Hewlett Packard
Enterprise

HPE Data Protector

Software Version: 10.00

Disaster Recovery Guide

Document Release Date: June 2017

Software Release Date: June 2017

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise Development LP 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. HPE 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 HPE 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.

This product includes an interface of the 'zlib' general purpose compression library, which is Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.

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 software updates, go to https://softwaresupport.hpe.com/patches.

To verify that you are using the most recent edition of a document, go to https://softwaresupport.hpe.com/manuals.

This site requires that you register for an HPE Passport and sign in. To register for an HPE Passport ID, go to https://hpp12.passport.hpe.com/hppcf/login.do.

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

Support

Visit the HPE Software Support Online web site at https://softwaresupport.hpe.com.

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

HPE 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
- Access product documentation
- Manage support contracts
- Look up HPE 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 HPE Passport user and sign in. Many also require a support contract.

To register for an HPE Passport ID, go to https://hpp12.passport.hpe.com/hppcf/login.do.

To find more information about access levels, go to https://softwaresupport.hpe.com/web/softwaresupport/access-levels.

Contents

Chapter 1: Introduction 9
Data Protector disaster recovery overview 9
Disaster recovery phasesprocess 11
Disaster Recovery Methods 11
Manual disaster recovery method 13
Disaster recovery using disk delivery 13
Enhanced Automated Disaster Recovery (EADR) 14
One Button Disaster Recovery (OBDR) 14
HPE Data Protector Integrations and Disaster Recovery 15

Chapter 2: How to Prepare for Disaster Recovery 16
Planning 16
Consistent and relevant backups 17
Creating a consistent and relevant backup 18
Encrypted backups 18
Updating and Editing the System Recovery Data 19

Chapter 3: Disaster recovery on Windows systems 20
Enhanced Automated Disaster Recovery (EADR) 20
Overview 20
Prerequisites 20
Preparation for Enhanced Automated Disaster Recovery (Windows and Linux) 21
Prerequisites 21
Limitations 23
General preparations 25
Additional preparations for the Cell Manager 26
Saving a Recovery Set to the Cell Manager 27
Saving the recovery set file to the Cell Manager for all clients in the backup specification 27
Steps 27
Saving the recovery set file to the Cell Manager for a particular client in the backup specification 28
Preparing the Encryption Keys 29
Preparing a DR OS image 29
Steps 29
Recovering Windows Systems Using Enhanced Automated Disaster Recovery 31
Steps 31
Phase 1 31
Phase 2 34
Phase 3 36

One Button Disaster Recovery (OBDR) 36
Overview 36
Requirements 37
Limitations 38
Preparation for One Button Disaster Recovery (Windows and Unix) 38
Preparatory steps 39
Creating the Backup Specification for One Button Disaster Recovery 40
Prerequisites 40
Limitations 40
Creating a backup specification for OBDR 41
Steps 41
Modifying an OBDR backup specification to use disk image backup 42
Steps 42
Preparing the Encryption Keys 43
Recovering Windows Systems Using One Button Disaster Recovery 43
Prerequisites 43
Steps 43
Phase 1 43
Phase 2 47
Phase 3 47
Advanced tasks 48
Disaster Recovery of Microsoft Cluster Server 48
About Disaster Recovery of a Microsoft Cluster Server 48
Possible scenarios 48
Preparation for Microsoft Cluster Server Disaster Recovery Specifics 48
EADR specifics 49
OBDR specifics 49
Recovering a Microsoft Cluster Server 49
At least one of the nodes is up and running 49
Prerequisites 49
All nodes in the cluster have experienced a disaster 50
Prerequisites 50
Steps 50
Merging P1S Files for Microsoft Cluster Server 51
Windows 51
UNIX 51
Steps 51
Restoring Original Hard Disk Signatures on Windows Systems 52
Restoring original hard disk signatures on Windows 52
Obtaining original hard disk signatures 52
Example of Hard Disk Signatures in the SRD File 53
Restoring the Data Protector Cell Manager specifics 53
Making IDB consistent (all recovery methods) 53
Enhanced Automated Disaster Recovery specifics 53
Restoring Internet Information Server Specifics 54
Requirements 54

Steps 54
Editing the kb.cfg File 54
Editing the SRD Files 55
EADR/OBDR 56
Steps 56
Windows systems 56
Linux systems 57
Example of Editing the SRD File 58
Changing the MA client 58
Changing the backup device 58
Windows BitLocker Drive Encryption 59
Limitation 59
Steps 59

Chapter 4: Disaster recovery on UNIX systems 61

Manual Disaster Recovery (MDR) 61
Overview 61
Preparation for Manual Disaster Recovery (HP-UX Cell Manager) 61
One-time preparation 62
HP-UX systems 62
Backing up the system 62
Installing and Configuring HP-UX Systems Manually (Cell Manager) 63
Steps 63
Phase 1 63
Restoring System Data Manually (HP-UX Cell Manager) 63
Prerequisites 63
Steps 63
Phase 2 63
Phase 3 63
Preparation for Manual Disaster Recovery (HP-UX Client) 64
Using custom installation medium (Golden Image) 64
Creating a Golden Image 64
Recovering an HP-UX Client 66
Recovery using a Golden Image 66
On the client 66
Steps 66
On the Ignite-UX Server 67
Steps 67
Recovery from the bootable backup tape 67
Steps 67
Recovery from the network 67
Using system recovery tools (make_tape_recovery, make_net_recovery) 67
Prerequisites 68
Creating an archive using make_tape_recovery 68
Creating an archive using make_net_recovery 69

