

HP Data Protector

Software Version: 8.10

Granular Recovery Extension User Guide for Microsoft Exchange Server

Document Release Date: November 2016

Software Release Date: November 2016



Legal Notices

Warranty

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

The information contained herein is subject to change without notice.

Restricted Rights Legend

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

Copyright Notice

© Copyright 2016 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

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

UNIX® is a registered trademark of The Open Group.

Documentation Updates

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

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

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

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

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

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

Support

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

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

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

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

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

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

To find more information about access levels, go to:

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

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

Contents

Contents	3
Chapter 1: Introduction	7
Granular Recovery Extension Documentation set	7
Backup	8
Restore and recovery	8
Chapter 2: Installation	11
Prerequisites	11
Supported environments	12
Installing the extension	13
Removing the extension	14
Chapter 3: Configuration	15
Meeting Data Protector configuration requirements for the Granular Recovery Extension	15
Configuring the Granular Recovery Web service port	15
Configuring user account for the Granular Recovery Extension	15
Data Protector user rights	15
Other necessary privileges	17
Privileges for executing Exchange Management cmdlet operations	17
Chapter 4: Backup	19
Chapter 5: Restore and recovery	20
Limitations	20
Considerations	20
Restore and recovery flow	22
Opening the HP Data Protector Granular Recovery Extension GUI	24
Remote powershell configuration	25
Importing mailbox databases	26
Mounting databases	32
Starting recovery	33
Dismounting databases	37
Removing databases	37

Changing settings	38
Changing the retention period	39
Chapter 6: Command line reference	41
SYNOPSIS	42
DESCRIPTION	43
OPTIONS	43
Examples	47
Changing Granular Recovery Extension settings	47
Restoring a mailbox database from Data Protector backup	47
Listing mailbox database information	48
Changing the retention period	49
Mounting a mailbox database	49
Dismounting a mailbox database	49
Searching a mailbox	49
Recovering items to the original location	50
Recovering items to another location	50
Removing sessions	51
Removing recovery databases	51
Chapter 7: Troubleshooting	53
Before you begin	53
Debugging	53
Enabling debugging option	53
Known issues and workarounds	54
Search Criteria Results page remains empty after at least one search keyword is entered	54
Manual removal of temporary mailboxes created by the extension	54
Search for mailbox items fails and reports an error	55
Mailboxes are missing from the list in the Import from Backup wizard	55
Mounting a restored database fails	56
Interprocess communication error being reported by the GUI	56
An Exchange GRE recovery or restore operation fails due to insufficient permission	56
The message Adding snap-in to console... is displayed for a long time	57

The About HP Data Protector Granular Recovery Extension for Microsoft Exchange Server does not display the product build number	58
Glossary	60
Index	101
We appreciate your feedback!	103

Chapter 1: Introduction

This guide describes the HP Data Protector Granular Recovery Extension for Microsoft Exchange Server 2010 and Microsoft Exchange Server 2013 (hereafter both referred to as **Microsoft Exchange Server** unless differences are pointed out).

The Granular Recovery Extension for Microsoft Exchange Server (**the extension**) does not provide you with any backup solution. Use the Data Protector Microsoft Exchange Server 2010 integration to back up Microsoft Exchange Server 2010 mailbox and public folder databases or Microsoft Exchange Server 2013 mailbox databases (**databases**). Use the extension to restore Microsoft Exchange Server mailbox database files and to recover Microsoft Exchange Server single items or complete mailboxes.

Thus, the extension enables you to recover individual mailbox items, such as e-mail folders, calendar, contacts, or notes, with no need to recover the whole Microsoft Exchange Server mailbox or the entire mailbox database.

