HP Data Protector

Software Version: 8.10

Disaster recovery guide

Document Release Date: November 2016

Software Release Date: November 2016

Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to: **http://h20230.www2.hp.com/selfsolve/manuals**

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to: **http://h20229.www2.hp.com/passport-registration.html**

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

Visit the HP Software Support Online web site at: **http://www.hp.com/go/hpsoftwaresupport**

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

http://h20229.www2.hp.com/passport-registration.html

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

HP Software Solutions Now accesses the HPSW Solution and Integration Portal Web site. This site enables you to explore HP Product Solutions to meet your business needs, includes a full list of Integrations between HP Products, as well as a listing of ITIL Processes. The URL for this Web site is **http://h20230.www2.hp.com/sc/solutions/index.jsp**

Contents

Chapter 1: Introduction

Data Protector disaster recovery overview

This chapter provides a general overview of the disaster recovery process, explains the basic terms used in the Disaster Recovery guide and provides an overview of disaster recovery methods.

A **computer disaster** refers to any event that renders a computer system unbootable, whether due to a human error, hardware failure, or natural disaster. In these cases, it is most likely that the boot partition or system partition of the computer is not available and the environment needs to be recovered before the normal restore operation can begin. The disaster recovery includes repartitioning and/or reformatting the boot partition and recovery of the operating system with all the configuration information that defines the environment. This step *must* be completed in order to recover other user data.

Original system refers to the system configuration backed up by Data Protector before a computer disaster hit the 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 affected and the target system is that the target system has all faulty hardware replaced.

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.

Hosting system is a working Data Protector client used for Disk Delivery Disaster Recovery with Disk Agent installed.

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.

Disaster recovery operating system (DR OS) is the 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 its own configuration data with the original configuration data.

Critical volumes are the volumes required to boot the system and Data Protector volumes. Regardless of the operating system, these include:

- Boot volume
- System volume
- the volume with Data Protector executables
- the volume where the IDB is located (for Cell Managers)

> **Note:** If the IDB is located on more than one volume then all volumes where the IDB resides are treated as critical.

Apart from the critical volumes stated above, CONFIGURATION is also a part of the critical volumes set for Windows and Linux systems. On Windows systems, services are backed up as a part of the CONFIGURATION backup.

On Windows systems, some items included in the CONFIGURATION object can be located on volumes other than system, boot, Data Protector, or IDB volumes. In this case these volumes are also a part of the critical volumes set:

- User profiles volume
- Certificate Server database volume on Windows Server systems
- Active Directory Service volume on domain controller on Windows Sever
- Quorum volume on Microsoft Cluster Server

On Linux systems, the CONFIGURATION object contains only data relevant for the automated disaster recovery methods, such as volumes, mount points, network settings, and similar.

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 the GUI, and so on).

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, and so on). Only standalone, SCSI Library, File Library and Backup to Disk (B2D) devices can be used for offline recovery. The Cell Manager can only be recovered offline.

Remote recovery is performed if all Media Agent systems 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 a severe event, however, the following factors can exacerbate the situation:

- The system has to be returned to online status as quickly and efficiently as possible.
- Disaster recovery is not a common event and administrators may not be familiar with the required

steps.

- The available personnel to perform the recovery may only have fundamental system knowledge.

Disaster recovery is not provided as an already-defined, easy-to-use solution. It is a complex process that involves extensive planning and preparation before execution. You have to thoroughly define a step-by-step process to be prepared for swift recovery from disastrous situations.

Disaster recovery process

The process of disaster recovery is split into four consecutive phases, regardless of the recovery method:

1. **Phase 0** (preparation) is the prerequisite for a successful disaster recovery. The planning and preparation must be done before a disaster occurs.

2. 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.

3. The operating system with all the configuration information that defines the environment with Data Protector (as it was) is restored in Phase 2.

4. 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.

Disaster Recovery Methods

This section provides a general overview of disaster recovery methods. For lists of disaster recovery methods that are supported on different operating systems, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.

Note: Each disaster recovery method has limitations you should consider before implementation.

"Overview of disaster recovery methods" below provides an overview of the Data Protector disaster recovery methods.

Figure 1: Overview of disaster recovery methods

Phase 0	Phase 1	Phase 2	Phase 3
Manual Disaster Recovery			

Full filesystem backup of the entire system, Internal Database backup (Cell Manager only). Update the SRD file (Windows systems only). Collect information on the original system to enable installation and configuration of the DR OS.	Install DR OS with network support. Repartition the disk and re-establish the original storage structure.	Execute the `drstart` command to automatically recover critical volumes.Additional steps are required to perform advanced recovery tasks.	Restore user and application data using the standard Data Protector restore procedure.
See "Assisted Manual Disaster Recovery (AMDR)" on page 24 or "Manual Disaster Recovery (MDR)" on page 92.			
Disk Delivery Disaster Recovery (DDDR) (UNIX systems only)			
Full filesystem backup of the entire system, internal Database backup (Cell Manager only), create the auxiliary disk.	Connect the auxiliary disk to the target system. Repartition the replacement disk and re-establish the original storage structure.	Restore the boot disk of the original system onto the replacement disk, remove the auxiliary boot disk. Restart the system. Additional steps are required to perform advanced recovery tasks.	Restore user and application data using the standard Data Protector restore procedure.
See "Disk Delivery Disaster Recovery (DDDR)" on page 101.			
Enhanced Automated Disaster Recovery (EADR)			
Full filesystem backup of the entire system, Internal Database backup (Cell Manager only). Prepare and update the SRD file. Prepare the DR OS image.	Boot the system from the disaster recovery CD, USB flash drive, or network and select the scope of recovery.	Automatic restore of critical volumes. Additional steps are required to perform advanced recovery tasks.	Restore user and application data using the standard Data Protector restore procedure.
See "Enhanced Automated Disaster Recovery (EADR)" on page 37 or "Enhanced Automated Disaster Recovery (EADR)" on page 107.			
One Button Disaster Recovery (OBDR)			

Full filesystem backup of the entire system using the OBDR wizard.Prepare and update the SRD file.	Boot the target system from the OBDR tape and select scope of recovery.	Automatic restore of critical volumes.	Restore user and application data using the standard Data Protector restore procedure.
See "One Button Disaster Recovery (OBDR)" on page 56 or "One Button Disaster Recovery (OBDR)" on page 119.			

The following has to be completed before you can proceed to the next phase:

- *Phase 0*:

  A full client backup and the IDB backup (on Cell Manager only) must be performed, and enough information must be collected by the administrator from the original system to enable installation and configuration of the DR OS. An auxiliary boot disk should be created for Disk Delivery Disaster Recovery of UNIX systems.

- *Phase 1*:

  DR OS must be installed and configured and the original storage structure must be re-established (all volumes are ready to be restored). The replacement disk for Disk Delivery Disaster Recovery on UNIX must be made bootable.

- *Phase 2*:

  Critical volumes are restored. Additional steps to perform advanced recovery tasks are required. See the section “Advanced recovery tasks”.

- *Phase 3*:

  Check if application data is restored correctly (for example, databases are consistent).

Manual disaster recovery method

This is a basic disaster recovery method that involves recovering the target system to the original system configuration.

First, you have to install and configure the DR OS. Then use Data Protector to restore data (including the operating system files) replacing the operating system files with the restored operating system files.

With manual recovery, it is important to collect the information regarding the storage structure, which is not kept in flat files (such as partition information, disk mirroring, and striping).

Disaster recovery using disk delivery

The Disk Delivery Disaster Recovery method (DDDR) is supported on UNIX clients. For details on supported operating systems, see the *HP Data Protector Product Announcements, Software Notes, and References*.

This method works without an additional client and requires a bootable auxiliary disk (which can be carried around) with a minimal operating system, networking, and a Data Protector Disk Agent installed. You need to collect enough information before the disaster to be able to correctly format and partition the disk.

This is a fast and simple method to recover clients.

> **Tip:** This method is especially 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.

See "Disk Delivery Disaster Recovery (DDDR)" on page 101.

Enhanced Automated Disaster Recovery (EADR)

Data Protector offers an enhanced disaster recovery procedure for Windows and Linux Data Protector clients and Cell Managers where user intervention is reduced to a minimum.

The EADR procedure collects all relevant environment data automatically at backup time. During a configuration backup, data required for temporary DR OS setup and configuration is packed in a single large **DR image (recovery set)** file is stored on the backup tape (and optionally on the 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, you can use the EADR wizard to restore the DR OS image from the backup medium (if it has not been saved on the Cell Manager during the full backup). You can either convert it to a **disaster recovery CD ISO image**, save it on a bootable USB drive, or create a bootable network image. You can then record the CD ISO image on a CD using any CD recording tool.

When you boot the target system from the CD, USB drive, or over the network, Data Protector automatically installs and configures the DR OS, formats and partitions the disks, and finally recovers the original system with Data Protector as it was at the time of backup.

The recovered volumes are:

- The boot volume
- The system volume
- The volume containing the Data Protector installation and configuration

Any remaining volumes can be recovered using the standard Data Protector restore procedure.

One Button Disaster Recovery (OBDR)

One Button Disaster Recovery (OBDR) is an automated Data Protector disaster recovery method for Windows and Linux Data Protector clients, where user intervention is reduced to minimum. It is based on the concept of using an OBDR device and copying an image file onto a tape. For details on supported operating systems, see the *HP Data Protector Product Announcements, Software Notes, and References*.

During OBDR backup, data required for the temporary DR OS installation and configuration is packed in a single large OBDR image file and stored on the backup tape. When a disaster occurs, the OBDR device 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 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.

The automatically-recovered volumes are:

- The boot volume
- The system volume
- The volume containing the Data Protector installation and configuration

The remaining volumes can be recovered using the standard Data Protector restore procedure.

> **Important:** You need to prepare a new OBDR boot tape locally on the client after each hardware, software or configuration change. This also applies to any network configuration changes, such as change of IP address or DNS server.
>
> HP recommends to restrict access to backup media, DR images, SRD files and disaster recovery CDs and USB drives storing DR OS data

Data Protector integrations and disaster recovery

Disaster recovery is a very complex process that involves products from several vendors. As such, successful disaster recovery depends on all the vendors involved. Use the information provided here only as a guideline.

Check the instructions of the database/application vendor on how to prepare for disaster recovery.

This is a general procedure on how to recover an application:

1. Perform Disaster Recovery.
2. Install, configure, and initialize the database/application so that data on Data Protector media can be loaded back to the system. Consult database/application vendor documentation for a detailed procedure and steps needed to prepare the database.
3. Ensure that the database/application server has the required Data Protector client software installed and is configured for the database/application. Follow the procedures in the appropriate *HP Data Protector Integration Guide*.
4. Start the restore. When the restore is complete, follow the instructions of the database/application vendor for any additional steps required to bring the database back online.

Chapter 2: How to Prepare for Disaster Recovery

Carefully follow the instructions below to prepare for disaster recovery and ensure a fast and efficient restore. The preparation procedure does not depend on the disaster recovery method, and includes developing a detailed disaster recovery plan, performing consistent and relevant backups, and updating the SRD file on Windows.

This chapter contains the general preparation procedure for disaster recovery for all disaster recovery methods. Additional preparation is required for each particular disaster recovery method. For additional preparation steps, see the corresponding topics.

Remember that preparing the Cell Manager for disaster recovery is critical and requires more attention.

> **Important:** Prepare for disaster recovery before a disaster occurs.

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 steps:

   - Make a list of the most important systems that should be recovered first. Critical systems are systems required for a network to function properly (DNS servers, domain controllers, gateways, and so on), Cell Managers, and Media Agent clients. They should be recovered prior to all other systems.

   - Select the disaster recovery methods that are appropriate for your systems. Based on these methods, consider which preparation steps are required for each system.

   - Determine a method to obtain the required information at recovery time, such as the media that stores the IDB, the location of the updated SRD file, and the location and labels of the Cell Manager backup media. Define the location of software libraries to enable the performance of new installations.

   - 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**

   Perform the following preparation steps before running the backup to guarantee environmental consistency during the backup:

   All systems:

- Perform regular and consistent backups.

- You need to understand volume groups and partition concepts. On UNIX systems, you should know where the information about the storage environment structure resides.

UNIX systems:

- Create pre-exec scripts, which collect the storage structure, and perform other client-specific preparations.

- Create tools, such as the auxiliary disk with the minimum operating system, network resources, and the Data ProtectorDisk Agent installed.

Windows systems*:*

- Ensure that you have a valid CONFIGURATION backup at your disposal.

- Update the SRD file and store it in a safe place. You should restrict access to SRD files due to security considerations.

3. **Perform recovery procedures**

   Follow the procedures and checklists you have tested to recover the affected system.

Caution: Do not change the default `Inet` listen port on systems that are prepared for disaster recovery. In the opposite case, if such systems are struck by a disaster, the disaster recovery process may fail.

Consistent and relevant backups

In case of a disaster, the target system should be returned to the original system configuration. Additionally, the system is expected to operate and function as it did just before the last valid backup was performed.

Note: On UNIX systems, some daemons or processes are active as soon as the system finishes booting, for various reasons (the run-level 2). Such processes may even read 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 systems, 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. Either the login account must be changed or the relevant service must be stopped.

Depending on what is active on the system when the backup runs, the data consistency of an application can be violated, causing re-start and execution issues after the recovery.

Creating a consistent and relevant backup

- Ideally, you would perform a backup with the relevant partition(s) set offline, which is often 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.
- Ensure minimal system activity. For example, only the core operating system, basic networking, and backup should be active. None of the low-level application services should be running. This can be achieved using an appropriate pre-exec script.

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 Server). See the topics pertaining to preparation for particular disaster recovery methods.

Encrypted backups

If your backups are encrypted, you must ensure that the encryption keys are safely stored and available when you start a disaster recovery. Without the access to the appropriate encryption key, the disaster recovery procedure aborts. Different disaster recovery methods have additional requirements.

The encryption keys are stored centralized on the Cell Manager; thus the disaster recovery client must be connected to the Cell Manager to get the encryption key. For details on encryption concepts, see the HP Data Protector Help index: "encryption".

Two disaster recovery scenarios are possible:

- Recovery of a client where you can establish a connection to the Cell Manager. No additional encryption related preparations are needed for such a scenario, as Data Protector automatically obtains the encryption keys.
- Disaster recovery of a Cell Manager or standalone client recovery, where you cannot establish a connection the Cell Manager.

  You must provide the encryption keys on removable media (for example a diskette) when prompted.

  The keys are not part of the disaster recovery OS image and are exported to the key file (`DR-`*`ClientName`*`-keys.csv`). You must manually store the keys to a separate removable media, such as a diskette or USB flash drive. Ensure that you have always an appropriate copy of the keys for each backup that is prepared for disaster recovery. If the encryption key is not available, disaster recovery is not possible.

Updating and Editing the System Recovery Data

The **System Recovery Data (SRD)** is a text file in the Unicode (UTF-16) format that contains information required to configure the target system. The SRD file is generated when a CONFIGURATION backup is performed on a Windows client and then stored on the Cell Manager into the directory:

Windows systems: *`Data_Protector_program_data`*`\Config\Server\DR\SRD`

UNIX systems: `/etc/opt/omni/server/dr/srd`.

> **Important:** When IDB is not available, information about objects and media is stored only in the SRD file.

The SRD filename on the Cell Manager is identical to the hostname of the computer on which it was generated (for example, `computer.company.com`).

After the CONFIGURATION backup, the SRD file 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 or Linux client. The name of the updated SRD file is `recovery.srd`.

There are three different methods possible for updating the SRD file:

- Update SRD File wizard (from Windows systems only)
- `omnisrdupdate` command as a standalone utility
- `omnisrdupdate` command as a backup session post-exec script

> **Important:** When you update the SRD file for Cell Manager, specify an IDB backup session which is newer than the filesystem backup session so that you can browse the file system backup sessions and data after a recovery.

For a detailed procedure on how to update the SRD file, see "Updating the SRD File (Windows Clients)" on page 29.

Chapter 3: Disaster recovery on Windows systems

Assisted Manual Disaster Recovery (AMDR)

While being recovered, Windows requires an installation of disaster recovery operating system (DR OS). The procedure of recovering the original operating system is automated by the `omnidr` command.

Windows systems provide additional possibilities for recovering a system before deciding for 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 the problems.

Overview

The general procedure for Assisted Manual Disaster Recovery of a Windows system is:

1. **Phase 0**

   a. Perform a full filesystem backup of the entire system including its CONFIGURATION object (client backup). If you are preparing for disaster recovery of a Cell Manager, also perform an Internal Database backup afterwards as soon as possible.

   b. Update the SRD file. Collect information on the original system to enable installation and configuration of the DR OS.

2. **Phase 1**

   a. Replace the faulty hardware.

   b. Reinstall the operating system (create and format the necessary volumes).

   c. Reinstall service packs.

   d. Manually re-partition the disk and re-establish the storage structure with original drive letter assignments.

   Tip: You can combine Phase 1 of Manual Disaster Recovery with automated deployment tools.

3. **Phase 2**

   a. Execute the Data Protector drstart command that will install the DR OS and start the restore of critical volumes.

   b. The system must be restarted after the drstart command finishes.

   c. Additional steps are required if you are recovering a Cell Manager or performing advanced recovery tasks. For more information, see “Advanced recovery tasks” (page 56).

4. **Phase 3**

   a. Use the Data Protector standard restore procedure to restore user and application data.

Requirements

- The partitions must be the same size or bigger than the partitions on the failed disk. This way the information stored on the crashed disk can be restored to the new one. Also, the type of the filesystem (FAT, NTFS) and compression attributes of the new volumes must match.
- The hardware configuration of the target system must be the same as of the original system. This includes SCSI BIOS settings (sector remapping).
- All hardware has to be the same.

Limitations

- The Internet Information Server 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.
- Using resumed object backups for recovery is not supported since the consistency of such backups cannot be guaranteed.

Steps

1. "Preparation for Assisted Manual Disaster Recovery (Windows Systems)" below.
2. "Installing and Configuring a Windows System Manually" on page 33.
3. "Restoring System Data Manually (Windows Systems)" on page 34.
4. "Restoring Vendor-Specific Partitions (Windows systems)" on page 35.
5. Restore user data.

Preparation for Assisted Manual Disaster Recovery (Windows Systems)

To prepare for a successful disaster recovery, follow the instructions related to the general preparation procedure for disaster recovery before performing the steps listed in this topic. Prepare in advance in order to perform a disaster recovery fast and efficiently. You should pay special attention to disaster recovery preparation for the Cell Manager.

Important: Prepare for disaster recovery before a disaster occurs.

General preparations

Before completing the steps listed in this section, see also "Planning" on page 20 for the general preparation procedure for all disaster recovery methods. 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. You can also use the Windows diskettes if you do not have a bootable CD-ROM drive.

2. Ensure that you have the drivers for the system that you want to recover. You may need to install some drivers, such as HBA and SCSI drivers, during Windows Setup.

3. To recover the affected system, you need the following information about the system before the disaster:

   - If DHCP was not used before the disaster, the TCP/IP properties (IP address, default gateway, subnet mask and DNS order (IPv4), subnet prefix length, preferred and alternate DNS server (IPv6))

   - Client properties (hostname, domain)

4. Ensure that the following is true:

   - You should have a valid full client backup image (including valid CONFIGURATION backup data). See the HP Data Protector Help index: “backup, Windows specific” and “backup, configuration”.

   - You should have an SRD file, updated with information about the objects from the backup sessions that you plan to use for the recovery.

   - For recovery of the Cell Manager, you should have a valid Internal Database backup image, created after the client backup image. For more information on how to configure and perform an IDB backup, see the HP Data Protector Help index: “IDB, configuration”.

   - In case of Microsoft Cluster Server, consistent backup also 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, also Cell Manager virtual server and IDB.

     For details, see "About Disaster Recovery of a Microsoft Cluster Server" on page 71.

   - The disk with the boot partition requires free disk space for the Data Protector disaster recovery installation (15 MB) and an DR OS installation. Additionally, you need as much free disk space as required for the restore of the original system.

5. Copy the drsetup images (“drsetup diskettes”) onto a USB flash drive or floppy disks. The number of diskettes depends on the platform and the version of the Windows operating system. The images are located in:

   - 32-bit Windows systems:

     Windows Vista and later releases: `*Data_Protector_program_data*\Depot\DRSetupX86`

     Windows XP, Windows Server 2003: `*Data_Protector_home*\Depot\DRSetupX86`

     Data Protector installation medium: `\i386\tools\DRSetupX86`

   - 64-bit Windows systems on the AMD64/Intel EM64T platform:

     Windows Vista and later releases: `*Data_Protector_program_data*\Depot\DRSetupX64`

     Windows XP, Windows Server 2003: `*Data_Protector_home*\Depot\DRSetupX64`

     Data Protector installation medium: `\i386\tools\DRSetupX64`

   - 64-bit Windows systems on the Itanium platform:

     Windows Vista and later releases: `*Data_Protector_program_data*\Depot\DRSetupIA64`

     Windows XP, Windows Server 2003: `*Data_Protector_home*\Depot\DRSetupIA64`

     Data Protector installation medium: `\i386\tools\DRSetupIA64`

   In case of a disaster, save the updated SRD file of the affected system to the first floppy disk (disk 1). Only one set of floppy disks is required per site for all Windows systems, but you must always copy the updated SRD file of the affected client onto 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 as they existed before the disaster, record the following information for each partition (it will be needed during the recovery process):

   - partition length and order

   - drive letter assigned to the partition

   - partition filesystem type

   This information is stored in the SRD file. The -type option in the diskinfo section of the SRD file shows the volume filesystem type for a particular volume:

   Table 1: How to determine the filesystem type from the SRD File

Type number	Filesystem
1	Fat12

Type number	Filesystem
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. For an empty template which can be used when preparing for the Assisted Manual Disaster Recovery, see "Example of the Disaster Recovery Preparation Table for Windows" on page 144.

Updating the recovery diskettes using the CLI

Data Protector does not offer a command to automatically create recovery images (diskettes). However, you can manually update the contents of the first diskette in the recovery set by executing the `omnisrdupdate` command. Insert the first diskette from the recovery set in to the floppy disk drive and specify `a:\` as the location, for example:

Data Protector client system:

```
omnisrdupdate -session 10/04/2011-1 -host clientsys.company.com -location a:\ -asr
```

Data Protector Cell Manager:

```
omnisrdupdate -session 10/04/2011-1 10/04/2011-2 -host cmsys.company.com -location
a:\ -asr
```

To manually create a recovery diskette, you also need to copy the `DRDiskNumber.cab` files from `Data_Protector_program_data\Depot\DRSetup\DiskDiskNumber` folders to the appropriate recovery diskette.

Example of the Disaster Recovery Preparation Table for Windows

Client properties	Computer name	ANAPURNA
	Hostname	anapurna.company.com
Drivers		tatpi.sys, aic78xx.sys
Windows Service Pack		Windows Vista

TCP/IP properties for IPv4	IP address	10.17.2.61
	Default gateway	10.17.250.250
	Subnet mask	255.255.0.0
	DNS order	10.17.3.108, 10.17.100.100
TCP/IP properties for IPv6	IP address	td10:1234:5678:abba::6:1600
	Subnet prefix length	64
	Default gateway	td10:1234:5678:abba::6:1603
	Preferred DNS server	td10:1234:5678:abba::6:1603
	Alternate DNS server	td10:1234:5678:abba::6:1604
Medium label/barcode number		"anapurna - 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	

Updating the SRD File (Windows Clients)

After the CONFIGURATION backup the SRD file contains only system information required for installation of the DR OS. It is located on the Cell Manager:

Windows systems: `Data_Protector_program_data\Config\Server\DR\SRD`

UNIX systems: `/etc/opt/omni/server/dr/srd`

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 SRD filename on the Cell Manager is identical to the hostname of the computer where it was generated - for example `computer.company.com`. The name of the updated SRD file is `recovery.srd`.

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.

Important: Store the Cell Manager SRD file in a safe place (not on the Cell Manager). It is recommended to restrict access to the SRD files.

Updating the SRD file using the Data Protector Disaster Recovery Wizard on Windows systems

Steps

1. In the Data Protector Context List, click **Restore**.
2. In the Scoping Pane, click **Task**, and then click **Disaster Recovery** to open the Disaster Recovery Wizard.
3. In the Host drop-down list, select the system for which you want to update the SRD file.
4. In the Disaster Recovery Method list, select **SRD file update**. Click **Next**.

   Data Protector first searches the Cell Manager for the SRD file. If it is not found, Data Protector restores it from the last backup.
5. Select the objects and versions needed to restore the logical volumes and the system configuration. Click **Next** for each object.
6. Specify the destination for the SRD file. Click **Finish**.

Updating the SRD file using the omnisrdupdate command

You can use `omnisrdupdate` as a standalone command.

To update the SRD file, either modify an existing backup specification or create a new one with a specified post-exec script.

Steps

1. In the Data Protector Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Filesystem**. All saved backup specifications are displayed.
3. Click the backup specification that you want to modify.

4. In the Options property page, under the Backup Specification Options, click **Advanced**.

5. In the Backup Options window, type `omnisrdupdate` in the Post-exec text box.

6. In the On client drop-down list, select the client on which this post-exec script will be executed, then click **OK**.

7. Click **Apply** to save the change and exit the wizard.

Updating the SRD file 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` 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 the omnisrdupdate command is executed as a post-exec utility, the session IDs are obtained automatically, without needing to be specified.

All other options can be specified the same way as with the standalone utility (`-location` *`Path`*`, -host` *`ClientName`*).

Important: You cannot use omnisrdupdate in a post-exec script to update the SRD for a Cell Manager because the IDB is backed up in a separate session.

Example of Editing the SRD File

If the information in the SRD file is not up to date anymore (for example, you changed a backup device), modify the updated SRD file (`recovery.srd`) before performing Phase 2 of disaster recovery to update the wrong information and thus enable a successful recovery.

You can display some of the device configuration information using the `devbra -dev` command.

Changing the MA client

You performed a backup for disaster recovery purposes 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 also the value of the -devaddr option in the updated SRD file accordingly.

After you have edited the file, save it in Unicode (UTF-16) format to the original location.

Changing the backup device

To perform disaster recovery using another device than the one which was used for the backup, modify the following option values in the updated SRD file:

-dev, -devaddr, -devtype, -devpolicy, -devioctl, and -physloc

Where:

-dev	specifies the logical name of the backup device or drive (library) to be used for the backup,
-devaddr	specifies its SCSI address,
-devtype	specifies the Data Protector device type,
-devpolicy	specifies the device policy, which 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),
-devioctl	specifies the robotics SCSI address.
-physloc	specifies the library slot
-storname	specifies the logical library name

For example, you performed a backup for disaster recovery purposes using an HP Ultrium standalone device with the device name Ultrium_dagnja, connected to the MA host dagnja (Windows systems). However, for the disaster recovery you would like to use an HP Ultrium robotics library with the logical library name Autoldr_kerala with drive Ultrium_kerala connected to the MA client kerala (Linux systems).