Disk Delivery Disaster Recovery (DDDR) 69
Overview 69
Limitations 70
Preparation for Disk Delivery Disaster Recovery of UNIX Clients 70
One-time preparation 70
HP-UX Example 71
Solaris Example 71
AIX 71
Preparing the auxiliary disk 71
Backing up the system 71
Creating the Backup Specification for Disaster Recovery of a UNIX Client 72
Steps 72
Installing and Configuring a UNIX Client Using DDDR 73
Prerequisites 73
Steps 73
Restoring System Data Using DDDR (UNIX Client) 73
Prerequisites 74
Steps 74
Phase 2 74
Phase 3 74
Enhanced Automated Disaster Recovery (EADR) 74
Overview 75
Requirements 75
Limitations 76
Disk and partition configuration 77
Preparation for Enhanced Automated Disaster Recovery 77
General preparations 78
Additional preparations for the Cell Manager 78
Saving a Recovery Set to the Cell Manager 78
Saving the recovery set to the Cell Manager for all clients in the backup specification 79
Steps 79
Saving the recovery set to the Cell Manager for a particular client in the backup specification 80
Preparing the Encryption Keys 80
Preparing a DR OS image 80
Steps 81
Recovering Linux Systems Using EADR 82
Prerequisites 82
Steps 82
Phase 1 82
Phase 2 84
Phase 3 84
One Button Disaster Recovery (OBDR) 84
Overview 85
Requirements 85

Limitations ... 86
Disk and partition configuration ... 87
Preparation for One Button Disaster Recovery ... 87
Preparatory steps ... 87
Creating the Backup Specification for One Button Disaster Recovery ... 87
Prerequisites ... 87
Limitations ... 88
Creating a backup specification for OBDR ... 88
Steps ... 88
Preparing the Encryption Keys ... 89
Recovering Linux Systems Using OBDR ... 89
Prerequisites ... 89
Steps ... 90
Phase 1 ... 90
Phase 2 ... 91
Phase 3 ... 92

Appendix A: Example Preparation Tasks ... 93
Example of Moving Kill Links on HP-UX 11.x ... 93
Example of the Disaster Recovery Preparation Table for Windows ... 93

Send documentation feedback ... 95

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.

For detailed information on disaster recovery, see the *HPE Data Protector Disaster Recovery Guide*.

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 phasesprocess

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

1. Phase 0
2. Phase 1
3. Phase 2
4. Phase 3

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 https://softwaresupport.hpe.com/.

> **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.

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	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.

information on the original system to enable installation and configuration of the DR OS.			
See Manual Disaster Recovery (MDR) , on page 61.			
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 69.			
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 20 or Enhanced Automated Disaster Recovery (EADR) , on page 74.			
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 36 or One Button Disaster Recovery (OBDR) , on page 84.			

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 *HPE 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 69.

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 *HPE 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.

HPE recommends to restrict access to backup media, DR images, SRD files and disaster recovery CDs and USB drives storing DR OS data

HPE 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 *HPE 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 Protector Disk 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.

Disaster recovery uses data from the btrfs sub volumes and volumes backed up through the root of the file system (cross file system boundary) to create a Disaster Recovery ISO image, and perform recovery and restore. This implies that all system, profile, and relevant user data must be included in the backup of the `/ (root)` file system object. All separately backed up data (one using the `OB2_SHOW_BTRFS_MOUNTS`) can be used only for the regular Disk Agent file system restore operations and not for the recovery process. This applies to the Linux operating system only.

> **NOTE:**
> Data Protector includes data from the manually created btrfs snapshots.

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 HPE 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.

Chapter 3: Disaster recovery on Windows systems

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).

This section describes the steps/tasks that you need to perform after you encounter a disaster recovery situation.

Overview

Ensure that you have performed all the general preparation steps that are mentioned in the preparation chapter. The general steps using the Enhanced Automated Disaster Recovery method for a Windows client are:

1. **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.
2. **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.
3. **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:

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

- For disaster recovery of the Cell Manager, you should have a valid Internal Database backup image that is newer than the filesystem backup image.
- 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).
- Ensure that you have enabled the Automount feature. The Automount feature ensures that all the volumes (without a mountpoint) are online. When the Automount is disabled, all the volumes without the drive letter are offline during the booting process. Therefore the System Reserve partition will not have access to the drive letter, and this may result in the failure of the disaster recovery procedure.

  If you need to disable the Automount feature, then ensure that you have mounted the System Reserve partition.
- On Windows Server 2008 and later releases, at least one volume must be an NTFS volume.
- For a remote restore, the network must be available when you boot DR OS image.

Preparation for Enhanced Automated Disaster Recovery (Windows and Linux)

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.

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 *HPE Data Protector Installation Guide*.
- On Windows Server 2008 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.
- Ensure that you have enabled the Automount feature. The Automount feature ensures that all the volumes (without a mountpoint) are online. When the Automount is disabled, all the volumes without the drive letter are offline during the booting process. Therefore the System Reserve partition will not have access to the drive letter, and this may result in the failure of the disaster recovery procedure.

  If you need to disable the Automount feature, then ensure that you have mounted the System Reserve partition.
- 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 Server 2008 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 Server 2008:

  Automated Installation Kit (AIK) for Windows Server 2008

  Windows Server 2008 R2:

  - Windows Automated Installation Kit (AIK) Supplement for Windows Server 2008 R2 SP1)

  Windows Server 2012:

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

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

  Windows Server 2012 R2:

  - Assessment and Deployment Kit (ADK 1.1) for 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 Server 2008 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 Server 2008 and later releases, install the IIS 6 Metabase Compatibility package.
- During the creation of recovery ISO image for RedHat 7 client, recovery media creation host must have **squashfs-tools** installed in order to create recovery ISO image successfully.

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 Server 2008, Windows Server 2008 R2 systems (on all supported platforms), Windows Server 2012, and later releases.
- VSS disk image backup of the logical volumes can be used for disaster recovery only on Windows Server 2008 and later releases.
- On Windows Server 2008 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 Server 2008 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.
- AUTODR does not support recovery of btrfs on multiple devices (various btrfs raid configurations) as they are not supported by SLES 11.3.
- The current btrfs tools on SLES 11.3 do not set the UUID on a newly created btrfs file system. Therefore, AUTODR cannot set the same UUID on btrfs file systems during recovery as done for backup.