Granular Recovery Extension Documentation set

- **Electronic PDF format**

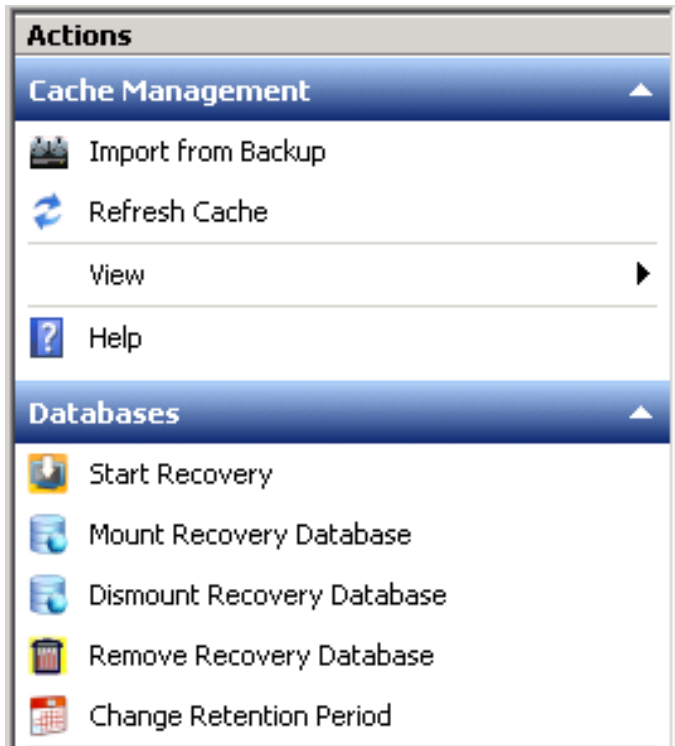
The HP Data Protector Granular Recovery Extension User Guide for Microsoft Exchange Server provides information specific to this extension:

- For detailed information about Data Protector specifics, see the Data Protector "[Documentation set](#)".
- For detailed information about Microsoft Exchange Server specifics, refer to the official Microsoft Exchange Server documentation.

- **Help**

To complete the information presented in this guide available in the electronic PDF format, the Granular Recovery Extension for Microsoft Exchange Server provides the context-sensitive (F1) Help integrated in the Microsoft Management Console (MMC). The Help explains the pages and options available in the Granular Recovery Extension Graphical User Interface (GUI). You can access the Help by pressing F1, or by clicking the question mark (? Help) in the action pane.

Figure 1: Accessing the Help



Backup

The Granular Recovery Extension for Microsoft Exchange Server does not provide you any backup solution. Back up your Microsoft Exchange Server databases using the HP Data Protector Microsoft Exchange Server 2010 integration.

- the HP Data Protector Microsoft Exchange Server 2007 integration
- the HP Data Protector Microsoft Exchange Server 2010 integration
- the HP Data Protector Microsoft Volume Shadow Copy Service integration

For more information on the HP Data Protector Microsoft Exchange Server backup solution, see the *HP Data Protector Integration Guide for Microsoft Applications: SQL Server, SharePoint Server, and Exchange Server*.

Restore and recovery

The extension offers you the following benefits:

- **Recovery granularity**

The smallest Microsoft Exchange Server object that you can restore is a database. After the restore you can browse individual Microsoft Exchange Server mailbox items, such as e-mail folders, calendar, contacts, or notes. Thus, you can select to recover the entire database or the desired mailbox items only.

- **Multiple restore requests**

You can receive multiple restore requests concurrently.

- **Recovery of multiple mailboxes**

You can recover multiple mailboxes concurrently.

- **Recovery to different locations**

You can recover Microsoft Exchange Server items to:

- the original location in a mailbox
- a different location:
 - a different mailbox
 - a Personal Folders file (.pst)

You can recover your Microsoft Exchange Server items to a Microsoft Office Outlook client located on a different Microsoft Exchange Mailbox Server without the extension's component installed, by using a Personal Folders file (.pst).

- a different Mailbox Server node without the extension's component installed

- **Easy to search**

You can filter your Microsoft Exchange Server items by specifying the e-mail subject, author, date, terms in the attachments name, or even terms in the message body of e-mail messages. The Microsoft Exchange Server items can be searched before the recovery process is started. This way you can preview all the Microsoft Exchange Server items which will be recovered.