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" -storname "AutoLdr_kerala" -mahost kerala.company.com.
```

After you have edited the file, save it in Unicode (UTF-16) format to the original location.

Installing and Configuring a Windows System Manually

After a disaster happens, you should first install and configure the operating system (Phase 1). After installing the operating system, you can proceed with the system data recovery.

Steps

Phase 1

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 that it was installed on before the disaster. Do not install the Internet Information Server (IIS) during the installation of the system.

   > **Important:** If the Windows operating system has been installed using the unattended setup, use the same script now to re-install Windows to make sure that both the *%SystemRoot%* and *%SystemDrive%*`\Documents and Settings` folders are installed to the same location.

2. When the Windows Partition Setup screen appears, proceed as follows:

   - If an EISA Utility Partition (EUP) existed on the system before the disaster, create (if it does not exist due to the disaster) and format a “dummy” FAT partition using the EUP information stored in the SRD file. The EUP will later on be recovered to the space occupied by the “dummy” partition. Create and format a temporary boot partition immediately after the “dummy” partition.

   - If an EUP did not exist on the system before the disaster, create (if the boot partition does not exist due to the disaster) and format the boot partition as it existed on the disk before the disaster.

     When the Windows Setup prompts you for the Windows installation directory, specify a new directory on the boot partition equal to the directory where the original Windows installation resided.

     > **Note:** During the installation, do not add the system to a Windows domain where it had previously resided, but to a workgroup. If you are restoring a Primary Domain Controller (PDC), ensure that the target restore system is not located in the domain that had been controlled by the affected PDC.

3. Install the TCP/IP protocol. If DHCP was not used before the disaster, configure the TCP/IP protocol as it was before the disaster by providing the following information: the hostname of the affected client, its IP address, default gateway, subnet mask and DNS server. This information can be obtained from the SRD file. 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 (for example, `DRAdmin`) and add it to the Data Protector Admin group on the Cell Manager. See the *HP Data Protector Help* index: “adding Data Protector users”.

   The user account must not have existed on the system before the disaster. The temporary Windows user account will be removed later on during this procedure.

5. Log off and log on to the system using the newly created account.

6. Create and format all unformatted partitions (including the “dummy” EISA Utility Partition, if used), as they existed on the disk before the disaster. Use the vendor specific procedure for creating utility partitions. The “dummy” EISA Utility Partition has to be formatted as a FAT filesystem. Assign drive letters as they were assigned before the disaster.

Restoring System Data Manually (Windows Systems)

After you have installed and configured the operating system (Phase 1), you can use Data Protector to recover the Data Protector client or Cell Manager. Disaster recovery of the Cell Manager and Internet Information Server (IIS) requires additional steps.

Restoring the Windows system

Steps

Phase 2

1. 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.

2. Run `drstart` from the *`Data_Protector_home`*`\Depot\drsetup\disk1` (Cell Manager) or `\i386\tools\drsetup\disk1` (Data Protector installation medium) directory.

   If you have prepared the `drsetup` diskettes, you can also run `drstart` from the first diskette.

3. `drstart` first scans the current working directory, floppy disk drive, and CD-ROM drive 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. If these files are not found, browse for them or enter their path in the `DR Installation Source` text box.

4. If the SRD file (`recovery.srd`) is found in the same directory as `dr1.cab` and `omnicab.ini`,

drstart copies recovery.srd to the %SystemRoot%\system32\OB2DR\bin directory and the omnidr utility is started automatically. Otherwise, you can enter the location of the SRD file (recovery.srd) in the SRD Path text box 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.

5. Remove the temporary Data Protector user account (added during the Phase 1) from the Data Protector Admin group on the Cell Manager, unless it existed on the Cell Manager before the disaster recovery.
6. Restart the system, log on and verify that the restored applications are running.

Phase 3

6. Additional steps are required if you are recovering a Cell Manager or performing advanced recovery tasks (such as restoring MSCS or IIS, editing the kb.cfg and SRD files). For more information, see "Restoring the Data Protector Cell Manager specifics" on page 77 and the “Advanced recovery tasks” section.
7. Restore user and application data using the standard Data Protector restore procedure.

The temporary DR OS will be deleted after the first login, except in the following cases:

- You interrupt the Disaster Recovery Wizard during the 10 second pause (after it has found the DR installation and the SRD file on the backup medium) and select the **Debugs** option.
- You manually execute the omnidr command with the -no_reset or -debug option.
- Disaster recovery fails.

Restoring the Data Protector Cell Manager specifics

After you have performed the general Manual Disaster Recovery procedure of a Windows system, perform additional steps to restore the Cell Manager using Data Protector.

To make the IDB recovery consistent, you restore the information about the backed up objects that was not restored during the disaster recovery. To do this, update the IDB by importing the media with the Cell Manager full client backup used for the disaster recovery.

Restoring Vendor-Specific Partitions (Windows systems)

If needed, conclude the general Manual Disaster Recovery procedure with the recovery of vendor-specific partitions (VSP).

Disclaimer

Recovering a VSP may be a complex procedure, requiring advanced skills and knowledge of the Windows operating system. The information provided here is meant for your convenience only. You should use this information at your *own risk*. If the partition order has changed after the restore of a VSP, the `boot.ini` file will need to be modified. A wrong `boot.ini` file results in the system not being bootable.

Preparation for the Disaster Recovery

This information applies only for Assisted Manual Disaster Recovery (AMDR), because Enhanced Automated Disaster Recovery (EADR) and One Button Disaster Recovery (OBDR) will automatically recover the VSP and are therefore the recommended methods to be used to recover VSPs.

When performing AMDR, you will have to manually recreate the previous storage structure (including the VSP).

ASR will recreate the previous storage structure automatically and preserve unallocated space on disk for VSP. Then you will have to recreate the VSP on the unallocated disk space using vendor-specific tools and procedures.

In order to enable Data Protector access to the VSP, you will have to map it in Windows using the Data Protector `omnipm` utility.

Steps

1. Run `*Data_Protector_home*\bin\utilns\omnipm`to start the Data Protector Partition Mapper.

2. In the Partition Mapper window, select the partition identified with vendor-specific ID under the Type column.

3. Click **Map** to assign a drive letter to the selected partition. In the dialog window, specify a drive letter and click **OK**.

4. Use the standard Data Protector restore procedure to restore the backed up data on the mapped EISA utility partition.

5. Unmap the partition you have mapped in step 3.

Caution: Do not overwrite the operating system files (usually `*.sys` files) in the root of the VSP during the recovery, because this could cause the system to be unbootable. It is therefore recommended to add these files to the exclude list.

Restoring an Eisa Utility Partition

Steps

1. If you do not maintain the Eisa utility partition (EUP), you must manually create it. Note that the EUP should reside on the first disk as seen by the system BIOS. Since Disk Manager cannot create an EUP, create a normal FAT16 partition and assign it a drive letter.

2. Restore its contents using Data Protector. Select the **Restore As** option for the Eisa Utility Partition Configuration object. The drive letter assigned must be the one that was assigned during the EUP creation and the directory to restore into must be the root directory (\).

3. Rearrange the root directory entries if necessary.

   a. Run `omnipm`, select the EUP and click **Root...**. The root directory of the EUP is displayed.

   b. Reorder the root directory entries to their original positions. Use either drag and drop or right click an entry to display the options menu.

4. Change the FAT16 partition into a real EUP.

   a. Select the EUP and click **Unmap**. The drive letter is removed.

   b. Click **Type**. A dialog windows displays. Select **Eisa Utility Partition**.

Enhanced Automated Disaster Recovery (EADR)

Enhanced Automated Disaster Recovery is used to recover ordinary Data Protector Cell Managers and clients as well as Data Protector Cell Managers and clients that are part of the Microsoft Cluster Server (MSCS).

Overview

The general steps using the Enhanced Automated Disaster Recovery method for a Windows client are:

1. **Phase 0**

   a. Perform a full backup of the system (client backup) that includes at least all critical volumes. If you are preparing for disaster recovery of a Cell Manager, also perform an Internal Database backup afterwards as soon as possible.

   b. Use the Enhanced Automated Disaster Recovery Wizard to prepare a DR OS image from the recovery set file of the affected system and record it on a CD. On Windows Vista and later releases, you can create a bootable network image or a bootable USB drive with the DR OS image instead of a disaster recovery CD. If the recovery set has not been saved on the Cell Manager during the full backup, the Disaster Recovery Wizard will restore it from the backup medium.

> **Important:** You need to perform a new backup and prepare a new DR OS image after each hardware, software, or configuration change. This also applies to any network changes, such as a change of IP address or DNS server.

c. If the full client backup was encrypted, store the encryption key on a removable medium so that it is available for disaster recovery. You will need the key for a Cell Manager recovery or if the connection to the Cell Manager cannot be established.

2. **Phase 1**

a. Replace the faulty hardware.

b. Start the target system from the disaster recovery CD, USB drive, or through the network and select the scope of recovery. This is a completely unattended recovery.

> **Important:** Windows Server 2003: If you are recovering a domain controller, before the Disaster Recovery Wizard is launched, a standard Windows logon dialog box prompts you to enter the username (Administrator) and password of the Directory Services Restore Mode administrator account.

3. **Phase 2**

a. Depending on the recovery scope you select, the selected volumes are automatically restored. Critical volumes (the boot partition and the operating system) are always restored.

4. **Phase 3**

a. Use the standard Data Protector restore procedure to restore user and application data.

> **Important:** Prepare a disaster recovery CD, a bootable USB drive, or a network bootable image with the recovery set in advance for any critical systems that must be restored first (especially DNS servers, Cell Managers, Media Agent clients, file servers, and so on.).
>
> Prepare removable media containing encryption keys in advance for Cell Manager recovery.

The following sections explain the limitations, preparation, and recovery that pertains to EADR of the Windows clients. See also the “Advanced recovery tasks” section for details.

Prerequisites

Before selecting this method of disaster recovery, consider the following requirements and limitations:

- The Data Protector Automatic Disaster Recovery component must be installed on systems for which you want to enable recovery using this method and on systems where the DR OS image will be prepared. For details, see the *HP Data Protector Installation and Licensing Guide*.

- The hardware configuration of the target system must be the same as that of the original system. This includes the SCSI BIOS settings (sector remapping).

- The new disk have to be the same size or bigger than the affected disk. If it is larger than the original disk, the difference will remain unallocated.
- The replacement disks have to be attached to the same host bus adapter on the same bus.
- On Windows XP and Windows Server 2003 systems, the boot partition (on which the DR OS is installed) must be larger than 200 MB or disaster recovery will fail. 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.
- On Windows Vista and later releases, at least one volume must be an NTFS volume.
- A backup of all necessary data for disaster recovery may require a significant amount of free space. While normally 500 MB is enough, up to 1 GB may be required depending on the operating system.
- During the DR OS image creation, the partition on which Data Protector is installed should have at least 500 MB of temporary free space. This space is required to create a temporary image.
- On Windows Server 2003 systems, all drivers required for booting must be installed under the *%SystemRoot%*folder.
- For a remote restore, the network must be available when you boot DR OS image.
- In a cluster environment, a cluster node can be successfully backed up if the bus address enumeration on each cluster node is the same. This means that you need:
  - Equal cluster node motherboard hardware
  - The same OS version on both nodes (service packs and updates)
  - The same number and type of bus controllers
  - Bus controllers must be inserted in the same PCI mother board slots.
- The operating system should be activated at the time of the backup. Otherwise, when the activation period expires, disaster recovery fails.
- To create a DR OS image for Windows Vista and later releases, you must install the appropriate version of Windows Automated Installation Kit (WAIK) or Assessment and Deployment Kit (ADK) on the system on which you will create the image:

  Windows Vista and Windows Server 2008:

  Automated Installation Kit (AIK) for Windows Vista SP1 and Windows Server 2008

  Windows 7 and Windows Server 2008 R2:

  - Windows Automated Installation Kit (AIK) for Windows 7
  - Windows Automated Installation Kit (AIK) Supplement for Windows 7 SP1 (optional, for Microsoft

Windows 7 SP1 and Windows Server 2008 R2 SP1)

Windows 8 and Windows Server 2012:

- Assessment and Deployment Kit (ADK 1.0) for Windows 8 and Windows Server 2012

Data Protector checks the WAIK/ADK version and aborts the image creation if no appropriate version is available.

Windows 8.1 and Windows Server 2012 R2:

- Assessment and Deployment Kit (ADK 1.1) for Windows 8.1 and Windows Server 2012 R2

- For a disaster recovery from a bootable USB device, make sure that:
  - the size of the USB storage device is at least 1 GB
  - the target system supports booting from the USB device. Older systems may require a BIOS update or might not be able to boot from an USB storage device at all.
- To create a bootable network image for Windows Vista and later Windows systems versions, the following requirements must be met:
  - On the target system, the network adapter is enabled to communicate through the PXE protocol. The BIOS of this system should be compliant with the PXE protocol.
  - Windows Deployment Services (WDS) server is installed and configured on the Windows Server 2008 and later Windows systems. WDS server must be either a member of an Active Directory domain or a domain controller for an Active Directory domain.
  - A DNS server and a DHCP server with an active scope are running in the network.
- To back up the IIS configuration object residing on a Windows Vista and later releases, install the IIS 6 Metabase Compatibility package.

Limitations

- Multiboot systems that do not use Microsoft's boot loader are not supported.
- The Internet Information Server database, Terminal Services database and Certificate Server database are not restored automatically during Phase 2. They can be restored to the target system using the standard Data Protector restore procedure.
- You can create a bootable USB drive on Windows 7, Windows 8, Windows Server 2008 R2 systems (on all supported platforms), Windows Server 2008 systems (on the Itanium platform), Windows Server 2012, and later releases.
- On Windows XP and Windows Server 2003, recovery of a SAN boot configuration is not supported.
- VSS disk image backup of the logical volumes can be used for disaster recovery only on Windows

Vista and later releases.

- On Windows XP and Windows Server 2003, you cannot boot the target system over the network.
- On Windows XP and Windows Server 2003, a console interface is available instead of the HP Data Protector Disaster Recovery GUI.
- On Windows Vista and later releases, originally encrypted folders can only be restored as unencrypted.
- Do not select backup object versions which belong to a checkpoint restart backup session.
- When selecting an object copy as the source for the recovery, the following applies:
  - Only copies of full backup objects can be selected for recovery.
  - Object copies can be selected only if you create a volume recovery set from a list of volumes. Sessions are not supported.
  - Media copies are not supported.
- Using resumed object backups for recovery is not supported since the consistency of such backups cannot be guaranteed.
- The DRM restore monitor monitors the overall bytes written to a disk by the VRDA process. The overall bytes written to a disk do not always match what is displayed in the Data Protector session manager.

  > **Note:** The new Recovery Session monitor is implemented only on Windows Vista and later releases.

- Sparse files are restored to their full size during offline restore. This may result in the target volume running out of space.

Disk and partition configuration

- Dynamic disks are not supported (including mirror sets upgraded from Windows NT).
- A 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 EADR.
- On Windows XP and Windows Server 2003 systems, disaster recovery ISO images cannot be created on systems where Data Protector is installed on FAT/FAT32 partitions. You need at least one client in the cell where Data Protector is installed on an NTFS volume to be able to create disaster recovery images
- Recovery of operating systems which have been deployed using the HP Intelligent Provisioning tool

(v.1.4 and v.1.5) may fail because of incorrect MBR partition information.

- Storage Spaces configurations where physical disks do not entirely belong to a storage pool are not supported.

Preparation for Enhanced Automated Disaster Recovery

To prepare for a successful disaster recovery, follow the instructions related to the general preparation procedure for all disaster recovery methods before completing the steps listed in this topic. You have to prepare in advance in order to perform a disaster recovery fast and efficiently. You should pay special attention to disaster recovery preparation for the Cell Manager.

Important: Prepare for disaster recovery before a disaster occurs.

General preparations

1. Perform a full backup of the client system. It is recommended that you back up the whole client, however, you need to select at least the following critical volumes and objects:

   - the boot and system volumes
   - the Data Protector installation volume
   - the volume where the CONFIGURATION object is located
   - the Active Directory database volume (in case of an Active Directory controller)
   - the quorum volume (in case of a Microsoft Cluster Server)

   For a *Data Protector Cell Manager system*, see "Additional preparations for the Cell Manager" on page 44..

   See the *HP Data Protector Help* index: “backup, Windows specific” and “backup, configuration”

   During a full client backup, the recovery set and P1S file are stored on the backup medium and (recovery set optionally) on the Cell Manager.

   Considerations:

   Windows Vista and later releases:

   - Make sure that you back up also the system volume if present.
   - You can back up logical volumes using disk image backup that uses VSS writers. VSS disk image backup ensures that the volume remains unlocked during the backup and can be accessed by other applications. The IDB and CONFIGURATION objects, as well as volumes that are not mounted or are mounted as NTFS folders, must be backed up using regular filesystem backup.

Windows Server 2012 (R2):

- Use disk image backup to back up volumes in the following cases:
  - Deduplicated volumes

    During a filesystem restore, the volume is rehydrated and you might run of space on the destination volume during recovery. A disk image restore keeps the size of the volume.

  - Volumes with Resilient File System (ReFS)

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.

  The above items should be included in the same backup session.

  For details, see "About Disaster Recovery of a Microsoft Cluster Server" on page 71.

- *Cluster Shared Volumes:* Before performing a full backup of the client system, back up the Virtual Hard Drive (VHD) files and CSV configuration data using the Data Protector Virtual Environment first. See the *HP Data Protector Integration Guide for Virtualization Environments*.

  Virtual Hard Drives (VHD) must be dismounted to ensure consistency.

- 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.

  If the full client backup was encrypted, store the encryption key on a removable medium so that it is available for disaster recovery. You will need the key if you are recovering a Cell Manager or if the connection to the Cell Manager cannot be established.

Active Directory on Windows Server 2008 and later Windows Server versions:

- If your Windows Server is a domain controller whose Active Directory size exceeds 512 MB, the backup specification for the client backup needs to be modified: in the source page, expand the `CONFIGURATION` object, and clear the checkboxes for the `ActiveDirectoryService` and `SYSVOL` items.

  > **Note:** The Active Directory and SYSVOL will still be backed up as part of the system volume (`C:/`) backup. By default, they are located in `C:/Windows/NTDS` and `C:/Windows/SYSVOL` respectively.

2. After a disaster occurs, use the EADR Wizard to convert the DR image into a disaster recovery CD ISO image.

   Windows Vista and later releases: Alternatively, create a bootable network image or a bootable USB drive with the DR OS image instead of a disaster recovery CD.

3. Record the disaster recovery CD ISO image on a CD using any CD recording tool that supports the ISO9660 format. This disaster recovery CD can then be used to boot the target system and automatically restore critical volumes.

4. Execute a disaster recovery test plan.

5. On Windows systems, if some service or driver is not operational after the boot, you may have to manually edit the `kb.cfg` file.

Additional preparations for the Cell Manager

Successful disaster recovery of the Cell Manager requires additional preparation.

- Regularly back up the IDB. The IDB session should not be older than the file system session.
- Store the Cell Manager's SRD file at a safe location (not on the Cell Manager).
- Prepare a disaster recovery OS image for the Cell Manager in advance.

Saving a Recovery Set to the Cell Manager

A recovery set is packed in a single large file and stored on the backup medium and optionally on the Cell Manager during a full client backup. Saving the recovery set file to the Cell Manager is useful if you plan to record the disaster recovery CD on the Cell Manager, because it is much faster to obtain the recovery set from the hard disk than to restore it from a backup medium.

If the recovery set is saved to the Cell Manager during backup, it is saved to the default Data Protector `P1S` files location.

To change the default location, specify a new global option `EADRImagePath = `*`valid_path`* (for example, `EADRImagePath = /home/images` or `EADRImagePath = C:\temp`).

See the HP Data Protector Help index: “Global Options, modifying”.

Tip: If you do not have enough free disk space in the destination directory, you can create a mount point (Windows systems) or a link to another volume (UNIX systems).

Saving the recovery set file to the Cell Manager for all clients in the backup specification

Steps

1. In the Context List, click **Backup**.

2. In the Scoping Pane, expand **Backup Specifications** and then **Filesystem**.

3. Select the backup specification you will use for a full client backup (create it if you have not done so already). For details, see the HP Data Protector Help index: “creating, backup specifications”.

4. In the Results Area, click **Options**.

5. Under **Filesystem Options** click **Advanced**.

6. In the **Other** page, select **Copy Recovery Set to disk**.

7. **Windows Vista and later releases:** In the **WinFS Options** page, select the **Detect NTFS hardlinks** and leave the **Use Shadow Copy** option selected, and leave **Allow Fallback** cleared. Note that the **Detect NTFS hardlinks** option is not automatically selected if you manually add objects or update existing backup specifications.

Figure 2: WinFS options tab

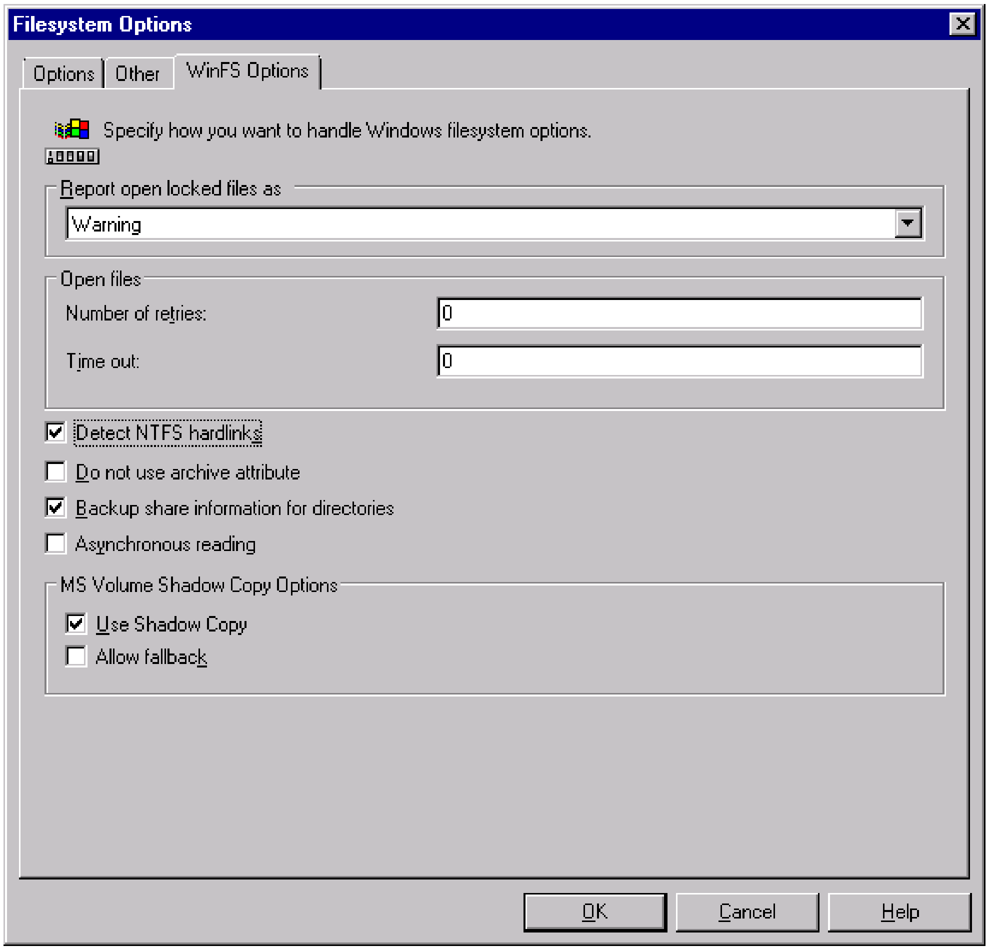

Specify how you want to handle Windows filesystem options.

Report open locked files as: Warning

Open files

Number of retries: 0

Time out: 0

☑ Detect NTFS hardlinks

☐ Do not use archive attribute

☑ Backup share information for directories

☐ Asynchronous reading

MS Volume Shadow Copy Options

☑ Use Shadow Copy

☐ Allow fallback

Saving the recovery set file to the Cell Manager for a particular client in the backup specification

To copy the recovery set files only for particular clients in the backup specification, perform the following steps:

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Filesystem**.
3. Select the backup specification you will use for a full client backup (create it if you have not done so already). For details, see the HP Data Protector Help index: “creating, backup specifications”.
4. In the Results Area, click **Backup Object Summary**.
5. Select the client for which you would like to store its recovery set file onto the Cell Manager and click **Properties**.

6. In the **Other** page, select **Copy Recovery Set to disk**.

7. **Windows Vista and later releases:** In the **WinFS Options** page, leave the **Detect NTFS hardlinks** and **Use Shadow Copy** options selected, and leave **Allow Fallback** cleared. Note that the **Detect NTFS hardlinks** option is not automatically selected if you manually add objects or update existing backup specifications.

Preparing the Encryption Keys

For a Cell Manager recovery or an offline client recovery, you must ensure that the encryption keys are available during the disaster recovery by storing them on a removable medium. For a Cell Manager recovery, prepare the removable medium in advance, before the disaster occurs.

The encryption keys are not part of the DR OS image file. During the disaster recovery image creation, the keys are automatically exported to the Cell Manager to the file `*Data_Protector_program_data*\Config\Server\export\keys\DR-*ClientName*-keys.csv` (Windows systems) or `/var/opt/omni/server/export/keys/DR-*ClientName*-keys.csv` (UNIX systems), where *`ClientName`* is the name of the client for which the image is being created.

Ensure that you have the correct encryption key for each backup that is prepared for a disaster recovery.

Preparing a DR OS image

After a disaster occurs, you should prepare a DR OS image to be recorded on a disaster recovery CD or saved to a bootable USB drive, which can then be used for Enhanced Automated Disaster Recovery. Alternatively, you can prepare a bootable network image.

Note that the Data Protector Automatic Disaster Recovery component must be installed on the system where a DR OS image will be prepared.

A new disaster recovery OS image has to be prepared after each hardware, software or configuration change from a new recovery set.

Prepare a DR OS image in advance for any critical systems that must be restored first, especially systems required for the network to function properly (DNS servers, domain controllers, gateways, and so on), Cell Managers, Media Agent clients, file servers, and so on.

It is recommended to restrict access to backup media and disaster recovery CDs or USB drives containing the OS image.

Steps

1. In the Data Protector Context List, click **Restore**.

2. In the Scoping Pane, click **Tasks**, and then click **Disaster Recovery** to start the Disaster Recovery Wizard.

3. In the Results Area, from the **Host to be recovered** drop down list, select the client you would like to prepare the DR OS image for.

4. From the **Recovery media creation host** drop down list, select the client on which you will prepare the DR OS image. By default, this is the same client for which the DR OS image is prepared for. The client on which you prepare the DR OS image must have the same OS type installed (Windows, Linux) and must have a Disk Agent installed.

5. Keep the **Enhanced Automated Disaster Recovery** selected and select whether the volume recovery set will be built from a backup session or a list of volumes. By default, **Backup session** is selected.

   Click **Next**.

6. Depending on the recovery set build method select:

   - If you selected Backup session, select the host backup session and in case of a Cell Manager, the IDB session.

   - If you selected Volume list, for each critical object select an appropriate object version.

   Click **Next**.

7. Select the location of the recovery set file. By default, **Restore recovery set file from a backup** is selected.

   If you have saved the recovery set file on the Cell Manager during backup, select **Path to the recovery set file** and specify its location. Click **Next**.

8. Select the image format. The following options are available:

   - **Create bootable ISO image**: a DR ISO image (by default, `recovery.iso`)

   - **Create bootable USB drive**: a DR OS image on a bootable USB drive

   - **Create bootable network image**: a DR OS image that can be used for the network boot (by default, `recovery.wim`)

9. If you are creating a bootable ISO image or a bootable network image, select the destination directory, where you would like to place the created image.

   If you are creating a bootable USB drive, select the destination USB drive or disk number, where you would like to place the created image.

   Important: During the creation of the bootable USB drive, all data stored on the drive will be lost.

10. Optionally, set a password to protect the DR OS image from unauthorized use. The lock icon indicates whether a password has been set.

    Click **Password** to open the Password Protect Image dialog window and enter the password. To remove the password, clear the fields.

11. **Windows Vista and later releases:**:

Review and if necessary, modify the list of drivers that are inserted into the DR OS image.

You can use this option to add missing drivers to the DR OS. Add or remove drivers manually by clicking **Add** or **Remove**. To reload the original drivers, click **Reload**. The drivers from the `%Drivers%` part of the recovery set are automatically injected into the DR OS image.

> **Important:** The drivers collected during the backup procedure and stored within the recovery set's `%Drivers%` directory may not always be appropriate for use in the DR OS. In some cases, Windows Preinstallation Environment (WinPE) specific drivers may need to be injected to ensure that the hardware is functioning properly during the recovery.

12. Click **Finish** to exit the wizard and create the DR OS image.

13. If you are creating a bootable CD or DVD, record the ISO image on a CD or DVD using a recording tool that supports the ISO9660 format.

Recovering Windows Systems Using Enhanced Automated Disaster Recovery

You can successfully perform the Enhanced Automated Disaster Recovery of a Windows system only if all preparation steps were fulfilled. If you are recovering a Cell Manager, first the Internal Database is restored from its backup image, and restore of the volumes and the CONFIGURATION object from their backup image follows afterwards. For details on supported operating systems, see the *HP Data Protector Product Announcements, Software Notes, and References*.

Prerequisites

- You need a new hard disk to replace your affected disk.

- You should have a valid full filesystem backup of the entire system that you want to recover (client backup).

- For disaster recovery of the Cell Manager, you should have a valid Internal Database backup image that is newer than the filesystem backup image.

- You need a disaster recovery CD or a bootable USB drive with the DR OS image or a bootable network image.

- **Windows Server 2003:** If the affected system is a domain controller, you need the password of the Directory Services Restore Mode administrator account.

Steps

Phase 1

1. Unless you are performing an offline disaster recovery, add a Data Protector account with the following properties to the Data Protector admin user group on the Cell Manager, depending on the operating system of the target system:

   Windows Vista and later releases:

   - Type: `Windows`
   - Name: `SYSTEM`
   - Group/Domain: `NT AUTHORITY`
   - Client: the temporary hostname of the system being recovered

     A temporary hostname is assigned to the system by the Windows Preinstallation Environment (WinPE). You can retrieve it by running the `hostname` command in the Command Prompt window of the WinPE.

   Windows XP, Windows Server 2003:

   - Type: `Windows`
   - Name: `DRM$Admin`
   - Group/Domain: hostname of the target system
   - Client: fully qualified domain name (FQDN) of the target system

   For more information on adding users, see the HP Data Protector Help index: “adding Data Protector users”.

Figure 3: Adding a user account

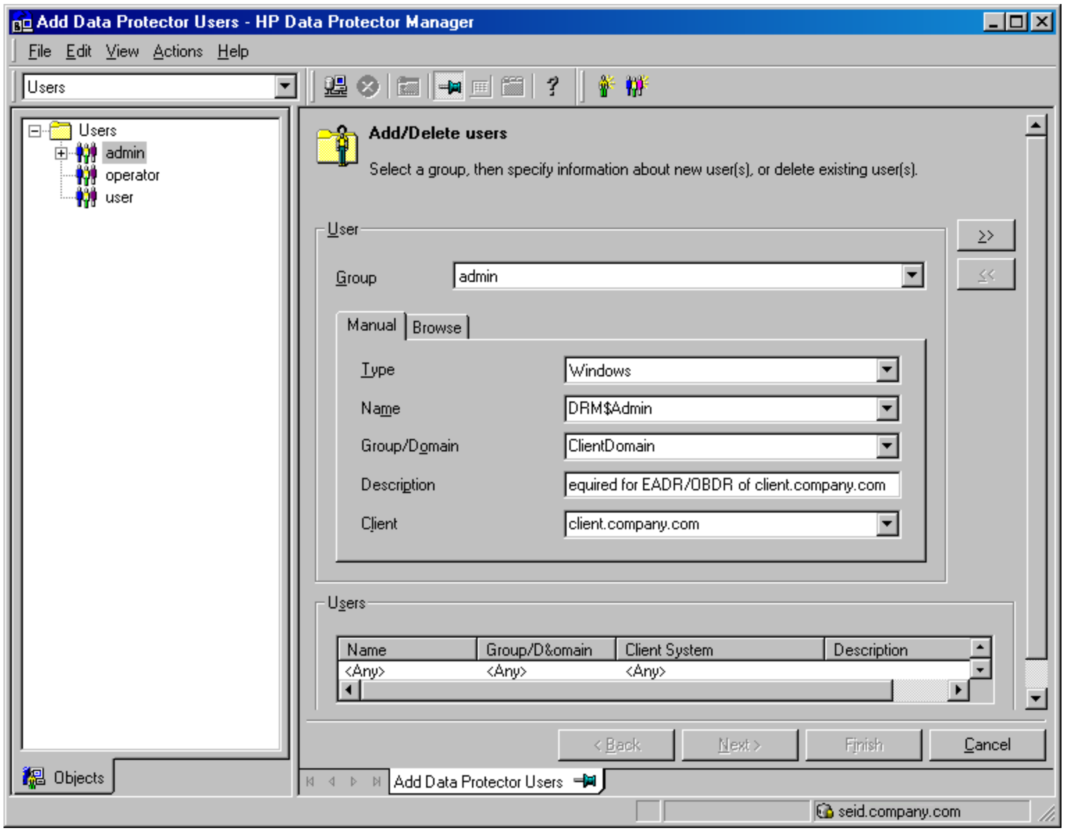

Add/Delete users

Select a group, then specify information about new user(s), or delete existing user(s).

User

Group: admin

Type: Windows

Name: DRM$Admin

Group/Domain: ClientDomain

Description: equired for EADR/OBDR of client.company.com

Client: client.company.com

Users

Name	Group/Domain	Client System	Description
<Any>	<Any>	<Any>	

Note: If you are using encrypted control communication between the clients in a Data Protector cell, you must add the client to the Security Exceptions list on the Cell Manager before you start the recovery. Unless you are using a local device, the Media Agent client must be added to the Security Exceptions list on the Cell Manager as well.

2. Boot the client system from the disaster recovery CD, the bootable USB drive, or the bootable network image of the original system. If you are starting the target system from a disaster recovery CD, ensure that no external USB disks (including USB keys) are connected to the system before you start the recovery procedure.

Note: If the screen is locked during a recovery, you can log on with following credentials:

User: `DRM$ADMIN`

Password: `Dr8$ad81n$pa55wD`

3. **Windows Server 2003:** If you are recovering a domain controller, when the Welcome to Windows dialog box appears, press **Ctrl+Alt+Delete**, enter the password of the Directory Services Restore Mode administrator account, and then click **OK**.

4. Select the scope of the recovery and recovery options. The following steps differ depending on the

operating system:

Windows Vista and later releases:

a. The Disaster Recovery GUI (the Installer Wizard) appears and displays the original system information. Click **Next**.

> **Tip:** There are some keyboard options available when the progress bar appears. You can check which options are available and their description by hovering over progress bar.

b. In the Recovery scope page, select the scope of the recovery:

- `Default Recovery`: Critical volumes (system disks, boot disk, and the Data Protector installation volume) are recovered. All other disks are partitioned and formatted and remain empty and ready for Phase 3.
- `Minimal Recovery`: Only system disks and boot disk are recovered.
- `Full Recovery`: All volumes in the Restore Set are recovered, not only the critical ones.
- `Full with Shared Volumes`: Available for Microsoft Cluster Server (MSCS). This option should be used if all nodes in the MSCS have been struck by a disaster and you are performing EADR 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, then shared volumes will not be restored because the node keeps them locked. In this case, you should use `Default Recovery`.

c. Optionally, to modify the recovery settings, click **Settings** to open the Recovery settings page.

The following additional recovery options are available, some of them are used in cases where the disaster recovery does not finish completely or requires additional steps:

- `Use original network settings`: Select this option if you need to restore the original network configuration (for example, due to a missing DHCP server). By default, this option is not selected and the DR OS recovery environment uses a DHCP network configuration.
- `Restore BCD`: If selected, Data Protector also restores the Boot Configuration Data (BCD) store in advance during the disaster recovery session, before it is restored in the Data Protector restore session. The option is selected by default.
- `Restore DAT`: If selected, the Data Protector disaster recovery module also restores Microsoft VSS writers' data. By default, the DR module skips the restore of VSS writer's data. You can use this option if Data Protector fails to back up critical writers during a non-VSS backup. To restore the data before a DR module restore, select `Pre`. To restore the data after a Data Protector, select `Post`.
- `Initialize Disks Manually`: This option enables you to manually map the original and current system disks and initialize them to match the original configuration. By default, this option is not selected.

If selected, a new disk mapping and initialization page is displayed when the recovery process starts. The disaster recovery module will provide the initial disk mapping and display the result of the initial mapping attempt. Use the provided options to change the disk mapping. Once the mapping is completed, the volumes are initialized and the system restarts.

- `Restore Storage Spaces`: By default, Storage Spaces are restored. You can deselect the option and restore the virtual disks directly to physical ones, at recovery time, if the storage configuration permits this. Note that you need to manually initialize the disks if you restore Storage Spaces to dissimilar hardware or USB disks.
- `Enable Dissimilar Hardware Restore`: If enabled, Data Protector scans the system for missing drivers during the recovery. The option is enabled by selecting one of the following methods from the drop-down list:
  - `Unattend` (default): This mode automatically configures the operating system to various hardware platforms using a predefined configuration file. This is the primary mode of recovery with dissimilar hardware. Use it in the first instance.
  - `Generic`: Select this if Unattend mode fails (perhaps because of a misconfiguration of the restored operating system). It is based on adapting the restored OS registry and its drivers and services to the dissimilar hardware.
- `Remove Devices`: Available if the `Dissimilar Hardware` option is enabled. If selected, Data Protector removes original devices from the registry of the restored operating system.
- `Connect iSCSI Devices`: This option is enabled and selected if the original machine was using iSCSI. By selecting this option Data Protector automatically restores the basic iSCSI configuration as it was at backup time. If not selected, the iSCSI configuration will be skipped.

  You can also use the native Microsoft iSCSI configuration wizard to manage a more complex iSCSI configuration. If the DR GUI detects certain iSCSI features (for example, security options) which require a manual configuration, it offers the option to run the Microsoft iSCSI configuration wizard.
- `Map Cluster Disks Manually`: Available on Windows Server 2008 and later releases. If selected, you can map cluster volumes manually. If not selected, the volumes will be mapped automatically. It is recommended to check that all volumes are mapped appropriately after automatic mapping.