  If you mount the btrfs file systems by UUID instead of a device name, you need to manually edit the `/etc/fstab` file after restore. This needs to be done to reflect the new and correct UUIDs of the recovered btrfs devices. The same is applicable for the GRUB configuration, so avoid the UUID for the root device and replace the device by name.

  After a system recovery, the btrfs has different UUIDs than the ones during backup. If another recovery is performed from backups created before the last recovery of the system, the AUTODR tries to identify healthy btrfs file systems and skips recreating them.

- AUTODR can only map the btrfs device configurations in backup to btrfs devices in the present system being recovered by UUID. It can skip recovering wrong devices or recreated ones.

  To avoid this, recover btrfs file systems only from backups created after the last system recovery or destroy manually present btrfs file systems before a system recovery. The same is applicable for btrfs file systems manually recreated by users after the last backup.

  > **NOTE:**
  > Data Protector warns users of this before starting the recovery process.

- btrfs snapshots can be backed up but restored only as oridinary sub volumes. During such an instance, none of the data will be shared between the snapshot and sub volume from where the snapshot is created. The overall Copy On Write (COW) relationship between the parent and its snapshot is lost. Therefore, in some cases, restore of complete data set is not possible, as data from the snapshot is duplicated and runs out-of-space on the underlying device during restore.
- Only data from the mounted btrfs sub volumes are protected. Consider child sub volumes accessible from an OS file system interface and parent sub volume being mounted. In such a case, the sub volumes are not protected, as Disk Agent (DA) detects them as a different file system and skips them because they do not have a dedicated mount point.
- Sub volumes mounted using the `subvolid` (refer to the *btrfs documentation*) mount option in `/etc/fstab` file can be skipped from mount in the recovered system or mounted on a wrong mount point, as `subvolid` of recovered sub volume need not be the same as the one during backup. Even though all sub volumes are recreated, the HPE Data Protector skips restore in such sub volumes or data can be restored in wrong ones.

  > **NOTE:**
  > Use the `subvol` option in `fstab` instead of `subvolid`.

Disk and partition configuration

- EADR is not supported for shared dynamic disks residing in Windows clusters.
- If the System Reserved volume resides on the dynamic disk, the volume will not be indicated by the yellow colored icon, instead it will be indicated as a green colored icon in the Data Protector GUI.
- When performing a disaster recovery with dynamic disks, all the disks need to be cleaned up before starting the EADR.
- After the EADR session, all the volumes will be recreated, but only volumes that are inside the recovery scope will be restored.
- 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.
- Recovery of operating systems which have been deployed using the HPE Intelligent Provisioning tool (v.1.4 and v.1.5) may fail because of incorrect MBR partition information.
- Sparse files are restored to their full size. This may result in that the target volume runs out of space.
- Storage Spaces configurations where physical disks do not entirely belong to a storage pool are not supported.

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 the next page..

   See the *HPE 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 Server 2008 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 48.

- *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 *HPE Data Protector Integration Guide*.

  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. Before performing a disaster recovery of a client, run the following command on the Cell Manager for an online recovery, and on the media hosts for an offline recovery:
   `omnicc -secure_comm -configure_for_dr <hostname_of_client_being_recovered>`
3. Post online recovery of a client, run the following command on the Cell Manager:
   `omnicc -secure_comm -configure_peer <client_host_name> -overwrite`
4. After a disaster occurs, use the EADR Wizard to convert the DR image into a disaster recovery CD ISO image.
   ***Windows Server 2008 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.
5. 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.
6. Execute a disaster recovery test plan.
7. 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.

- Before performing disaster recovery for the cell manager, run the following command on the media host used for the disaster recovery:
  `omnicc -secure_comm -configure_for_dr <cell_manager_hostname>`
- After the recovery is complete, run the following command on the media hosts:
  `omnicc -secure_comm -configure_peer <cell_manager_hostname>`

- 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 HPE 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 HPE 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 Server 2008 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.