- **Secure operation of the extension**

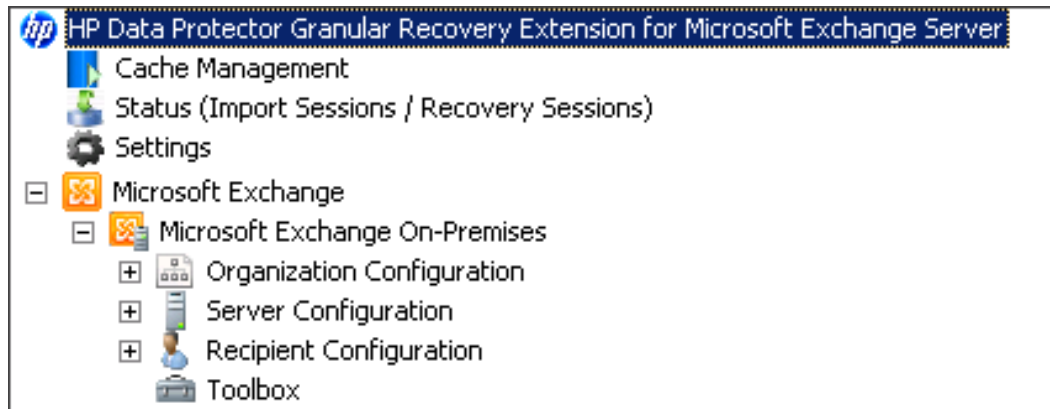
To restore and recover Microsoft Exchange Server items, you (as a Microsoft Exchange Server administrator) must be assigned the `start_restore` user right by a Data Protector backup administrator.

For detailed information, see the ["Configuring user account for the Granular Recovery Extension" on page 15](#).

- **Microsoft Management Console (MMC) snap-in**

The extension's Graphical User Interface (GUI) is a Microsoft Management Console (MMC) snap-in integrated with the Exchange Management Console (EMC). You can find the extension's entry point above the EMC entry point (Microsoft Exchange icon) in the console tree.

Figure 2: The Granular Recovery Extension entry point



The integration eases switching between managing Exchange tasks and performing Granular Recovery Extension tasks such as requesting restores, starting recovery sessions, and so on.

Chapter 2: Installation

This chapter describes how to install the Data Protector Granular Recovery Extension for Microsoft Exchange Server.

Prerequisites

Microsoft Exchange Server 2007

Install the following:

- Microsoft Management Console (MMC) 3.0 or newer
- .NET Framework 3.0 or newer
- Windows PowerShell 1.0 or newer
- Internet Information Services (IIS) 6.0 or newer

Make sure you have:

- Exchange Organization Administrator role assigned.

When you access the extension for the first time, you have to enter the Exchange Organization Administrator user credentials: user name, password, and domain.

- Local Administrator account permission

Microsoft Exchange Server software

Install the following:

- Microsoft Exchange Server

Make sure that you have correctly installed and configured the Microsoft Exchange Server environment.

For supported versions, platforms, devices, and other information, see the latest support matrices at <http://www.hp.com/support/manuals>.

For information on installing, configuring, and using Microsoft Exchange Server, see the Microsoft Exchange Server documentation.

- Microsoft Management Console (MMC) 3.0 or later
- .NET Framework 3.5.1
- Internet Information Services (IIS) 6.0 or later

Data Protector software

Install the following Data Protector components:

- The Data Protector User Interface component
- The Data Protector MS Exchange Server 2010+ Integration component on all Microsoft Exchange Server systems

Make sure that you installed and configured your Data Protector backup solution as described in the *HP Data Protector Installation and Licensing Guide* and in the *HP Data Protector Integration Guide for Microsoft Applications: SQL Server, SharePoint Server, and Exchange Server*.

Other non-Data Protector software and services

- Install Windows PowerShell 1.0 or later (a Windows Management Framework Core package)
- Keep the TCP/IP port 60000 (default) free for the Granular Recovery Web service.

For detailed procedure on configuring an alternative port number for this service, see ["Configuring the Granular Recovery Web service port" on page 15](#).

Configure a firewall to allow new ports.

Supported environments

The extension can be integrated with Microsoft Exchange Server in different Microsoft Exchange Server environments:

- a standalone Microsoft Exchange Server system (a standalone environment)
- multiple Microsoft Exchange Mailbox Server systems (multiple Server systems)
- Microsoft Exchange Server Database Availability Group environments (DAG environments)

Depending on your Microsoft Exchange Server environment, install the extension as follows:

Standalone environment

All Microsoft Exchange Server services and data are installed on a single Microsoft Exchange Mailbox Server, which is adequate in small scale environments. Install the MS Exchange Granular Recovery Extension component to the Exchange Mailbox Server system.

