The activities are specified in a backup specification or an interactive session. All clients configured in one backup specification are backed up together in one backup session using the same backup type (full or incremental). The result of a backup session is a set of media, which was written to, also called the backup or media set.
See also **incremental backup** and **full backup.**

backup set
A complete set of integration objects associated with a backup.

backup set *(Oracle specific term)*
A logical grouping of backed up files that are created using the RMAN backup command. A backup set is a complete set of files associated with a backup. The files can be multiplexed to increase performance. A backup set contains either datafiles or archived logs, but not both together.

backup specification
A list of objects to be backed up, together with a set of devices or drives to be used, backup options for all objects in the specification, days and time that you want backups to be performed. The objects are entire disks/volumes or parts of them such as files, directories, or even the Windows

Registry for example. File selection lists such as include-lists and exclude-lists can be specified.

backup system *(ZDB specific term)*
A system connected to target volumes of one or multiple application systems. The backup system is typically connected to a backup device to perform the backup of the data in a replica.
See also **application system**, **target volume**, and **replica**.

backup types
See **incremental backup**, **differential backup**, **transaction backup**, **full backup** and **delta backup**.

backup view
Data Protector provides different views for backup specifications: By Type - according to the type of data available for backups/templates. Default view. By Group - according to the group to which backup specifications/templates belong. By Name - according to the name of backup specifications/templates. By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.

BC *(EMC Symmetrix specific term)*
Business Continuance are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.
See also **BCV**.

BC EVA
See **Business Copy EVA**.

BC VA
See **Business Copy VA**.

BC XP
See **Business Copy XP**.

BC Process *(EMC Symmetrix specific term)*
A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices.
See also **BCV**.

Business Copy EVA *(HP StorageWorks EVA specific term)*
For local replication, the HP StorageWorks Business Copy (BC) EVA configuration is used, which allows data replication within the same array. With this, large replica sets can be used, the number of members being limited primarily by the available space on the array. Once established, BC operations continue unattended, providing local data replication.

Business Copy VA *(HP StorageWorks Virtual Array specific term)*
Business Copy VA allows you to maintain internal copies of HP StorageWorks Virtual Array LUNs for data backup or data duplication within the same virtual array. The copies (child or Business Copy LUNs) can be used for various purposes, such as backup, data analysis or development. When used for backup purposes, the original (parent) LUNs are connected to the application system and the Business Copy (child) LUNs are connected to the backup system.
See also **HP StorageWorks Virtual Array LUN**, **application system**, and **backup system**.

Business Copy XP *(HP StorageWorks Disk Array XP specific term)*
The Business Copy XP allows to maintain internal copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data backup or data duplication. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system for various purposes, such as backup and development. For backup purposes, P-VOLs should be connected to the application system, and one of the S-VOL mirror sets to the backup system.
See also **HP StorageWorks Disk Array XP LDEV**, **Continuous Access XP**, **Main Control Unit**, **application system**, and **backup system**.

BCV *(EMC Symmetrix specific term)*
Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected.
See also **BC** and **BC Process**.

Boolean operators
The Boolean operators for the full text search functionality of the online Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition
A volume/disk/partition with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.

BRARCHIVE *(SAP R/3 specific term)*
An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process.
See also **SAPDBA**, **BRBACKUP** and **BRRESTORE**.

BRBACKUP *(SAP R/3 specific term)*
An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files.
See also **SAPDBA**, **BRARCHIVE** and **BRRESTORE**.

BRRESTORE *(SAP R/3 specific term)*
An SAP R/3 tool that can be used to restore files of the following type:

- Database data files, control files, and online redo log files saved with BRBACKUP
- Redo log files archived with BRARCHIVE
- Non-database files saved with BRBACKUP

You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup.
See also **SAPDBA**, **BRBACKUP** and **BRARCHIVE**.

BSM
The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.

CA EVA *(HP StorageWorks Enterprise Virtual Array specific term)*
See **Continuous Access EVA**.

CA XP *(HP StorageWorks Disk Array XP specific term)*
See **Continuous Access XP**.

BSM
The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.

Continuous Access EVA *(HP StorageWorks Enterprise Virtual Array specific term)*
In order to provide high availability and disaster-tolerance of the stored data, two identical EVA disk arrays are interconnected in a so-called Continuous Access configuration (CA EVA).
In this configuration, Data Replication groups are configured in such a way that the data located on the source virtual disks on the local disk array (visible to the application system) is constantly synchronized with the destination

virtual disks on the remote disk array. *See also* **Business Copy EVA** *(HP StorageWorks EVA specific term).*

Continuous Access XP *(HP StorageWorks Disk Array XP specific term)*
Continuous Access XP allows you to create and maintain remote copies of HP StorageWorks Disk Array XP LDEVs for purposes such as data duplication, backup, and disaster recovery. CA operations involve the main (primary) disk arrays and the remote (secondary) disk arrays. The main disk arrays contain the CA primary volumes (P-VOLs), which contain the original data and are connected to the application system. The remote disk arrays contain the CA secondary volumes (S-VOLs) connected to the backup system.
See also **Business Copy XP** *(HP StorageWorks Disk Array XP specific term)*, **Main Control Unit** and **HP StorageWorks Disk Array XP LDEV**.

CAP *(StorageTek specific term)*
Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

catalog protection
Defines how long information about backed up data (such as file names and file versions) is kept in the IDB.
See also **data protection**.

CDB
The Catalog Database is a part of the IDB that contains information about backups, object copies, restores, media management sessions, and backed up data. Depending on the selected logging level, it also contains file names and file versions. This part of the database is always local to the cell.
See also **MMDB**.

CDF file *(UNIX specific term)*
A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.

cell
A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN. Central control is available to administer the backup and restore policies and tasks.

Cell Manager
The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI

used for management tasks can be located on a different system. Each cell has one Cell Manager system.

centralized licensing
Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then allocate licenses to specific cells to suit your needs.
See also **MoM**.

Centralized Media Management Database (CMMDB)
See **CMMDB**.

channel *(Oracle specific term)*
An Oracle Recovery Manager resource allocation. Every allocated channel starts a new Oracle process, which performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used:

- type “disk”
- type ‘SBT_TAPE’

If the specified channel is type ‘SBT_TAPE’ and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.

circular logging *(Microsoft Exchange Server and Lotus Domino Server specific term)*
Circular logging is a Microsoft Exchange Server database and Lotus Domino Server database mode, in which the transaction log file contents are periodically overwritten after the corresponding data is committed to the database. Circular logging reduces disk storage space requirements.

client backup
A backup of all filesystems mounted on a client. Filesystems mounted on the client after the backup specification was created are not automatically detected.

client backup with disk discovery
A backup of all filesystems mounted on a client. When the backup starts, Data Protector discovers the disks on the clients. Client backup with disk discovery simplifies backup configuration and improves backup coverage of systems that often mount or dismount disks.

client or **client system**
Any system configured with any Data Protector functionality and configured in a cell.

cluster-aware application
It is an application that supports the cluster Application Programming Interface. Each cluster-aware

application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on MC/ServiceGuard), application services, IP names and addresses ...).