WinFS options tab

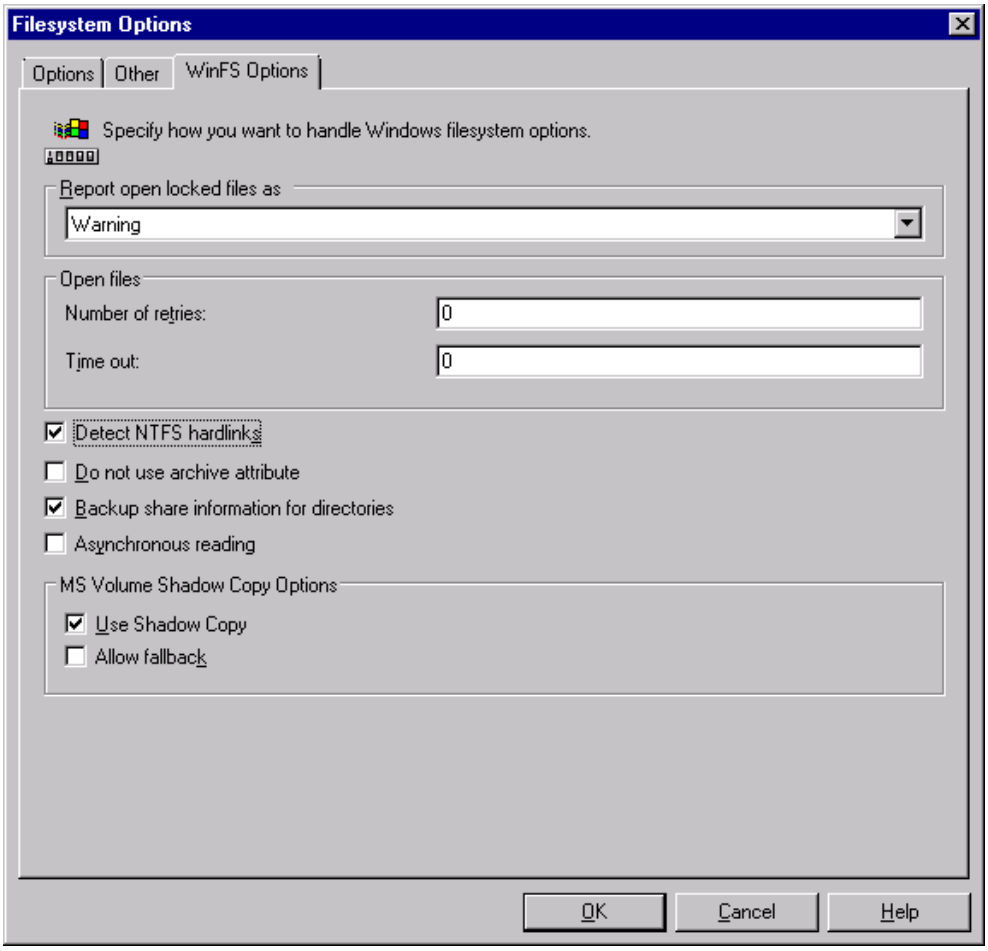

Saving the recovery set file to the Cell Manager for a particular client in the backup specification

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 HPE 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 Server 2008 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

Before 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, select the client for which you would like to prepare the DR OS image from the **Host to be recovered** drop down list and click **Validate** to validate the client.

   > **NOTE:**
   > The validated client gets added to the **Host to be recovered** drop down list.

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 Server 2008 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 *HPE Data Protector Product Announcements, Software Notes, and References*.

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 Server 2008 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.

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

   Adding a user account

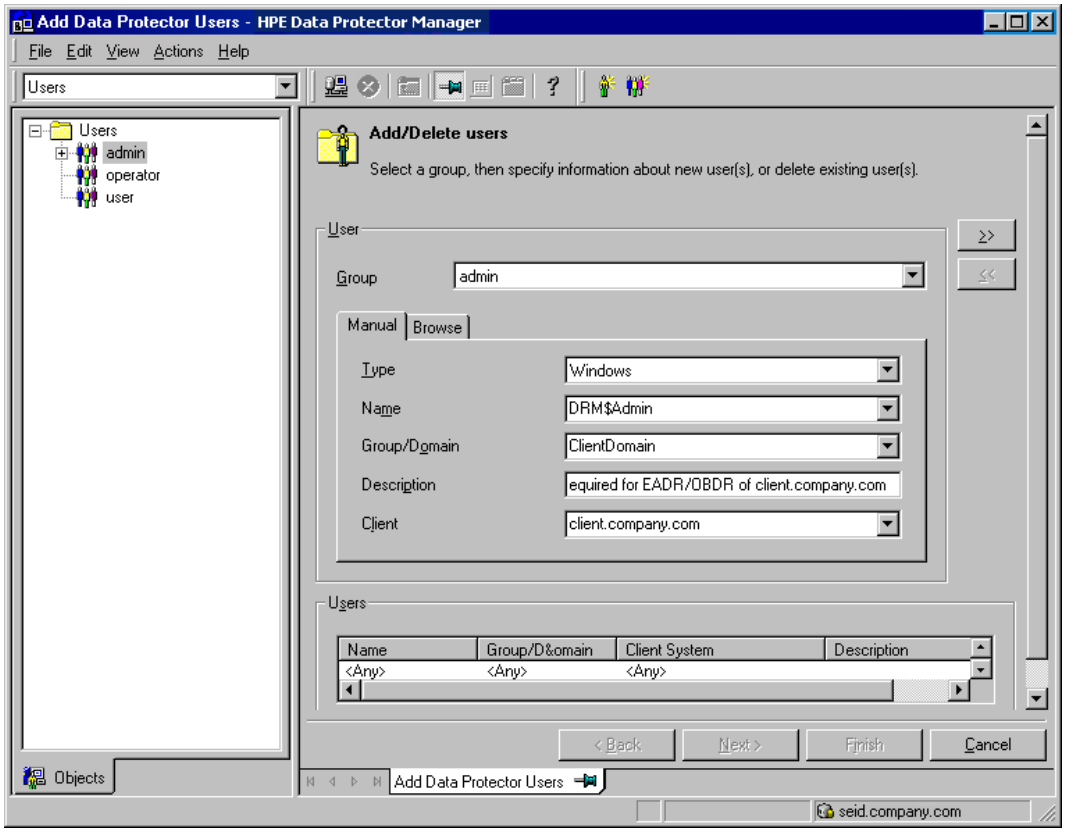

2. Boot the client system from the disaster recovery CD, the bootable USB drive, or the bootable network image of the original system. If you are starting the target system from a disaster recovery CD, ensure that no external USB disks (including USB keys) are connected to the system before you start 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. Select the scope of the recovery and recovery options. The following steps differ depending on the operating system:

   Windows Server 2008 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 disk, 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

Phase 2

3. 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 Server 2008 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.

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.
- `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.

4. If the DR OS image is password protected, provide the password and continue the recovery.
5. 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.
6. If the information in the SRD file is not up to date (for example, because you changed the backup device after the disaster) and you are performing an offline recovery, edit the SRD file before continuing with this procedure.
7. 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 Server 2008 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.
8. 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.