- `Remove Boot Descriptor`: Available on Intel Itanium systems. Removes all Boot Descriptors left over by the disaster recovery processes.
- `Manual disk selection`: Available on Intel Itanium systems. If the disk setup has changed significantly, the disaster recovery module may not be able to find the boot disk(s). Use this option to select the correct boot disk.

To reset the options to the default settings, click **Reset default settings**.

Click **Save >** to save the changes.

d. Click **Finish** to start the recovery. The recovery process starts and you can monitor the progress.

If the volumes are encrypted using BitLocker Drive Encryption, you are prompted to unlock the encrypted drives.

> **Tip:** In the Disaster Recovery GUI, you can click **Tasks** to perform the following:
>
> - run Command Prompt, Task Manager, or Disk Administrator
> - access the `Map Network Drives` and `Load Drivers` tools
> - view log files specific to the disaster recovery process
> - enable or disable the DRM configuration file, view this file in text editor, and edit it
> - edit the hosts file of the WinPE recovery environment
> - access Help and view the legends to GUI icons

Windows XP and Windows Server 2003 systems:

a. Press **F12** when the following message is displayed: `To start recovery of the machine` *`Hostname`* `press F12`.

b. The scope selection menu is displayed at the beginning of the boot process. Select the scope of recovery and press **Enter**. There are five different scopes of recovery:

   - `Reboot`: Disaster recovery is not performed and the system is restarted.
   - `Default Recovery`: Critical volumes (system disks, boot disk, and the Data Protector volume) are recovered. All other disks are partitioned and formatted and remain empty and ready for Phase 3.
   - `Minimal Recovery`: Only system disks and boot disk are recovered.
   - `Full Recovery`: All volumes in the Restore Set are recovered, not only the critical ones.
   - `Full with Shared Volumes`: Available for Microsoft Cluster Server (MSCS). This option should be used if all nodes in the MSCS have been struck by a disaster and you are performing EADR 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, then shared volumes will not be restored because the node keeps them locked. In this case, you should use `Default Recovery`.

   The following additional recovery options are available, some of them are used in cases where the disaster recovery does not finish completely or requires additional steps:

   - `Remove Boot Descriptor`: Available on Intel Itanium systems. Removes all Boot Descriptors left over by the disaster recovery processes.
   - `Manual disk selection`: Available on Intel Itanium systems. If the disk setup has changed

significantly, the disaster recovery module may not be able to find the boot disk(s). Use this option to select the correct boot disk.

Phase 2

4. After you have selected the scope of the recovery, Data Protector starts setting up the DR OS. You can monitor the progress and, when the DR OS is set up, the system restarts. On Windows Vista and later releases, the system restart is not performed.

   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.

   On Windows XP and Windows Server 2003, if the DR OS does not boot normally or cannot access network, you may need to edit the kb.cfg file.

   The Disaster Recovery Wizard appears. To modify the disaster recovery options, press any key to stop the wizard during the countdown and modify the options.

   The following options are available:

   - `Debugs...`: Enables debugging. See "Debugging disaster recovery sessions" on page 129.
   - `Omit deleted files`: Files, deleted between successive incremental backups, are not restored. This may slow down the recovery.
   - `Install only`: This option will install only the temporary operating system to the target system and thus finish the Phase 1 of disaster recovery. Phase 2 of disaster recovery will not start automatically. You can use this option for example if you need to edit the SRD file.

   Additionally, you can start the Registry Editor, the command line, or the Task manager using the appropriate buttons.

   Click **Finish** to continue with the disaster recovery.

5. If the DR OS image is password protected, provide the password and continue the recovery.

6. If the disaster recovery backup is encrypted and you are either recovering the Cell Manager or a client where the Cell Manager is not accessible, the following prompt will appear:

   ```
   Do you want to use AES key file for decryption [y/n]?
   ```

   Press **y**.

   Ensure that the keystore (`DR-`*`ClientName`*`-keys.csv`) is available on the client (for example, by inserting a CD-ROM, floppy disk, or USB flash drive) and enter the full path to the keystore file. The keystore file is copied to the default location on the DR OS and is used by the Disk Agents. Disaster recovery now continues without further interruption.

7. 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.

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 interrupt the Disaster Recovery Wizard during the 10 second pause (after it has found the DR installation and SRD file on the backup medium) and select the **Debugs** option.
   - You manually execute the `omnidr` command with the `-no_reset` or `-debug` option.
   - Disaster recovery fails.

   On Windows Vista and later releases, the temporary DR OS is never retained.

   Note that Data Protector will first try to perform online recovery. If the online recovery fails for any reason (for example, the Cell Manager or network services are not available, the firewall is preventing access to the Cell Manager) Data Protector will then try to perform a remote offline recovery. If even the remote offline restore fails (for example, because the Media Agent host accepts only requests from the Cell Manager), Data Protector will perform a local offline restore.

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. If you are recovering a Cell Manager, make the IDB consistent.

Phase 3

10. 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 new files created to be compressed as well.

11. Additional steps are necessary if you are performing disaster recovery of all nodes in a Microsoft Cluster Server.

One Button Disaster Recovery (OBDR)

One Button Disaster Recovery (OBDR) is a automated Data Protector recovery method for Windows Data Protector clients, where user intervention is reduced to minimum. For details on supported operating systems, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.

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, the 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.

Once DR OS Image is booted, Data Protector automatically formats and partitions the disk 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 storing the Data Protector installation data

Any remaining partitions can be recovered using the standard Data Protector recovery procedure.

Overview

The general steps using the One Button Disaster Recovery method for a Windows client are:

1. **Phase 0**

   a. You need an OBDR backup image (create the backup specification using the Data Protector One Button Disaster Recovery Wizard).

   b. If you are using encrypted backups, store the encryption key on a removable medium so that it is available for disaster recovery.

2. **Phase 1**

   Boot from the recovery tape and select the scope of recovery.

3. **Phase 2**

   Depending on the recovery scope you select, the selected volumes are automatically restored.

   Critical volumes (the boot partition and the operating system) are always restored.

4. **Phase 3**

   Restore any remaining partitions using the standard Data Protector restore procedure.

> **Important:** HP recommends to restrict access to OBDR boot media.

The following sections explain the requirements, limitations, preparation and recovery pertaining to One Button Disaster Recovery on Windows systems. See also the section “Advanced recovery tasks”.

Requirements

- The Data Protector Automatic Disaster Recovery must be installed on systems for which you want to enable recovery using this method. For details, see the *HP Data Protector Installation and Licensing Guide*.

- The client system must support booting from the tape device that will be used for OBDR.

  For more information about supported systems, devices and media, see the HP Tape Hardware Compatibility Table and the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.

- The hardware configuration of the target system must be the same as that of the original system. This includes the SCSI BIOS settings (sector remapping).

- The new disk have to be the same size or bigger than the affected disk. If it is larger than the original disk, the difference will remain unallocated.

- Replacement disks have to be attached to the same host bus adapter on the same bus.

- Windows XP, Windows Server 2003: 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.

- During OBDR backup,, the partition on which Data Protector is installed should have at least 500 MB of temporary free space. This space is required to create a temporary image.

- Windows Server 2003: All drivers required for booting must be installed under the `%SystemRoot%`folder.

- A media pool with a Non-appendable media usage policy and a Loose media allocation policy has to be created for the OBDR capable device. Only media from this pool can be used for disaster recovery.

- Windows XP, Windows Server 2003: The operating system should be activated at the time of the backup. Otherwise, when the activation period expires, disaster recovery fails.

- To create a DR OS image for Windows Vista and later releases, you must install the appropriate version of Windows Automated Installation Kit (WAIK) or Assessment and Deployment Kit on the system on which you will create the image:

  Windows Vista and Windows Server 2008:

  Automated Installation Kit (AIK) for Windows Vista SP1 and Windows Server 2008

  Windows 7 and Windows Server 2008 R2:

  - Windows Automated Installation Kit (AIK) for Windows 7

  - Windows Automated Installation Kit (AIK) Supplement for Windows 7 SP1 (optional, for Microsoft Windows 7 SP1 and Windows Server 2008 R2 SP1)

Windows 8 and Windows Server 2012:

- Assessment and Deployment Kit (ADK 1.0) for Windows 8 and Windows Server 2012

Windows 8.1 and Windows Server 2012 R2:

- Assessment and Deployment Kit (ADK 1.1) for Windows 8.1 and Windows Server 2012 R2

- To back up the IIS configuration object residing on a Windows Vista, Windows 7, or Windows Server 2008 system, install the IIS 6 Metabase Compatibility package.

Limitations

- One Button Disaster Recovery (OBDR) is not available for Data Protector Cell Managers.
- Multiboot systems that do not use Microsoft's boot loader are not supported.
- On Windows XP and Windows Server 2003, recovery of a SAN boot configuration is not supported.
- VSS disk image backup of the logical volumes can be used for disaster recovery only on Windows Vista and later releases.
- On Windows XP and Windows Server 2003, a console interface is available instead of the HP Data Protector Disaster Recovery GUI.
- On Windows XP and Windows Server 2003, recovery of a configuration with Network Teaming adapters is not supported.
- On Windows Vista and later releases, originally encrypted folders can only be restored as unencrypted.
- The Internet Information Server 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.
- The DRM restore monitor monitors the overall bytes written to a disk by the VRDA process. The overall bytes written to a disk do not always match what is displayed in the Data Protector session manager.

  > **Note:** The new Recovery Session monitor is implemented only on Windows Vista and later releases.

- Sparse files are restored to their full size during offline restore. This may result in the target volume running out of space.

Disk and partition configuration

- Dynamic disks are not supported (including mirror sets upgraded from Windows NT).
- A 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.
- OBDR is supported on systems where Data Protector is installed on an NTFS volume.
- On Intel Itanium systems, recovery of a boot disk is supported only for local SCSI disks.

Preparation for One Button Disaster Recovery

To prepare for a successful disaster recovery, follow the instructions related to the general preparation procedure for disaster recovery before completing the steps listed in this topic. Prepare in advance in order to perform a disaster recovery fast and efficiently.

Important: Prepare for disaster recovery before a disaster occurs.

Preparatory steps

After you have completed the general preparation for disaster recovery, perform the following specific steps to prepare for OBDR.

1. Create a media pool for DDS or LTO media with the Non-appendable media usage policy and the Loose media allocation policy (because the backup media is formatted during OBDR backup). In addition,specify this media pool as the default media pool for the OBDR device. See the *HP Data Protector Help* index: “creating media pool”. Only media from such pool can be used for OBDR.

2. Perform the OBDR backup locally on the system for which you want to enable recovery using OBDR.

   Considerations

   Windows Vista and later releases: Make sure that you back up system volumes (such as boot volumes) if present.

   Windows Server 2012 (R2): Use disk image backup to back up volumes in the following cases:

   - Deduplicated volumes

     During a filesystem restore, the volume is rehydrated and you might run of space on the destination volume during recovery. A disk image restore keeps the size of the volume.

   - Volumes with Resilient File System (ReFS)

   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, the client system's virtual server.

  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.

  Cluster Shared Volumes: Before performing a full backup of the client system, back up the Virtual Hard Drive (VHD) files and CSV configuration data using the Data Protector Virtual Environment first. See the *HP Data Protector Integration Guide for Virtualization Environments*. The backup must be performed on a separate device, because an OBDR backup can be performed only on non-appendable media.

  Virtual Hard Drives (VHD) must be dismounted to ensure consistency.

  If the full client backup was encrypted, store the encryption key on a removable medium so that it is available for disaster recovery. You will need the key if the connection to the Cell Manager cannot be established.

3. Execute a disaster recovery test plan.

4. On Windows systems, if some service or driver is not operational after the system startup, you may have to manually edit the `kb.cfg` file.

Creating the Backup Specification for One Button Disaster Recovery

You need to create a One Button Disaster Recovery (OBDR) backup specification in order to prepare the OBDR boot tape.

Prerequisites

- Before adding an OBDR device, create a media pool for DDS or LTO media with the Non-appendable media usage policy and the Loose media allocation policy. The created media pool must be selected as the default media pool for the OBDR device.

- This device has to be connected locally to the system, for which you want to enable recovery using OBDR.

- The Data Protector Automatic Disaster Recovery and User Interface components must be installed on systems for which you want to enable recovery using the OBDR method.

- This backup specification has to be created locally on the system, for which you want to enable recovery using OBDR.

Tip: To enable an automatic restore of all shared disk volumes in the MS Cluster using the OBDR method, move all volumes temporarily to the node for which you are preparing the OBDR boot tape. It is practically impossible to collect enough information to configure the disk in Phase 1 for shared disk volumes that are locked by another node.

Limitations

- One Button Disaster Recovery (OBDR) is not available for Data Protector Cell Managers.

This backup specification is unique to the One Button Disaster Recovery method. By default, the required volumes are backed up as filesystems. However, on Windows Vista and later releases, you can choose to back up logical volumes as disk images by using the VSS writers. This ensures that the volumes remain unlocked during the backup and can be accessed by other applications. To back up logical volumes as disk images, you must modify the backup specification created for OBDR.

Creating a backup specification for OBDR

Modifying an OBDR backup specification to use disk image backup

Creating a backup specification for OBDR

Steps

1. In the Data Protector Context List, click **Backup**.
2. In the Scoping Pane, click **Tasks**, and then click **One Button Disaster Recovery Wizard**.
3. In the Results Area, select the client for which you would like to perform an OBDR backup (locally on the client) from the drop-down list and click **Next**.
4. The critical volumes that you need to back up are already selected. Click **Next**.

   Important: Important volumes are selected automatically and cannot be deselected. Select any other partitions you want to keep, because during the recovery procedure Data Protector deletes all partitions from your system.

5. Select the local device or drive to be used for the backup. Only one device or drive can be selected. Click **Next**.
6. **Windows Vista or later releases:**

   Review and if necessary, modify the list of drivers that are inserted into the DR OS image.

You can use this option to add missing drivers to the DR ISO image. Add or remove drivers manually by clicking **Add** or **Remove**. To reload the original drivers, click **Reload**. The drivers from the `%Drivers%` part of the recovery set are automatically injected into the DR OS image.

> **Important:** The drivers collected during the backup procedure and stored within the recovery set's `%Drivers%` directory may not always be appropriate for use in the DR OS. In some cases, Windows Preinstallation Environment (WinPE) specific drivers may need to be added to ensure that the hardware is functioning properly during the recovery.

Optionally, select the backup options.

Click **Next**.

7. Optionally, schedule a backup. Click **Next**.
8. In the Backup Summary page, review the backup specification settings, and then click **Next**.

   You cannot change a previously selected backup device or the order in which the backup specifications follow one another. Only OBDR non-essential backup objects can be deleted and only general object properties can be viewed. You can also change a backup object description.
9. Save the modified backup specification as an OBDR backup specification to keep it in the original One Button Disaster Recovery format.

Modifying an OBDR backup specification to use disk image backup

Steps

1. In the Scoping Pane, click the created OBDR backup specification. When you are asked, whether you want to treat it as an OBDR backup specification or as an ordinary backup specification, click **No**.

   > **Note:** When an OBDR backup specification is saved as an ordinary backup specification, it can be still used for the OBDR.
2. In the Backup Object Summary page, select the logical volumes that you want to back up as disk images and click **Delete**.

   > **Note:** You can back up only logical volumes. The configuration objects, as well as volumes that are not mounted or are mounted as NTFS folders, should be backed up with filesystem backup.
3. Click **Manual add** to open the wizard.
4. In the Select Backup Object page, click the **Disk image object** option, and then click **Next**.

5. In the General Selection page, select a client with the disk image you want to back up and provide an appropriate description. Click **Next**.

    > **Note:** Description must be unique for each disk image object. Use a descriptive name, for example, `[Disk Image C] for C: volume`.

6. In the General Object Options property page, set data protection to **None**. Click **Next**.

    > **Note:** When you set data protection to **None**, the content of the tape can be overwritten by the newer OBDR backups.

7. In the Advanced Object Options property page, you can specify advanced backup options for the disk image object. Click **Next**.

8. In the Disk Image Object Options property page, specify the disk image sections to back up. Use the following format:

    `\\.\`*`DriveLetter`*`:`, for example: `\\.\E:`

    > **Note:** When the volume name is specified as a drive letter, the volume is not being locked during the backup. A volume that is not mounted or is mounted as an NTFS folder cannot be used for the disk image backup.

9. Click **Finish** to exit the wizard.

10. In the Backup Object Summary page, review the summary of the backup specification. The logical volumes that you specified as disk images should be of a Disk Image type. Click **Apply**.

Preparing the Encryption Keys

For a Cell Manager recovery or an offline client recovery, you must ensure that the encryption keys are available during the disaster recovery by storing them on a removable medium. For a Cell Manager recovery, prepare the removable medium in advance, before the disaster occurs.

The encryption keys are not part of the DR OS image file. During the disaster recovery image creation, the keys are automatically exported to the Cell Manager to the file *`Data_Protector_program_data`*`\Config\Server\export\keys\DR-`*`ClientName`*`-keys.csv` (Windows systems) or `/var/opt/omni/server/export/keys/DR-`*`ClientName`*`-keys.csv` (UNIX systems), where *`ClientName`* is the name of the client for which the image is being created.

Ensure that you have the correct encryption key for each backup that is prepared for a disaster recovery.

Recovering Windows Systems Using One Button Disaster Recovery

You can successfully perform the One Button Disaster Recovery (OBDR) of a Windows system only if all preparation steps were fulfilled.

For details on supported operating systems for OBDR, see the *HP Data Protector Product Announcements, Software Notes, and References*.

Prerequisites

- You need a new hard disk to replace your affected disk.
- You should have a bootable OBDR backup medium with all critical objects of the client that you want to recover. The OBDR backup has to be performed locally on the client.
- You need an OBDR device connected locally to the target system.

Steps

Phase 1

1. Unless you are performing an offline disaster recovery, add the account with the following properties to the Data Protector `admin` user group on the Cell Manager, depending on the operating system of the target system:

   Windows Vista and later releases:

   - Type: `Windows`
   - Name: `SYSTEM`
   - Group/Domain: `NT AUTHORITY`
   - Client: the temporary hostname of the system being recovered

     A temporary hostname is assigned to the system by the Windows Preinstallation Environment (WinPE). You can retrieve it by running the `hostname` command in the Command Prompt window of the WinPE.

   Windows XP, Windows Server 2003:

   - Type: `Windows`
   - Name: `DRM$Admin`

- Group/Domain: hostname of the target system
- Client: fully qualified domain name (FQDN) of the target system

For more information on adding users, see the HP Data Protector Help index: “adding Data Protector users”.

Figure 4: Adding a user account

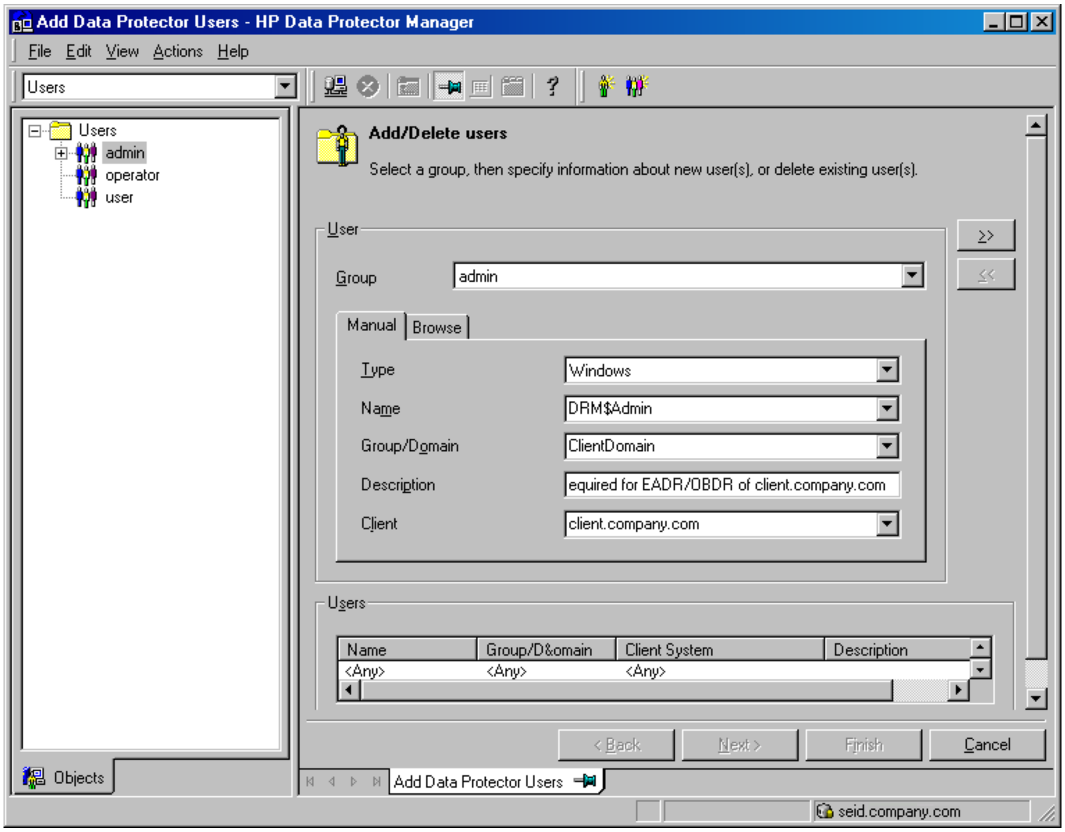

Add/Delete users

Select a group, then specify information about new user(s), or delete existing user(s).

User

Group: admin

Type: Windows

Name: DRM$Admin

Group/Domain: ClientDomain

Description: equired for EADR/OBDR of client.company.com

Client: client.company.com

Users

Name	Group/Domain	Client System	Description
<Any>	<Any>	<Any>	

> **Note:** If you are using encrypted control communication between the clients in a Data Protector cell, you must add the client to the Security Exceptions list on the Cell Manager before you start the recovery. Unless you are using a local device, the Media Agent client must be added to the Security Exceptions list on the Cell Manager as well.

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. Ensure that no external USB disks (including USB keys) are connected to the system before you start the recovery procedure.

4. Power the target system on and, while it is being initialized, press the **Eject** button on the tape device and power it on. For details, see the device documentation.

5. Select the scope of the recovery and recovery options. The following steps differ depending on the operating system:

Windows Vista and later releases::

a. The Disaster Recovery GUI (the Installer Wizard) appears and displays the original system information. Click **Next**.

> **Tip:** There are some keyboard options available when the progress bar appears. You can check which options are available and their description by hovering over progress bar.

b. In the Recovery scope page, select the scope of the recovery:

   - `Default Recovery`: Critical volumes (system disks, boot disk, and the Data Protector installation volume) are recovered. All other disks are partitioned and formatted and remain empty and ready for Phase 3.

   - `Minimal Recovery`: Only system disks and boot disk are recovered.

   - `Full Recovery`: All volumes in the Restore Set are recovered, not only the critical ones.

   - `Full with Shared Volumes`: Available for Microsoft Cluster Server (MSCS). This option should be used if all nodes in the MSCS have been struck by a disaster and you are performing EADR 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, then shared volumes will not be restored because the node keeps them locked. In this case, you should use `Default Recovery`.

c. Optionally, to modify the recovery settings, click **Settings** to open the Recovery settings page.

   The following additional recovery options are available, some of them are used in cases where the disaster recovery does not finish completely or requires additional steps:

   - `Use original network settings`: Select this option if you need to restore the original network configuration (for example, due to a missing DHCP server). By default, this option is not selected and the DR OS recovery environment uses a DHCP network configuration.

   - `Restore BCD`: If selected, Data Protector also restores the Boot Configuration Data (BCD) store in advance during the disaster recovery session, before it is restored in the Data Protector restore session. The option is selected by default.

   - `Restore DAT`: If selected, the Data Protector disaster recovery module also restores Microsoft VSS writers' data. By default, the DR module skips the restore of VSS writer's data. You can use this option if Data Protector fails to back up critical writers during a non-VSS backup. To restore the data before a DR module restore, select `Pre`. To restore the data after a Data Protector, select `Post`.

   - `Initialize Disks Manually`: This option enables you to manually map the original and current system disks and initialize them to match the original configuration. By default, this option is not selected.

If selected, a new disk mapping and initialization page is displayed when the recovery process starts. The disaster recovery module will provide the initial disk mapping and display the result of the initial mapping attempt. Use the provided options to change the disk mapping. Once the mapping is completed, the volumes are initialized and the system restarts.

- `Restore Storage Spaces`: By default, Storage Spaces are restored. You can deselect the option and restore the virtual disks directly to physical ones, at recovery time, if the storage configuration permits this. Note that you need to manually initialize the disks if you restore Storage Spaces to dissimilar hardware or USB disks.

- `Enable Dissimilar Hardware Restore`: If enabled, Data Protector scans the system for missing drivers during the recovery. The option is enabled by selecting one of the following methods from the drop-down list:

  - `Unattend` (default): This mode automatically configures the operating system to various hardware platforms using a predefined configuration file. This is the primary mode of recovery with dissimilar hardware. Use it in the first instance.

  - `Generic`: Select this if Unattend mode fails (perhaps because of a misconfiguration of the restored operating system). It is based on adapting the restored OS registry and its drivers and services to the dissimilar hardware.

- `Remove Devices`: Available if the `Dissimilar Hardware` option is enabled. If selected, Data Protector removes original devices from the registry of the restored operating system.

- `Connect iSCSI Devices`: This option is enabled and selected if the original machine was using iSCSI. By selecting this option Data Protector automatically restores the basic iSCSI configuration as it was at backup time. If not selected, the iSCSI configuration will be skipped.

  You can also use the native Microsoft iSCSI configuration wizard to manage a more complex iSCSI configuration. If the DR GUI detects certain iSCSI features (for example, security options) which require a manual configuration, it offers the option to run the Microsoft iSCSI configuration wizard.

- `Map Cluster Disks Manually`: Available on Windows Server 2008 and later releases. If selected, you can map cluster volumes manually. If not selected, the volumes will be mapped automatically. It is recommended to check that all volumes are mapped appropriately after automatic mapping.

- `Remove Boot Descriptor`: Available on Intel Itanium systems. Removes all Boot Descriptors left over by the disaster recovery processes.

- `Manual disk selection`: Available on Intel Itanium systems. If the disk setup has changed significantly, the disaster recovery module may not be able to find the boot disk(s). Use this option to select the correct boot disk.

To reset the options to the default settings, click **Reset default settings**.

Click **Save >** to save the changes.

d. The recovery process starts and you can monitor the progress.

If the volumes are encrypted using BitLocker Drive Encryption, you are prompted to unlock the encrypted drives.

> **Tip:** In the Disaster Recovery GUI, you can click **Tasks** to perform the following:
>
> - run the Command Prompt, Task Manager, or Disk Administrator
> - access the `Map Network Drives` and `Load Drivers` tools
> - view log files specific to the disaster recovery process
> - enable or disable the DRM configuration file, view this file in text editor, and edit it
> - edit the hosts file of the WinPE recovery environment
> - access Help and view the legends to GUI icons

Windows XP, Windows Server 2003:

a. Press **F12** when the following message is displayed: `To start recovery of the machine HOSTNAME press F12`.

b. The scope selection menu is displayed at the beginning of the boot process. Select the scope of recovery and press **Enter**. There are five different scopes of recovery:

   - `Reboot`: Disaster recovery is not performed and the system is restarted.
   - `Default Recovery`: Critical volumes (system disks, boot disk, and the `OBInstall` volume) are recovered. All other disks are partitioned and formatted and remain empty and ready for Phase 3.
   - `Minimal Recovery`: Only system disks and boot disk are recovered.
   - `Full Recovery`: All volumes in the Restore Set are recovered, not only the critical ones.
   - `Full with Shared Volumes`: Available for Microsoft Cluster Server (MSCS). This option should be used if all nodes in the MSCS have been struck by a disaster and you are performing OBDR 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, then shared volumes will not be restored because the node keeps them locked. In this case, you should use `Default Recovery`.

   The following additional recovery options are available, some of them are used in cases where the disaster recovery does not finish completely or requires additional steps:

   - `Remove Boot Descriptor`: Available on Intel Itanium systems. Removes all Boot Descriptors left over by the disaster recovery processes.
   - `Manual disk selection`: Available on Intel Itanium systems. If the disk setup has changed

significantly, the disaster recovery module may not be able to find the boot disk(s). Use this option to select the correct boot disk.

Phase 2

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 restarts. If the DR OS does not boot normally or cannot access network, then you may need to edit the kb.cfg file. On Windows Vista and later releases, the DR OS is not installed and the system restart is not performed.

7. If the disaster recovery backup is encrypted and you are recovering a client whose Cell Manager is not accessible, the following prompt is displayed:

   ```
   Do you want to use AES key file for decryption [y/n]?
   ```

   Press **y**.

   Ensure that the keystore (`DR-`*`ClientName`*`-keys.csv`) is available on the client (for example, by inserting a CD-ROM, floppy disk, or USB flash drive) and enter the full path to the keystore file. The keystore file is copied to the default location on the DR OS and is used by the Disk Agents. Disaster recovery now continues without further interruption.

8. 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.

9. 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 interrupt the Disaster Recovery Wizard during the 10 second pause (after it has found the DR installation and SRD file on the backup medium) and select the **Debugs** option.
   - You manually execute the `omnidr` command with the `-no_reset` or `-debug` option.
   - Disaster recovery fails.

   Note that Data Protector first tries to perform an online restore. If the online restore fails for any reason (for example, the Cell Manager or network service is not available or firewall is preventing access to the Cell Manager) Data Protector tries to perform remote offline recovery. If the remote offline restore fails (for example, because the Media Agent host accepts requests only from the Cell Manager), Data Protector performs a local offline restore.

10. 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.

Phase 3

12. 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 have to manually set the volume compression if you want any new files to be compressed as well.

13. Additional steps are necessary if you are performing disaster recovery of all nodes in a Microsoft Cluster Server.

Advanced tasks

Disaster Recovery of Microsoft Cluster Server

About Disaster Recovery of a Microsoft Cluster Server

Microsoft Cluster Server (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 also apply for the disaster recovery of the MSCS. 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.

For details on supported operating systems, see the *HP Data Protector Product Announcements, Software Notes, and References*.

All prerequisites for disaster recovery (for example, a consistent and up-to-date backup, an updated SRD file, replaced faulty hardware, and so on) must be met to recover the MSCS.

Possible scenarios

There are two possible scenarios for disaster recovery of an MSCS:

- a disaster occurred to a non-active(s) node
- all nodes in the cluster have experienced a disaster

Preparation for Microsoft Cluster Server Disaster Recovery Specifics

All prerequisites for disaster recovery (such as consistent and up-to-date backup images, an updated SRD file, replaced faulty hardware, ...) must be met to recover the Microsoft Cluster Server (MSCS). All specifics, limitations, and requirements pertaining a particular disaster recovery method will also apply for the disaster recovery of an MSCS.

Consistent backup image for an MSCS includes:

- all nodes
- the virtual server
- if Data Protector is configured as a cluster-aware application, the Cell Manager should be included in the backup specification

EADR specifics

It is practically 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.

OBDR specifics

To enable faster restore, use the `omnisrdupdate` command as a post-exec command to update the SRD file after the OBDR backup. Insert the diskette with an updated SRD file in the floppy disk drive when performing OBDR to provide Data Protector with information on the location of backed up objects on the tape. Restoring the MSCS database will be faster because Data Protector will not search the tape for the location of the MSCS database.

To enable the automatic restore of all shared disk volumes in the MSCS, temporarily move all volumes to the node for which you are preparing the OBDR boot tape. 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.

Recovering a Microsoft Cluster Server

There are two possible scenarios for disaster recovery of a Microsoft Cluster Server (MSCS):

At least one of the nodes is up and running

All nodes in the cluster have experienced a disaster

At least one of the nodes is up and running

This is the basic scenario for disaster recovery of an MSCS. The prerequisites that follow must be fulfilled in addition to other prerequisites for disaster recovery.

Prerequisites

- At least one of the cluster nodes is functioning properly (active node).
- The cluster service is running on this node.
- All physical disk resources must be on-line (that is, owned by the cluster).
- All normal cluster functionality is available (the cluster administration group is on-line).
- 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 affected non-active node.

Only local disks are restored, because all shared disks are moved to the working node after the disaster 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 a part of the CONFIGURATION object on Windows systems.

All nodes in the cluster have experienced a disaster

In this case, all nodes in the MSCS are unavailable and the cluster service is not running.

The prerequisites that follow must be fulfilled in addition to other prerequisites for disaster recovery.

Prerequisites

- The primary node must have write access to the quorum disk (the quorum disk must not be locked).
- The primary node must have access to all IDB volumes, when recovering the Cell Manager.

In this case, you have to 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.

For AMDR, 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 MSCS Service will not recognize the replaced disks as valid cluster resources, and cluster groups depending on these resources will fail. To prevent this, restore the original hard disk signatures, if you replaced the shared cluster disks.

Steps

1. Perform disaster recovery of the primary node (including the quorum disk).

   Assisted Manual Disaster Recovery (AMDR):

   All user and application data on the quorum disk will be restored automatically by the `drstart -full_clus` command.

   Enhanced Automatic Disaster Recovery (EADR), One Button Disaster Recovery (OBDR):

   When you are asked to select the scope of recovery, select **Full with Shared Volumes** to restore the quorum disk.

2. Restart the system.

3. Restore the MSCS database, which is a part of the CONFIGURATION object on Windows systems. MSCS service must be running in order to be able to restore the MSCS database, therefore it cannot be restored automatically during Phase 2 of disaster recovery. However, the cluster database can be restored manually at the end of Phase 2 using the standard Data Protector restore procedure.

4. **Methods other than One Button Disaster Recovery (OBDR):**

   If you are recovering a Cell Manager, make the IDB consistent.

5. The quorum and IDB 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 perform the following steps:

   a. Disable the cluster service and cluster disk driver (the steps are described in MSDN Q176970).

   b. Restart the system.

   c. Reestablish the previous storage structure.

   d. Enable the cluster disk driver and cluster service.

   e. Restart the system and restore user and application data.

6. Restore the remaining nodes.

Merging P1S Files for Microsoft Cluster Server

After a backup has been performed, another step is required for Enhanced Automated Disaster Recovery (EADR) to restore the active node. Information on shared cluster volumes in P1S files for all nodes in the Microsoft Cluster Server (MSCS) has to be merged so that the P1S file of each node contains information on the shared cluster volumes configuration. This is necessary to enable restore of all shared cluster volumes. You can avoid merging P1S files after backup by moving all shared cluster volumes temporarily to the node which you are going to back up. In this case, all required information about all shared cluster volumes can be collected. This means that only that node can be the primary node.

Windows

To merge the P1S files of all nodes, execute the merge.exe command from the *Data_Protector_home*\bin\drim\bin directory:

```
merge 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.

Filenames of updated P1S files have .merged appended (for example, computer.company.com.merged). Rename the merged P1S files back to their original names (delete the .merged extension).

For example, to merge the P1S files for an MSCS with 2 nodes, type:

```
merge Data_Protector_program_data\Config\server\dr\p1s\node1.company.com Data_Protector_program_data\Config\server\dr\p1s\node2.company.com.
```

The merged files will be node1.company.com.merged and node2.company.com.merged.

UNIX

The merge.exe command works only on Windows systems with the Data Protector Automatic Disaster Recovery component installed. On a UNIX Cell Manager, perform the procedure below.

Steps

1. Copy the P1S files to a Windows client which has an Automatic Disaster Recovery component installed.
2. Merge the files.
3. Rename the merged P1S files back to their original names.
4. Copy the merged P1S files back to the UNIX Cell Manager.

Restoring Original Hard Disk Signatures on Windows Systems

The Microsoft Cluster Server (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 (that is, if all nodes in the cluster have experienced a disaster), 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 and 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, then 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.

Restoring original hard disk signatures on Windows

On Windows systems, 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.

For example, to restore a MSCS Database from `C:\temp\ClusterDatabase`, type the following at the command prompt:

```
clubar r C:\temp\ClusterDatabase force q:.
```

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` utility included in the Windows Resource Kit. For details on using `dumpcfg`, run `dumpcfg /?` or see the Windows Resource Kit documentation. For more information on problems with hard disk signatures on Windows systems, see MSDN article Q280425.

Obtaining original hard disk signatures

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.

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.

Example of Hard Disk Signatures in the SRD File

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. The following is an example of a hard disk signature in the SRD file:

```
-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 letter `C`) and a quorum disk (with drive letter `Q`).

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 recovery 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 into the IDB. 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 volumes that remain to be restored for enabling the medium or media to be imported in the IDB. For more information on recycling media, see the HP Data Protector Help index: “recycling media”.

   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 executing commands:

   a. `omnisv -stop`

   b. `del Data_Protector_program_data\tmp\*.*`

   c. `omnisv -start`

2. Using the Data Protector GUI, export the medium or media with the backup of the volumes that remain to be restored. For more information on exporting media, see the HP Data Protector Help index: “exporting, media”.

3. Using the Data Protector GUI, import the medium or media with the backup of the partitions that remain to be restored. For more information on importing media, see the HP Data Protector Help index: “importing, media”.

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:

- A disaster recovery CD or an USB drive containing the DR OS image or a network bootable image for the Cell Manager should be prepared in advance.

  > **Important:** Perform a new backup and prepare a new DR OS image 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 safe locations as a part of the disaster recovery preparation policy, because the SRD file is the only

Data Protector file where information about objects and media is stored when the 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” (page 27).

- If your backups are encrypted, you must save the encryption key to a removable medium before a disaster occurs. If the encryption key is saved only on the Cell Manager, it is not accessible if the Cell Manager fails. Without the encryption key, disaster recovery is not possible. See “Preparation” (page 27).

Important: HP recommends to restrict access to backup media, recovery set files, SRD files, removable media with encryption keys, disaster recovery CDs, and USB drives storing DR OS data.

Restoring Internet Information Server Specifics

Internet Information Server (IIS) is not supported for disaster recovery. To recover the IIS, the following requirements must be met (in addition to the requirements required for Assisted Manual Disaster Recovery):

Requirements

- Do not install the IIS during the clean installation of the system.

Perform the following steps (in addition to the steps required for Assisted Manual disaster recovery):

Steps

1. Stop or uninstall the IIS Admin Service, if it is running.

2. Run the `drstart` command.

The IIS Database is restored as a plain file (with the filename `DisasterRecovery`) into the default IIS location (*`%SystemRoot%`*`\system32\inetsrv`).

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.

Editing the kb.cfg File

The `kb.cfg` file is located in the *`Data_Protector_home`*`\bin\drim\config` directory and stores information on the location of driver files from the *`%SystemRoot%`* directory. 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 `kb.cfg` file already contains all files necessary for industry standard hardware configurations.

For example, functionality of some drivers is split into several separate files, all required for the driver to function properly. Sometimes, Data Protector cannot identify all driver files, if they are not listed in the `kb.cfg` file on a case-by-case basis. In this case, they will not be included in the DR OS. Create and execute a test plan using the default version of the `kb.cfg` file. If the DR OS does not boot normally or cannot access network, then you may need to modify the file.

If you want to back up these drivers, add information about dependent files to the `kb.cfg` file in the appropriate format as described in the instructions at the beginning of the `kb.cfg` file. The easiest way to edit the file is to copy and paste an existing line and 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 span several lines. You can also add comment lines that start with a “#” (pound) sign.

After you finished editing the `kb.cfg` file, save it to the original location. Then perform another full client backup to include the added files in the recovery set.

Important: 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 HP.

Caution: It is recommended to create and execute a test plan to be sure the disaster recovery will work after you have edited the `kb.cfg` file.

Editing the SRD Files

The information about backup devices or media stored in the updated SRD file (`recovery.srd`) 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. However, 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 stored in the SRD file will be wrong and the recovery will fail. In this case, edit the updated SRD file before performing the Phase 2 of the disaster recovery to update the wrong information and thus enable a successful recovery.

To edit the SRD file, open it (for the location of the SRD file see specifics for particular method below) 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 target system has changed, replace the value of the `-host` option. You can also edit the information regarding the:

- Cell Manager client name (`-cm`),
- Media Agent client (`-mahost`),
- device name (`-dev`),
- device type (`-type`),
- address (`-devaddr`),

- policy (`-devpolicy`),
- robotics SCSI address (`-devioctl`)
- library slot (`-physloc`), and so on.

After you have edited the file, save it in Unicode (UTF-16) format to the original location.

The procedure on using the edited SRD file for disaster recovery differs between some disaster recovery methods and operating systems. Specific details for particular disaster recovery methods are explained below.

Important: You should restrict access to the SRD files due to security reasons.

AMDR

EADR/OBDR

AMDR

Perform the following steps before performing the regular AMDR recovery procedure, if the information in the SRD file is out-of-date.

Steps

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 (UTF-16) format.

EADR/OBDR

If the information in the SRD file is out-of-date, perform the following additional steps before proceeding with the regular EADR/OBDR procedure.

Steps

Windows systems

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 the Phase 1 of disaster recovery. Phase 2 of disaster recovery will not start automatically if the **Install only** option is selected.

Figure 5: The Install Only option in the Disaster Recovery Wizard

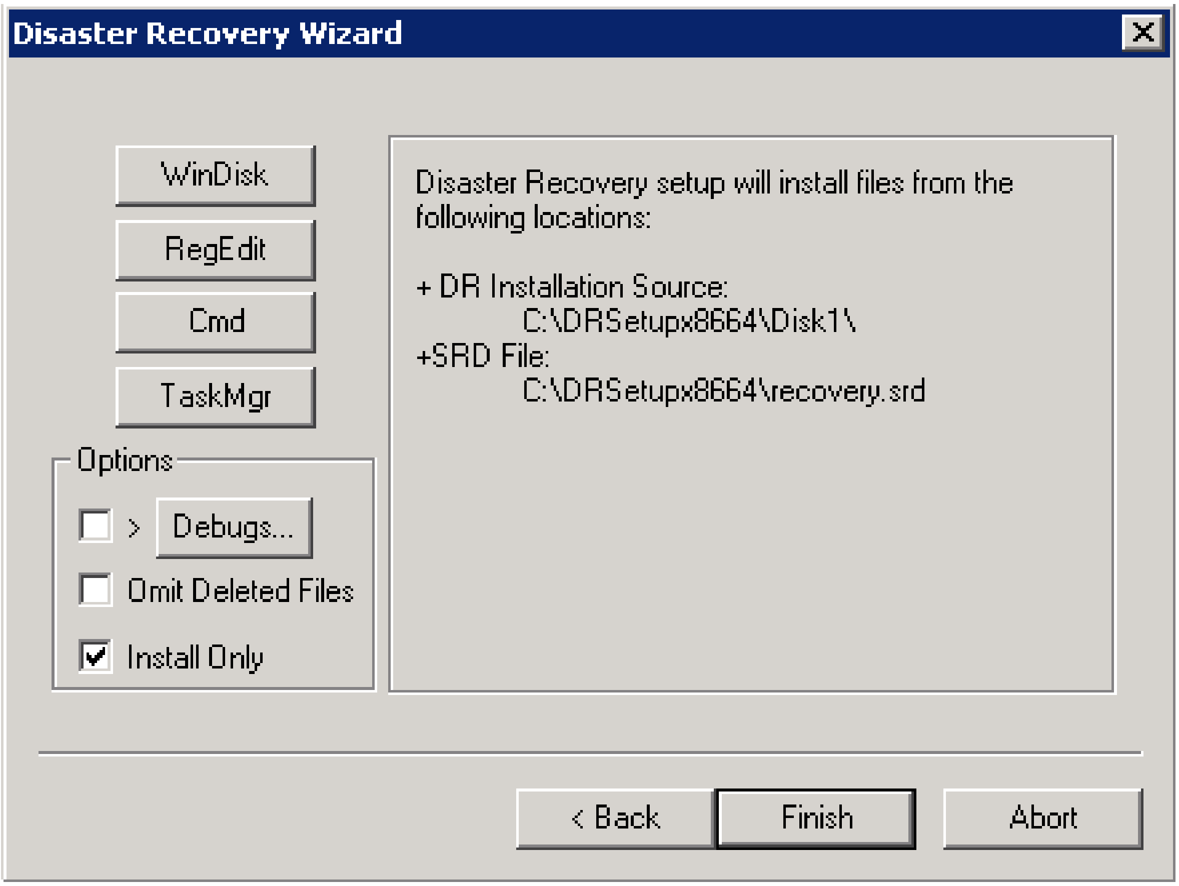

☐ > Debugs...
☐ Omit Deleted Files
☑ Install Only

Disaster Recovery setup will install files from the following locations:

+ DR Installation Source:
C:\DRSetupx8664\Disk1\
+SRD File:
C:\DRSetupx8664\recovery.srd

2. Select **Omit Deleted Files** option. This option enables removal of deleted files between successive incremental backups at restore time. If specified the `omnidr` binary will forward the same option to Data Protector restore tools (`omnir` and `omniofflr`) in case of incremental backup. The option has no effect on the restore of full backup object versions. However, selecting this option can significantly prolong the time of restore.

3. Run **Windows Task Manager** (press **Ctrl+Alt+Del** and select **Task Manager**).

4. In the Windows Task Manager, click **File** and then **New Task (Run...)**.

5. Run the following command from the Run dialog: `notepad C:\DRSYS\System32\OB2DR\bin\recovery.srd` and press **Enter**. The SRD file will be opened in the Notepad.

6. Edit the SRD file.

7. After you have edited and saved the SRD file to the original location, run the following command from `C:\DRSYS\System32\OB2DR\bin`

   ```
   omnidr -drimini C:\$DRIM$.OB2\OBRecovery.ini
   ```

8. Proceed with the next step in the regular EADR/OBDR recovery procedure.

Linux systems

1. When the Disaster Recovery Wizard appears, press **q** to stop the wizard during the countdown and select the **Install Only** option. This option will install only a minimal version of Data Protector to the target system. Phase 2 of disaster recovery will not start automatically if the Install Only option is selected.

2. Switch to another shell.

   Edit the SRD file `/opt/omni/bin/recovery.srd`. For details, see the *HP Data Protector Disaster Recovery Guide*.

3. After you have edited and saved the SRD file, execute:

   ```
   omnidr -srd recovery.srd -drimini /opt/omni/bin/drim/drecovery.ini
   ```

4. Once the recovery finishes, return to the previous shell and proceed with the next step in the ordinary EADR/OBDR recovery procedure.

Example of Editing the SRD File

If the information in the SRD file is not up to date anymore (for example, you changed a backup device), modify the updated SRD file `(recovery.srd)` before performing Phase 2 of disaster recovery to update the wrong information and thus enable a successful recovery.

You can display some of the device configuration information using the `devbra -dev` command.

Changing the MA client

You performed a backup for disaster recovery purposes 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 also the value of the `-devaddr` option in the updated SRD file accordingly.

After you have edited the file, save it in Unicode (UTF-16) format to the original location.

Changing the backup device

To perform disaster recovery using another device than the one which was used for the backup, modify the following option values in the updated SRD file:

`-dev`, `-devaddr`, `-devtype`, `-devpolicy`, `-devioctl`, and `-physloc`

Where:

`-dev`	specifies the logical name of the backup device or drive (library) to be used for the backup,
`-devaddr`	specifies its SCSI address,
`-devtype`	specifies the Data Protector device type,
`-devpolicy`	specifies the device policy, which 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),
`-devioctl`	specifies the robotics SCSI address.
`-physloc`	specifies the library slot
`-storname`	specifies the logical library name

For example, you performed a backup for disaster recovery purposes using an HP Ultrium standalone device with the device name `Ultrium_dagnja`, connected to the MA host `dagnja` (Windows systems). However, for the disaster recovery you would like to use an HP Ultrium robotics library with the logical library name `Autoldr_kerala` with drive `Ultrium_kerala` connected to the MA client `kerala` (Linux systems).

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" -storname "AutoLdr_kerala" -mahost kerala.company.com.
```

After you have edited the file, save it in Unicode (UTF-16) format to the original location.

Windows BitLocker Drive Encryption

During the disaster recovery process on Windows Vista and later releases, you can unlock volumes that are encrypted using BitLocker Drive Encryption.

Limitation

If you do not unlock a specific volume or if the volume is damaged, cannot be unlocked, and must therefore be formatted, the volume is no longer encrypted after disaster recovery. In such circumstances, you need to encrypt the volume again.

Note that the system volume is always restored unencrypted.

Steps

1. When the disaster recovery module detects an encrypted volume, you are prompted to unlock it.

   Click **Yes** to start the Unlocker wizard. Note that if you click **No**, the encrypted volumes will remain locked.

2. In the Select Locked Volumes page, the detected encrypted volumes are listed. Select the volumes you want to unlock and then click **Next**.

3. In the Unlock Volume pages (one page for each selected volume), you are requested to specify the unlock method. The following unlock methods are available:

   - Password *(available on Windows 7 and later releases)*

     A string of characters that was used when you encrypted the volume.

   - Passphrase

     A string of characters longer than the usual password that you used when you encrypted the volume.

   - Recovery key

     A special hidden key you created on each volume that you encrypted. The recovery key has a `BEK` extension, it is saved in the recovery key text file. You can click **Browse** to locate the recovery key file.

   Type the requested information in the text box and then click **Next**.

4. Check whether the volumes were unlocked successfully and then click **Finish**.

   > **Note:** If the unlocking process failed, you can review the error information and retry or skip the unlocking procedure.

Recovery to Dissimilar Hardware

> **Note:** Recovery to dissimilar hardware is an extension of . You should refer to that as well as the information here.

After hardware failure or a similar disaster, you may need to restore a backup to a system where some or all of the hardware is different from the original (**dissimilar hardware**).

Dissimilar hardware restore adds the following to the standard EADR and OBDR procedures:

1. At backup time, the disaster recovery module also collects network configuration information and hardware information.

2. It enables the injection of critical device drivers into the DR OS image, so that they are available during restore. You can also inject missing drivers manually at restore time if some are missing.

3. During restore, the network and hardware information is used to configure and map the network properly for the restored OS, and to detect critical hardware that is missing.

When dissimilar hardware restore might be needed

- **Hardware failure**

  Dissimilar hardware restore is needed when some of the boot-critical hardware (such as the storage controller, processor, or motherboard) fails and must be replaced with non-identical hardware.

- **Disaster**

  Dissimilar hardware restore is needed after total machine disaster where:

  - No matching machine can be found (because of limited budget, the failing machine’s age or other causes).

  - Down-time cannot be afforded; the system must be up and running immediately.

  In these situations, the use of dissimilar hardware restore can mean lower budget cost since exact clones of the original systems are not needed.

- **Migration**

  Dissimilar hardware restore is needed in the following situations:

  - Moving to another machine (for example, to faster or newer hardware) where OS reinstallation and reconfiguration is not an option.

  - Moving from a physical system to a virtual environment or the other way round.

    From the disaster recovery module's point of view, a virtual environment is another hardware platform for which you need to provide critical drivers in order to restore a system backup taken on some other virtual or physical platform. Limitations and requirements listed below also apply to virtual environments.

Overview

The phases of restore to dissimilar hardware are the standard disaster recovery phases *with the following differences*:

- **Phase 0**: Additional information is collected about the network configuration and the hardware.

- **Phase 1**: The machine is brought into a state where disaster recovery executables have access to disks, file systems, the network and WIN32 API. Restore critical devices are checked. If any drivers are missing, you are prompted to provide them.

- **Phase 2**, restoring the OS, is the same, but an extra sub-phase occurs after it:
    - **Phase 2a**: The restored operating system is prepared and adapted to the hardware, through injecting critical drivers, updating the registry and mapping the network.
- **Phase 3** is the same, in which data not restored in Phase 2 is restored.

Requirements

- You must provide at least all boot-critical drivers (including network drivers) for the target machine. These drivers can be added directly to the image at image creation time (recommended) or loaded at restore time (during Phase 1). In addition, drivers of locally attached backup devices (such as a tape device) must also be available if a local restore is attempted.

  For more information see "Drivers" on the next page.
- Automatic network configuration restore for the restored OS requires network drivers to be present at restore time.
- The restore system must have at least the same number of disks (with the same or greater size) as the backup system did.
- The original OS should be supported on the target machine (server or workstation) by the hardware manufacturer.
- It is advisable for the system firmware of the target machine to be up-to-date before a dissimilar hardware restore.
- If you need to disable dissimilar hardware support during backup, edit the file `drm.cfg` on the system you want to back up and set the `enable_disshw` option to `0`.
- The system must include at least one NTFS volume, which serves as a storage point for VSS purposes during the backup phase.

Limitations

The disaster recovery module only supports dissimilar hardware restores if the backup was performed with the **Use Shadow Copy** option (selected by default for supported platforms).

- Dissimilar hardware support is provided only for EADR and OBDR on the following operating system releases:
    - Windows Vista
    - Windows 7
    - Windows Server 2008
    - Windows Server 2008 R2

- Windows 8
- Windows 8.1
- Windows Server 2012
- Windows Server 2012 R2

For details, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.

- The following cross-platform restore combinations are supported:

From	To
64-bit (x64) operating system	64-bit (x64) hardware architecture
32-bit operating system	32-bit or 64-bit (x64) hardware architecture

Dissimilar hardware restore of upgraded operating systems is only supported with the “Generic” recovery mode option (see "Recovery procedure" on the next page).

- Network card teaming configurations are not supported. If you need them, you must reconfigure them after the OS is restored. The disaster recovery module only restores physical network card configurations.
- The disaster recovery module can only inject drivers for which an INF file is provided. Drivers that have their own installation procedures (such as graphics drivers) are not supported and cannot be injected during Phase 1 or Phase 2a. However, for boot-critical device drivers, manufacturers typically provide INF files.
- The target machine’s disks should be kept attached to the same type of host adapter buses (such as SCSI or SAS), otherwise the restore may fail.
- When restoring Domain Controllers, using the “Unattend” mode, you must login manually in order to complete sysprep cleanup. Once the cleanup is completed the OS will reboot automatically and the system will be ready for usage.

Recommendations

The system firmware of the target machine should be up-to-date before a dissimilar hardware restore is attempted.

Drivers

Note: The DR OS image includes a large database of generic critical drivers, especially for storage controllers. If you cannot find original drivers to inject, there is a good chance that generic ones already exist in the DR OS image.

To enable restore onto dissimilar hardware, drivers vital for the restore and boot of the new system must be available. You will need to provide the following drivers:

- For all storage controllers of the target system. This will enable the detection of the underlying storage at restore/boot time.
- Network card drivers to enable network restore and access to existing driver store locations, along with drivers for locally attached backup devices (such as tape drives) if a local restore is attempted.

Drivers for the original hardware can be included in the DR OS image during backup in the preparation phase (Phase 0), and you can add drivers for new hardware during the creation of the image. You also have the option of adding them manually during the restore process.

Although the disaster recovery module searches only for boot-critical drivers during the restore process, you can add additional non-boot-critical drivers in the DR OS image, which you can then inject during the restore using the “Load Drivers” Tasks menu option.

When the operating system has been booted you should install other missing hardware drivers.

You can inject drivers from a CD-ROM, DVD-ROM, or USB drive, a network share, or a local folder.

Preparation

Note: You need to perform this preparation after each hardware configuration change to the system.

Preparation is the same as for EADR (see) and OBDR (see) with the following changes:

- The disaster recovery module also collects network configuration and hardware information.
- Critical device drivers (such as for storage, network or tape) should be present, so the disaster recovery module can inject the drivers into the DR OS image at the image creation time. See "Drivers" on the previous page.

Recovery procedure

If you enabled the dissimilar hardware restore in the Recovery Options page of the HP Data Protector Disaster Recovery GUI, the system is scanned for the missing drivers during the recovery process. If any critical driver (such as storage, tape, network drivers, or disk controllers) is missing, you are prompted to load the missing driver.

Steps

1. When you are prompted to load missing drivers during the disaster recovery procedure, click **Yes** to start the Dissimilar Hardware wizard. If you click **No**, the driver injection procedure is skipped.
2. In the Select Devices page, select the devices, for which you want to load drivers. Click **Next**.
3. In the Driver Search Locations page, specify the locations on the running system where you keep

your drivers. Browse for the device driver or type the location in the Driver path text box and then click **Add path** to add the specified path to the list. You can use the **Search tree depth** option to adjust the search to your system specifics.

> **Note:** You can remove the specified location from the search list by right-clicking this location and then selecting **Remove**.

The specified locations will be searched for the missing drivers. Click **Next**.

4. After the specified locations are searched for the missing drivers, the following results are possible:

   - The device driver is found: the full path to the corresponding driver information file (`*.inf`) is specified in the Driver path text box. Verify, if this driver is appropriate and click **Next** to load it.

   - The device driver is not found: the Driver path text box is empty. Do one of the following:

     If you want to search for a different driver, click **Browse**. In the Browse file dialog, select the device driver path and then click **Next**.

     If you do not want to load a driver to this device, you can leave the Driver path text box empty and click **Next** to proceed to the next page, or, you can click **Skip** to exit the wizard.

     > **Note:** If you specify a driver that does not correspond to the device, this driver is indicated as being invalid and you are not able to load it. If the driver is not appropriate, you can change it or skip loading.

5. In the Driver Installation Progress page, you can view whether the device drivers were loaded successfully. If any errors are reported, you can retry to load the drivers by clicking **Retry**. Click **Finish**.

Restoring and preparing the OS

The process of restoring the OS is the same as in the standard EADR (from Step 5) and OBDR (from Step 6) processes. After it, the recovery process prepares and adapts the restored OS to the dissimilar hardware to prepare the OS for the restore of applications and files. This includes injecting boot-critical drivers, updating the registry of the restored OS and mapping the network.

Since all boot-critical drivers should exist (loaded into the running DR OS image during Phase 0 or added manually during the restore of the OS), injecting them occurs automatically. However your intervention may be required to correct the network mappings.

Correcting network mappings

After you finished with restoring to dissimilar hardware, disaster recovery module checks, whether the network adapters on the system you are recovering are different from the network adapters of the original system. The disaster recovery module cannot always map the network configuration of the original system to the network configuration of the target system on its own. This happens, for example, when the target system has one network card and the original system has two or more network cards, or when you add additional network adapters to the target system. When such difference is detected or if the correct network mappings cannot be determined automatically, you have an option to map the original network adapters to the network adapters discovered on the target system.

> **Note:** The network mapping occurs only for available network adapters. Network adapters without drivers cannot be mapped. Because of this, you should load network card drivers before the restore process begins.

Steps

1. In the Network Adapter Mapping page, select the network adapters of the original system in the Original network adapters drop-down list. In the current network adapters drop-down list, select one of the network adapters available on the target system. Click **Add mapping**. The mapping you created is added to the list.

   > **Note:** You can remove a mapping from the list by right-clicking the mapping and then selecting **Remove**.

2. When you mapped all the network adapters you wanted, click **Finish**.

After the OS is successfully restored

Dissimilar hardware restore resets the OS activation. Once the OS is successfully restored, you should:

- Re-activate the OS.
- Check and, if needed, reinstall missing system drivers.

Restoring user and application data

This phase is the same as for EADR. See "Enhanced Automated Disaster Recovery (EADR)" on page 37.

> **Note:** Third-party application services and drivers may fail to load once the OS is booted. These applications will probably need to be reinstalled, reconfigured or removed from the current system if they are not needed.

Recovery of a physical system to a virtual machine (P2V)

Data Protector supports recovery to virtualization environments which provide support for the original operating system, such as VMware vSphere, Microsoft Hyper-V, or Citrix XenServer.

Prerequisites

The target virtual machine must meet the following requirements:

- The guest operating system must be of the same type as the original one (Windows, Linux).
- The virtual machine must have the same or larger number of disks than the original system.
- The disks must have the same or lager size as their original counterparts.
- The disk order must be the same as on the original system.
- The amount of memory assigned to a virtual machine may have an impact on the recovery process, therefore it is recommended to allocate at least 1 GB of memory to the virtual machine.
- The virtual video card memory size must meet the requirement of the original system based on the display resolution of the original system. If possible, use automatic settings.
- Add the same number of network adapters as on the original machine. The adapters must be connected to the same network as the original ones.

Procedure

Boot the virtual machine using the DR OS image and follow the standard disaster recovery procedure to dissimilar hardware.

Recovery of a virtual machine to a physical system (V2P)

Disaster recovery of a virtual machine to a physical system is performed using the standard disaster recovery to dissimilar hardware.

Chapter 4: Disaster recovery on UNIX systems

Manual Disaster Recovery (MDR)

Manual Disaster Recovery is a basic recovery method. This method involves recovering the system by reinstalling it in the same way that it was initially installed. Data Protector is used to restore all files, including the operating system.