CMD Script for OnLine Server
(Informix specific term)
Windows CMD script that is created in INFORMIXDIR when Informix OnLine Server is configured. The CMD script is a set of system commands that export environment variables for OnLine Server.

CMMDB
The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the MoM Manager. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended
See also **MoM**.

COM+ Registration Database
(Windows specific term)
The COM+ Registration Database and the Windows Registry store COM+ application attributes, class attributes, and computer-level attributes. This guarantees consistency among these attributes and provides common operation on top of these attributes.

command-line interface
A set of DOS and UNIX like commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.

Command View (CV) EVA *(HP StorageWorks EVA specific term)*
The user interface that allows you to configure, manage, and monitor your HP StorageWorks EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, creating snapclones and snapshots of virtual disks. The Command View EVA software runs on the HP OpenView Storage Management Appliance, and is accessed by a Web browser.
See also **HP StorageWorks EVA Agent (legacy)** and **HP StorageWorks EVA SMI-S Agent**.

concurrency
See **Disk Agent concurrency**.

control file *(Oracle and SAP R/3 specific term)*
An Oracle data file that contains entries specifying the physical structure of the

database. It provides database consistency information used for recovery.

CRS
The Cell Request Server process (service) runs on the Data Protector Cell Manager. It starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager.
CRS runs under the account root on UNIX systems, and under any Windows account. By default, it runs under the account of the user, specified at installation time.

CSM
The Data Protector Copy Session Manager process controls the object copy session and runs on the Cell Manager system.

data file *(Oracle and SAP R/3* specific term)
A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.

data protection
Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions.
See also **catalog protection**.

Data Protector Event Log
A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The Event Log is accessible only to Data Protector users in the Admin group and to Data Protector users who are granted the Reporting and notifications user rights. You can view or delete all events in the Event Log.

Data Protector user account
You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

data stream
Sequence of data transferred over the communication channel.

database library
A Data Protector set of routines that enables data transfer between Data

Protector and a server of an online database integration, for example, the Oracle Server.

database parallelism
More than one database is backed up at a time if the number of available devices allows you to perform backups in parallel.

database server
A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

Dbobject *(Informix specific term)*
An Informix physical database object. It can be a blobspace, dbspace, or logical-log file.

DC directory
The Detail Catalog (DC) directory consists of DC binary files, which store information about file versions. It represents the DCBF part of the IDB, which occupies approximately 80% of the IDB. The default DC directory is called the dcbf directory and is located in the <*Data_Protector_home*>\db40 directory on a Windows Cell Manager and in the /var/opt/omni/server/db40 directory on a UNIX Cell Manager. You can create more DC directories and locate them as appropriate to you. Up to 10 DC directories are supported per cell. The default maximum size of a DC directory is 4 GB.

DCBF
The Detail Catalog Binary Files (DCBF) part of the IDB stores information about file versions and attributes. It occupies approximately 80% of the IDB. One DC binary file is created for each Data Protector medium used for backup.

delta backup
A delta backup is a backup containing all the changes made to the database from the last backup of any type.
See also **backup types**

device
A physical unit which contains either just a drive or a more complex unit such as a library.

device chain
A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group *(EMC Symmetrix specific term*)
A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group

must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.

device streaming
A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.

DHCP server
A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic configuration of IP addresses and related information.

differential backup
An incremental backup (incr) based on any previous Data Protector backup (full or any incremental), which must still be protected.
See **incremental backup**.

differential backup *(MS SQL specific term)*
A database backup that records only the data changes made to the database after the last full database backup.
See also **backup types**.

differential database backup
A differential database backup records only those data changes made to the database after the last full database backup.

direct backup
A SAN-based backup solution in which data movement directly from disk to tape (or to other secondary storage) is facilitated by the SCSI Extended Copy (Xcopy) command. Direct backup lessens the backup I/O load on systems in a SAN environment. The data movement is facilitated directly from disk to tape (or to other secondary storage) by the SCSI Extended Copy (XCopy) command. The command is provided by any element of the infrastructure including bridges, switches, tape libraries, and disk subsystems.
See also **XCopy engine**.

directory junction *(Windows specific term)*
Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.

Directory Store (DS) *(Microsoft Exchange specific term)*
A part of the Microsoft Exchange Server directory. The Microsoft Exchange Server directory contains objects used by Microsoft Exchange applications in order to find and access services, mailboxes, recipients, public folders, and other addressable objects within the messaging system.
See also **Information Store (MDB)**.

disaster recovery
A process to restore a client's main system disk to a state close to the time when a (full) backup was performed.

Disk Agent
A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk.

Disk Agent concurrency
The number of Disk Agents that are allowed to send data to one Media Agent concurrently.

disk discovery
The detection of disks during client backup with disk discovery. During this backup, Data Protector discovers (detects) the disks that are present on the client — even though they might not have been present on the system when the backup was configured — and backs them up. This is particularly useful in dynamic environments, where configurations change rapidly. After the disks are expanded, each inherits all options from its master client object. Even if pre- and post-exec commands are specified once, they are started many times, once per each object.

disk group *(Veritas Volume Manager specific term)*
The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.

disk image (rawdisk) backup
A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.

disk quota
A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.

disk staging
The process of backing up data in several phases to improve the performance of backups and restores, reduce costs of storing the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).

Distributed File System (DFS)
A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ
The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server
In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

domain controller
A server in a network that is responsible for user security and verifying passwords within a group of other servers.

DR image
Data required for temporary disaster recovery operating system (DR OS) installation and configuration.

DR OS
A disaster recovery operating system is an operating system environment in which disaster recovery runs. It provides Data Protector a basic runtime environment (disk, network, tape, and filesystem access). The OS has to be installed and configured before the Data Protector disaster recovery can be performed. DR OS not only hosts the Data Protector disaster recovery process but is also a part of the restored system because it replaces its own configuration data with the original configuration data.

drive
A physical unit that receives data from a computer system and can write it onto a magnetic medium (typically a tape drive). It can also read the data from the medium and send it to the computer system.

drive index
A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.

dynamic client
See **client backup with disk discovery**.

EMC Symmetrix Agent (SYMA)

(EMC Symmetrix specific term)
See **Symmetrix Agent (SYMA)**

emergency boot file *(Informix specific term)*
An Informix configuration file that resides in the <*INFORMIXDIR*>\etc directory (on HP-UX) or <*INFORMIXDIR*>/etc directory (on Windows) and is called ixbar.<*server_id*>, where <*INFORMIXDIR*> is the OnLine Server home directory and <*server_id*> is the value of the SERVERNUM configuration parameter. Each line of the emergency boot file corresponds to one backup object.

Enterprise Backup Environment
Several cells can be grouped together and managed from a central cell. The enterprise backup environment includes all clients located in several Data Protector cells which are managed and administered from a central cell using the Manager-of-Managers concept.
See also **MoM**.