9. 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 https://softwaresupport.hpe.com/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

Ensure that you have performed all the general preparation steps that are mentioned in the preparation chapter.The general steps using the One Button Disaster Recovery method for a Windows client are:

1. **Phase 1**

   Boot from the recovery tape and select the scope of recovery.
2. **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.
3. **Phase 3**

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

> **IMPORTANT:**
> HPE 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 *HPE Data Protector Installation 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 HPE Tape Hardware Compatibility Table and the latest support matrices at https://softwaresupport.hpe.com/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.
- 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.
- 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.
- To create a DR OS image for Windows Server 2008 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 Server 2008:

  Automated Installation Kit (AIK) for Windows Server 2008

  Windows Server 2008 R2:

  - Windows Automated Installation Kit (AIK) Supplement for Windows Server 2008 R2 SP1)

  Windows Server 2012:

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

  Windows Server 2012 R2:

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

- To back up the IIS configuration object residing on a 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.
- VSS disk image backup of the logical volumes can be used for disaster recovery only on Windows Server 2008 and later releases.
- On Windows Server 2008 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 Server 2008 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 (Windows and Unix)

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 *HPE 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 Server 2008 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 *HPE Data Protector Integration Guide*. 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. Before performing a disaster recovery of a client, run the following command on the cell manager for an online recovery and on the media hosts for an offline recovery:
   ```
   omnicc –secure_comm –configure_for_dr <hostname_of_client_being_recovered>
   ```

4. Post online recovery of a client, run the following command on the Cell Manager:
   `omnicc –secure_comm –configure_peer <client_host_name> -overwrite`
5. Execute a disaster recovery test plan.
6. 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 Server 2008 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 Server 2008 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.

   Optionally, select the backup options.

   > **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.

   Linux: Select backup options. For more details on available options, see the *HPE Data Protector Help* index: “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.
10. a. 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: 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.

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 *HPE 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 Server 2008 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.

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

Adding a user account

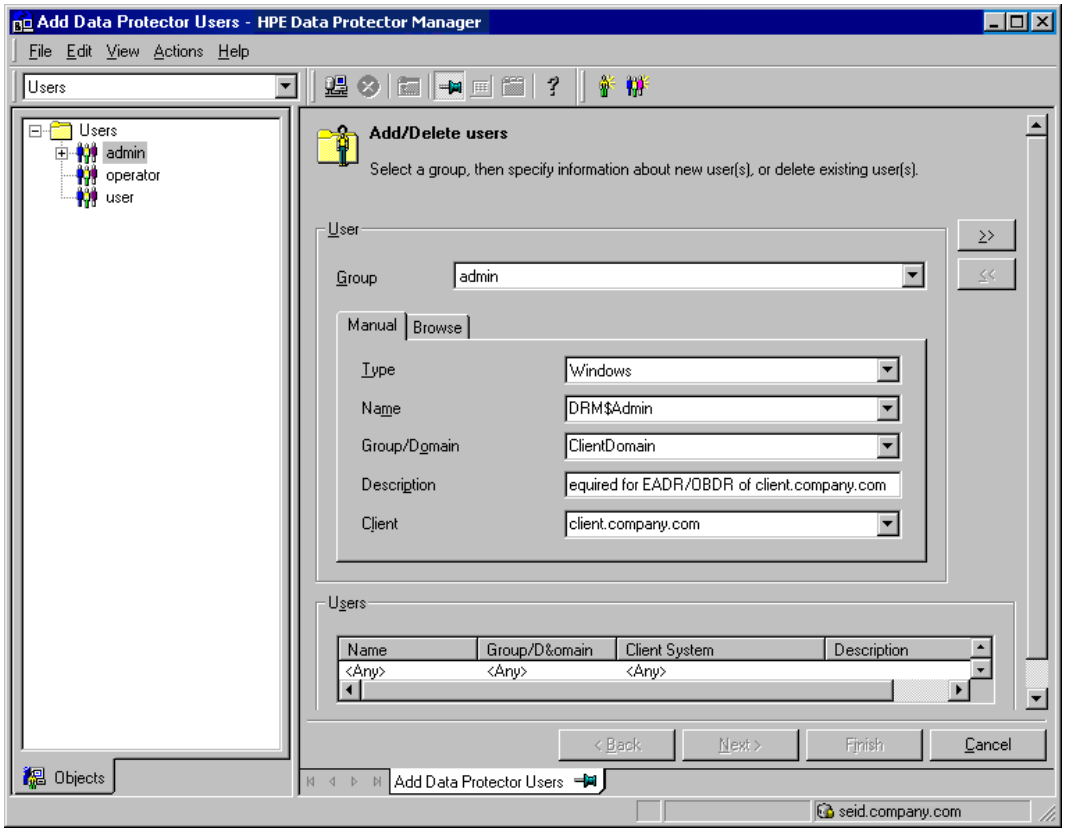

2. Insert the tape containing the image file and your backed up data into an OBDR device.
3. Shut down the target system and power off the tape device. Ensure that no external USB disks (including USB keys) are connected to the system before you start the recovery procedure.
4. Power the target system 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 Server 2008 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

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 Server 2008 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 *HPE 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.

Steps

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

   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 HPE 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 HPE 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 HPE 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:
HPE 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:

Requirements

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

Perform the following steps:

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 HPE.

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.

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.