MDR of an HP-UX client 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 must be used to restore the user and application data in order to complete Phase 3 of disaster recovery.

> **Note:** This section does not cover the full functionality of Ignite-UX. For detailed information, see the *Ignite-UX administration guide*.

Overview

Ignite-UX offers 2 different approaches to preparing a system for and recovering a system from a disaster:

- Using a custom installation medium (Golden Image)
- Using system recovery tools (make_tape_recovery, make_net_recovery)

While the usage of a custom installation medium is most suitable for IT environments with a large number of basically identical hardware configurations and OS releases, the usage of system recovery tools supports the creation of recovery archives, which are customized for individual systems.

Both methods allow the creation of bootable installation media like DDS-Tapes or CDs. 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 must be located on an NFS share on the network.

Use Ignite-UX GUI where it is supported.

Preparation for Manual Disaster Recovery (HP-UX Cell Manager)

To prepare for a successful disaster recovery, you should follow the instructions related to the general preparation procedure, together with the specific method requirements. You have to prepare in advance in order to perform a disaster recovery fast and efficiently.

Preparation for a Manual Disaster Recovery of the Cell Manager includes:

- Gathering information for your backup specification
- Preparing your backup specification (using a pre-exec script)
- Executing a backup
- Executing Internal Database backup sessions regularly

All of these preparatory steps are necessary before executing disaster recovery on the Cell Manager.

One-time preparation

You should document the location of these files in the disaster recovery plan so that you can find the information when disaster strikes. Also you should consider version administration (there is a collection of the “auxiliary information” per backup).

If the system to be backed up has application processes active at low run-levels, you should establish a state of `minimal activity` (modified `init 1 run-level`) to prepare the Cell Manager for a consistent backup.

HP-UX systems

- 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 which would otherwise be suspended by moving to run-level 1, and they are needed for the backup.
- Ensure that `rpcd` is configured on the system (configure the option `RPCD=1` within the `/etc/rc.config.d/dce` file).

This prepares the system so that it enters a state of minimal activity that can be characterized as follows:

- Init-1 (`FS_mounted, hostname_set, date_set, syncer_running`)
- Running processes: `network`, `inetd`, `rpcd`, `swagentd`

Backing up the system

After you have prepared the backup specification, you should execute the backup procedure. Repeat it on a regular basis, or at least after every major system configuration change, especially after any change in the physical or logical volume structure. Pay special attention to the IDB and filesystem backup:

- Back up the IDB regularly, ideally in a separate backup specification, and scheduled after the backup of the Cell Manager itself.
- Run the IDB and filesystem backup on a specific device attached to the Cell Manager system so you know that the medium in the device contains the most recent backup version of the IDB.

Installing and Configuring HP-UX Systems Manually (Cell Manager)

After a disaster happens, you should first install and configure the operating system (Phase 1). Then you can recover the Cell Manager.

Steps

Phase 1

1. Replace the affected disk.

2. Boot your system from the operating system installation medium.

3. Reinstall the operating system. During the installation, use the data gathered during the preparation phase (using a pre-exec script) to re-create and configure the physical and logical storage/volume structure, filesystem, mount points, network settings, and so on.

Restoring System Data Manually (HP-UX Cell Manager)

After you have installed and configured the operating system (Phase 1), you can use Data Protector to recover the Cell Manager.

Prerequisites

- You need media containing the latest backup image of the root volume of the Cell Manager system and a newer latest backup image of the IDB.

- You need a device connected to the Cell Manager system.

Steps

Phase 2

1. Reinstall the Data Protector software on the Cell Manager.

2. Restore the IDB and the `/etc/opt/omni` directory from their respective latest backup images to a temporary directory. This simplifies the restore of all other files from backup media. Remove the `/etc/opt/omni/` directory and replace it with the `/etc/opt/omni` directory from the temporary directory. This re-creates the previous configuration.

3. Start Data Protector processes with the `omnisv -start` command.

Phase 3

4. Start the Data Protector GUI and restore the needed files from your backup images.

5. Restart the system.

Your Cell Manager should now be successfully recovered.

Preparation for Manual Disaster Recovery (HP-UX Client)

Ignite-UX offers 2 different approaches to preparing a system for and recovering a system from a disaster:

Using custom installation Medium (Golden Image)

Using system recovery tools (make_tape_recovery, make_net_recovery)

Using custom installation medium (Golden Image)

Large IT environments often consist of a large number of systems that are based on identical hardware and software. The installation time for the OS, applications and required patches for a new system can be significantly reduced if a complete snapshot of an installed system is used to install other systems. Ignite-UX includes a feature that allows you to modify parameters like networking or filesystem settings, as well as add software like Data Protector to the image (with the 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.

The general steps using a custom installation medium are:

1. **Phase 0**

   a. Create a Golden Image of a client system.

2. **Phase 1** and **2**

   a. Replace the faulty disk with a replacement disk.

   b. Boot the HP-UX client from the Ignite-UX server and configure the network.

   c. Install the Golden Image from the Ignite-UX server.

3. **Phase 3**

   a. Use the standard Data Protector restore procedure to restore user and application data.

Creating a Golden Image

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 password-free 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.11.31/archive_HPUX11_31_DP70_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_31_DP70_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 newly-created depot, add a `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_31_DP70_Client" {

description "HPUX 11.i OS incl Patches and DP70 Client"

"/opt/ignite/data/Rel_B.11.31/config"

"/var/opt/ignite/data/Rel_B.11.31/archive_HPUX11_31_DP70_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 will 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.

You need to repeat these steps to create a Golden Image for all systems with different hardware and software configurations.

Ignite-UX enables you to create a bootable tape or CD based on the created Golden Image. See the *Ignite-UX Administration Guide* for more information.

Using system recovery tools (make_tape_recovery, make_net_recovery)

The usage of the system recovery tools bundled with the Ignite-UX enables 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 starting up 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 starting up either from a bootable tape created by the Ignite-UX `make_boot_tape` command or the system boots directly from the Ignite-UX server. Starting up directly from the Ignite-UX server can be automated with the Ignite-UX bootsys command or interactively specified on the boot console.

The general steps using system recovery tools are:

1. **Phase 0**

   a. Create a recovery archive of an HP-UX client using the Ignite-UX GUI on the Ignite-UX server.

2. **Phase 1** and **2**

   a. Replace the faulty disk with a replacement disk.

   b. For local restore, boot from the prepared recovery tape.

   c. In case of a local restore, the recovery process starts automatically.

      For network restore, boot from the Ignite-UX client and configure the network and UI.

      In case of a network restore, install the Golden Image from the Ignite-UX server.

3. **Phase 3**

   a. Use the standard Data Protector restore procedure to restore user and application data.

Prerequisites

Before you can to prepare your system for disaster, the Ignite-UX fileset must 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

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 tapes will work 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 that you want to include in 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 a bootable archive tape out of a compressed archive file. See the chapter `Create a Bootable Archive Tape via the Network` in the `Ignite-UX Administration Guide`.

Recovering an HP-UX Client

There are 3 different methods to recover HP-UX clients using Manual Disaster Recovery (MDR):

Recovery using a Golden Image

Recovery from the bootable backup tape

Recovery from the network

Recovery using a Golden Image

You can recover an HP-UX client by applying the Golden Image, which is located on an NFS share on your network.

On the client

Steps

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 GUI 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

Steps

1. Right-click the client icon in the Ignite-UX GUI and select **Install Client - New Install**.
2. Select the Golden Image you want to install, check the settings (network, filesystem, time zone, ...) and click **Go!**.
3. You can check the installation progress by right-clicking the client icon and choosing **Client Status**.
4. After the installation has finished, restore additional user and application data using the standard Data Protector restore procedure.

Recovery from the bootable backup tape

A bootable backup tape is created using the make_tape_recovery command.

Steps

1. Replace the faulty hardware.
2. Make sure that the tape device is locally connected to the affected 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` *`Pnumber`*.
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

You can boot the target system over the network from the recovery archive file located on the Ignite-UX server. Follow the instructions on how to perform a recovery using a Golden Image and make sure you have selected the desired archive for the installation.

Disk Delivery Disaster Recovery (DDDR)

There are two possible methods for Disk Delivery Disaster Recovery. You can use a working Data Protector client system and create the new disk while connected to this client. Alternatively, you can use an auxiliary disk without an additional working client. You need to collect enough data before the disaster to be able to correctly format and partition the disk.

Overview

Disk Delivery of a UNIX client is performed using an auxiliary disk (which can be carried around), with a minimal operating system with networking and a Data Protector agent installed on it.

The general steps using an auxiliary disk for a UNIX client are:

1. **Phase 0**
    a. Perform a full filesystem backup of the entire system (client backup).
    b. Create an auxiliary disk.
2. **Phase 1**
    a. Replace the faulty disk with a replacement disk, connect the auxiliary disk to the target system and restart the system with the minimal operating system installed on the auxiliary disk.
    b. Manually re-partition the replacement disk and re-establish the storage structure and make the replacement disk bootable.
3. **Phase 2**
    a. Use the standard Data Protector restore procedure to restore the boot disk of the original system onto the replacement disk (use the **Restore into** option).
    b. Shut down the system and remove the auxiliary disk. You do not need to shut down the system if you are using a hot-swappable hard disk drive.
    c. Restart the system.
4. **Phase 3**
    a. Use the standard Data Protector restore procedure to restore user and application data.

Limitations

- An auxiliary disk should be prepared on a system of the same hardware class as the target system.
- The cluster environment recovery may differ from the standard procedure. Depending on the configuration of the cluster environment, additional steps and modification to the environment may be necessary.
- RAID is not supported.

Preparation for Disk Delivery Disaster Recovery of UNIX Clients

To prepare for a successful disaster recovery, you should follow the instructions related to the general preparation procedure, together with the specific method requirements. You have to prepare in advance in order to perform a disaster recovery fast and efficiently. For details on supported operating systems, see the *HP Data Protector Product Announcements, Software Notes, and References*.

Preparation for a Disk Delivery Disaster Recovery includes:

- gathering information for your backup specification
- preparing an auxiliary disk
- preparing your backup specification (using a pre-exec script)
- executing the backup

All of these preparatory procedures are necessary before executing a disaster recovery on the client system.

One-time preparation

If the information is collected as part of a pre-exec command, you should document the location of these files in the disaster recovery plan so that you can find the information when disaster strikes. Also, you should consider version administration (there is a collection of the “auxiliary information” per backup).

You should also establish a state of `minimal activity` (modified `init 1 run-level`) on each client system to prepare it for a consistent backup and thus avoid problems after recovery. Consult your operating system documentation for details.

HP-UX Example

- 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 which would otherwise be suspended by moving to run-level 1, and they are needed for the backup.
- Ensure that `rpcd` is configured on the system (configure the option `RPCD=1` within the `/etc/rc.config.d/dce` file).

This prepares the system so that it enters a state of minimal activity that can be characterized as follows:

- Init-1 (`FS_mounted, hostname_set, date_set, syncer_running`)
- Network must be running
- Running processes: `network`, `inetd`, `rpcd`, `swagentd`

Solaris Example

- Move some kill links from `/etc/rc1.d to /etc/rc0.d` and complement the changes for the boot-up section. The kill links include the basic services which would otherwise be suspended by moving to run-level 1, and they are needed for the backup.

- Ensure that `rpcbind` is configured on the system.

This prepares the system so that it enters a state of minimal activity that can be characterized as follows:

- Init-1

- Network must be running

- Running processes: `network, inetd, rpcbind`

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.

Preparing the auxiliary disk

If you want to work with an auxiliary disk, you need to prepare it first. Only one bootable auxiliary disk is required per cell and platform. This disk has to contain the operating system and network configuration, and has to be bootable.

Backing up the system

After you have prepared the backup specification, you should execute the backup procedure. Repeat it on a regular basis, or at least after every major system configuration change, especially after any change in the physical or logical volume structure.

Creating the Backup Specification for Disaster Recovery of a UNIX Client

To configure a backup specification for Disaster Recovery of a UNIX client, either modify an existing specification or create a new one with specified pre- and post-exec scripts. For details on supported operating systems, see the *HP Data Protector Product Announcements, Software Notes, and References*.

Steps

1. Provide a Pre-exec script that will perform the following:
   - Collect all necessary information about the environment and store it where it is available in case a disaster recovery is needed. The information includes:
     - The physical and logical storage structure of the system
     - The current logical volume structure (for example, on HP-UX systems, using `vgcfgbackup` and `vgdisplay -v`)
     - Cluster configuration data, disk-mirroring, and striping
     - A filesystem and mountpoint overview (for example, on HP-UX systems, using `bdf` or copy of `/etc/fstab`)
     - System paging space information (for example, on HP-UX systems, the output of the `swapinfo` command)
     - An I/O-structure overview (for example, on HP-UX systems, using `ioscan -fun` and `ioscan -fkn` on HP-UX systems)
     - Client network settings

     An emergency copy of the data can also be put into the backup itself. If so, extract the information prior to the actual recovery.
   - Log out all users from the system.
   - Shut down all applications, unless the application data gets backed up separately, for example, using online database backup.
   - Optionally, restrict network access to the system, so that nobody can log on to the system while the backup is running (for example, on HP-UX systems, overwrite `inetd.sec` and use `inetd -c`).
   - If needed, enter a state of minimal system activity (for example, on HP-UX systems, use `sbin/init 1`; `wait 60`; check if `run-level 1` is reached). Note that this is a modified "init 1" state.
2. Provide a post-exec script that will restore the system to the standard run-level, restart applications, and so on.
3. Configure a backup specification for the client on the Data Protector Cell Manager using pre- and post-exec scripts. It should include all the disks.
4. 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).

Installing and Configuring a UNIX Client Using DDDR

After a disaster occurs, you should first install and configure a new disk for the faulty client (Phase 1).

Prerequisites

- You need a new hard disk to replace your affected disk.
- An auxiliary disk should be prepared on a system of the same hardware class as the target system.
- An auxiliary disk should contain the relevant UNIX operating system and the Data Protector agents.
- You should have a valid full backup of the client that you want to recover.

Steps

1. Replace the faulty disk with a new disk of comparable size.
2. Attach the auxiliary disk (which contains the required 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 systems). Use the backed-up data for the non-root volume groups (for example, with `vgcfgrestore` or SAM on HP-UX systems).
5. Additionally, create the root volume group to be restored on the repaired disk (for example, using `vgimport` on HP-UX systems). It will not look like a root volume group during the restore process, because the operating system from the auxiliary disk will be running.
6. Make the new disk bootable using the relevant UNIX commands.
7. Reconstruct any other storage structures like mirror, striping, HP Serviceguard, 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 (for example, `/etc_restore` for `/etc`, and so on).
9. Remove any files in the mountpoints to be restored; they must be empty.
10. Proceed with restoring the system data.

Restoring System Data Using DDDR (UNIX Client)

You can restore a system to the state when the last successful backup was performed. You should first install and configure the UNIX client (Phase 1). For details on supported operating systems, see the *HP Data Protector Product Announcements, Software Notes, and References*.

Prerequisites

- The relevant operating system should be installed and configured.
- Data Protector should be installed.
- You should have a valid full backup of the client that you want to recover.
- The media required for the restore should be available.

Steps

Phase 2

1. Start the Data Protector user interface and open a connection to the Data Protector Cell Manager.
2. Import the system with the auxiliary disk into the cell.
3. Select the backup version from which you want to restore.
4. 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.
5. Shut down and restart the system that was just restored.
6. Disconnect the auxiliary disk from the system.
7. Restart the system from the new (or repaired) disk.

Phase 3

8. Restore user and application data using the standard Data Protector restore procedure.

Enhanced Automated Disaster Recovery (EADR)

Data Protector offers an enhanced disaster recovery procedure for Linux Data Protector Cell Manager and clients. For details on supported operating systems, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.

EADR collects all relevant environment data automatically at backup time. During a full backup of the entire client system, data required for the temporary DR OS setup and configuration is packed in a single large recovery set file and stored on the backup tape (and optionally on the Cell Manager) for each backed up client in the cell.

In addition to this image file, a Phase 1 Startup file (P1S file), required for correct partitioning and formatting of the disk is stored on a backup medium and on the Cell Manager. When a disaster occurs, the Enhanced Automated Disaster Recovery Wizard is used to restore the recovery set from the backup medium (if it has not been saved on the Cell Manager during the full backup) and convert it into a disaster recovery CD ISO image. The CD ISO image can be recorded on a CD using any CD burning tool and used to boot the target system.

Once DR OS Image is booted, Data Protector automatically formats and partitions the disks, and finally recovers the original system with Data Protector as it was at the time of the backup.

Important: HP recommends to restrict access to backup media, recovery set files, SRD files, and disaster recovery CDs.

Overview

The general steps using the Enhanced Automated Disaster Recovery method for a Linux client are:

1. **Phase 0**

    a. Perform a full backup of the system (client backup) that includes at least all critical volumes. If you are preparing for disaster recovery of a Cell Manager, also perform an Internal Database backup afterwards as soon as possible.

    b. Use the Enhanced Automated Disaster Recovery Wizard to prepare a disaster recovery OS image (DR OS image) from the recovery set file of the affected system and record it on a CD. If the recovery set has not been saved on the Cell Manager during the full backup, the Enhanced Automated Disaster Recovery Wizard will restore it from the backup medium.

    Important: You need to perform a new backup and prepare a new DR OS image after each hardware, software, or configuration change. This also applies to any network changes, such as a change of IP address or DNS server.

    c. If the full client backup was encrypted, store the encryption key on a removable medium so that it is available for disaster recovery. You will need the key for a Cell Manager recovery or if the connection to the Cell Manager cannot be established.

2. **Phase 1**

    a. Replace the faulty hardware.

    b. Boot the target system from the disaster recovery CD or USB flash drive and select the scope of recovery. This is a completely unattended recovery.

3. **Phase 2**

    a. Depending on the recovery scope you select, the selected volumes are automatically restored. Critical volumes (the boot and root volumes and the volumes containing the Data Protector installation and configuration) are always restored.

4. **Phase 3**

    a. Use the standard Data Protector restore procedure to restore user and application data.

Important: Prepare a DR OS image in advance for any critical systems that must be restored first (especially DNS servers, Cell Managers, Media Agent clients, file servers, and so on).

Prepare removable media containing encryption keys in advance for Cell Manager recovery.

The following sections explain the limitations, preparation steps, and the recovery procedure that pertains to EADR of the Linux clients.

Requirements

- The Data Protector Automatic Disaster Recovery component must be installed on systems for which you want to enable recovery using this method and on systems where the DR OS image will be prepared. For details, see the *HP Data Protector Installation and Licensing Guide*.
- The hardware configuration of the target system must be the same as that of the original system. This includes the 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.
- During the EADR preparation, the volume on which Data Protector is installed should have at least 800 MB of temporary free space. This space is required to create a temporary image.
- The system’s BIOS must support bootable CD extensions, as defined in the El-Torito standard, and read/write access to hard disk drives using LBA addressing via INT13h function XXh. The BIOS options can either be checked in the user manuals of the system or by inspecting the system setup before the boot.

Limitations

- Enhanced Automated Disaster Recovery (EADR) and One Button Disaster Recovery (OBDR) are available on Linux systems only.
- You must create DR ISO images for Linux systems on Linux systems. You cannot create DR ISO images for on other systems (Windows systems, HP-UX systems, Solaris systems). The limitation does not apply for updating the SRD file or other tasks.
- If you have a mount point with the name `CONFIGURATION` and it contains the directory `SystemRecoveryData`, data in the directory `SystemRecoveryData` will not be backed up.
- Do not mount disks using the disk ID, because the ID is unique and depends on the disk serial number. In case of a disaster, the disk may be replaced and the new disk will have a new ID. As a result, the disaster recovery fails.
- A custom kernel installation or configuration is not supported, only the original kernels provided with the distributions are supported.

- When restoring a Linux client with SELINUX enforcing mode enabled, the system has to relabel all system files after recovery which, depending on system configuration, can take some time to complete. If permissive mode is used, the system log will contain a large number of SELINUX warning messages.

- When you create a backup specification with the CONFIGRATION/SYSTEMRECOVERYDATA object selected, the folders `/opt/omni/bin/drim/log` and `/opt/omni/bin/drim/tmp` are by default excluded from the backup. However, this exclusion is not set if you manually update existing backup specifications.

- Using resumed object backups for recovery is not supported since the consistency of such backups cannot be guaranteed.

- Fusion IO disks that do not automatically attach at MiniOS boot time need to be manually attached prior to recovery. This is required when replacing an old Fusion IO disk with a new one or when an internal Fusion IO disk error occurs. Those disks need to be formatted using specific tools before being attached in the MiniOS. To format and attach a Fusion IO disk manually to the system, you need to run the following commands in Linux shell present in MiniOS before the recovery starts:

  - `fio-status` – List the status of all the Fusion IO disks.

  - `fio-format [path]` – Perform low-level format of the Fusion IO disk.

  - `fio-attach [path]` – Attach the Fusion IO disk to the system.

- Sparse files are restored to their full size during offline restore. This may result in the target volume running out of space.

Disk and partition configuration

- A 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 EADR.

Preparation for Enhanced Automated Disaster Recovery

To prepare for a successful disaster recovery, follow the instructions related to the general preparation procedure for all disaster recovery methods before completing the steps listed in this topic. You have to prepare in advance in order to perform a disaster recovery fast and efficiently. You should pay special attention to disaster recovery preparation for the Cell Manager.

Important: Prepare for disaster recovery before a disaster occurs.

General preparations

1. Perform a full backup of the client system. It is recommended that you back up the whole client, however, you need to select at least the following critical volumes and objects:

   - the boot and system volumes
   - the Data Protector installation volume
   - the volume where the CONFIGURATION object is located

   For a *Data Protector Cell Manager system*, see "Additional preparations for the Cell Manager" below..

   See the *HP Data Protector Help* index: “backup, UNIX specific” and “backup, configuration”

   During a full client backup, the recovery set and P1S file are stored on the backup medium and (optionally) on the Cell Manager.

2. After a disaster occurs, use the EADR Wizard to convert the DR image into a disaster recovery CD ISO image.

3. Record the disaster recovery CD ISO image on a CD using any CD recording tool that supports the ISO9660 format. This disaster recovery CD can then be used to boot the target system and automatically restore critical volumes.

4. Execute a disaster recovery test plan.

Additional preparations for the Cell Manager

Successful disaster recovery of the Cell Manager requires additional preparation.

- Regularly back up the IDB. The IDB session should not be older than the file system session.
- Store the Cell Manager's SRD file at a safe location (not on the Cell Manager).
- Prepare a disaster recovery OS image for the Cell Manager in advance.

Saving a Recovery Set to the Cell Manager

A recovery set is packed in a single large file and stored on the backup medium and optionally on the Cell Manager during a full client backup. Saving the recovery set file to the Cell Manager is useful if you plan to record the disaster recovery CD on the Cell Manager, because it is much faster to obtain the recovery set file from the hard disk than to restore it from a backup medium.

If the recovery set is saved to the Cell Manager during backup, it is saved to the default Data Protector `P1S` files location.

To change the default location, specify a new global option `EADRImagePath = *valid_path*` (for example, `EADRImagePath = /home/images` or `EADRImagePath = C:\temp`).

See the HP Data Protector Help index: “Global Options, modifying”.

Tip: If you do not have enough free disk space in the destination directory, you can create a mount point (Windows systems) or a link to another volume (UNIX systems).

Saving the recovery set to the Cell Manager for all clients in the backup specification

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Filesystem**.
3. Select the backup specification you will use for a full client backup (create it if you have not done so already). For details, see the HP Data Protector Help index: “creating, backup specifications”.
4. In the Results Area, click **Options**.
5. Under **Filesystem Options** click **Advanced**.
6. In the **Other** page, select **Copy Recovery Set to disk**.

Figure 6: Other options tab

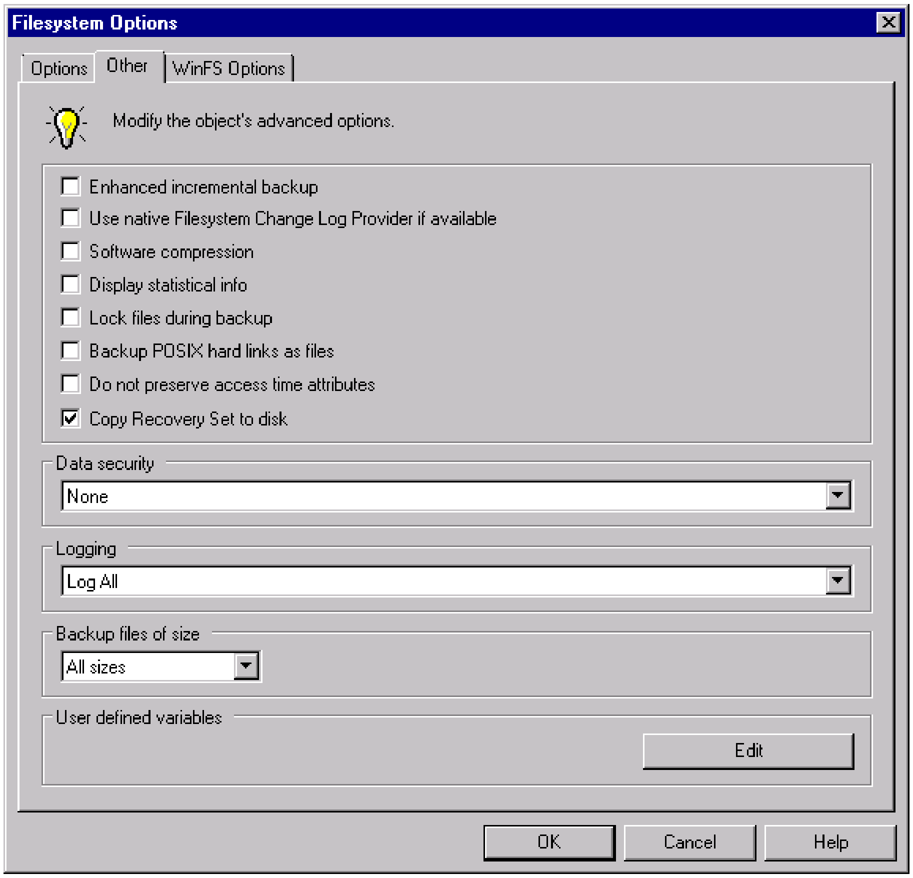

Modify the object's advanced options.

☐ Enhanced incremental backup
☐ Use native Filesystem Change Log Provider if available
☐ Software compression
☐ Display statistical info
☐ Lock files during backup
☐ Backup POSIX hard links as files
☐ Do not preserve access time attributes
☑ Copy Recovery Set to disk

Data security: None

Logging: Log All

Backup files of size: All sizes

User defined variables

Saving the recovery set to the Cell Manager for a particular client in the backup specification

To copy the recovery set files only for particular clients in the backup specification, perform the following steps:

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Filesystem**.
3. Select the backup specification you will use for a full client backup (create it if you have not done so already). For details, see the HP Data Protector Help index: “creating, backup specifications”.
4. In the Results Area, click **Backup Object Summary**.
5. Select the client for which you would like to store its recovery set file onto the Cell Manager and click **Properties**.
6. In the **Other** page, select **Copy Recovery Set to disk**.

Preparing the Encryption Keys

For a Cell Manager recovery or an offline client recovery, you must ensure that the encryption keys are available during the disaster recovery by storing them on a removable medium. For a Cell Manager recovery, prepare the removable medium in advance, before the disaster occurs.

The encryption keys are not part of the DR OS image file. During the disaster recovery image creation, the keys are automatically exported to the Cell Manager to the file `*Data_Protector_program_ data*\Config\Server\export\keys\DR-*ClientName*-keys.csv` (Windows systems) or `/var/opt/omni/server/export/keys/DR-*ClientName*-keys.csv` (UNIX systems), where *ClientName* is the name of the client for which the image is being created.

Ensure that you have the correct encryption key for each backup that is prepared for a disaster recovery.

Preparing a DR OS image

After a disaster occurs, you should prepare a DR OS image to be recorded on a disaster recovery CD or saved to a bootable USB drive, which can then be used for Enhanced Automated Disaster Recovery. Alternatively, you can prepare a bootable network image.

Note that the Data Protector Automatic Disaster Recovery component must be installed on the system where a DR OS image will be prepared.

A new disaster recovery OS image has to be prepared after each hardware, software or configuration change.

Prepare a DR OS image in advance for any critical systems that must be restored first, especially systems required for the network to function properly (DNS servers, domain controllers, gateways, and so on), Cell Managers, Media Agent clients, file servers, and so on.

It is recommended to restrict access to backup media and disaster recovery CDs or USB drives containing the OS image.

Steps

1. In the Data Protector Context List, click **Restore**.

2. In the Scoping Pane, click **Tasks**, and then click **Disaster Recovery** to start the Disaster Recovery Wizard.

3. In the Results Area, from the **Host to be recovered** drop down list, select the client you would like to prepare the DR OS image for.

4. From the **Recovery media creation host** drop down list, select the client on which you will prepare the DR OS image. By default, this is the same client for which the DR OS image is prepared for. The client on which you prepare the DR OS image must have the same OS type installed (Windows, Linux) and must have a Disk Agent installed.

5. Keep the **Enhanced Automated Disaster Recovery** selected and select whether the volume

recovery set will be built from a backup session or a list of volumes. By default, **Backup session** is selected.

Click **Next**.

6. Depending on the recovery set build method select:

   - If you selected Backup session, select the host backup session and in case of a Cell Manager, the IDB session.
   - If you selected Volume list, for each critical object select an appropriate object version.

   Click **Next**.

7. Select the location of the recovery set file. By default, **Restore recovery set file from a backup** is selected.

   If you have saved the recovery set file on the Cell Manager during backup, select **Path to the recovery set file** and specify its location. Click **Next**.

8. Select the image format. The following options are available:

   - **Create bootable ISO image**: a DR ISO image (by default, `recovery.iso`)
   - **Create bootable USB drive**: a DR OS image on a bootable USB drive
   - **Create bootable network image**: a DR OS image that can be used for the network boot (by default, `recovery.wim`)

9. If you are creating a bootable ISO image or a bootable network image, select the destination directory, where you would like to place the created image.

   If you are creating a bootable USB drive, select the destination USB drive or disk number, where you would like to place the created image.

   > **Important:** During the creation of the bootable USB drive, all data stored on the drive will be lost.

10. Optionally, set a password to protect the DR OS image from unauthorized use. The lock icon indicates whether a password has been set.

    Click **Password** to open the Password Protect Image dialog window and enter the password. To remove the password, clear the fields.

11. Click **Finish** to exit the wizard and create the DR OS image.

12. If you are creating a bootable CD or DVD, record the ISO image on a CD or DVD using a recording tool that supports the ISO9660 format.

Recovering Linux Systems Using EADR

You can successfully perform the Enhanced Automated Disaster Recovery of a Linux system only if all preparation steps were fulfilled. If you are recovering a Cell Manager, first the Internal Database is restored from its backup image, and restore of the volumes and the CONFIGURATION object from their backup image follows afterwards. For details on supported operating systems, see the *HP Data Protector Product Announcements, Software Notes, and References*.

Prerequisites

- You need a new hard disk to replace your affected disk.
- You should have a valid full filesystem backup of the entire system that you want to recover (client backup).
- For disaster recovery of the Cell Manager, you should have a valid Internal Database backup image that is newer than the filesystem backup image.
- You need a disaster recovery CD.

Steps

Phase 1

1. Unless you are performing an offline disaster recovery, add a Data Protector `admin` account with the following properties to the Data Protector `admin` user group on the Cell Manager:

   - Start restore
   - Restore to other clients
   - Restore as root

   > **Note:** The disaster recovery procedure can only be performed by the root user.