Event Logs
Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.

exchanger
Also referred to as SCSI Exchanger.
See also **library**.

exporting media
A process that removes all data about backup sessions, such as systems, objects, and file names, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged.
See also **importing media.**

Extensible Storage Engine (ESE)

(Microsoft Exchange Server 2000/2003 specific term)
A database technology used as a storage system for information exchange in Microsoft Exchange Server 2000/2003.

failover
Transferring of the most important cluster data, called group (on Windows)

or package (on Unix) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.

FC bridge
See **Fibre Channel bridge**

Fibre Channel
An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed bidirectional transmission of large data files and can be deployed between sites kilometers apart.
Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.

Fibre Channel bridge
A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.

file depot
A file containing the data from a backup to a file library device.

file jukebox device
A device residing on disk consisting of multiple slots used to store file media.

file library device
A device which resides on a disk emulating a library with several media, hence containing multiple files, referred to as file depots.

File Replication Service (FRS)
A Windows service that replicates the domain controller store logon scripts and group policies. FRS also enables replication of Distributed File System (DFS) shares between systems and allows any server to perform replication activity.

file version
The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.

filesystem
The organization of files on a hard disk. A filesystem is backed up so that the file attributes and the file contents are stored on the backup media.

first level mirror (*HP StorageWorks Disk Array XP specific term*)
HP StorageWorks Disk Array XP allows up to three mirror copies of a Primary Volume and each of these copies can have additional two copies. The three mirror copies are called first level mirrors.
See also **Primary Volume**, and **MU numbers**.

fnames.dat
The fnames.dat files of the IDB contain information on the names of the backed up files. Typically, these files occupy about 20% of the IDB, if filenames are stored.

formatting
A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (media ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.

free pool
An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.

full backup
A backup in which all selected objects are backed up, whether or not they have been recently modified.
See also **backup types**.

full database backup
A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.

full mailbox backup
A full mailbox backup is a backup of the entire mailbox content.

global options file
A file that allows you to customize Data Protector. It explains the global options, which cover various aspects of Data Protector, typically time-outs and limits, and affect the entire Data Protector cell. The file is located in the /etc/opt/omni/server/options directory on HP-UX and Solaris systems and in the *<Data_Protector_home>*\Config\Server\Options directory on Windows systems.

group *(Microsoft Cluster Server specific term*)
A collection of resources (for example disk volumes, application services, IP names and addresses) that are needed to run a specific cluster-aware applications.

GUI
A cross-platform (HP-UX, Solaris, and Windows) graphical user interface, provided by Data Protector for easy access to all configuration, administration, and operation tasks.

hard recovery *(Microsoft Exchange Server specific term)*
A Microsoft Exchange Server database recovery that is performed after a restore by the database engine, using transaction log files.

heartbeat
A cluster data set with a time stamp carrying information about the operational status of a particular cluster node. This data set or packet is distributed among all cluster nodes.

Hierarchical Storage Management (HSM)
A method for optimizing the use of expensive hard disk storage by migrating less frequently used data to less expensive optical platters. When needed, the data is migrated back to hard disk storage. This balances the need for fast retrieval from hard disk with the lower cost of optical platters.

Holidays file
A file that contains information about holidays. You can set different holidays by editing the Holidays file: /etc/opt/omni/server/Holidays on the UNIX Cell Manager and *<Data_Protector_home>*\Config\Server\holidays on the Windows Cell Manager.

host backup
See **client backup with disk discovery**.

hosting system
A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.

HP ITO
See **OVO**.

HP OpC
See **OVO**.

HP OpenView SMART Plug-In (SPI)
A fully integrated, out-of-the-box solution which "plugs into" HP OpenView Operations, extending the managed domain. Through the Data Protector integration, which is implemented as an HP OpenView SMART Plug-In, a user can have an arbitrary number of Data Protector Cell Managers monitored as an extension to HP OpenView Operations (OVO).

HP OVO
See **OVO**.

HP StorageWorks Disk Array XP LDEV
A logical partition of a physical disk within an HP StorageWorks Disk Array XP. LDEVs are entities that can be replicated in the Continuous Access XP (CA) and Business Copy XP (BC) configurations, or can be used as standalone entities.
See also **BC** *(HP StorageWorks Disk Array XP specific term)*, **CA** *(HP StorageWorks Disk Array XP specific term)*, and **replica**.

HP StorageWorks EVA Agent (legacy)
A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration operating on HP StorageWorks EVA with Command View (CV) EVA software v3.1 or lower, and the EVA VCS firmware v3.01x or lower.
See also **Command View (CV) EVA** and **HP StorageWorks EVA SMI-S Agent**.

HP StorageWorks EVA SMI-S Agent
A Data Protector software module that executes all tasks required for the HP StorageWorks Enterprise Virtual Array integration operating on HP StorageWorks EVA with Command View (CV) EVA software starting with v3.2. With the EVA SMI-S Agent, the control over the array is established through HP StorageWorks SMI-S EVA provider, which directs communication between incoming requests and CV EVA.
See also **Command View (CV) EVA**, **HP StorageWorks SMI-S EVA provider**, and **HP StorageWorks EVA Agent (legacy)**.

HP StorageWorks SMI-S EVA provider
An interface used for controlling HP StorageWorks Enterprise Virtual Array. SMI-S EVA provider runs as a separate service on the HP OpenView Storage Management Appliance system and acts as a gateway between incoming requests and Command View EVA. With the Data Protector HP StorageWorks EVA integration, SMI-S EVA provider accepts standardized requests from the EVA SMI-S Agent, communicates with Command View EVA for information or method invocation, and returns standardized responses.
See also **HP StorageWorks EVA SMI-S Agent** and **Command View (CV) EVA**.

HP StorageWorks Virtual Array LUN
A logical partition of a physical disk within an HP StorageWorks Virtual Array. LUNs are entities that can be replicated in the HP StorageWorks

Business Copy VA configuration, or can be used as standalone entities.
See also **BC VA** and **replica**.

HP VPO
See **OVO**.

ICDA *(EMC Symmetrix specific term)*
EMC's Symmetrix Integrated Cached Disk Arrays (ICDA) is a disk array device that combines a set of physical disks, a number of FWD SCSI channels, an internal cache memory, and control and diagnostic software commonly referred to as the microcode.

IDB
The Data Protector Internal Database is an embedded database located on the Cell Manager that keeps information regarding which data is backed up, on which media it is backed up, how backup and restore sessions are run, and which devices and libraries are configured.

importing media
A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media.
See also **exporting media.**

incremental backup
A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, allowing selective backup of only files that have changed since the last incremental backup.
See also **backup types**.

incremental backup *(Microsoft Exchange Server specific term)*
A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up.
See also **backup types**.

incremental mailbox backup
An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.

incremental1 mailbox backup
An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.

incremental (re)-establish *(EMC Symmetrix specific term)*
A BCV or SRDF control operation. In BCV control operations, an incremental establish causes the BCV device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.
In SRDF control operations, an incremental establish causes the target

(R2) device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.

incremental restore *(EMC Symmetrix specific term)*
A BCV or SRDF control operation. In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was written to the target (R2) device during the time of the original pair split, and the data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

Inet
A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.