   The Install Only option in the Disaster Recovery Wizard

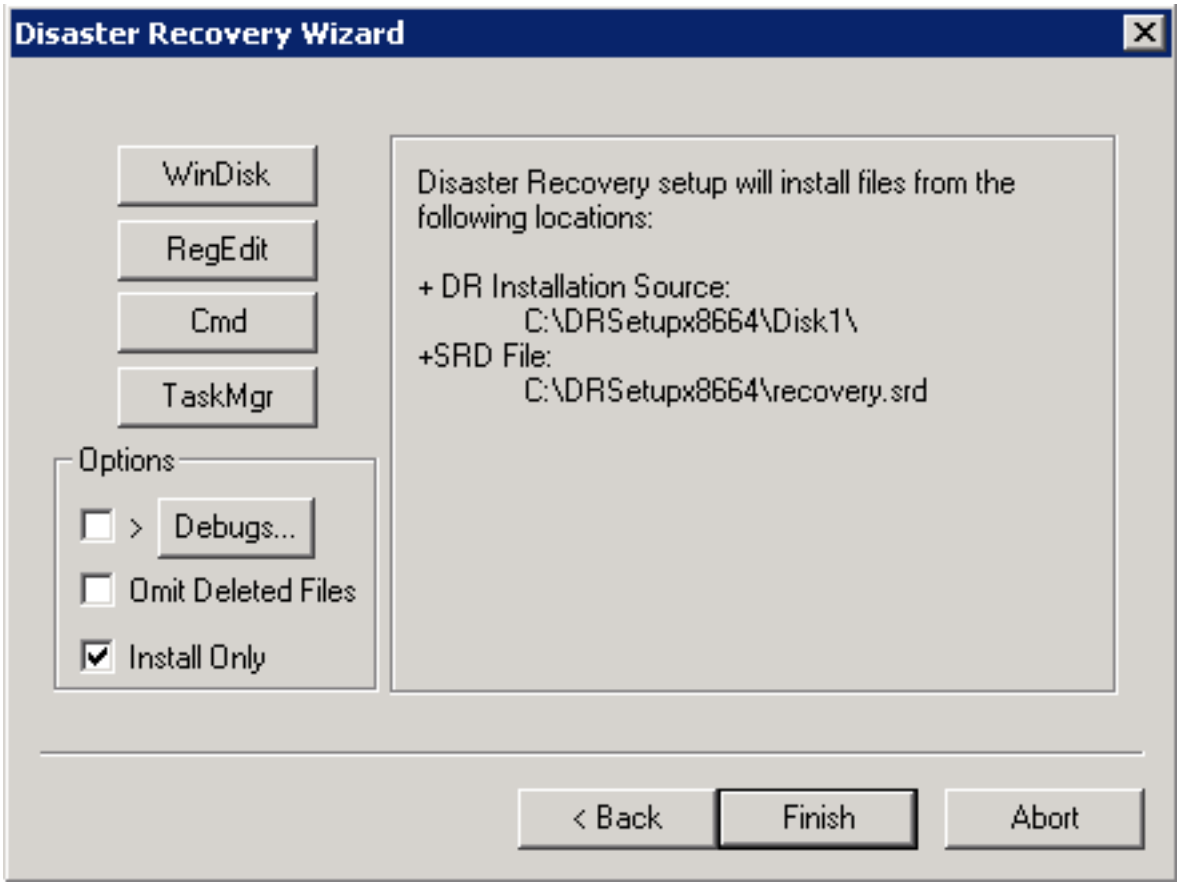

2. Select **Omit Deleted Files** option. This option enables removal of deleted files between successive incremental backups at restore time. If specified the `omnidr` binary will forward the same option to Data Protector restore tools (`omnir` and `omniofflr`) in case of incremental backup. The option 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 *HPE 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 HPE 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 HPE 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 Server 2008 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 Server 2008 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.

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_1"
     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.

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 P`*`number`*.
4. The recovery process starts automatically.
5. After the recovery has completed successfully, restore additional user and application data using the standard Data Protector restore procedure.

Recovery from the network

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.

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`.

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.

Ensure that you have performed all the general preparation steps that are mentioned in the preparation chapter. The general steps using an auxiliary disk for a UNIX client are:

1. **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.
2. **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.

3. **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 *HPE 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 *HPE 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, HPE 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 *HPE 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 https://softwaresupport.hpe.com/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:**
> HPE recommends to restrict access to backup media, recovery set files, SRD files, and disaster recovery CDs.

Overview

Ensure that you have performed all the general preparation steps that are mentioned in the preparation chapter. The general steps using the Enhanced Automated Disaster Recovery method for a Linux client are:

1. **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.
2. **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.
3. **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 *HPE Data Protector Installation 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.
- 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.
- AUTODR does not support recovery of btrfs on multiple devices (various btrfs raid configurations) as they are not supported by SLES 11.3.
- The current btrfs tools on SLES 11.3 do not set the UUID on a newly created btrfs file system. Therefore, AUTODR cannot set the same UUID on btrfs file systems during recovery as done for backup.

If you mount the btrfs file systems by UUID instead of a device name, you need to manually edit the `/etc/fstab` file after restore. This needs to be done to reflect the new and correct UUIDs of the recovered btrfs devices. The same is applicable for the GRUB configuration, so avoid the UUID.

After a system recovery, the btrfs has different UUIDs than the ones during backup. If another recovery is performed from backups created before the last recovery of the system, the AUTODR tries to identify healthy btrfs file systems and skips recreating them.

- AUTODR can only map the btrfs device configurations in backup to btrfs devices in the present system being recovered by UUID. It can skip recovering wrong devices or recreated ones.

  To avoid this, recover btrfs file systems only from backups created after the last system recovery or destroy manually present btrfs file systems before a system recovery. The same is applicable for btrfs file systems manually recreated by users after the last backup.

  > **NOTE:**
  > Data Protector warns users of this before starting the recovery process.

- btrfs snapshots can be backed up but restored only as oridinary sub volumes. During such an instance, none of the data will be shared between the snapshot and sub volume from where the snapshot is created. The overall Copy On Write (COW) relationship between the parent and its snapshot is lost. Therefore, in some cases, restore of complete data set is not possible, as data from the snapshot is duplicated and runs out-of-space on the underlying device during restore.
- Only data from the mounted btrfs sub volumes are protected. Consider child sub volumes accessible from an OS file system interface and parent sub volume being mounted. In such a case, the sub volumes are not protected, as Disk Agent (DA) detects them as a different file system and skips them because they do not have a dedicated mount point.
- Sub volumes mounted using the `subvolid` (refer to the *btrfs documentation*) mount option in `/etc/fstab` file can be skipped from mount in the recovered system or mounted on a wrong mount point, as `subvolid` of recovered sub volume need not be the same as the one during backup. Even though all sub volumes are recreated, the HPE Data Protector skips restore in such sub volumes or data can be restored in wrong ones.