   For more information on adding users, see the HP Data Protector Help index: “adding Data Protector users”.

Figure 7: Adding a user account

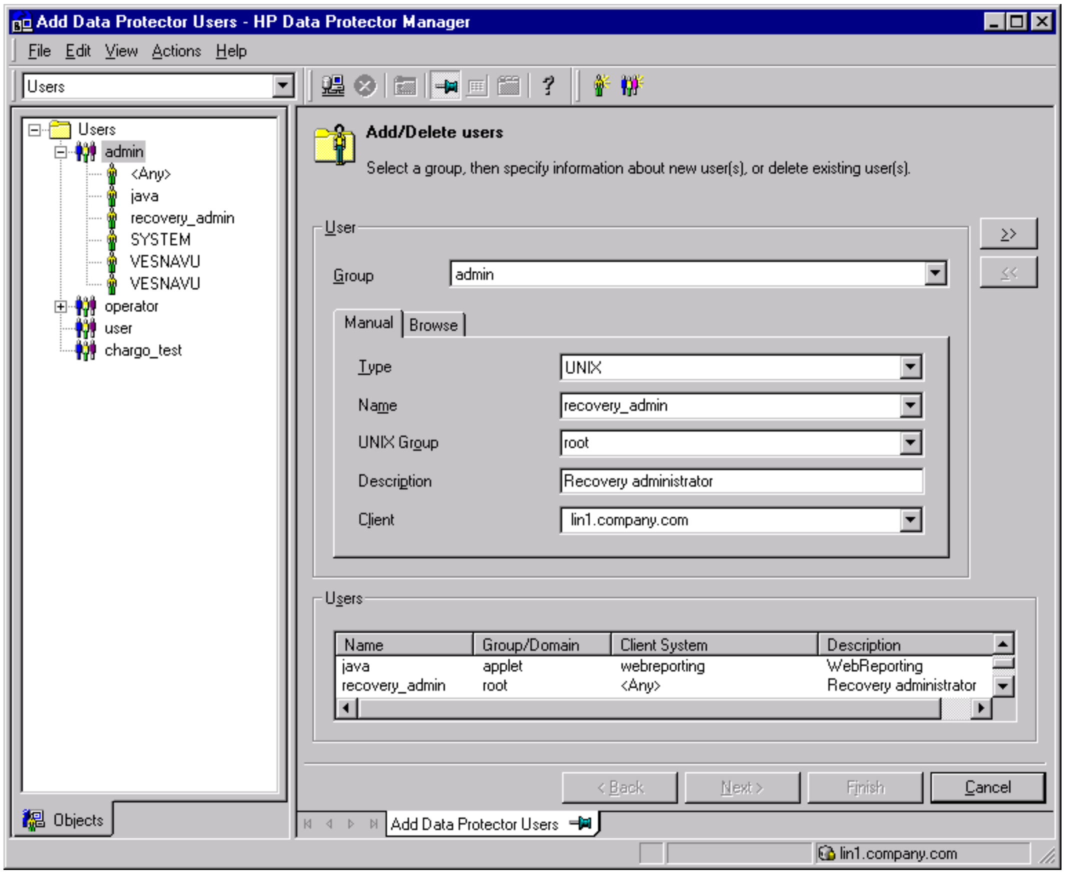

Add/Delete users

Select a group, then specify information about new user(s), or delete existing user(s).

User

Group: admin

Manual | Browse

Type: UNIX

Name: recovery_admin

UNIX Group: root

Description: Recovery administrator

Client: lin1.company.com

Users

Name	Group/Domain	Client System	Description
java	applet	webreporting	WebReporting
recovery_admin	root	<Any>	Recovery administrator

Note: If you are using encrypted control communication between the clients in a Data Protector cell, you must add the client to the Security Exceptions list on the Cell Manager before you start the recovery. Unless you are using a local device, the Media Agent client must be added to the Security Exceptions list on the Cell Manager as well.

2. Boot client system from the disaster recovery CD of the original system.

3. Press **Enter** when the following message is displayed: Press Enter to boot from Recovery CD.

4. The DR OS is loaded first into memory and then the scope menu is displayed. Select the scope of recovery. There are four different scopes of recovery and two additional options:

   - `Reboot`: Disaster recovery is not performed and the computer is restarted.
   - `Default Recovery`: Recovers the `/boot` and `/` (root) volumes and all volumes on which Data Protector installation and configuration files are located (`/opt`, `/etc`, and `/var`). All other disks are not partitioned and formatted and are ready for Phase 3.
   - `Minimal Recovery`: Recovers only the `/boot` and `/` (root) volumes.
   - `Full Recovery`: All volumes are recovered, not only the critical ones.

- `Full with Shared Volumes`: All volumes are recovered, including shared volumes that were locked at backup time.
- `Run shell`: Runs the Linux shell. You can use it for advanced configuration or recovery tasks.

Phase 2

5. The Disaster Recovery Wizard appears. To modify the disaster recovery options, press any key to stop the wizard during the countdown and modify the options. To continue with the disaster recovery, select **Proceed With Restore**.

6. If the disaster recovery backup is encrypted and you are either recovering the Cell Manager or a client where the Cell Manager is not accessible, the following prompt will appear:

   ```
   Do you want to use AES key file for decryption [y/n]?
   ```

   Press **y**.

   Ensure that the keystore (`DR-`*`ClientName`*`-keys.csv`) is available on the client (for example, by inserting a CD-ROM, floppy disk, or USB flash drive) and enter the full path to the keystore file. The keystore file is copied to the default location on the DR OS and is used by the Disk Agents. Disaster recovery now continues without further interruption.

7. If the information in 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.

8. Data Protector will then reestablish the previous storage structure within the selected scope of recovery and restore all critical volumes.

   Note that Data Protector will first try to perform online restore. If the online restore fails for any reason (for example, the Cell Manager or network service is not available or firewall is preventing access to the Cell Manager), Data Protector tries to perform remote offline recovery. If even the remote offline restore fails (for example, because the Media Agent host accepts only requests from the Cell Manager), Data Protector will perform a local offline restore.

9. Remove the client’s local Data Protector 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. If you are recovering a Cell Manager, make the IDB consistent.

Phase 3

11. Restore user and application data using the standard Data Protector restore procedure.

12. Additional steps are necessary if you are performing disaster recovery of all nodes in a cluster.

One Button Disaster Recovery (OBDR)

One Button Disaster Recovery (OBDR) is an automated Data Protector recovery method for Linux Data Protector clients, where user intervention is reduced to minimum. For details on supported operating systems, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.

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 (recovery set) 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 runs and configures the disaster recovery operating system (DR OS), partitions and formats 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 OBDR procedure recovers volumes depending on the selected scope of the recovery.

Any remaining volumes can be recovered using the standard Data Protector restore.

Overview

The general steps using the One Button Disaster Recovery method for a Windows client are:

1. **Phase 0**

   a. You need an OBDR backup image (create the backup specification using the Data Protector One Button Disaster Recovery Wizard).

   b. If you are using encrypted backups, store the encryption key on a removable medium so that it is available for disaster recovery.

2. **Phase 1**

   Boot from the recovery tape and select the scope of recovery.

3. **Phase 2**

   Depending on the recovery scope you select, the selected volumes are automatically restored.

   Critical volumes (the boot partition and the operating system) are always restored.

4. **Phase 3**

   Restore any remaining partitions using the standard Data Protector restore procedure.

Important: HP recommends to restrict access to OBDR boot media.

The following sections explain the requirements, limitations, preparation and recovery pertaining to One Button Disaster Recovery on Windows systems.

Requirements

- The Data Protector Automatic Disaster Recovery component must be installed on systems for which you want to enable recovery using this method. Additionally, the Automatic Disaster Recovery component must be installed on systems where the DR OS image will be prepared. For details, see the *HP Data Protector Installation and Licensing Guide*.
- The client system must support booting from the tape device that will be used for OBDR.

  For more information about supported systems, devices and media, see the HP Tape Hardware Compatibility Table and the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.
- The hardware configuration of the target system must be the same as that of the original system. This includes the SCSI BIOS settings (sector remapping).
- Replacement disks have to be attached to the same host bus adapter on the same bus.
- The volume on which Data Protector is installed should have at least 800 MB of free space. This space is required to create a temporary image.
- A media pool with a Non-appendable media usage policy and a Loose media allocation policy has to be created for the OBDR capable device. Only media from this pool can be used for disaster recovery.
- In a SAN boot configuration, make sure the following items on the target system are identical to the ones on the original system:
  - The local HBA's BIOS parameters
  - The SAN disks LUN numbers
- In multipath SAN disk configurations, the LUNs and WWIDs of the target system disks must be identical to the ones on the original system.

Limitations

- One Button Disaster Recovery (OBDR) is not available for Data Protector Cell Managers.
- A One Button Disaster Recovery backup session can only be run 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.
- USB tape storage devices are not supported.
- If you have a mount point with the name `CONFIGURATION` and it contains the directory `SystemRecoveryData`, data in the directory `SystemRecoveryData` will not be backed up.

- Do not mount disks using the disk ID, because the ID is unique and depends on the disk serial number. In case of a disaster, the disk may be replaced and the new disk will have a new ID. As a result, the disaster recovery fails.
- When restoring a Linux client with SELINUX enforcing mode enabled, the system has to relabel all system files after recovery which, depending on system configuration, can take some time to complete. If permissive mode is used, the system log will contain a large number of SELINUX warning messages.
- When you create a backup specification with the CONFIGRATION/SYSTEMRECOVERYDATA object selected, the folders `/opt/omni/bin/drim/log` and `/opt/omni/bin/drim/tmp` are by default excluded from the backup. However, this exclusion is not set if you manually update existing backup specifications.
- Fusion IO disks that do not automatically attach at MiniOS boot time need to be manually attached prior to recovery. This is required when replacing an old Fusion IO disk with a new one or when an internal Fusion IO disk error occurs. Those disks need to be formatted using specific tools before being attached in the MiniOS. To format and attach a Fusion IO disk manually to the system, you need to run the following commands in Linux shell present in MiniOS before the recovery starts:
  - `fio-status` – List the status of all the Fusion IO disks.
  - `fio-format [path]` – Perform low-level format of the Fusion IO disk.
  - `fio-attach [path]` – Attach the Fusion IO disk to the system.
- Sparse files are restored to their full size during offline restore. This may result in the target volume running out of space.

Disk and partition configuration

- A 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 for One Button Disaster Recovery

To prepare for a successful disaster recovery, follow the instructions related to the general preparation procedure for disaster recovery before completing the steps listed in this topic. Prepare in advance in order to perform a disaster recovery fast and efficiently.

> **Important:** Prepare for disaster recovery before a disaster occurs.

Preparatory steps

After you have completed the general preparation for disaster recovery, perform the following specific steps to prepare for OBDR.

1. Create a media pool for DDS or LTO media with the Non-appendable media usage policy and the Loose media allocation policy (because the backup media is formatted during OBDR backup). In addition,specify this media pool as the default media pool for the OBDR device. See the *HP Data Protector Help* index: “creating media pool”. Only media from such pool can be used for OBDR.

2. Perform the OBDR backup locally on the system for which you want to enable recovery using OBDR.

   If the full client backup was encrypted, store the encryption key on a removable medium so that it is available for disaster recovery. You will need the key if the connection to the Cell Manager cannot be established.

3. Execute a disaster recovery test plan.

Creating the Backup Specification for One Button Disaster Recovery

You need to create a One Button Disaster Recovery (OBDR) backup specification in order to prepare the OBDR boot tape.

Prerequisites

- Before adding an OBDR device, create a media pool for DDS or LTO media with the Non-appendable media usage policy and the Loose media allocation policy. The created media pool must be selected as the default media pool for the OBDR device.

- This device has to be connected locally to the system, for which you want to enable recovery using OBDR.

- The Data Protector Automatic Disaster Recovery and User Interface components must be installed on systems for which you want to enable recovery using the OBDR method.

- This backup specification has to be created locally on the system, for which you want to enable recovery using OBDR.

> **Tip:** To enable an automatic restore of all shared disk volumes in the MS Cluster using the OBDR method, move all volumes temporarily to the node for which you are preparing the OBDR boot tape. It is practically impossible to collect enough information to configure the disk in Phase 1 for shared disk volumes that are locked by another node.

Limitations

- One Button Disaster Recovery (OBDR) is not available for Data Protector Cell Managers.

Creating a backup specification for OBDR

Steps

1. In the Data Protector Context List, click **Backup**.
2. In the Scoping Pane, click **Tasks**, and then click **One Button Disaster Recovery Wizard**.
3. In the Results Area, select the client for which you would like to perform an OBDR backup (locally on the client) from the drop-down list and click **Next**.
4. The critical volumes that you need to back up are already selected. Click **Next**.

   > **Important:** Important volumes are selected automatically and cannot be deselected. Select any other partitions you want to keep, because during the recovery procedure Data Protector deletes all partitions from your system.

5. Select the local device or drive to be used for the backup. Only one device or drive can be selected. Click **Next**.
6. Select backup options. For more details on available options, see the *HP Data Protector Help* index: “backup options”.
7. Click Next to proceed to the Scheduler page, which can be used to schedule the backup. See the *HP Data Protector Help* index: “scheduling backups on specific dates and times”.
8. In the Backup Summary page, review the backup specification settings, and then click **Next**.

   > **Note:** You cannot change a previously selected backup device or the order in which the backup specifications follow one another. Only OBDR non-essential backup objects can be deleted and only general object properties can be viewed.
   >
   > You can also change a backup object description.

9. In the final page of the Backup wizard, you can save the backup specification, start the interactive backup, or preview the backup.

   HP recommends to save the backup specification so that you can schedule or modify it later.

   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 ensure that you do not override OBDR-specific options in it. If saved as a standard backup specification, it may not be 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.

If the backup is an encrypted, encryption IDs are exported automatically by the omnisrdupdate utility which is executed as a post-exec command.

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.

Preparing the Encryption Keys

For a Cell Manager recovery or an offline client recovery, you must ensure that the encryption keys are available during the disaster recovery by storing them on a removable medium. For a Cell Manager recovery, prepare the removable medium in advance, before the disaster occurs.

The encryption keys are not part of the DR OS image file. During the disaster recovery image creation, the keys are automatically exported to the Cell Manager to the file `*Data_Protector_program_data*\Config\Server\export\keys\DR-*ClientName*-keys.csv` (Windows systems) or `/var/opt/omni/server/export/keys/DR-*ClientName*-keys.csv` (UNIX systems), where *`ClientName`* is the name of the client for which the image is being created.

Ensure that you have the correct encryption key for each backup that is prepared for a disaster recovery.

Recovering Linux Systems Using OBDR

You can successfully perform the One Button Disaster Recovery (OBDR) of a Linux system only if all preparation steps were fulfilled.

For details on supported operating systems for OBDR, see the *HP Data Protector Product Announcements, Software Notes, and References*.

Prerequisites

- You need a new hard disk to replace your affected disk.
- You should have a bootable OBDR backup medium with all critical objects of the client that you want to recover. The OBDR backup has to be performed locally on the client.
- You need an OBDR device connected locally to the target system.

Steps

Phase 1

1. Unless you are performing an offline disaster recovery, add a Data Protector `admin` account with the following properties to the Data Protector `admin` user group on the Cell Manager, depending on the operating system of the target system:

   - Start restore
   - Restore to other clients
   - Restore as root

   Note: The disaster recovery procedure can only be performed by the root user.

   For more information on adding users, see the HP Data Protector Help index: “adding Data Protector users”.

   Figure 8: Adding a user account

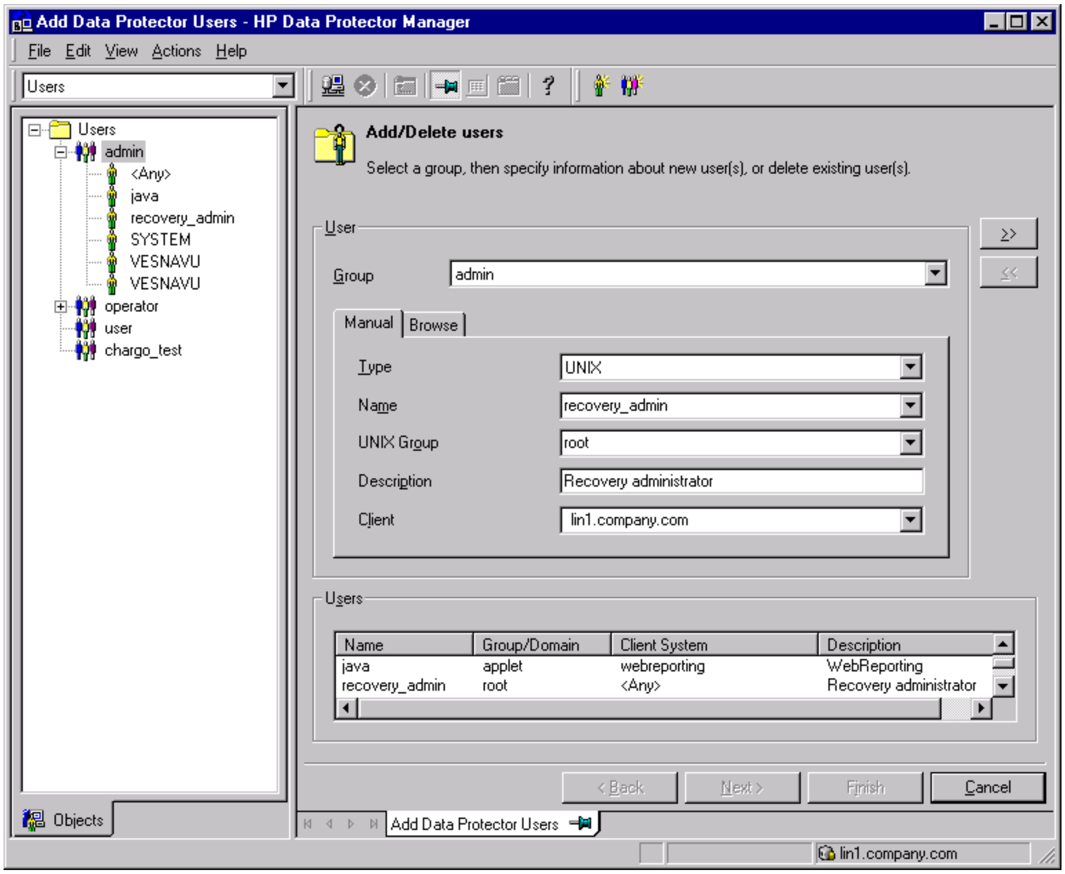

Select a group, then specify information about new user(s), or delete existing user(s).

User

Group: admin

Type: UNIX

Name: recovery_admin

UNIX Group: root

Description: Recovery administrator

Client: lin1.company.com

Users

Name	Group/Domain	Client System	Description
java	applet	webreporting	WebReporting
recovery_admin	root	<Any>	Recovery administrator

> **Note:** If you are using encrypted control communication between the clients in a Data Protector cell, you must add the client to the Security Exceptions list on the Cell Manager before you start the recovery. Unless you are using a local device, the Media Agent client must be added to the Security Exceptions list on the Cell Manager as well.

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 the target system on and, while it is being initialized, press the Eject button on the tape device and power it on. For details, see the device documentation.

5. The DR OS is loaded first into memory and then the scope menu is displayed. Select the scope of recovery. There are four different scopes of recovery and two additional options:

   - `Reboot`: Disaster recovery is not performed and the computer is restarted.

   - `Default Recovery`: Recovers the `/boot` and `/` (root) volumes and all volumes on which Data Protector installation and configuration are located (`/opt`, `/etc`, and `/var`). All other disks are not partitioned and formatted and are ready for Phase 3.

   - `Minimal Recovery`: Recovers only the `/boot` and `/` (root) volumes.

   - `Full Recovery`: All volumes are recovered, not only the critical ones.

   - `Full with Shared Volumes`: All volumes are recovered, including shared volumes that were locked at backup time.

   - `Run shell`: Runs the Linux shell. You can use it for advanced configuration or recovery tasks.

Phase 2

6. The Disaster Recovery Wizard appears. To modify the disaster recovery options, press any key to stop the wizard during the countdown and modify the options. Select Proceed With Restore to continue with the disaster recovery.

7. If the disaster recovery backup is encrypted and you are recovering a client whose Cell Manager is not accessible, the following prompt is displayed:

   ```
   Do you want to use AES key file for decryption [y/n]?
   ```

   Press **y**.

   Ensure that the keystore (`DR-`*`ClientName`*`-keys.csv`) is available on the client (for example, by inserting a CD-ROM, floppy disk, or USB flash drive) and enter the full path to the keystore file. The keystore file is copied to the default location on the DR OS and is used by the Disk Agents. Disaster recovery now continues without further interruption.

8. 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.

9. Data Protector will then reestablish the previous storage structure within the selected scope of recovery and restore all critical volumes.

   Note that Data Protector first tries to perform an online restore. If the online restore fails for any reason (for example, the Cell Manager or network service is not available or firewall is preventing access to the Cell Manager) Data Protector tries to perform remote offline recovery. If the remote offline restore fails (for example, because the Media Agent host accepts requests only from the Cell Manager), Data Protector performs a local offline restore.

10. Remove the client's local Data Protector 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.

Phase 3

11. Additional steps are required if you are recovering a Cell Manager or performing advanced recovery tasks (such as editing the SRD files).

12. Restore user and application data using the standard Data Protector restore procedure.

Chapter 5: Troubleshooting disaster recovery

This chapter contains descriptions of problems you might encounter while performing a disaster recovery. You can start with problems connected to a particular disaster recovery method and continue with general disaster recovery problems. For information where to find the error messages, see "Troubleshooting disaster recovery" above.

For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*e.

Before you begin

- Ensure that the latest official Data Protector patches are installed. For more information on how to verify this, see the *HP Data Protector Help* index: “patches”.
- For general Data Protector limitations, as well as known problems and workarounds, see the *HP Data Protector Product Announcements, Software Notes, and References*.
- For an up-to-date list of supported versions, platforms, and other information, see http://support.openview.hp.com/selfsolve/manuals.

Troubleshooting Automatic Disaster Recovery

The AUTODR.log file

Automatic Disaster Recovery includes two disaster recovery methods: EADR and OBDR. Messages relevant to these methods are logged in the `AUTODR.log` file, located in the default Data Protectortemporary files directory directory. 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 levels of messages (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) in the `AUTODR.log` file:

- `Critical error`: the error is so serious that the backup of the object cannot continue and will be aborted.
- `Error`: there is an error but if it is critical, depends on different factors.

  For example, `AUTODR.log` reports an error that some driver has not been included in the DR OS. The missing driver may be the reason why the recovered system is not operational after the recovery. This can also result in some non-critical service not to be running after the boot of the operating system. The seriousness of an error 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.

Two of the most common messages stated in the `AUTODR.log` file are:

- `unsupported location`: Data Protector notices that a certain file required by a service or a driver that will be included in the 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. The success of disaster recovery depends on which service or driver was affected. A possible solution to 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.

Debugging disaster recovery sessions

During a disaster recovery session, the debugging settings and the location of the debug logs depend on the disaster recovery phase:

- During the DR OS preparation, the debug logs are automatically saved to `X:\$DRM$\log` (Windows Vista and later releases), `c:\$DRM$\log` (Windows XP, Windows Server 2003), or `/opt/omni/bin/drim/log/Phase1.log` (Linux systems).

- During the data restore step, you must manually select the debugging options in the Disaster Recovery Wizard to enable debugging.

Windows

To enable creation of debug logs:

1. In the Disaster Recovery Wizard, press any key to stop the wizard during the countdown.

   Select the check box to the left of the Debugs button.

Figure 9: Enabling debugs during a disaster recovery session

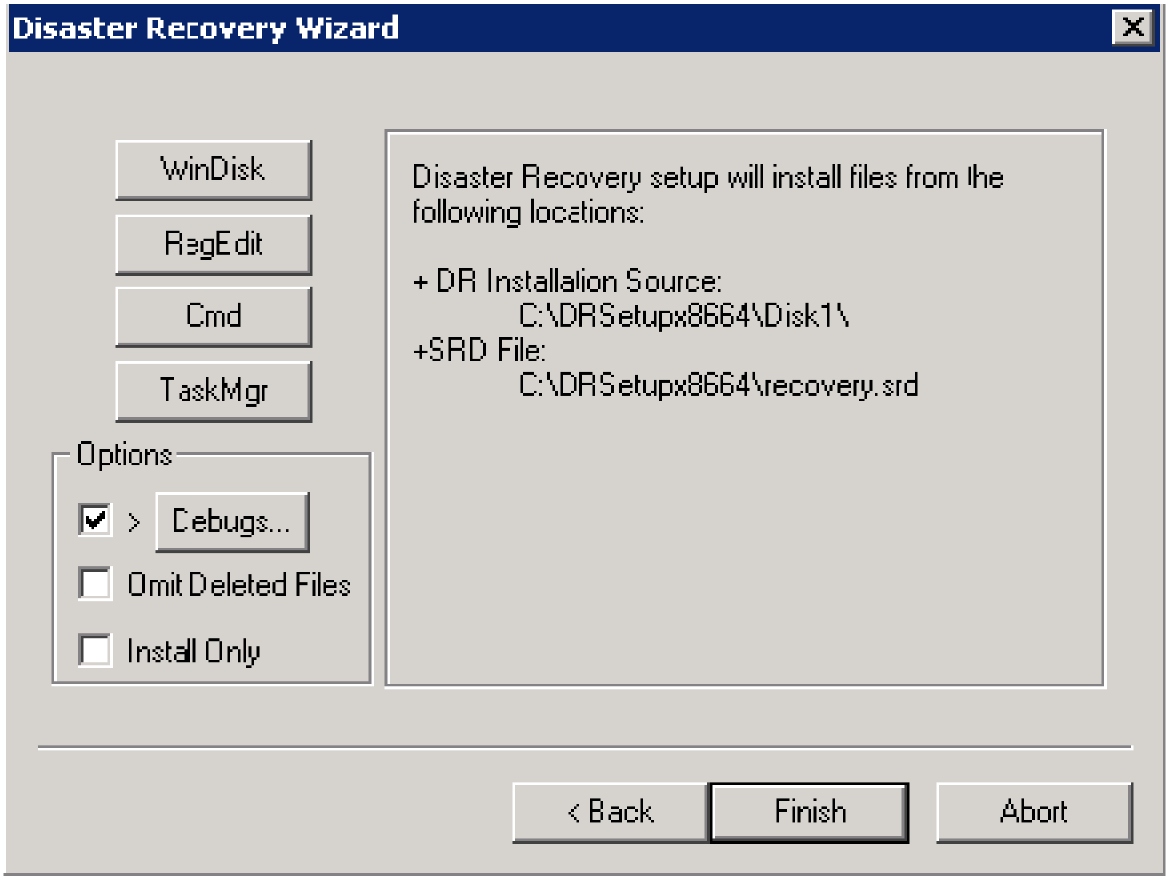

Disaster Recovery setup will install files from the following locations:

+ DR Installation Source:
C:\DRSetupx8664\Disk1\
+SRD File:
C:\DRSetupx8664\recovery.srd

☑ > Debugs...
□ Omit Deleted Files
□ Install Only

2. To specify the debug options, such as the location where the debugs are saved, click **Debugs...**. By default, the debugs are saved into the *%SystemRoot%*`\system32\OB2DR\tmp` directory.

> **Note:** On Windows Vista and later releases, the directory *%SystemRoot%*`\system32\OB2DR\tmp` resides on RAM disk. The RAM disk size is typically limited to less than 64 MB. Once the RAM disk usage reaches the limit, Data Protector may start to behave unpredictably. Thus, if you expect that the disaster recovery session will produce a large amount of debugs, you must change the location to which the debugs will be saved.

The Debug Options window appears.

Figure 10: Changing the debug logs location

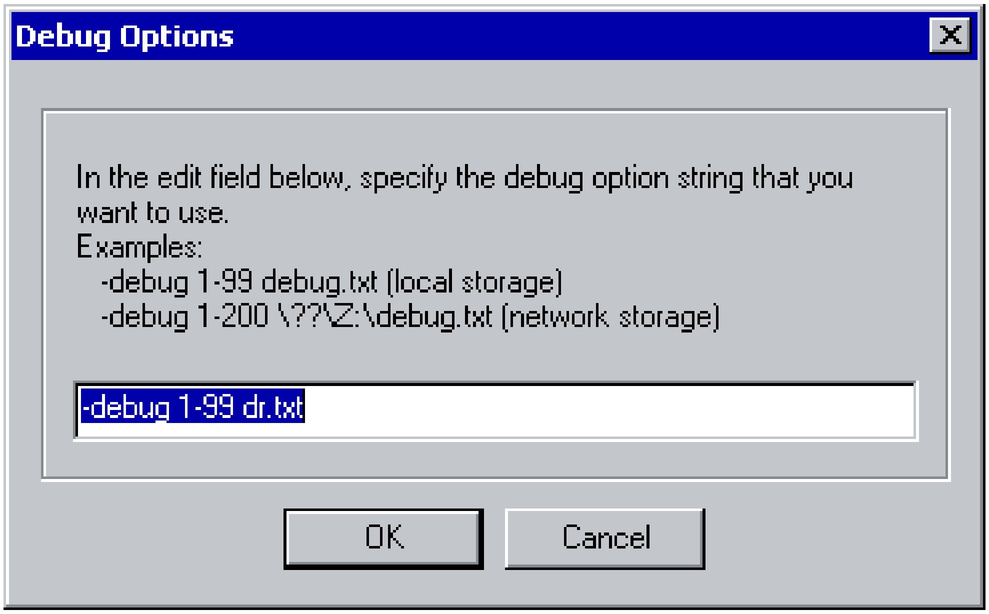

In the edit field below, specify the debug option string that you want to use.
Examples:
-debug 1-99 debug.txt (local storage)
-debug 1-200 \??\Z:\debug.txt (network storage)

-debug 1-99 dr.txt

3. Enter the location, where the debug logs are saved. The drive must be preceded by `\\?`, for example, `\\?\Z:\debug.txt`. If you choose to save the debugs on a network share, use the net use command to mount the share to which the debug logs are written. For example, `net use X: "\\client\debug_output_folder /user:username password"`.

Linux systems

To enable creation of debug logs:

1. In the Disaster Recovery Wizard, select **Use debugs**.

2. On the debug options screen, select to either use the default options or modify them.

   ```
   Select one of following options:

   1) Use Default Debug Option "-debug 1-200 dr.txt"

   2) Specify Different Debug Option

   3) Disable Debug option

   Command [1-3]:
   ```

   > **Note:** On Linux systems, the directory to which the debug logs are saved, resides on RAM disk. The RAM disk size is typically limited. Once the RAM disk usage reaches the limit, Data Protector may start to behave unpredictably. Thus, if you expect that the disaster recovery session will produce a large amount of debugs, you should change the location to which the debugs will be saved. To change the location, select **Specify Different Debug Option**.

3. A new screen will appear on which you can enter the debug parameters.

   ```
   Examples:

   -debug 1-200 debug.txt (local storage)

   -debug 1-200 //servername/sharename/debug.txt (windows share)

   -debug 1-200 servername:/sharename/debug.txt (nfs share)

   Specify the debug option string that you want to use:
   ```

   You can choose to save the debug files to a Windows shared disk or an NFS shared folder.

Setting omnirc options during disaster recovery

For general information on omnirc options, see the *HP Data Protector Troubleshooting Guide*.

If you need to set an omnirc option during the disaster recovery on a Windows or Linux system, perform the following steps:

Windows systems

1. When the Disaster Recovery Wizard appears, press any key to stop the wizard during the countdown.

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` option set to 1000 seconds.