Information Store *(Microsoft Exchange Server 2000/2003 specific term)*
The Microsoft Exchange Server 2000/2003 service that is responsible for storage management. Information Store in Microsoft Exchange Server 2000/2003 manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users.
See also **Key Management Service** and **Site Replication Service**.

Information Store *(Microsoft Exchange Server 5.5 specific term)*
This is the default message store provider for the Microsoft Exchange Server 5.5. Information Store consists of the following stores:

- public information store
- private information store
- personal folder store

- offline information store.

The public information store contains public folders and messages that can be shared among multiple users and applications. A single public store is shared by all users within an Exchange Server 5.5 organization, even if multiple Exchange Servers are used. The private information store consists of mail boxes that can belong to users or to applications. The mail boxes reside on the server running the Exchange Server 5.5.
See also **Directory Store (DS)**.

initializing
See **formatting**.

Installation Server
A computer system that holds a repository of the Data Protector software packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.

instant recovery (*ZDB specific term*)
A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape sessions, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore from tape. Depending on the application/database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery.
See also **replica**, **zero downtime backup (ZDB)**, **ZDB to disk**, and **ZDB to disk+tape**.

integrated security *(MS SQL specific term)*
Integrated security allows the Microsoft SQL Server to use Windows authentication mechanisms to validate Microsoft SQL Server logins for all connections. Using integrated security means that users have one password for both Windows and Microsoft SQL Server. Integrated security should be used in environments where all clients support trusted connections. Connections validated by Windows Server and accepted by Microsoft SQL Server are referred to as trusted connections. Only trusted connections are allowed.

integration object
A backup object of a Data Protector integration, such as Oracle or SAP DB.

Internet Information Server (IIS)

(Windows specific term)
Microsoft Internet Information Server is a network file and application server

that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).

IP address
Internet Protocol address is a numeric address of a system used to uniquely identify the system on the network. The IP address consists of four groups of numbers separated by periods (full stops).

ISQL *(Sybase specific term)*
A Sybase utility used to perform system administration tasks on Sybase SQL Server.

ITO
See **OVO**.

jukebox
See **library**.

jukebox device
A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the "file jukebox device".

Key Management Service *(Microsoft Exchange Server 2000/2003 specific term)*
The Microsoft Exchange Server 2000/2003 service that provides encryption functionality for enhanced security. *See also* **Information Store** and **Site Replication Service**.

LBO *(EMC Symmetrix specific term)*
A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.

library
Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.

lights-out operation or **unattended operation**
A backup or restore operation that takes place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.

LISTENER.ORA *(Oracle specific term)*
An Oracle configuration file that describes one or more Transparent Network Substrate (TNS) listeners on a server.

load balancing
By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If you do not want to use load balancing, you can select which device will be used for each object in the backup specification. Data Protector will access the devices in the specified order.

local and remote recovery
Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.

lock name
You can configure the same physical device several times with different characteristics, by using different device names. The lock name is a user specified string that is used for locking all such device configurations to prevent collision if several such devices (device names) are used concurrently. Use an identical lock name for all device definitions which use the same physical device.

log_full shell script *(Informix UNIX specific term)*
A script provided by ON-Bar that you can use to start backing up logical-log files when OnLine Server issues a log-full event alarm. The Informix ALARMPROGRAM configuration parameter defaults to the <*INFORMIXDIR*>/etc/log_full.sh, where <*INFORMIXDIR*> is the OnLine Server home directory. If you do not want logical logs to be backed up continuously, set the ALARMPROGRAM configuration parameter to <*INFORMIXDIR*>/etc/no_log.sh.

logging level
The logging level determines the amount of details on files and directories written to the IDB during backup or object copying. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings

influence the IDB growth, backup speed, and the convenience of browsing data for restore.

logical-log files
This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.

login ID *(MS SQL Server specific term)*
The name a user uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.

login information to the Oracle Target Database *(Oracle and SAP R/3 specific term)*
The format of the login information is <user_name>/<password>@<service>, where:

- <user_name> is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have been granted Oracle SYSDBA or SYSOPER rights.
- <password> is a string used for data security and known only to its owner. Passwords are entered to connect to an operating system or software application. The password has to be the same as the password specified in the Oracle password file (orapwd), which is used for authentication of users performing database administration.
- <service> is the name used to identify an SQL*Net server process for the target database.

login information to the Recovery Catalog Database *(Oracle specific term)*
The format of the login information to the Recovery (Oracle) Catalog Database is <user_name>/<password>@<service>, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database. In this case, <service> is the name of the service to the Recovery Catalog Database, not the Oracle target database.

Note that the Oracle user specified here has to be the owner of the Oracle Recovery (Oracle) Catalog.

Lotus C API *(Lotus Domino Server specific term)*
An interface for the exchange of backup

and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM
A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system consists of several volume groups, where each volume group has several volumes.

Magic Packet
See **Wake ONLAN**.

mailbox (*Microsoft Exchange Server specific term)*
The location to which e-mail is delivered, which is set up by the administrator for each user. If a set of personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.

Mailbox Store (*Microsoft Exchange Server 2000/2003 specific term)*
A part of the Information Store that maintains information about user mailboxes. A mailbox store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

Main Control Unit (MCU) *(HP StorageWorks Disk Array XP specific term)*
An HP StorageWorks XP disk array that contains the primary volumes for the Continuous Access configuration and acts as a master device.
See also **BC** *(HP StorageWorks Disk Array XP specific term)*, **CA** *(HP StorageWorks Disk Array XP specific term)*, and **HP StorageWorks Disk Array XP LDEV**.

Manager-of-Managers (MoM)
See **Enterprise Cell Manager**.

Media Agent
A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore session, a Media Agent locates data on the backup medium and sends it to the Disk Agent. The Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

MAPI *(Microsoft Exchange specific term)*
The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

media allocation policy
Determines in which sequence media are used for backup. The Strict allocation policy directs Data Protector to prompt for a specific medium. The Loose policy directs Data Protector to prompt for any suitable medium. The Formatted First policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library.

media condition
The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of read and write errors with tape media. Media need to be replaced when they are marked as POOR.

media condition factors
The user-assigned age threshold and overwrite threshold used to determine the state of a medium.

media ID
A unique identifier assigned to a medium by Data Protector.

media label
A user-defined identifier used to describe a medium.

media location
A user-defined physical location of a medium, such as "building 4" or "off-site storage".

media management session
A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.

media pool
A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.

media set
The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.

media type
The physical type of media, such as DDS or DLT.

media usage policy
The media usage policy controls how new backups are added to the already used media. It can be Appendable, Non-Appendable, or Appendable for incrementals only.

merging
This defines one mode to resolve file conflicts during restore. If the file to be restored already exists at the destination, the one with the more recent

modification date is kept. Files not present on the disk are always restored. *See also* **overwrite**.