  > **NOTE:**
  > Use the `subvol` option in `fstab` instead of `subvolid`.

- EADR of systems with Fibre Channel over Ethernet (FCoE) LUNs and Fibre Channel over Ethernet (FCoE) SAN boot are not supported..

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 *HPE 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 HPE 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 HPE 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**.

   Other options tab

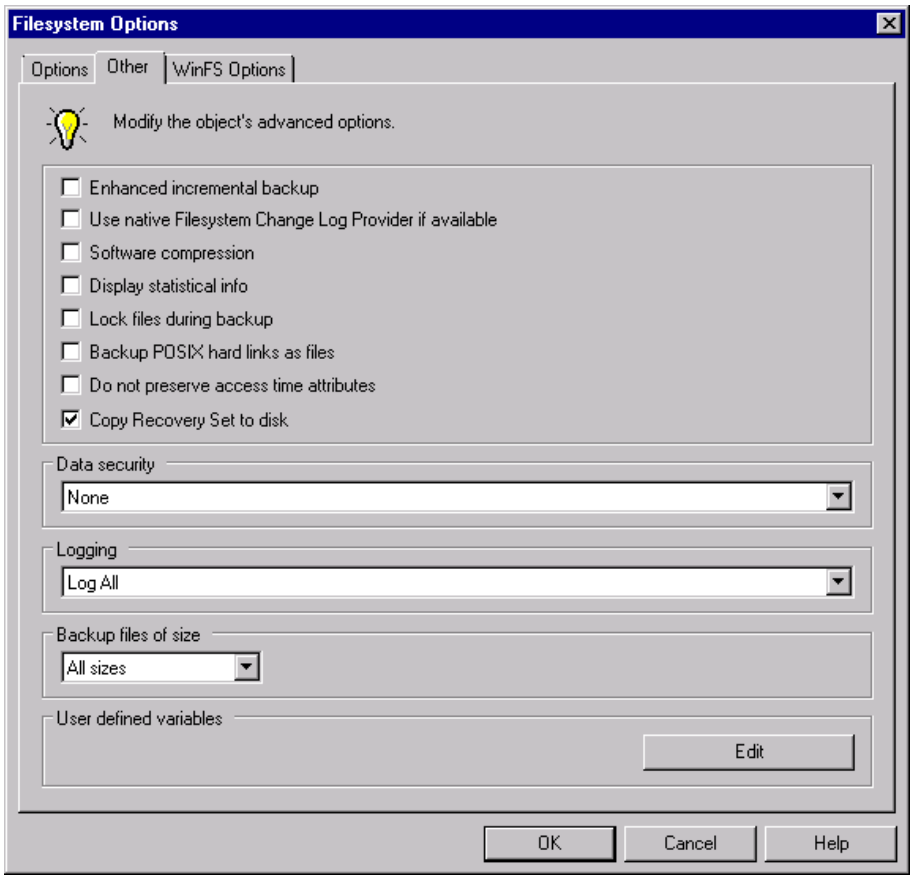

Saving the recovery set to the Cell Manager for a particular client in the backup specification

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 HPE 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

Before 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, select the client for which you would like to prepare the DR OS image from the **Host to be recovered** drop down list and click **Validate** to validate the client.

   > **NOTE:**
   > The validated client gets added to the **Host to be recovered** drop down list.

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 *HPE 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 HPE Data Protector Help index: “adding Data Protector users”.