4. Close the command prompt and click **Next** in the Disaster Recovery Wizard to proceed with disaster recovery.

Linux systems

1. In the Disaster Recovery Wizard, switch to another console by pressing **Alt F3**.

2. In the console, run the following command:

```
echo variable > /opt/omni/.omnirc
```

where *variable* is the omnirc option exactly as it should be written in the .omnirc file.

Example:

```
echo OB2RECONNECT_RETRY=1000 > /opt/omni/.omnirc
```

This command creates a .omnirc file in the disaster recovery operating system with the OB2RECONNECT_RETRY option set to 1000 seconds.

3. Type **exit** to exit the shell and proceed with disaster recovery in the Disaster Recovery Wizard.

The drm.cfg file on Windows

The Data Protector disaster recovery configuration is set up to cover a broad range of system configurations. However, in some cases, these settings may not be the most appropriate, or you may want to modify some of the settings in order to troubleshoot issues on your system.

The drm.cfg file contains several parameters that you can modify and which affect the disaster recovery process, along with a description of their impact. The file is available for EADR and OBDR.

To change the parameters:

1. Copy the template file drm.cfg.tmpl to drm.cfg. The template is created during an installation or upgrade in *Data_Protector_home*\bin\drim\config, with all parameters set to their default values.

2. Edit the drm.cfg file. Set the desired value for parameters. Follow the instructions in the file.

Disabling the automatic collection of EADR or OBDR

When running a full client backup, the CONFIGURATION backup may fail while collecting data needed for a certain backup method, even though this method will not be used for disaster recovery, because Data Protector by default collects data for all automatic disaster recovery methods. For example, this may happen while Data Protector collects data for EADR if the boot disks are LDM disks.

Disable automatic collecting of data for the disaster recovery method that failed. This will allow Data Protector to collect data needed for other methods.

Set the option OB2_TURNOFF_COLLECTING to one of the following values:

Value	Description
0	Default setting, data collection is turned on for all automatic methods (EADR, OBDR).
1	Turn off collecting of EADR/OBDR data
2	EADR/OBDR data is still collected.
3	Turn off collecting for all methods.

Common Problems (All Methods)

While performing disaster recovery, you may encounter the following problems:

You cannot perform a disaster recovery from a media copy or an object 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.

Action

- 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.
- 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.

You cannot log on after disaster recovery finishes

Problem

Problems occur logging on to the system after disaster recovery finishes.

You may receive the following message:

```
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.
```

Such a message may be caused by one of the following reasons:

- After collecting all information for disaster recovery, you reinstalled Windows and added it into the offending domain.
- After collecting all information for disaster recovery, you removed your system from the offending domain and later added it into the same or some other domain.

In cases like this, Windows generates new system security information, which is incompatible with information that is restored during disaster recovery.

Action

1. Log on to the system locally as an Administrator.
2. In the Control Panel, click **Network** and, using the **Identification** tab, remove the system from its current domain to a temporary workgroup.
3. Reinsert the system into the domain from which it was previously removed. You need a domain administrator password. Click **OK**.
4. Restart the system.

To update this new state, repeat all necessary disaster recovery preparation steps.

Disaster recovery fails due to inappropriate network settings

Problem

A disaster recovery session fails because Data Protector recovers a client with unsuitable network configuration.

The default settings that are used to configure the client's network depend on the client's operating system:

Windows XP, Windows Server 2003:

The original network configuration (network configuration at the time of backup), which is specified in the SRD file.

Windows Vista and later releases:

Network configuration that is defined by the DHCP settings.

Action

To switch to the non-default network configuration:

1. Start a disaster recovery session.
2. When Data Protector displays:

   Windows XP, Windows Server 2003:

   Press F8 in the next 10 seconds to switch network to DHCP...

   Windows Vista and later releases:

   Press F8 in the next 10 seconds to switch to the network setup at the time of backup...

   press **F8**.

Troubleshooting Assisted Manual Disaster Recovery

While performing Assisted Manual Disaster Recovery, you may encounter the following problems:

"Cannot copy file"

Problem

Drstart reports: “`Can not copy` *`filename`*`.`"

This error is reported because the *drstart* utility cannot 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 **Abort**.

Troubleshooting Enhanced Automated Disaster Recovery and One Button Disaster Recovery

You may run into the following problems during disaster recovery with the Enhanced Automated Disaster Recovery or One Button Disaster Recovery method:

Automatic DR information could not be collected

Problem

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

Possible reasons for this error are stored in `autodr.log` file located in the default Data Protectortemporary files directory directory:

1. Check if all storage devices are configured correctly. If Device Manager reports a device as “Unknown Device”, you have to install the proper device drivers before you can perform EADR/OBDR. A similar entry would appear in `autodr.log`, if improperly configured storage devices are attached to your system:

   ```
   DRIM_WIN_ERROR 13 SetupDiGetDeviceRegistryProperty
   ```

2. 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 `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 will work) and contact technical support.

Some non-critical errors were detected

Problem
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. Review the Automatic DR log file.`
Action
A non-critical error detected during the execution of the Automatic Disaster Recovery module means that the backup can most likely still be used for disaster-recovery purposes. Possible reason for non-critical errors are stored in `autodr.log` located in the default Data Protectortemporary files directory directory. For example: 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.

Network is not available during restore

Problem

This can be caused by various reasons, for example a damaged network cable or switch. Another possible reason for network failure is 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.

Action

1. Ensure that the problem is not in with switches, cables, and so on.

2. If the DNS server (as configured at backup time) is offline during the restore, you can either:

   - perform offline recovery and change the DNS settings after recovery.

   - edit the registry before Phase 2 is started. In this case you have to restart the system before Phase 2 for the changes to take effect. After Phase 2 finishes, you must correct the settings before Phase 3 can be started.

     > **Caution:** Editing the registry incorrectly may result in failed disaster recovery.

Network is not available due to missing network drivers

Problem

On a Windows Vista or Windows Server 2008 system, during a disaster recovery, the network is not available because the DR OS does not support the network cards.

Action

Inject the missing drivers in to the DR OS image.

EADR and OBDR online recovery fails when the Cell Manager and a client are in the different domains

Problem

This can be caused by incorrect network configuration.

Action

1. Update the `host` files on both Cell Manager and client systems. These files must contain host names of the Cell Manager and of the client and their IP addresses.

2. Check whether the `ping` request between the Cell Manager and the client returns the correct value. In case of a problem, contact your network administrator.

3. Check whether the DNS resolution between the Cell Manager and the client is correct with the `omnicheck -dns` command. For more details, see the `omnicheck` man page or the *HP Data Protector Command Line Interface Reference*. In case of a problem, contact your network administrator.

When encrypted control communication is enabled, Cell Manager does not respond during an online recovery of a client

Problem
On Windows Vista and later releases, when you perform an online disaster recovery of a client in a DHCP environment with encrypted control communication enabled on the Cell Manager and the client added as an exception, online restore fails, disaster recovery continues with the offline restore. The reason is that a new temporary hostname is generated for the DR OS by default.
Action
During disaster recovery, switch to the original network configuration by selecting the **Restore Network Configuration** option. Alternatively, check the hostname of the system after the DR OS is started and add this name as an exception on the Cell Manager before starting the restore.

Auto logon does not work

Problem
Sometimes auto logon does not work.
Action
Manually log on using an administrator account with a blank password.

Computer stops responding during EADR

Problem
This can be caused by problems with the disaster recovery CD.
Action
• Check if the CD is readable. • Do not reuse CD-RWs too many times.

Cannot create a CD ISO image for EADR of Microsoft Cluster Server

Problem
The quorum disk has to be backed up in order to be able to create a CD ISO image.
Action
Back up the quorum disk.

Creating a CD ISO image on a Microsoft Cluster Server client fails

Problem
In a Microsoft Cluster Server environment, you cannot create an ISO image on a cluster client. The file system restore works as expected. The issue arises because Data Protector tries to use the cluster IP (which is a virtual one) instead of the domain name (which is resolved to the IP of the physical client).
Action
Change the connection order for network services so that the `Local Area Connection` is on top.

Volume is not re-mounted during phase 1

Problem
On some systems (depending on the disk controller and its configuration) a volume (without a drive letter assigned) associated with a mount point on a different volume may not be re-mounted properly during phase 1 of the disaster recovery. This may occur if the volume containing the mount point is recreated or reformatted (for example the System Volume with DR OS), causing the operating system to boot in “Safe Mode” and to miss the detection of the file system present on the original mount point's target volume. Consequently, the disaster recovery module does not recognize this volume and reports it as MISSING in the `drecovery.ini` file. The contents of such a volume are intact, even if it is not recognized.
Action
• Mount the volume with a drive letter and verify it with the `chkdsk /v /f` command or wait until the system is completely restored and then recreate the original mount point. • Manually restart the system directly to MiniOS (do not start the system from the recovery CD). The previously dismounted volume will be automatically mounted to a drive letter.

After a failed or aborted disaster recovery, Boot Descriptors are left

Problem
On Intel Itanium systems, after a failed or aborted disaster recovery session, Boot Descriptors (named `DRM Temporary OS`) may be left in the EFI environment. This can cause unwanted behavior when restarting the disaster recovery process.
Action
Remove the boot descriptor using the option **Remove Boot Descriptor** from the scope selection menu. After the boot descriptor is removed, you can proceed with disaster recovery, by selecting the scope.

A wrong or no boot disk is selected on an Intel Itanium system

Problem
On an Intel Itanium system, the wrong boot disk (or no boot disk at all) is selected.
Action
1. Select **Manual Disk Selection** from the scope selection menu. A new menu, listing all available disks, will display. 2. Determine the correct boot disk. Press **o** to view information about the original disk and **d** to see details about the selected one. 3. Select the disk from the list using cursor keys and press **b**. You can remove a selection by pressing **c**. If the boot disk is not the same as the system disk (by default, both disks are the same), you must select the system disk as well. Select **Back**. 4. Select the scope of the recovery and disaster recovery will continue.

Disaster recovery fails with "There is not enough space" message

Problem

A disaster recovery of a Windows Server 2008 R2 domain controller fails with an error similar to the following:

```
[Major] From: VRDA@computer.company.com "Dev1" [/CONFIGURATION]" Time:

07.12.2012 15:33:58 X:\windows\System32\OB2DR\tmp\config\

ActiveDirectoryService\D$\ Windows\NTDS\ntds.dit Cannot write:

([112] There is not enough space on the disk. ) => not restored.
```

Action

1. Modify the backup specification for the client backup: in the source page, expand the `CONFIGURATION` object and clear the checkboxes for the `ActiveDirectoryService` and `SYSVOL` items.

   > **Note:** The Active Directory and SYSVOL will still be backed up as part of the system volume (`C:/`) backup. By default, they are located in `C:/Windows/NTDS` and `C:/Windows/SYSVOL` respectively.

2. Repeat the disaster recovery procedure.

Minor errors or warnings are displayed during a client backup

Problem

During a client backup, minor errors may be reported:

```
Cannot perform stat(): ([2] No such file or directory)

File is shorter than it was when it was opened
```

Such warnings and errors may appear due to changed files inside temporary Data Protector directories. This can happen for example if the `/CONFIGURATION` mount point and the `/` (root) mount points are backed up simultaneously.

Action

Exclude the `/opt/omni/bin/drim/tmp` and `/opt/omni/bin/drim/log` directories from your backup specifications.

In backup specifications created with 8.10 or later versions, these files are automatically excluded.

Troubleshooting Disaster Recovery of Internet Information Server

Problems with disaster recovery of Internet Information Server (IIS) are usually the result of either services not running or services not being installed.

IIS dependent services do not start automatically

Problem
Any of the IIS dependent services (for example, SMTP, NNTP) does not start automatically after the recovery of IIS.
Action
1. Start the services manually. 2. If this fails, stop the IIS Admin Service and restore the `%SystemRoot%\system32\inetsrv\MetaBase.bin` file using the Overwrite option. **Note:** The `%SystemRoot%\system32\inetsrv` directory is the default location of the IIS Service. If you have installed the service into another location, use this location as a destination for restore of the `MetaBase.bin` file. 3. Start the IIS Admin Service and all dependent services.

Appendix A: Example Preparation Tasks

Example of Moving Kill Links on HP-UX 11.x

```
# The system will go from "run-level" 4 to "run-level 1"

# retaining the (rpcd), 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
```

Example of the Disaster Recovery Preparation Table for Windows

Client properties	Computer name	ANAPURNA
	Hostname	anapurna.company.com
Drivers		tatpi.sys, aic78xx.sys
Windows Service Pack		Windows Vista

TCP/IP properties for IPv4	IP address	10.17.2.61
	Default gateway	10.17.250.250
	Subnet mask	255.255.0.0
	DNS order	10.17.3.108, 10.17.100.100
TCP/IP properties for IPv6	IP address	td10:1234:5678:abba::6:1600
	Subnet prefix length	64
	Default gateway	td10:1234:5678:abba::6:1603
	Preferred DNS server	td10:1234:5678:abba::6:1603
	Alternate DNS server	td10:1234:5678:abba::6:1604
Medium label/barcode number		"anapurna - 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	

Glossary

[

[%=DP.DP_AbbCompanyName%] [%=DP.HW_SW_P9000_XP_special%] Agent

A [%=DP.DP_BriefProductName%] software component that executes all tasks needed by the [%=DP.DP_BriefProductName%] [%=DP.HW_SW_P9000_XP_full%] integration. It communicates with the [%=DP.HW_SW_P9000_XP_abbrev%] storage system via the RAID Manager Library.

[%=DP.DP_AbbCompanyName%] Business Copy (BC) [%=DP.HW_SW_P6000_EVA_special%] ([%=DP.HW_SW_P6000_EVA_full%] specific term)

A local replication software solution that enables creation of point-in-time copies (replicas) of the source volumes using the snapshot and clone capabilities of the [%=DP.HW_SW_P6000_EVA_special%] firmware. See also replica, source volume, snapshot, and [%=DP.DP_AbbCompanyName%] Continuous Access + Business Copy (CA+BC) [%=DP.HW_SW_P6000_EVA_special%].

[%=DP.DP_AbbCompanyName%] Business Copy (BC) [%=DP.HW_SW_P9000_XP_special%] ([%=DP.HW_SW_P9000_XP_full%] specific term)

An [%=DP.HW_SW_P4000_LH_full%] configuration that enables creation and maintenance of internal copies of LDEVs for various purposes, such as data duplication and backup. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system. For [%=DP.DP_BriefProductName%] zero downtime backup purposes, P-VOLs should be available to the application system, and one of the S-VOL sets should be available to the backup system. See also LDEV, [%=DP.DP_AbbCompanyName%] Continuous Access (CA) [%=DP.HW_SW_P9000_XP_special%], Main Control Unit (MCU), application system, and backup system.

[%=DP.DP_AbbCompanyName%] Command View (CV) EVA ([%=DP.HW_SW_P6000_EVA_full%] specific term)

The user interface that enables you to configure, manage, and monitor your [%=DP.HW_SW_P6000_EVA_special%] storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, and creating snapshots, snapclones, or mirrorcloens of virtual disks. The [%=DP.DP_AbbCompanyName%] Command View EVA software runs on the [%=DP.DP_AbbCompanyName%] Storage Management Appliance, and is accessed by a Web browser. See also [%=DP.DP_AbbCompanyName%] P6000 / [%=DP.DP_AbbCompanyName%] 3PAR SMI-S Agent and [%=DP.DP_AbbCompanyName%] SMI-S [%=DP.HW_SW_P6000_EVA_abbrev%] provider.

[%=DP.DP_AbbCompanyName%] Continuous Access (CA) [%=DP.HW_SW_P9000_XP_special%] ([%=DP.HW_SW_P9000_XP_full%] specific term)

An [%=DP.HW_SW_P9000_XP_full%] confgiuration that enables creation and maintenance of remote copies of LDEVs for purposes such as data duplication, backup, and disaster recovery. [%=DP.DP_AbbCompanyName%] CA [%=DP.HW_SW_P9000_XP_special%] operations involve main (primary) disk array units and remote (secondary) disk array units. The main disk array units are connected to the application system and contain primary volumes (P-VOLs), which store original data. The remote disk aray units are connected to the backup system and contain secondary volumes (S-VOLs). See also [%=DP.DP_AbbCompanyName%] BC [%=DP.HW_SW_P9000_XP_special%] ([%=DP.HW_SW_P9000_XP_full%] specific term), Main Control Unit (MCU), and LDEV.

[%=DP.DP_AbbCompanyName%] Continuous Access + Business Copy (CA+BC) [%=DP.HW_SW_P6000_EVA_special%] ([%=DP.HW_SW_P6000_EVA_full%] specific term)

An [%=DP.HW_SW_P6000_EVA_full%] configuration that enables creation and maintainance of copies (replicas) of the source volumes on a remote [%=DP.HW_SW_P6000_EVA_special%], and later use of these copies as the source for local replication on this remote array. See also [%=DP.DP_AbbCompanyName%] BC [%=DP.HW_SW_P6000_EVA_special%], replica, and source volume.

[%=DP.DP_AbbCompanyName%] P6000 / [%=DP.DP_AbbCompanyName%] 3PAR SMI-S Agent

A [%=DP.DP_BriefProductName%] software module that executes all tasks required for the [%=DP.HW_SW_P6000_EVA_full%] integration. With the P6000 / 3PAR SMI-S Agent, the control over the array is established through an appropriate SMI-S provider, which directs communication between incoming requests and the storage system's native interface. See also [%=DP.DP_AbbCompanyName%] Command View (CV) EVA and [%=DP.DP_AbbCompanyName%] SMI-S [%=DP.HW_SW_P6000_EVA_abbrev%] provider.

[%=DP.DP_AbbCompanyName%] SMI-S [%=DP.HW_SW_P6000_EVA_abbrev%] provider

An interface used for controlling [%=DP.HW_SW_P6000_EVA_full%]. SMI-S [%=DP.HW_SW_P6000_EVA_abbrev%] provider runs as a separate service on the [%=DP.DP_AbbCompanyName%] Storage Management Appliance system and acts as a gateway between incoming requests and [%=DP.DP_AbbCompanyName%] Command View EVA. With the [%=DP.DP_BriefProductName%] [%=DP.HW_SW_P6000_EVA_full%] integration, SMI-S [%=DP.HW_SW_P6000_EVA_abbrev%] provider accepts standardized requests from the [%=DP.DP_AbbCompanyName%] P6000 / [%=DP.DP_AbbCompanyName%] 3PAR SMI-S Agent, communicates with [%=DP.DP_AbbCompanyName%] Command View EVA for information or method invocation, and returns standardized responses. See also [%=DP.DP_AbbCompanyName%] P6000 / [%=DP.DP_AbbCompanyName%] 3PAR SMI-S Agent and [%=DP.DP_AbbCompanyName%] Command View (CV) EVA.

[%=DP.DP_BriefProductName%] user account

You can use [%=DP.DP_BriefProductName%] only if you have a [%=DP.DP_BriefProductName%] user account, which restricts unauthorized access to [%=DP.DP_BriefProductName%] and to backed up data. [%=DP.DP_BriefProductName%] administrators create this account specifying a user logon name, the systems from which the user can log on, and a [%=DP.DP_BriefProductName%]

user group membership. This is checked whenever the user starts the [%=DP.DP_BriefProductName%] user interface or performs specific tasks.

[%=DP.PROD_HomeDir%]

A reference to the directory containing [%=DP.DP_BriefProductName%] program files (on Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012) or the directory containing [%=DP.DP_BriefProductName%] program files and data files (on other Windows operating systems). Its default path is %ProgramFiles%\OmniBack, but the path can be changed in the [%=DP.DP_BriefProductName%] Setup Wizard at installation time. See also [%=DP.PROD_ProgDataDir%].

[%=DP.PROD_ProgDataDir%]

A reference to the directory containing [%=DP.DP_BriefProductName%] data files on Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012. Its default path is %ProgramData%\OmniBack, but the path can be changed in the [%=DP.DP_BriefProductName%] Setup Wizard at installation time. See also [%=DP.PROD_HomeDir%].

A

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.

AES 256-bit encryption

The [%=DP.DP_BriefProductName%] software encryption, based on the AES-CTR (Advanced Encryption Standard in Counter Mode) encryption algorithm that uses random keys of 256-bit length. The same key is used for both encryption and decryption. With AES 256-bit encryption, data is encrypted before it is transferred over a network and before it is written to media.

AML (ADIC/GRAU specific term)

Automated Mixed-Media library.

AMU (ADIC/GRAU specific term)

Archive Management Unit.

application agent

A component needed on a client system 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.

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.

archived log files ([%=DP.DP_BriefProductName%] specific term)

Files that keep track of changes made to the [%=DP.DP_BriefProductName%]

Internal Database (IDB). They are used for online or offline IDB restore and recovery where the IDB needs to be recreated either in its latest possible state, beyond the time of the most recent IDB backup session, or in a state between the times of two consecutive IDB backup sessions.

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 an archived log destination. This copy is the archived redo log. The presence or absence of an archived redo log is determined by the mode the database is using: ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered if an instance or a disk fails. A "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.

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 as well as on the backup medium. After a disaster occurs, the ASR archive file is extracted to diskettes which you need to perform ASR.

audit logs

Data files to which auditing information is stored.

audit report

User-readable output of auditing information created from data stored in audit log files.

auditing information

Data about every backup session that was performed over an extended, user-defined period for the whole [%=DP.DP_BriefProductName%] cell.

autochanger

See library.

autoloader

See library.

Automatic Storage Management (ASM) (Oracle specific term)

A filesystem and volume manager integrated into Oracle which manages Oracle database files. It eliminates complexity associated with data and disk management and optimizes performance by providing striping and mirroring capabilities.

auxiliary disk

A bootable disk that has a minimal operating system with networking and [%=DP.DP_BriefProductName%] 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.

B

BACKINT (SAP R/3 specific term)

A [%=DP.DP_BriefProductName%] interface program that lets the SAP R/3 backup programs communicate with the [%=DP.DP_BriefProductName%] software via calls to an open interface. For backup and restore, SAP R/3 programs issue commands through the

[%=DP.DP_BriefProductName%] backint interface.

backup API (Oracle specific term)

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 reading and writing of data to media and also to allow creation, searching, and removing of backup files.

backup chain

See restore chain.

backup device

A device configured for use with [%=DP.DP_BriefProductName%] that 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/application entity or a disk image (rawdisk). A backup object is defined by: client name (hostname of the [%=DP.DP_BriefProductName%] client where the backup object resides), mount point (for filesystem objects - the access point in a directory structure on the client where the backup object is located (drive on Windows systems and mount point on UNIX systems), for integration objects - backup stream identification, indicating the backed up database/application items), description (for filesystem objects - uniquely defines objects with identical client name and mount point, for integration objects - displays the integration type), and type (for filesystem objects - filesystem type, for integration objects - "Bar").

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. The result of a backup session is a set of media, which was written to, also called the backup or media set. See also backup specification, 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; and 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. All clients configured in one backup specification are backed up at the same time in one backup session using the same backup type (full or incremental).

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

[%=DP.DP_BriefProductName%] provides different views of your 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 Continuances are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.

BC Process (EMC Symmetrix specific term)

A protected storage environment solution, which uses specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices. See also BCV.

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 splitable 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 [%=DP.DP_BriefProductName%] 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 containing 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 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 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 BRARCHIVE and BRBACKUP.

BSM

The [%=DP.DP_BriefProductName%] Backup Session Manager, which controls the backup session. This process always runs on the Cell Manager system.

C

CAP (StorageTek specific term)

The Cartridge Access Port built into the door panel of a library. The purpose is to enter or eject media.

Catalog Database (CDB)

A part of the [%=DP.DP_BriefProductName%] Internal Database (IDB) that contains information about backup, restore, object copy, object consolidation, object verification, and media management sessions. This part of the IDB is always local to the cell. It is stored in the embedded database. See also MMDB.

catalog protection

Defines how long information about backed up data (such as filenames and file attributes) is kept in the IDB. See also data protection.

CDB

See Catalog Database (CDB).

CDF file (UNIX systems 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 [%=DP.DP_BriefProductName%] 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

[%=DP.DP_BriefProductName%] allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All

[%=DP.DP_BriefProductName%] 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.

Certificate Server

A Windows certificate server can be installed and configured to provide certificates for clients. It provides customizable services for issuing and managing certificates for the enterprise. These services issue, revoke, and manage certificates employed in public key-based cryptography technologies.

Change Journal (Windows specific term)

A Windows filesystem feature that logs a record of each change as it occurs to the files and directories on a local NTFS volume.

Change Log Provider

A module that can be queried to determine which objects on a filesystem have been created, modified, or deleted.

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' or type 'sbt_tape'. If the specified channel is of type 'sbt_tape' and Oracle is integrated with [%=DP.DP_BriefProductName%], the server process will attempt to read backups from or write data files to [%=DP.DP_BriefProductName%].

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 volumes (filesystems) mounted on a [%=DP.DP_BriefProductName%] client. What is actually backed up depends on how you select objects in a backup specification. If you select the check box next to the client system name, a single backup object of the Client System type is created. As a result, at the time of the backup, [%=DP.DP_BriefProductName%] first detects all volumes that are mounted on the selected client and then backs them up. On Windows clients, CONFIGURATION is also backed up. If you individually select all volumes that are mounted on the client system, a separate backup object of the Filesystem type is created for each volume. As a result, at the time of the backup, only the selected volumes are backed up. Volumes that have been potentially mounted on the client after the backup specification was created are not backed up.

client or client system

Any system configured with any [%=DP.DP_BriefProductName%] functionality and configured in a cell.

cluster continuous replication (Microsoft Exchange Server specific term)

Cluster continuous replication (CCR) is a high availability solution that uses cluster management and failover options to create and maintain an exact copy (CCR copy) of a storage group. A storage group is replicated to a separate server. CCR removes any single point of failure in your Exchange back-end servers. You can perform backups using VSS on your passive Exchange Server node where a CCR copy is located and thus reducing the load on the active node. A CCR copy is used for disaster recovery since you can switch to the CCR copy in a few seconds. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) like an ordinary storage group. See also Exchange Replication Service and local continuous replication.

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 HP Serviceguard), application services, IP names and addresses ...).

CMD script for Informix Server (Informix Server specific term)

A Windows CMD script that is created in INFORMIXDIR when an Informix Server database is configured. The CMD script is a set of system commands that export environment variables for Informix Server.

CMMDB

The [%=DP.DP_BriefProductName%] 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 Manager-of-Managers. A reliable network connection between the MoM cell and the other [%=DP.DP_BriefProductName%] cells is highly recommended. See also MoM.

COM+ Class Registration Database (Windows specific term)

The COM+ Class Registration Database and the Windows Registry store COM+ application attributes, class attributes, and computer-level attributes. This guaranties consistency among these attributes and provides common operation on top of these attributes.

command device ([%=DP.HW_SW_P9000_XP_full%] specific term)

A dedicated volume in the disk array which acts as the interface between a management application and the disk array's storage system. It cannot be used for data storage and only accepts requests for operations that are then executed by the disk array.

command-line interface (CLI)

A set commands that you can use in shell scripts to perform [%=DP.DP_BriefProductName%] configuration, backup, restore, and management tasks.

concurrency

See Disk Agent concurrency.

container ([%=DP.HW_SW_P6000_EVA_full%] specific term)

Space on a disk array, which is pre-allocated for later use as a standard snapshot, vsnap, or snapclone.

control file (Oracle and SAP R/3 specific term)

A data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

copy set ([%=DP.HW_SW_P6000_EVA_full%] specific term)

A pair that consists of the source volumes on a local [%=DP.HW_SW_P6000_EVA_special%] and their replica on a remote [%=DP.HW_SW_P6000_EVA_special%]. See also source volume, replica, and [%=DP.DP_AbbCompanyName%] Continuous Access + Business Copy (CA+BC) [%=DP.HW_SW_P6000_EVA_special%].

CRS

The [%=DP.DP_BriefProductName%] Cell Request Server process (service), which runs on the Cell Manager, starts and controls the backup and restore sessions. The service is started as soon as [%=DP.DP_BriefProductName%] is installed on the Cell Manager. The CRS runs under the account root on UNIX systems. On Windows systems it runs under the account of the user, specified at installation time.

CSM

The [%=DP.DP_BriefProductName%] Copy and Consolidation Session Manager process controls the object copy and object consolidation sessions and runs on the Cell Manager system.

D

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, [%=DP.DP_BriefProductName%] will not overwrite it. When the protection expires, [%=DP.DP_BriefProductName%] will be able to reuse the media in one of the next backup sessions. See also catalog protection.

data replication (DR) group ([%=DP.HW_SW_P6000_EVA_full%] specific term)

A logical grouping of [%=DP.HW_SW_P6000_EVA_full%] virtual disks. It can contain up to eight copy sets provided they have common characteristics and share a common [%=DP.DP_AbbCompanyName%] CA [%=DP.HW_SW_P6000_EVA_special%] log. See also copy set.

data stream

Sequence of data transferred over the communication channel.

database library

A [%=DP.DP_BriefProductName%] set of routines that enables data transfer between [%=DP.DP_BriefProductName%] and a server of an online database integration, for example, 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 Server specific term)

An Informix Server physical database object. It can be a blobspace, dbspace, or logical log file.

DC directory

A directory that contains DC binary files, one for each configured Data Protector backup medium. DC directories constitute the Detail Catalog Binary Files part of the Data Protector Internal Database. See also Detail Catalog Binary Files (DCBF) and Internal Database (IDB).

DCBF

See Detail Catalog Binary Files (DCBF).

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.

Detail Catalog Binary Files (DCBF)

A part of the Data Protector Internal Database that stores names, versions, and metadata of the backed up items. It consists of DC directories with DC binary files. See also DC directory and Internal Database (IDB).

device

See backup device.

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. -Get configuration, status, and performance statistics by device group. -Issue control operations that apply to all devices in the device group.

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 writing to the tape, and so on. If data is written to the tape slower than it is delivered to the device then the device is streaming. Streaming significantly improves the use of space and the performance of the device.

DHCP server

A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic IP address assignment and network configuration for DHCP clients.

differential backup

An incremental backup that backs up changes made since the last full backup. To perform this type of backup, specify the Incr1 backup type. See also incremental backup.

differential backup (Microsfot SQL Server 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.

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.

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.

disaster recovery operating system

See DR OS.

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. During an object verification session the Disk Agent receives data from the Media Agent and performs the verification process, but no data is written to disk.

Disk Agent concurrency

The number of Disk Agents that are allowed to send data to one Media Agent concurrently.

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 [%=DP.DP_BriefProductName%] 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 media format

A media format, available with the file library, which supports a space efficient type of synthetic backup called virtual full backup. Using this format is a prerequisite for virtual full backup. See also virtual full backup.

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.

DR image

Data required for temporary disaster recovery operating system (DR OS) installation and configuration.

DR OS

An operating system environment in which disaster recovery runs. It provides [%=DP.DP_BriefProductName%] with a basic runtime environment (disk, network, tape, and filesystem access). It has to be installed on disk or loaded into memory and configured before the [%=DP.DP_BriefProductName%] disaster recovery can be performed. DR OS can be either temporary or active. A 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. An active DR OS not only hosts the [%=DP.DP_BriefProductName%] disaster recovery process but can also be 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. 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.

drive-based encryption

The [%=DP.DP_BriefProductName%] drive-based encryption uses the encryption functionality of the drive. While performing the backup, the drive encrypts both the data and the metadata that is written to the medium.