MFS
The Migrating File System enables a standard JFS filesystem with migration capabilities (on HP-UX 11.00). The MFS is accessed via a standard filesystem interface (DMAPI), it is mounted to a directory the same way as any HP-UX filesystem. In an MFS, only the superblock, the inode and the 'extended attribute' information remain permanently on the hard disk and are never migrated.
See also **VBFS**.

Microsoft Exchange Server
A “client-server” messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.

Microsoft Management Console (MMC) *(Windows specific term)*
An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.

Microsoft SQL Server 7.0/2000
A database management system designed to meet the requirements of distributed "client-server" computing.

Microsoft Volume Shadow Copy service (VSS)
A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets.
See also **shadow copy, shadow copy provider, writer**.

mirror *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
See **target volume**.

mirror rotation *(HP StorageWorks Disk Array XP specific term)*
See **replica set rotation**.

MMD
The Media Management Daemon process (service) runs on the Data

Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.

MMDB
The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells.
See also **CMMDB**, **CDB**.

MoM
Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The MoM allows you to configure and manage multiple cells from a central point.

mount request
A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.

mount point
The access point in a directory structure for a disk or logical volume, for example /opt or d:. On UNIX, the mount points are displayed using the bdf or df command.

MSM
The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.

MU number (*HP StorageWorks Disk Array XP specific term*)
A Mirror Unit number is an integer number (0, 1 or 2), used to indicate a first level mirror.
See also **first level mirror**.

multi-drive server
A license that allows you to run an unlimited number of Media Agents on a single system. This license, which is bound to the IP address of the Cell Manager, is no longer available.

obdrindex.dat
An IDB file with information about IDB backups, media, and devices used for the backup. This information can significantly simplify IDB recovery. It is recommended to relocate the file, together with IDB transaction logs, on a separate physical disk from other IDB

directories, and, additionally, to make a copy of the file and locate it where you want.

OBDR capable device
A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.

object
See **backup object**

object copy
A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.

object copy session
A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.

object copying
The process of copying selected object versions to a specific media set. You can select object versions from one or several backup sessions to be copied.

Object ID *(Windows specific term)*
The object IDs (OIDs) enable access to NTFS 5 files no matter where in the system the files reside. Data Protector treats the OIDs as alternate streams of the files.

object mirror
A copy of a backup object created using object mirroring. Object mirrors are often referred to as object copies.

object mirroring
The process of writing the same data to several media sets during a backup session. Data Protector enables you to mirror all or some backup objects to one or more media sets.

offline backup
A backup during which an application database cannot be used by the application.

- For simple backup methods (non ZDB), the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the whole backup period (~minutes/hours). For instance, for backup to tape, until streaming of data to the tape is finished.

- For ZDB methods, the database is also put into the quiescent state, but for the period of the data replication process only (~seconds). Normal

database operation can then be resumed for the rest of the backup process.

See also **zero downtime backup (ZDB)** and **online backup**.

offline recovery
Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI library devices can be used for offline recovery. Recovery of the Cell Manager is always offline.

offline redo log
See **archived redo log**

OmniStorage
Software providing transparent migration of less frequently used data to the optical library while keeping more frequently used data on the hard disk. HP OmniStorage runs on HP-UX systems.

On-Bar *(Informix specific term)*
A backup and restore system for OnLine Server. ON-Bar enables you to create a copy of your OnLine Server data and later restore the data. The ON-Bar backup and restore system involves the following components:

- onbar utility
- Data Protector, as the backup solution
- XBSA interface
- ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.

onbar utility *(Informix specific term)*
The Informix utility that communicates backup and restore requests to OnLine Server. The utility uses XBSA to exchange control data and back up and restore data with Data Protector.

ONCONFIG *(Informix specific term)*
An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, OnLine uses the configuration values from the file *<INFORMIXDIR>*/etc/onconfig (on HP-UX) or *<INFORMIXDIR>*\etc\onconfig (on Windows).

online backup
A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period that the backup application requires access to the original data objects. During this period, the database

is fully operational, but there may be a small performance impact and log files may grow very quickly.

- For simple backup methods (non ZDB), backup mode is required for the whole backup period (~minutes/ hours). For instance, for backup to tape, until streaming of data to tape is finished.

- For ZDB methods, backup mode is required for the short period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process.

In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored.
See also **zero downtime backup (ZDB)** and **offline backup**.

online redo log *(Oracle specific term)*
Redo logs that have not been archived, but are either available to the instance for recording database activity or are filled and waiting to be archived or reused.
See also **archived redo log**.

OnLine Server *(Informix specific term)*
Refers to INFORMIX-OnLine Dynamic Server.

OpC
See **OVO**.

Oracle instance *(Oracle specific term)*
Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.

ORACLE_SID *(Oracle specific term)*
A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired *<ORACLE_SID>*. The *<ORACLE_SID>* is included in the CONNECT DATA parts of the connect descriptor in a TNSNAMES.ORA file and in the definition of the TNS listener in the LISTENER.ORA file.

original system
The system configuration backed up by Data Protector before a computer disaster hits the system.

overwrite
An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files.
See also **merging**.

OVO
HP OpenView Operations for Unix provides powerful capabilities for operations management of a large number of systems and applications on

in a network. Data Protector provides an integration into this management product. This integration is implemented as a SMART Plug-In for OVO management servers on HP-UX and Solaris. Earlier versions of OVO were called IT/Operation, Operations Center and Vantage Point Operations.
See also **merging**.

ownership
The ownership of a backup determines who can restore from the backup. The user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the session is not considered interactive. In that case, if the backup owner has been defined in the backup specification, they remain the session owner. Otherwise, the session owner becomes the user who started the backup in question. For the scheduled backups, by default, the session owner is for the UNIX Cell Manager: root.sys@<Cell Manager>, and for the Windows Cell Manager, the user that was specified during the installation of the Cell Manager. It is possible to modify the ownership, so that the specific user becomes the session owner.

P1S file
P1S file contains information on how to format and partition all disks installed in the system during Enhanced Automated Disaster Recovery (EADR). It is created during full backup and is saved on backup medium and on Cell Manager into *<Data_Protector_home>*\Config\Server\dr\p1s directory on a Windows Cell Manager or in /etc/opt/omni/server/dr/p1s directory on a UNIX Cell Manager with the filename recovery.p1s.

package *(MC/ServiceGuard and Veritas Cluster specific term)*
A collection of resources (for example volume groups, application services, IP names and addresses) that are needed to run a specific cluster-aware application.

pair status *(HP StorageWorks Disk Array XP specific term)*
A mirrored pair of disks can have various status values depending on the action performed on it. The three most important status values are:

- COPY - The mirrored pair is currently resynchronizing. Data is transferred from one disk to the other. The disks do not contain the same data.
- PAIR - The mirrored pair is completely synchronized and both disks (the primary volume and the mirrored volume) contain identical data.