   Adding a user account

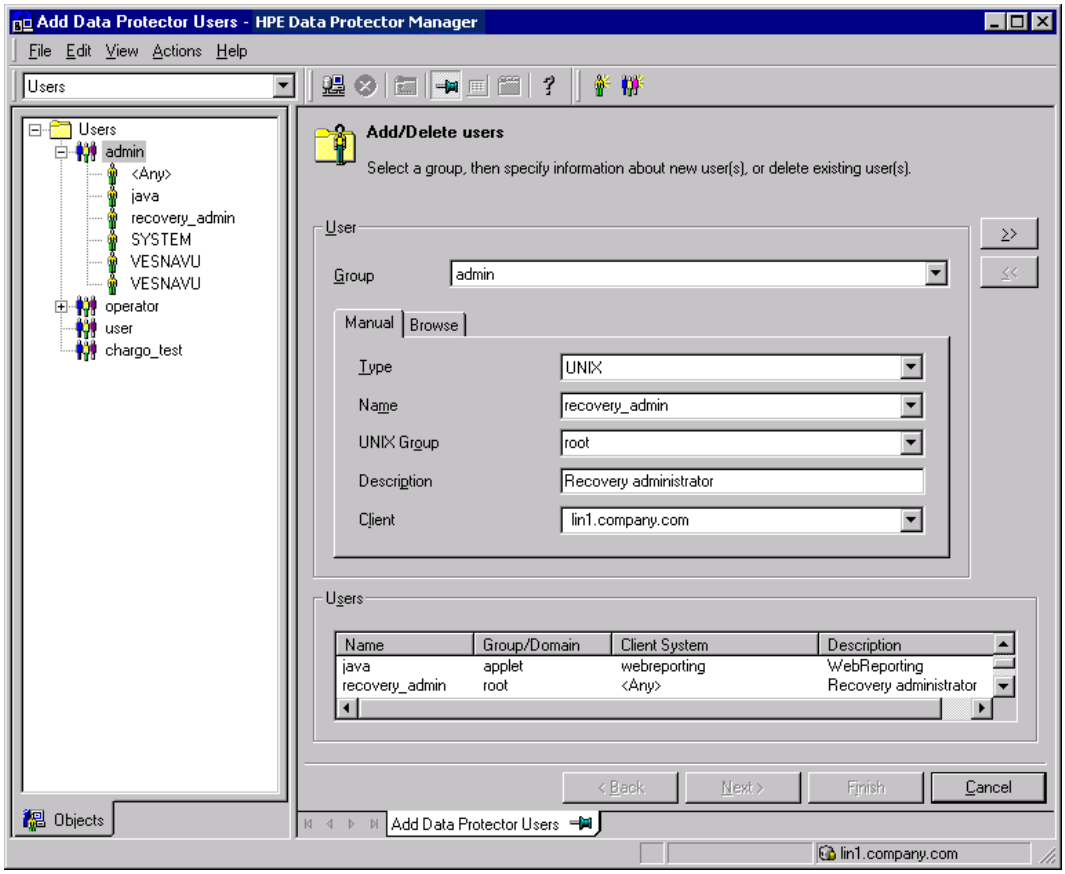

2. Boot client system from the disaster recovery CD of the original system.
3. Press **Enter** when the following message is displayed: Press Enter to boot from Recovery CD.
4. The DR OS is loaded first into memory and then the scope menu is displayed. Select the scope of recovery. There are four different 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.

> **NOTE:**
> All BTRFS volumes and sub volumes are recovered by Disaster Recovery irrespective of the selected recovery scope (default, minimal, or full recovery).

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**.

   > **NOTE: Note**: Ensure that the Cell Manager and Media (backup) host is reachable. If not, you may have to modify the NIC and MAC addresses. For more information, see Cell Manager and RMA hosts are not responding.

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 https://softwaresupport.hpe.com/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

Ensure that you have performed all the general preparation steps that are mentioned in the preparation chapter. The general steps using the One Button Disaster Recovery method for a Windows client are:

1. **Phase 1**

   Boot from the recovery tape and select the scope of recovery.
2. **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.
3. **Phase 3**

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

> **IMPORTANT:**
> HPE 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 *HPE Data Protector Installation 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 HPE Tape Hardware Compatibility Table and the latest support matrices at https://softwaresupport.hpe.com/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.
- 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 *HPE 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 *HPE 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 *HPE 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.

   HPE 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 *HPE 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 HPE Data Protector Help index: “adding Data Protector users”.

   Adding a user account

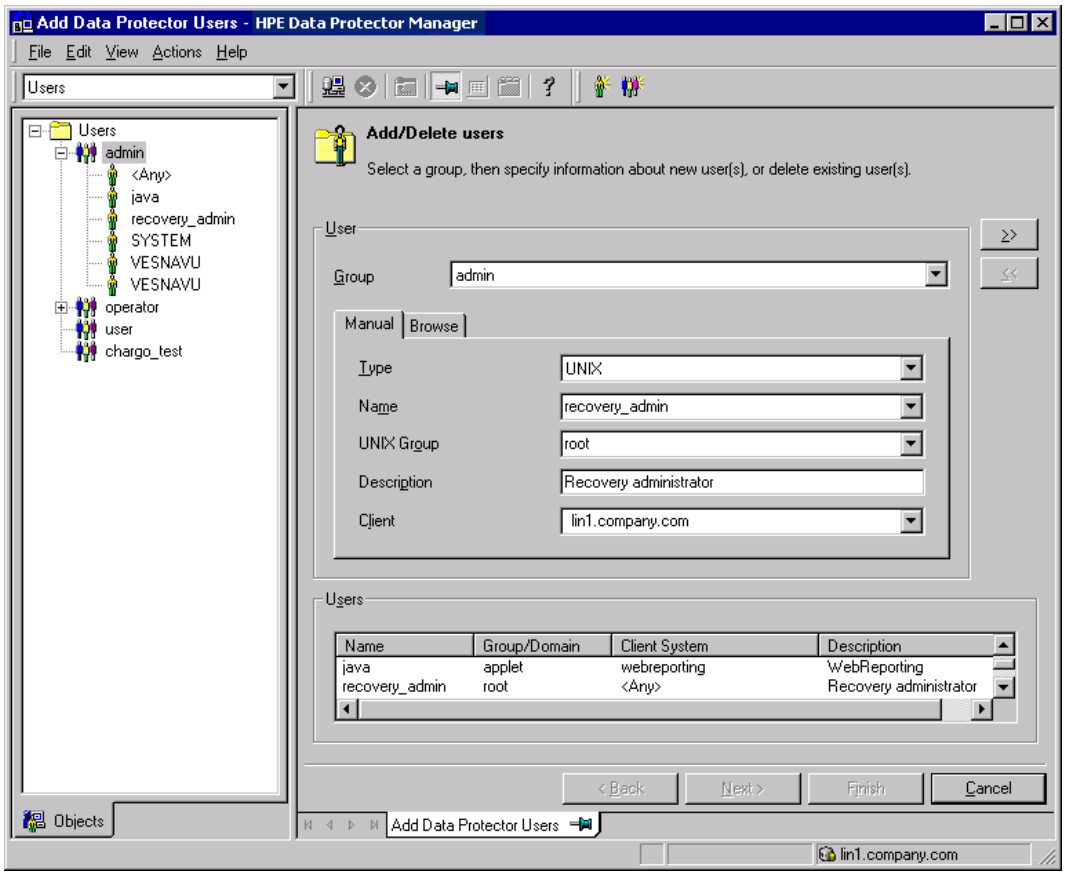

2. Insert the tape containing the image file and your backed up data into an OBDR device.
3. Shut down the target system and power off the tape device.

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.

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	

Send documentation feedback

If you have comments about this document, you can contact the documentation team by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Disaster Recovery Guide (HPE Data Protector 10.00)

Add your feedback to the email and click **Send**.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to AutonomyTPFeedback@hpe.com.

We appreciate your feedback!