E

EMC Symmetrix Agent

A [%=DP.DP_BriefProductName%] software module that prepares the EMC Symmetrix environment for backup and restore operations.

emergency boot file (Informix Server specific term)

The Informix Server configuration file ixbar.<server_id> that resides in the directory <INFORMIXDIR>/etc (on Windows systems) or <INFORMIXDIR>\etc (on UNIX systems). <INFORMIXDIR> is the Informix 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.

encrypted control communication

[%=DP.DP_BriefProductName%] secure communication between the clients in the [%=DP.DP_BriefProductName%] cell is based on a Secure Socket Layer (SSL) that uses export grade SSLv3 algorithms to encrypt control communication. Control communication in a [%=DP.DP_BriefProductName%] cell is all communication between [%=DP.DP_BriefProductName%] processes, except the data transfer from Disk Agent (and Integrations) to Media Agent, and the other way round.

encryption key

A 256-bit randomly generated number used by the [%=DP.DP_BriefProductName%] encryption algorithm to encode information during backups for which AES 256-bit software encryption or drive-based encryption has been specified. The same key is used for subsequent decryption of the information. Encryption keys for a [%=DP.DP_BriefProductName%] cell are stored in a central keystore on the Cell Manager.

encryption KeyID-StoreID

Combined identifier used by the [%=DP.DP_BriefProductName%] Key Management Server to identify and administer encryption keys used by [%=DP.DP_BriefProductName%]. KeyID identifies the key within the keystore. StoreID identifies the keystore on the Cell Manager. If [%=DP.DP_BriefProductName%] has been upgraded from an earlier version with encryption functionality, there may be several StoreIDs used on the same Cell Manager.

enhanced incremental backup

Conventional incremental backup backs up files that have changed since a previous backup, but has certain limitations in detection of changes. Unlike conventional incremental backup, enhanced incremental backup reliably detects and backs up also renamed and moved files, as well as files with changes in attributes.

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 [%=DP.DP_BriefProductName%] cells which are managed and administered from a central cell using the Manager-of-Managers concept. See also MoM.

Event Log ([%=DP.DP_BriefProductName%] Event Log)

A central repository of all [%=DP.DP_BriefProductName%]-related notifications. By default, all notifications are sent to the Event Log. The events are logged on the Cell Manager into the Ob2EventLog.txt file residing in the default [%=DP.DP_BriefProductName%] server log files directory. The Event Log is accessible only to users of the [%=DP.DP_BriefProductName%] admin user group and to users who are granted the [%=DP.DP_BriefProductName%] Reporting and notifications user rights. You can view or delete all events in the Event Log.

Event Logs (Windows specific term)

Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. [%=DP.DP_BriefProductName%] can back up Windows Event Logs as part of the Windows configuration backup.

Exchange Replication Service (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that represents storage groups that were replicated using either local continuous replication (LCR) or cluster continuous replication (CCR) technology. See also cluster continuous replication and local continuous replication.

exchanger

See library.

exporting media

A process that removes all data about backup sessions, such as systems, objects, and filenames, 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 specific term)

A database technology used as a storage system for information exchange in Microsoft Exchange Server.

F

failover

Transferring of the most important cluster data, called group (on Windows systems) or package (on UNIX systems) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.

failover ([%=DP.HW_SW_P6000_EVA_full%] specific term)

An operation that reverses the roles of source and destination in [%=DP.DP_AbbCompanyName%] Continuous Access + Business Copy (CA+BC) [%=DP.HW_SW_P6000_EVA_special%] configurations. See also [%=DP.DP_AbbCompanyName%] Continuous Access + Business Copy (CA+BC) [%=DP.HW_SW_P6000_EVA_special%].

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 bi-directional 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 tree walk

The process of traversing a filesystem to determine which objects have been created, modified, or deleted.

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, [%=DP.DP_BriefProductName%] 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 ([%=DP.HW_SW_P9000_XP_full%] specific term)

A mirror of an internal disk (LDEV) of a disk array of the [%=DP.HW_SW_P9000_XP_full%] which can be further mirrored itself, producing second-level mirrors. For [%=DP.DP_BriefProductName%] zero downtime backup and instant recovery purposes, only first-level mirrors can be used. See also primary volume and mirror unit (MU) number.

flash recovery area (Oracle specific term)

A directory, filesystem, or Automatic Storage Management (ASM) disk group managed by Oracle that serves as a centralized storage area for files related to backup, restore, and database recovery (recovery files). See also recovery files.

formatting

A process that erases any data contained on a medium and prepares it for use with [%=DP.DP_BriefProductName%]. Information about media (media ID, description, and location) is saved in the IDB as well as on the respective media (media header). Use the Force operation option to reformat [%=DP.DP_BriefProductName%] media with non-protected data. [%=DP.DP_BriefProductName%] 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.

full ZDB

A ZDB-to-tape or ZDB-to-disk+tape session in which all selected objects are streamed to tape, even if there are no changes from the previous backup. See also incremental ZDB.

G

global options

A set of options that define behavior of the entire [%=DP.DP_BriefProductName%] cell. The options are stored in a plain text file on the Cell Manager.

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 graphical user interface provided by [%=DP.DP_BriefProductName%] for easy access to all configuration, administration, and operation tasks. It is available for Microsoft Windows operating systems.

H

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 that resides on the Cell Manager in the default [%=DP.DP_BriefProductName%] server configuration directory.

hosting system

A working [%=DP.DP_BriefProductName%] client used for Disk Delivery Disaster Recovery with a [%=DP.DP_BriefProductName%] Disk Agent installed.

I

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

See Internal Database (IDB).

IDB recovery file

A file that maintains information about completed IDB backup sessions and the backup media and backup devices used in them. If available, it significantly simplifies and speeds up offline recovery of the Internal Database in case of a Cell Manager disaster. Its filename is obdrindex.dat.

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 (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 backup

A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, which enables detailed control of restore chain length. 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.

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.

incremental ZDB

A filesystem ZDB-to-tape or ZDB-to-disk+tape session in which only changes from the last protected full or incremental backup are streamed to tape. See also full ZDB.

Incremental1 Mailbox Backup

An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.

Inet

A process that runs on each UNIX system or service that runs on each Windows system in the [%=DP.DP_ BriefProductName%] 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 [%=DP.DP_ BriefProductName%] is installed on a system. The Inet process is started by the inetd daemon.

Information Store (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that is responsible for storage management. Information Store in Microsoft Exchange Server 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.

Informix Server (Informix Server specific term)

Refers to Informix Dynamic Server.

initializing

See formatting.

Installation Server

A computer system that holds a repository of the [%=DP.DP_ BriefProductName%] installation packages for a specific architecture. The Installation Server is used for remote installation of [%=DP.DP_ BriefProductName%] 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.

integration object

A backup object of a [%=DP.DP_ BriefProductName%] integration, such as Oracle or SAP MaxDB.

Internal Database (IDB)

An entity in [%=DP.DP_ BriefProductName%] that keeps information regarding which data was backed up, to which media it was backed up, how and when backup, restore, and other sessions were run, which devices, libraries, and disk arrays are configured, and so on. It is implemented with an embedded database and a collection of proprietary data files which reside on the Cell Manager. See also DC directory and Detail Catalog Binary Files (DCBF).

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).

ISQL (Sybase specific term)

A Sybase utility used to perform system administration tasks on Sybase SQL Server.

J

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'.

K

Key Management Service (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that provides encryption functionality for enhanced security. See also Information Store and Site Replication Service.

keychain

A tool that eliminates the supply of a passphrase manually when decrypting the private key. It needs to be installed and configured on the Installation Server if you perform remote installation using secure shell.

keystore

All encryption keys are centrally stored in the keystore on the Cell Manager and administered by the Key Management Server (KMS).

KMS

Key Management Server (KMS) is a centralized service that runs on the Cell Manager and provides key management for the [%=DP.DP_BriefProductName%] encryption functionality. The service is started as soon as [%=DP.DP_BriefProductName%] is installed on the Cell Manager.

L

LBO (Symmetric specific term)

A 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.

LDEV ([%=DP.HW_SW_P9000_XP_full%] specific term)

A logical partition of a physical disk of a disk array of the [%=DP.HW_SW_P9000_XP_full%]. An LDEV is the entity that can be replicated using the split-mirror or snapshot functionality of such disk array. See also [%=DP.DP_AbbCompanyName%] Business Copy (BC) [%=DP.HW_SW_P9000_XP_special%], [%=DP.DP_AbbCompanyName%] Continuous Access (CA) [%=DP.HW_SW_P9000_XP_special%], and replica.

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, [%=DP.DP_BriefProductName%] 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. [%=DP.DP_BriefProductName%] 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, [%=DP.DP_BriefProductName%] prompts you to select the device, which will be used for restore.

local continuous replication (Microsoft Exchange Server specific term)

Local continuous replication (LCR) is a single-server solution that creates and maintains an exact copy (LCR copy) of a storage group. An LCR copy is located on the same server as the original storage group. When an LCR copy is created, it is kept up to date through change propagation (log replay) technology. The replication feature in LCR guarantees that logs that have not been replicated are not deleted. The implication of this behavior is that running backups in a mode that deletes logs may not actually free space if replication is sufficiently far behind in its log copying. An LCR copy is used for disaster recovery because you can switch to the LCR copy in a few seconds. If an LCR copy is used for backup and if it is located on a different disk than the original data, then the I/O load on a production database is minimal. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) as a normal storage group. See also cluster continuous replication and Exchange Replication Service.

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 Server UNIX systems specific term)

A script provided by ON-Bar that you can use to start backing up logical log files when Informix Server issues a logfull event alarm. The Informix Server ALARMPROGRAM configuration parameter defaults to the <INFORMIXDIR>/etc/log_full.sh, where <INFORMIXDIR> is the Informix 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

An optino that determines the amount of details on files and directories written to the IDB during backup, object copying, or object consolidation. You can always restore your data, regardless of the logging level used during backup. [%=DP.DP_BriefProductName%] provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings mainly influence the IDB growth 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 (Microsoft SQL Server specific term)

The name a user needs 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 Oracle SYSDBA or SYSOPER rights. <password> must 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 must be the owner of the Oracle Recovery 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 [%=DP.DP_BriefProductName%].

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.

M

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 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 and a streaming native internet content .stm file.

Main Control Unit (MCU) ([%=DP.HW_SW_P9000_XP_full%] specific term)

An [%=DP.HW_SW_P9000_XP_full%] unit that contains primary volumes (P-VOLs) for the [%=DP.DP_AbbCompanyName%] CA [%=DP.HW_SW_P9000_XP_abbrev%] or [%=DP.DP_AbbCompanyName%] CA+BC [%=DP.HW_SW_P9000_XP_abbrev%] configuration and acts as a master device. See also [%=DP.DP_AbbCompanyName%] Business Copy (BC) [%=DP.HW_SW_P9000_XP_special%], [%=DP.DP_AbbCompanyName%] Continuous Access (CA) [%=DP.HW_SW_P9000_XP_special%], and LDEV.

maintenance mode

An operating mode you can initiate on the Cell Manager to prevent changes in the Internal Database. It enables you to perform various maintenance tasks, such as upgrading and patching the installation.

make_net_recovery

make_net_recovery is an Ignite-UX command, which 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 either 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.

make_tape_recovery

make_tape_recovery is a command on Ignite-UX which 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.

Manager-of-Managers

See MoM.

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.

MCU

See Main Control Unit (MCU).

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 or object verification session, a Media Agent locates data on the backup medium and sends it to the Disk Agent for processing. For a restore session, the Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

media allocation policy

Determines in which sequence media are used for backup. The Strict allocation policy directs [%=DP.DP_BriefProductName%] to prompt for a specific medium. The Loose policy directs [%=DP.DP_BriefProductName%] to prompt for any suitable medium. The Formatted First policy directs [%=DP.DP_BriefProductName%] 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 assert the condition of a medium.

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.

medium ID

A unique identifier assigned to a medium by [%=DP.DP_BriefProductName%].

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.

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

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, replica, and writer.

mirror (EMC Symmetrix and [%=DP.HW_SW_P9000_XP_full%] specific term)

See target volume.

mirror rotation ([%=DP.HW_SW_P9000_XP_full%] specific term)

See replica set rotation.

mirror unit (MU) number ([%=DP.HW_SW_P9000_XP_full%] specific term)

A non-negative integer number that determines a secondary volume (S-VOL) of an internal disk (LDEV) located on a disk array of the [%=DP.HW_SW_P9000_XP_full%]. See also first-level mirror.

mirrorclone ([%=DP.HW_SW_P6000_EVA_full%] specific term)

A dynamic replica of a storage volume, which is kept updated with changes made to the original storage volume via a local replication link. Replication between the original storage volume and its mirrorclone can be suspended. For each storage volume, a single mirrorclone can be created on the disk array.

MMD

The Media Management Daemon process (service) (MMD) runs on the [%=DP.DP_BriefProductName%] Cell Manager and controls media management and device operations. The process is started when [%=DP.DP_BriefProductName%] 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 [%=DP.DP_BriefProductName%] media used for backup. In an enterprise backup environment, this part of the database can be common to all cells. See also CMMDB and Catalog Database (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 cells are called MoM clients. The MoM enables you to configure and manage multiple cells from a central point.

mount point

The access point in a directory structure for a disk or logical volume, for example /opt or d:. On UNIX systems, the mount points are displayed using the bdf or df command.

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.

MSM

The [%=DP.DP_BriefProductName%] Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.

multisnapping ([%=DP.HW_SW_P6000_EVA_full%] specific term)

Simultaneous creation of target volumes so that the backup data is consistent not only on each individual target volume, but also across all the volumes that constitute a snapshot. See also snapshot.

O

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.

obdrindex.dat

See IDB recovery file.

object

See backup object.

object consolidation

The process of merging a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object. The process is a part of the synthetic backup procedure. The result is a synthetic full backup of the specified backup object.

object consolidation session

A process that merges a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this 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. [%=DP.DP_BriefProductName%] 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. [%=DP.DP_BriefProductName%] enables you to mirror all or some backup objects to one or more media sets.

object verification

The process of verifying the data integrity of backup objects, from the [%=DP.DP_BriefProductName%]point of view, and the ability of [%=DP.DP_BriefProductName%] to deliver them to the required destination. The process can be used to provide a level of confidence in the ability to restore object versions created by backup, object copy, or object consolidation sessions.

object verification session

A process that verifies the data integrity of specified backup objects or object versions and the ability of selected [%=DP.DP_BriefProductName%] network components to deliver them to a specified host. Object verification sessions can be run interactively, or as specified in automated post-backup, or scheduled specifications.

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. Cell Manager can only be recovered offline.

offline redo log

See archived redo log.

ON-Bar (Informix Server specific term)

A backup and restore system for Informix Server. ON-Bar enables you to create a copy of your Informix Server data and later restore the data. The ON-Bar backup and restore system involves the following components: the onbar command, [%=DP.DP_BriefProductName%] as the backup solution, the XBSA interface, and ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.

ONCONFIG (Informix Server specific term)

An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the onconfig file located in the directory <INFORMIXDIR>\etc (on Windows systems) or <INFORMIXDIR>/etc/ (on UNIX systems).

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 recovery

A type of the Internal Database recovery that you can use when the Cell Manager is accessible. In this case, the Cell Manager runs sessions, the sessions are logged into the IDB, and you can monitor the session progress using the GUI.

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.

OpenSSH

A set of network connectivity tools used to access remote machines securely, by using a variety of authentication and encryption methods. It needs to be installed and configured on the Installation Server and the client if you perform remote installation using secure shell.

Oracle Data Guard (Oracle specific term)

Oracle Data Guard is Oracle’s primary disaster recovery solution. Oracle Data Guard is able to maintain up to nine standby databases, each of which is a real-time copy of the production (primary) database, to protect against corruptions, data failures, human errors, and disasters. If a failure occurs on the production (primary) database, then a failover to one of the standby databases which becomes the new primary database is possible. In addition, planned downtime for maintenance can be reduced because the production processing can be moved from the current primary database to a standby database and back quickly.

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 [%=DP.DP_BriefProductName%] 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.

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. When copying or consolidating backup objects, the owner of the resulting objects is the user who started the original backup session.

P

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 with the filename recovery.p1s.

package (HP 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 ([%=DP.HW_SW_P9000_XP_full%] specific term)

The status of a disk pair (secondary volume and its corresponding primary volume) of a disk array of the [%=DP.HW_SW_P9000_XP_full%]. Depending on the circumstances, the paired disks can be in various states. The following states are particularly important for the operation of the [%=DP.DP_BriefProductName%] [%=DP.DP_AbbCompanyName%] [%=DP.HW_SW_P9000_XP_special%] Agent: PAIR - The secondary volume is prepared for zero downtime backup. If it is a mirror, it is completely synchronized, and if it is a volume to be used for snapshot storage, it is empty. SUSPENDED - The link between the disks is suspended. However, the pair relationship is still maintained, and the secondary disk can be prepared for zero downtime backup again at a later time. COPY - The disk pair is currently busy and making a transition into the PAIR state. If the secondary volume is a mirror, it is re-synchronizing with the primary volume, and if it is a volume to be used for snapshot storage, its contents are getting cleared.

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.

phase 0 of disaster recovery

Preparation for disaster recovery - the prerequisite condition for a successful disaster recovery.

phase 1 of disaster recovery

Installation and configuration of DR OS, establishing previous storage structure.

phase 2 of disaster recovery

Restoration of operating system (with all the configuration information that defines the environment) and [%=DP.DP_BriefProductName%].

phase 3 of disaster recovery

Restoration of user and application data.

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 [%=DP.DP_BriefProductName%]. You need to create your own. They can be written as executables or batch files on Windows systems and as shell scripts on UNIX systems. 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 [%=DP.DP_BriefProductName%]. You need to create your own commands. They can be written as executables or batch files on Windows systems and as shell scripts on UNIX systems.

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 [%=DP.DP_BriefProductName%]. You need to create your own. They can be written as executables or batch files on Windows systems and as shell scripts on UNIX systems. See also post-exec.

primary volume (P-VOL) ([%=DP.HW_SW_P9000_XP_full%] specific term)

An internal disk (LDEV) of a disk array of the [%=DP.HW_SW_P9000_XP_full%] for which a secondary volume (S-VOL), either its mirror or a volume to be used for its snapshot storage, exists. In the [%=DP.DP_AbbCompanyName%] CA [%=DP.HW_SW_P9000_XP_special%] and [%=DP.DP_AbbCompanyName%] CA+BC [%=DP.HW_SW_P9000_XP_special%] configurations, primary volumes are located in the Main Control Unit (MCU). See also secondary volume (S-VOL) and Main Control Unit (MCU).

protection

See data protection and catalog protection.

public folder store (Microsoft Exchange Server 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 [%=DP.DP_BriefProductName%] users private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

R

RAID

Redundant Array of Independent Disks.

RAID Manager [%=DP.HW_SW_P9000_XP_special%] ([%=DP.HW_SW_P9000_XP_full%] specific term)

A software application that provides a command-line interface to disk arrays of the [%=DP.HW_SW_P9000_XP_full%]. It offers an extensive set of commands for reporting and controlling the status of a [%=DP.HW_SW_P9000_XP_abbrev%] storage system, and for performing various operations on the disk array.

RAID Manager Library ([%=DP.HW_SW_P9000_XP_full%] specific term)

A software library that is used for accessing the configuration, status, and performance measurement data of a [%=DP.HW_SW_P9000_XP_abbrev%] storage system, and for invoking operations on the disk array. It translates function calls into sequences of low-level SCSI commands. See also [%=DP.DP_AbbCompanyName%] [%=DP.HW_SW_P9000_XP_special%] Agent.

rawdisk backup

See disk image backup.

RCU

See Remote Control Unit (RCU).

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.

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, and 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.

recovery files (Oracle specific term)

Recovery files are Oracle specific files that reside in the flash recovery area: the current control file, online redo logs, archived redo logs, flashback logs, control file autobackups, datafile copies, and backup pieces. See also flash recovery area.

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.

RecoveryInfo

When backing up Windows configuration files, [%=DP.DP_BriefProductName%] collects the information about the current system configuration (information on disk layout, volume, and network configuration). This information is needed for disaster recovery.

recycle or unprotect

A process that removes the data protection from all backed up data on a medium, allowing [%=DP.DP_BriefProductName%] 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 (RCU) ([%=DP.HW_SW_P9000_XP_full%] specific term)

An [%=DP.HW_SW_P9000_XP_full%] unit that acts as a slave device to the Main Control Unit (MCU) in the [%=DP.DP_AbbCompanyName%] CA [%=DP.HW_SW_P9000_XP_special%] or [%=DP.DP_AbbCompanyName%] CA+BC [%=DP.HW_SW_P9000_XP_special%] configuration. In bidirectional configurations, the RCU can also 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 or 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 perspective of a basic operating system, the complete physical disk containing a backup object is replicated. However, if a volume manager is used on a UNIX system, the whole volume/disk group containing a backup object is replicated. If partitions are used on a Windows system, the whole physical volume containing the selected partition 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 chain

Backup images that are needed to restore a backed up object to the state it was in at the selected point in time. In general, a restore chain of an object consists of its full backup image and one or more related incremental backup images.

restore session

A process that copies data from backup media to a client.

resync mode ([%=DP.HW_SW_P9000_XP_full%] VSS provider specific term)

One of two [%=DP.HW_SW_P9000_XP_abbrev%] VSS hardware provider operation modes. When the [%=DP.HW_SW_P9000_XP_abbrev%] provider is in the resync mode, the source volume (P-VOL) and its replica (S-VOL) are in the suspended mirror relationship after a backup. The maximum number of replicas (S-VOLs per a P-VOL) rotated is three provided that MU range is 0-2 or 0, 1, 2. Restore from a backup in such a configuration is possible only by re-synchronization of an S-VOL with its P-VOL. See also VSS compliant mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), MU number, and replica set rotation.

RMAN (Oracle specific term)

See Recovery Manager.

RSM

The [%=DP.DP_BriefProductName%] Restore Session Manager controls restore and object verification sessions. 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.

S

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 [%=DP.DP_BriefProductName%] 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) ([%=DP.HW_SW_P9000_XP_full%] specific term)

An internal disk (LDEV) of a disk array of the [%=DP.HW_SW_P9000_XP_full%] which is paired with another LDEV: a primary volume (P-VOL). It can act as a mirror of the P-VOL or as a volume to be used for the P-VOL's snapshot storage. An S-VOL is assigned a SCSI address different from the one used for the P-VOL. In an [%=DP.DP_AbbCompanyName%] CA [%=DP.HW_SW_P9000_XP_special%] configuration, the S-VOLs acting as mirrors can be used as failover devices in a MetroCluster configuration. See also primary volume (P-VOL) and Main Control Unit (MCU).

session

See backup session, media management session, and restore session.

session ID

An identifier of a backup, restore, object copy, object consolidation, object verification, 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 [%=DP.DP_BriefProductName%] 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 (Microsoft 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 and replica.

shadow copy provider (Microsoft 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 (Microsoft VSS specific term)

A collection of shadow copies created at the same point in time. See also shadow copy and replica set.

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 [%=DP.DP_BriefProductName%] Disk Agent installed.

Site Replication Service (Microsoft Exchange Server specific term)

The Microsoft Exchange Server 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. [%=DP.DP_BriefProductName%] 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, restore, object copy, object consolidation, object verification, and media management sessions. One binary file is created per session. The files are grouped by year and month.

SMI-S Agent (SMISA)

See [%=DP.DP_AbbCompanyName%] P6000 / [%=DP.DP_AbbCompanyName%] 3PAR SMI-S Agent.

snapshot ([%=DP.HW_SW_P4000_LH_full%], [%=DP.HW_SW_P6000_EVA_full%], [%=DP.HW_SW_P9000_XP_full%], and [%=DP.DP_AbbCompanyName%] 3PAR StoreServ Storage specific term)

A type of target volumes created using a specific replication technology. Depending on the disk array model and the chosen replication technique, a range of snapshot types with different characteristics is available. Basically, each snapshot may be either a virtual copy, still reliant upon the contents of the source volume, or an independent duplicate (clone) of the source volume. See also replica and snapshot creation.

snapshot backup

See ZDB to tape, ZDB to disk, and ZDB to disk+tape.

snapshot creation ([%=DP.HW_SW_P4000_LH_full%], [%=DP.HW_SW_P6000_EVA_full%], [%=DP.HW_SW_P9000_XP_full%], and [%=DP.DP_AbbCompanyName%] 3PAR StoreServ Storage specific term)

A replica creation process in which copies of the selected source volumes are created using storage virtualization technology. Such a replica is considered to be created at a particular point in time, and is immediately available for use. However, with certain snapshot types, a background data copying process continues to run on the disk array after the moment of the replica 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 [%=DP.HW_SW_P9000_XP_full%] specific term)

A type of target volumes created using a specific replication technology. A split-mirror replica provides independent duplicates (clones) of the source volumes. See also replica and split mirror creation.

split mirror backup ([%=DP.HW_SW_P9000_XP_full%] specific term)

See ZDB to tape, ZDB to disk, and ZDB to disk+tape.

split mirror backup (EMC Symmetrix specific term)

See ZDB to tape.

split mirror creation (EMC Symmetrix and [%=DP.HW_SW_P9000_XP_full%] 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 [%=DP.HW_SW_P9000_XP_full%] specific term)

A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is first copied from the backup media to a replica, and from the replica to the source volumes afterwards. 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 or registry (Informix Server specific term)

An Informix Server connectivity information file (on UNIX systems) or registry (on Windows systems) 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

A text file in the Unicode (UTF-16) format, generated during CONFIGURATION backup of a Windows system and stored on the Cell Manager. It contains system information required for installing and configuring the operating system on the target system in the case of a disaster. See also target system.

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 (SSEA)

See [%=DP.DP_AbbCompanyName%] [%=DP.HW_SW_P9000_XP_special%] Agent.

sst.conf file

The file /usr/kernel/drv/sst.conf is required on each [%=DP.DP_BriefProductName%] 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 [%=DP.DP_BriefProductName%] 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.

Storage Group (Microsoft Exchange Server specific term)

A collection of mailbox stores and public folder stores that share a set of transaction log files. Exchange Server manages each storage group with a separate server process.

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. Typically, these can be created or exist within a storage system such as a disk array.

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.

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 [%=DP.DP_BriefProductName%].

Sybase SQL Server (Sybase specific term)

The server in the Sybase “client-server” architecture. The 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.

SYMA

See EMC Symmetrix Agent.

synthetic backup

A backup solution that produces a synthetic full backup, an equivalent to a conventional full backup in terms of data, without putting stress on the production servers or the network. A synthetic full backup is created from a previous full backup and any number of incremental backups.

synthetic full backup

The result of an object consolidation operation, where a restore chain of a backup objects is merged into a new, synthetic full version of this object. A synthetic full backup is equivalent to a conventional full backup in terms of restore speed.

System Backup to Tape (SBT) (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 Recovery Data file

See SRD file.

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 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.

T

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 (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 database (Oracle specific term)

In RMAN, the target database is the database that you are backing up or restoring.

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 faulty 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 (Microsoft SQL Server 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 EMC 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

See 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 log table (Sybase specific term)

A system table in which all changes to the database are automatically recorded.

transportable snapshot (Microsoft 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).

U

unattended operation

See lights-out operation.

user account ([%=DP.DP_BriefProductName%] user account)

You can use [%=DP.DP_BriefProductName%] only if you have a [%=DP.DP_BriefProductName%] user account, which restricts unauthorized access to [%=DP.DP_BriefProductName%] and to backed up data. [%=DP.DP_BriefProductName%]

administrators create this account specifying a user logon name, the systems from which the user can log on, and a [%=DP.DP_BriefProductName%] user group membership. This is checked whenever the user starts the [%=DP.DP_BriefProductName%] user interface or performs specific tasks.

User Account Control (UAC)

A security component in Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012 that limits application software to standard user privileges until an administrator authorizes an increase in privilege level.

user disk quotas

NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. [%=DP.DP_BriefProductName%] backs up user disk quotas on the whole system and for all configured users at a time.

user group

Each [%=DP.DP_BriefProductName%] 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. [%=DP.DP_BriefProductName%] 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 [%=DP.DP_BriefProductName%] 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.

user_restrictions file

A file that restricts specific user actions, which are available to [%=DP.DP_BriefProductName%] user groups according to the user rights assigned to them, to be performed only on specific systems of the [%=DP.DP_BriefProductName%] cell. Such restrictions apply only to [%=DP.DP_BriefProductName%] user groups other than admin and operator.

V

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.

verify

A function that lets you check whether the [%=DP.DP_BriefProductName%] 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) ([%=DP.HW_SW_P6000_EVA_full%] specific term)

The firmware that manages all aspects of storage system operation, including communication with [%=DP.DP_AbbCompanyName%] Command View EVA through the HSV controllers. See also [%=DP.DP_AbbCompanyName%] Command View (CV) EVA.

Virtual Device Interface (Microsoft SQL Server specific term)

This is a Microsoft SQL Server programming interface that allows fast backup and restore of large databases.

virtual disk ([%=DP.HW_SW_P6000_EVA_full%] specific term)

A unit of storage allocated from a storage pool of a disk array of the [%=DP.HW_SW_P6000_EVA_full%]. A virtual disk is the entity that can be replicated using the snapshot functionality of such disk array. See also source volume and target volume.

virtual full backup

An efficient type of synthetic backup where data is consolidated using pointers instead of being copied. It is performed if all the backups (the full backup, incremental backups, and the resulting virtual full backup) are written to a single file library that uses distributed file medium format.

Virtual Library System (VLS)

A disk-based data storage device hosting one or more virtual tape libraries (VTLs).

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.

virtual tape

An archival storage technology that backs up data to disk drives in the same way as if it were being stored on tape. Benefits of virtual tape systems include improved backup and recovery speed and lower operating costs. See also Virtual Library System (VLS) and virtual tape library (VTL).

virtual tape library (VTL)

Data storage virtualization software that is able to emulate tape devices and libraries thus providing the functionality of traditional tape-based storage. See also Virtual Library System (VLS).

volser

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 (VSS).

VSS

See Microsoft Volume Shadow Copy Service (VSS).

VSS compliant mode ([%=DP.HW_SW_P9000_XP_full%] VSS provider specific term)

One of two [%=DP.HW_SW_P9000_XP_abbrev%] VSS hardware provider operation modes. When the [%=DP.HW_SW_P9000_XP_abbrev%] provider is in the VSS compliant mode, the source volume (P-VOL) and its replica (S-VOL) are in simplex, unpaired state after a backup. Therefore the number of replicas (S-VOLs per a P-VOL) rotated is not limited. Restore from a backup in such a configuration is possible only by switching of the disks. See also resync mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), and replica set rotation.

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.

W

Wake ONLAN

Remote power-up support for systems running in power-save mode from some other system on the same LAN.

Web reporting

The [%=DP.DP_BriefProductName%] functionality that allows you to view reports on backup, object copy, and object consolidation status and [%=DP.DP_BriefProductName%] 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

[%=DP.DP_BriefProductName%] 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. [%=DP.DP_BriefProductName%] can back up WINS server data as part of the Windows configuration.

writer (Microsoft 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.

X

XBSA Interface (Informix Server specific term)

ON-Bar and [%=DP.DP_ BriefProductName%] communicate with each other through the X/Open Backup Services Application Programmer's Interface (XBSA).

Z

ZDB

See zero downtime backup.

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 [%=DP.DP_ BriefProductName%] restore from tape, or with specific disk array families, 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 [%=DP.DP_ BriefProductName%] restore from tape. With specific disk array families, split mirror restore can also be used. See also zero downtime backup (ZDB), ZDB to disk, ZDB to disk+tape, instant recovery, 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.

Index