- SUSPENDED - The link between the mirrored disks is suspended. That means that both disks are accessed and updated independently. However, the mirror relationship is still maintained and the pair can be resynchronized without transferring the complete disk.

parallel restore
Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

parallelism
The concept of reading multiple data streams from an online database.

physical device
A physical unit that contains either a drive or a more complex unit such as a library.

post-exec
A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
See also **pre-exec**.

pre- and post-exec commands
Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows and as shell scripts on UNIX.

prealloc list
A subset of media in a media pool that specifies the order in which media are used for backup.

pre-exec
A backup option that executes a command or script before the backup of an object or before the entire session is started. Pre-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows and as shell scripts on UNIX.
See also **post-exec**.

Primary Volume (P-VOL) *(HP StorageWorks Disk Array XP specific term)*

Standard HP StorageWorks Disk Array XP LDEVs that act as a primary volume for the CA and BC configurations. The P-VOL is located in the MCU.
See also **Secondary Volume (S-VOL)**.

Private Information Store *(Microsoft Exchange Server 5.5 specific term)*
A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text .edb file.

protection
See **data protection** and also **catalog protection**.

public folder store (*Microsoft Exchange Server 2000/2003 specific term)*
The part of the Information Store that maintains information in public folders. A public folder store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

public/private backed up data
When configuring a backup, you can select whether the backed up data will be:

- public, that is visible (and accessible for restore) to all Data Protector users
- private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

RAID
Redundant Array of Inexpensive Disks.

RAID Manager Library *(HP StorageWorks Disk Array XP specific term)*
The RAID Manager Library is used internally by Data Protector on Solaris systems to allow access to HP StorageWorks Disk Array XP configuration, status, and performance data and to key HP StorageWorks Disk Array XP features through the use of function calls translated into a sequence of low level SCSI commands.

RAID Manager XP *(HP StorageWorks Disk Array XP specific term)*
The RAID Manager XP application provides an extensive list of commands to report and control the status of the CA and BC applications. The commands communicate through a RAID Manager instance with the HP StorageWorks Disk Array XP Disk Control Unit. This instance translates the commands into a sequence of low level SCSI commands.

rawdisk backup
See **disk image backup**.

RCU *(HP StorageWorks specific term)*
The Remote Control Unit acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

RDBMS
Relational Database Management System.

RDF1/RDF2 *(EMC Symmetrix specific term)*
A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.

RDS
The Raima Database Server process (service) runs on the Data Protector Cell Manager and manages the IDB. The process is started when Data Protector is installed on the Cell Manager.

Recovery Catalog *(Oracle specific term)*
A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about:

- The physical schema of the Oracle target database
- Data file and archived log backup sets
- Data file copies
- Archived Redo Logs
- Stored scripts.

Recovery Catalog Database *(Oracle specific term)*
An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

RecoveryInfo
When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume, and network configuration). This information is needed for disaster recovery.

Recovery Manager (RMAN) *(Oracle specific term)*
An Oracle command-line interface that directs an Oracle Server process to back up, restore, or recover the database it is connected to. RMAN uses either the recovery catalog or the control file to

store information about backups. This information can be used later in restore sessions.

recycle
A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.

redo log *(Oracle specific term)*
Every Oracle database has a set of two or more redo log files. The set of redo log files for the database is known as the database's redo log. Oracle uses the redo log to record all changes made to data.

Remote Control Unit *(HP StorageWorks Disk Array XP specific term)*
The Remote Control Unit (RCU) acts as a slave of an MCU in a CA configuration. In bidirectional configurations, the RCU can act as an MCU.

Removable Storage Management Database *(Windows specific term)*
A Windows service used for managing removable media (such as tapes and disks) and storage devices (libraries). Removable Storage allows applications to access and share the same media resources.

reparse point *(Windows specific term)*
A system-controlled attribute that can be associated with any directory or file. The value of a reparse attribute can have user-controlled data. The format of the data is understood by the application that stored the data and a filesystem filter that was installed to interpret the data and process such files. Whenever the filesystem encounters a file with a reparse point, it attempts to find the filesystem filter associated with the data format.

replica *(ZDB specific term)*
An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware/software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror), or a virtual copy (for example, a snapshot). From a host's perspective, on a basic UNIX or Windows system, the complete physical disk containing a backup object is replicated. However, if a volume manager is used on UNIX, the whole volume/disk group containing a backup object is replicated.
See also **snapshot**, **snapshot creation**, **split mirror**, and **split mirror creation**.

replica set *(ZDB specific term)*
A group of replicas, all created using the same backup specification.
See also **replica** and **replica set rotation**.

replica set rotation *(ZDB specific term)*
The use of a replica set for regular backup production: Each time the same backup specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set.
See also **replica** and **replica set**.

restore session
A process that copies data from backup media to a client.

RMAN *(Oracle specific term)*
See **Recovery Manager**.

RSM
The Data Protector Restore Session Manager controls the restore session. This process always runs on the Cell Manager system.

RSM *(Windows specific term)*
Removable Storage Manager (RSM) includes a media management service that facilitates communication among applications, robotic changers, and media libraries. It enables multiple applications to share local robotic media libraries and tape or disk drives and to manage removable media.

SAPDBA *(SAP R/3 specific term)*
An SAP R/3 user interface that integrates the BRBACKUP, BRARCHIVE, and BRRESTORE tools.

scan
A function that identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library).

scanning
A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the device if someone has manually manipulated media without using Data Protector to eject or enter, for example.

Scheduler
A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.

Secondary Volume (S-VOL) *(HP StorageWorks Disk Array XP specific term*)
Secondary Volumes, or S-VOLs, are XP

LDEVs that act as a secondary CA or BC mirror of another LDEV (a P-VOL). In the case of CA, S-VOLs can be used as failover devices in a MetroCluster configuration. The S-VOLs are assigned separate SCSI addresses, different from the addresses used by the P-VOLs.
See also **Primary Volume (P-VOL).**

session
See **backup session**, **media management session,** and **restore session**.

session ID
An identifier of a backup, restore, object copy, or media management session, consisting of the date when the session ran and a unique number.

session key
This environment variable for the Pre- and Post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and it is used for specifying options for the omnimnt, omnistat and omniabort CLI commands.

shadow copy *(MS VSS specific term)*
A volume that represents a duplicate of the original volume at a certain point in time. The data is then backed up from the shadow copy and not from the original volume. The original volume continues to changes as the backup process continues, but the shadow copy of the volume remains constant.
See also **Microsoft Volume Shadow Copy service**.

shadow copy provider *(MS VSS specific term)*
An entity that performs the work on creating and representing the volume shadow copies. Providers own the shadow copy data and expose the shadow copies. Providers can be software (for example, system providers) or hardware (local disks, disk arrays).
See also **shadow copy.**

shadow copy set *(MS VSS specific term)*
A collection of shadow copies created at the same point in time.
See also **shadow copy**.

shared disks
A Windows disk on another system that has been made available to other users on the network. Systems with shared disks can be backed up without a Data Protector Disk Agent installed.

SIBF
The Serverless Integrations Binary Files (SIBF) is a part of the IDB that stores raw NDMP meta data. This data is necessary to perform restore of NDMP objects.

Site Replication Service *(Microsoft Exchange Server 2000/2003 specific term)*
The Microsoft Exchange Server 2000/ 2003 service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service.
See also **Information Store** and **Key Management Service**.

slot
A mechanical position in a library. Each slot can hold a medium, such as a DLT tape. Data Protector references each slot with a number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.

SMB
See **split mirror backup**.

SMBF
The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, object copy, restore, and media management sessions. One binary file is created per session. The files are grouped by year and month.

snapshot *(HP StorageWorks VA and HP StorageWorks EVA specific term)*
A form of replica produced using snapshot creation techniques. A range of snapshot types is available, with different characteristics, depending on the arrays/techniques used. Such replicas are dynamic and may be either virtual copies, still reliant upon the contents of the source volumes, or independent exact duplicates (clones), depending on the snapshot type and the time since creation.
See also **replica** and **snapshot creation**.

snapshot backup *(HP StorageWorks VA and HP StorageWorks EVA specific term)*
See **ZDB to tape**, **ZDB to disk**, and **ZDB to disk+tape**.

snapshot creation *(HP StorageWorks VA and HP StorageWorks EVA specific term)*
A replica creation technique, in which copies of source volumes are created using storage virtualization techniques. The replicas are considered to be created at one particular point-in-time, without pre-configuration, and are immediately available for use. However background copying processes normally continue after creation.
See also **snapshot**.

source (R1) device *(EMC Symmetrix specific term)*
An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix

unit. An R1 device must be assigned to an RDF1 group type.
See also **target (R2) device**.

source volume *(ZDB specific term)*
A storage volume containing data to be replicated.

sparse file A file that contains data with portions of empty blocks. Examples are: -A matrix in which some or much of the data contains zeros -files from image applications -high-speed databases If sparse file processing is not enabled during restore, it might be impossible to restore this file.

split mirror *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
A replica created using split mirror techniques. Such a replica provides an independent, exact duplicate, or clone, of the contents of the source volumes.
See also **replica** and **split mirror creation**.

split mirror backup *(EMC Symmetrix specific term)*
See **ZDB to tape**.

split mirror backup *(HP StorageWorks Disk Array XP specific term)*
See **ZDB to tape**, **ZDB to disk**, and **ZDB to disk+tape**.

split mirror creation *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes.
See also **split mirror**.

split mirror restore *(EMC Symmetrix and HP StorageWorks Disk Array XP specific term)*
A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is restored from tape media to a split mirror replica, which is then synchronized to the source volumes. Individual backup objects or complete sessions can be restored using this method.
See also **ZDB to tape**, **ZDB to disk+tape**, and **replica**.

sqlhosts file *(Informix specific term)*
An Informix connectivity-information file that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.

SRD file
The Data Protector System Recovery Data (SRD) file contains system information required for installing and configuring the operating system in case of a disaster. The SRD file is an ASCII file, generated when a CONFIGURATION backup is performed on a Windows client and stored on the Cell Manager.

SRDF *(EMC Symmetrix specific term)*
The EMC Symmetrix Remote Data Facility is a business continuation process that enables effective, real-time data replication of SLDs between dislocated processing environments. These environments could be situated within the same root computer environment or separated by long distances.

SSE Agent *(HP StorageWorks Disk Array XP specific term)*
A Data Protector software module that executes all tasks required for a split mirror backup integration. It communicates with the HP StorageWorks Disk Array XP storing system using the RAID Manager XP utility (HP-UX and Windows systems) or RAID Manager Library (Solaris systems).

sst.conf file
The file /usr/kernel/drv/sst.conf is required on each Data Protector Sun Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.

st.conf file
The file /kernel/drv/st.conf is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.

stackers
Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.

standalone file device
A file device is a file in a specified directory to which you back up data.

standard security *(MS SQL specific term)*
Standard security uses the login validation process of the Microsoft SQL Server for all connections. Standard security is useful in network environments with a variety of clients, some of which may not support trusted

connections. It also provides backward compatibility for older versions of the Microsoft SQL Server.
See also **integrated security**.

Storage Group
(Microsoft Exchange Server 2000/2003 specific term)
A collection of databases (stores) that share a common set of transaction log files. Exchange manages each storage group with a separate server process.

StorageTek ACS library
(StorageTek specific term)
Automated Cartridge System is a library system (also known as Silo) consisting of one Library Management Unit (LMU) and one to 24 Library Storage Modules (LSM) connected to the unit.

storage volume *(ZDB specific term)*
A storage volume represents an object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, file systems, or other objects may exist. The volume management systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array.

switchover
See **failover**

Sybase Backup Server API *(Sybase specific term)*
An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.

Sybase SQL Server *(Sybase specific term)*
The server in the Sybase "client-server" architecture. Sybase SQL Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.

Symmetrix Agent (SYMA) *(EMC Symmetrix specific term)*
The Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

System Backup to Tape *(Oracle specific term)*
An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.

system databases *(Sybase specific term)*
The four system databases on a newly installed Sybase SQL Server are the:

- master database (master)
- temporary database (tempdb)
- system procedure database (sybsystemprocs)
- model database (model).

system disk
A system disk is a disk containing operating system files. Microsoft terminology defines the system disk as a disk containing the files required for initial step of boot process.

system partition
A system partition is a partition containing operating system files. Microsoft terminology defines a system partition as a partition containing the files required for initial step of boot process.

System State *(Windows specific term)*
The System State data comprises the Registry, COM+ Class Registration database, system startup files, and the Certificate Services database (if the server is a certificate server). If the server is a domain controller, Active Directory directory services and the Sysvol directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.

system volume/disk/partition
A volume/disk/partition containing operating system files. Microsoft terminology defines the system volume/disk/partition as the volume/disk/partition containing files required for the initial step of the boot process.

SysVol (*Windows specific term)*
A shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.

tablespace
A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.

tapeless backup *(ZDB specific term)*
See **ZDB to disk**.

target database *(Oracle specific term)*
In RMAN, the target database is the database that you are backing up or restoring.

target (R2) device *(EMC Symmetrix specific term)*
An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type.
See also **source (R1) device**

target system *(Disaster Recovery specific term)*
A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore this system to the original system configuration. The difference between a crashed system and a target system is that a target system has all faulty hardware replaced.

target volume *(ZDB specific term)*
A storage volume to which data is replicated.

Terminal Services *(Windows specific term)*
Windows Terminal Services provide a multi-session environment that allows clients to access a virtual Windows desktop session and Windows-based programs running on the server.

thread *(MS SQL Server 7.0/2000 specific term)*
An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.

TimeFinder *(EMC Symmetrix specific term)*
A business continuation process that creates an instant copy of single or multiple Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).

TLU
Tape Library Unit.

TNSNAMES.ORA *(Oracle and SAP R/3 specific term)*
A network configuration file that contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.

transaction
A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes.

transaction backup
Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred.

transaction backup *(Sybase and SQL specific term*)
A backup of the transaction log providing a record of changes made since the last full or transaction backup.

transaction log backup
Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.

transaction log files
Files that record transactions of the database modifications, and provide fault tolerance in case of a database disaster.

transaction logs (*Data Protector specific term)*
Keeps track of IDB changes. The archiving of transaction logs should be enabled to prevent you from losing the transaction log files that are created after the last IDB backup and are necessary for IDB recovery.

transaction log table *(Sybase specific term)*
A system table in which all changes to the database are automatically recorded.

transportable snapshot *(MS VSS specific term)*
A shadow copy that is created on the application system and can be presented to the backup system which performs the backup.
See also **Microsoft Volume Shadow Copy service (VSS)**.

TSANDS.CFG file *(Novell NetWare specific term)*
A file that allows you to specify the names of containers where you want backups to begin. It is text file located in the SYS:SYSTEM\TSA directory on the server where TSANDS.NLM is loaded.

unattended operation
See **lights-out operation**.

user account
You can use Data Protector only if you have a Data Protector user account,

which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

user disk quotas
NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. Data Protector backs up user disk quotas on the whole system and for all configured users at a time.

user group
Each Data Protector user is member of a User Group. Each User Group has a set of user rights that are given to every user in that User Group. The number of User Groups with their associated user rights can be defined as desired. Data Protector provides three default user groups: admin, operator, and user.

user profile *(Windows specific term)*
Configuration information retained on a user basis. This information includes desktop settings, screen colors, network connections, and so on. When the user logs on, the user profile is loaded and the Windows environment is set accordingly.

user rights
User rights or access rights are the permissions needed to perform specific Data Protector tasks. Configuring a backup, starting a backup session, or starting a restore session are typical user rights. Users have the access rights of the user group to which they belong.

vaulting media
The process of storing media to a safe and remote place. The media are brought back to the data center when they are needed for restore or are ready for reuse in subsequent backups. The vaulting procedure depends on your company's backup strategy and policies for data protection/reliability.

VBFS *(OmniStorage specific term)*
A Very Big File System is an extension of the standard HP-UX file system on HP-UX 9.x. It is mounted to a directory the same way as any HP-UX file system. In a VBFS, only the superblock, the inode and the 'extended attribute' information remain permanently on the hard disk and are never migrated.
See also **MFS**.

verify
A function that lets you check whether the Data Protector data on a specified medium is readable. Additionally, consistency within each block can be

checked if the backup was performed with the cyclic redundancy check (CRC) option ON.

Virtual Controller Software (VCS)

(HP StorageWorks EVA specific term)
The firmware that manages all aspects of storage system operation, including communication with Command View EVA through the HSV controllers.
See also **Command View (CV) EVA**.

Virtual Device Interface *(MS SQL Server 7.0/2000 specific term)*
This is a SQL Server 7.0/2000 programming interface that allows fast backup and restore of large databases.

virtual disk *(HP StorageWorks EVA specific term)*
A unit of storage allocated from an HP StorageWorks Enterprise Virtual Array storage pool. Virtual disks are the entities that are replicated using the HP StorageWorks Enterprise Virtual Array snapshot functionality.
See also **source volume** and **target volume**.

virtual server
A virtual machine in a cluster environment defined in a domain by a network IP name and address. Its address is cached by the cluster software and mapped to the cluster node that is currently running the virtual server resources. This way all requests for a particular virtual server are cached by a specific cluster node.

volser *(ADIC and STK specific term)*
A VOLume SERial number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/GRAU and StorageTek devices.

volume group
A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.

volume mountpoint (*Windows specific term*)
An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to the target volume. Provided that the volume is mounted, users and applications can refer to the data on the mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

Volume Shadow Copy service
See **Microsoft Volume Shadow Copy service**.

VPO
See **OVO**.

VSS
See **Microsoft Volume Shadow Copy service**.

VxFS
Veritas Journal Filesystem.

VxVM (Veritas Volume Manager)
A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

Wake ONLAN
Remote power-up support for systems running in power-save mode from some other system on the same LAN.

Web reporting
The Data Protector functionality that allows you to view reports on backup status and Data Protector configuration using the Web interface.

wildcard character
A keyboard character that can be used to represent one or many characters. The asterisk (*), for example, typically represents one or more characters, and the question mark (?) typically represents a single character. Wildcard characters are often used in operating systems as a means of specifying more than one file by name.

Windows CONFIGURATION backup
Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.

Windows Registry
A centralized database used by Windows to store configuration information for the operating system and the installed applications.

WINS server A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.

writer
(MS VSS specific term)
A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

XBSA interface *(Informix specific term)*
The onbar utility and Data Protector communicate with each other through

the X/Open Backup Specification Services Programmer's Interface (XBSA).

XCopy engine *(direct backup specific term)*
A SCSI-3 copy command that allows you to copy data from a storage device having a SCSI source address to a backup device having a SCSI destination address, thus enabling direct backup. The data flows from a source device (either block or streaming, that is, disk or tape) to the destination device (either block or streaming) through XCopy. This releases the controlling server of reading the data from the storage device into memory and then writing the information to the destination device.
See also **direct backup**.

ZDB
See **zero downtime backup (ZDB)**.

ZDB database *(ZDB specific term)*
A part of the IDB, storing ZDB related information such as source volumes, replicas and security information. The ZDB database is used for ZDB, instant recovery, and split mirror restore.
See also **zero downtime backup (ZDB)**.

ZDB to disk *(ZDB specific term)*
A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process.
See also **zero downtime backup (ZDB)**, **ZDB to tape**, **ZDB to disk+tape**, **instant recovery**, and **replica set rotation**.

ZDB to disk+tape *(ZDB specific term)*
A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using the instant recovery process, the standard Data Protector restore from tape, or on split mirror arrays, split mirror restore.
See also **zero downtime backup (ZDB)**, **ZDB to disk**, **ZDB to tape**, **instant recovery**, **replica**, and **replica set rotation**.

ZDB to tape *(ZDB specific term)*
A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be

retained on the disk array after backup completion. The backed-up data can be restored using standard Data Protector restore from tape. On split mirror arrays, split mirror restore can also be used.
See also **zero downtime backup (ZDB)**, **ZDB to disk**, **instant recovery**, **ZDB to disk+tape**, and **replica**.

zero downtime backup (ZDB)
A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation.
See also **ZDB to disk**, **ZDB to tape**, **ZDB to disk+tape**, and **instant recovery**.

C

F

P

Q

R

S

T

U

V

W

X

Index