Multiple Exchange Server systems environment

Your environment contains more than one Microsoft Exchange Server database. Install the MS Exchange Granular Recovery Extension component to the Exchange Mailbox Server system of which single items you want to recover.

DAG environment

Your environment can contain up to 16 Microsoft Exchange Mailbox Server systems. Install the MS Exchange Granular Recovery Extension component to any Microsoft Exchange Server mailbox role system node. Once the component is installed, the Granular Recovery Extension Graphical User Interface (GUI) displays all mailbox database objects of all the Mailbox Server nodes in the DAG environment. The extension considers the dynamic behavior of the DAG environment automatically.

CCR environment

Install the MS Exchange 2010 Granular Recovery Extension component on a mailbox server node.

LCR environment

Install the MS Exchange 2010 Granular Recovery Extension component on the server where the active and passive mailbox databases are located.

For detailed information on Microsoft Exchange Server concepts, see the Microsoft Exchange Server documentation.

Installing the extension

The HP Data Protector Granular Recovery Extension for Microsoft Exchange Server is provided as a Data Protector component. The MS Exchange Granular Recovery Extension component contains the Granular Recovery Extension graphical user interface, the command line options, the Web Service components, and the context-sensitive (F1) Help. All the content is installed together.

Note: The extension must be installed on the Microsoft Exchange Server mailbox role systems in the Microsoft Exchange Organization only. These systems contain Microsoft Exchange Server mailbox database and recovery technology such as Recovery Databases (RDB) required for restoring complete Microsoft Exchange Server databases and recovering mailbox items.

Procedure

Install the extension using the Data Protector Graphical User Interface (GUI):

Important: Make sure you have the Windows local user account SYSTEM or a Windows domain user account administrative privileges granted on the Microsoft Exchange Server mailbox role system. You need to be allowed to create registry entries and to install files or folders to the Program Files directory.

1. Perform a remote installation of a client by:
 - adding a client
 - importing a client
2. Add the MS Exchange Granular Recovery Extension component to the Data Protector client system.

For detailed information on the Data Protector installation, see:

- The *HP Data Protector Help* index: “clients systems, installing”.
- The *HP Data Protector Help* index: “adding, Data Protector components”.
- The *HP Data Protector Help* index: “clients systems, importing”.
- The *HP Data Protector Installation and Licensing Guide*.

Removing the extension

Perform one of the following:

- Using the Data Protector GUI, remotely remove the client which has the extension component installed.

For details on removing the Data Protector clients, see the *HP Data Protector Help* index: “uninstalling, client”.

- Manually remove the MS Exchange Granular Recovery Extension component.

For details on removing the Data Protector software components, see the *HP Data Protector Help* index: “uninstalling, Data Protector software”.

Chapter 3: Configuration

This chapter describes the configuration steps that you need to follow.

Meeting Data Protector configuration requirements for the Granular Recovery Extension

Configuring the Granular Recovery Web service port

The Granular Recovery Web service establishes communication using the TCP/IP 60000 port number. If other service is using this port number, configure the Granular Recovery Web service to use an alternative port number:

1. Without starting the extension, search for the following Windows Registry key: HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\OmniBackII\Plugins\exchgre.
2. Edit the key `Client port` and enter the new port number.
3. Update the IIS configuration to have the new port value for the Granular Recovery Web service by running the command:

```
IISWeb /create Web Service web site pathwebsite name/b new port number
```

Where:

Web Service web site path is the root path for the Granular Recovery Web service web site. The default path `C:\inetpub\wwwroot`.

website name is the web site hosted by the Granular Recovery Web service.

new port number is the new port number on which Granular Recovery Web service establishes communication.

For example:

```
IISWeb /create c:\inetpub\wwwroot "HP MS Exchange GRE" /b 8000
```

Configuring user account for the Granular Recovery Extension

Configure your Granular Recovery Extension (GRE) user account with the following user rights and privileges:

Data Protector user rights

Make sure you have the following Data Protector user rights assigned:

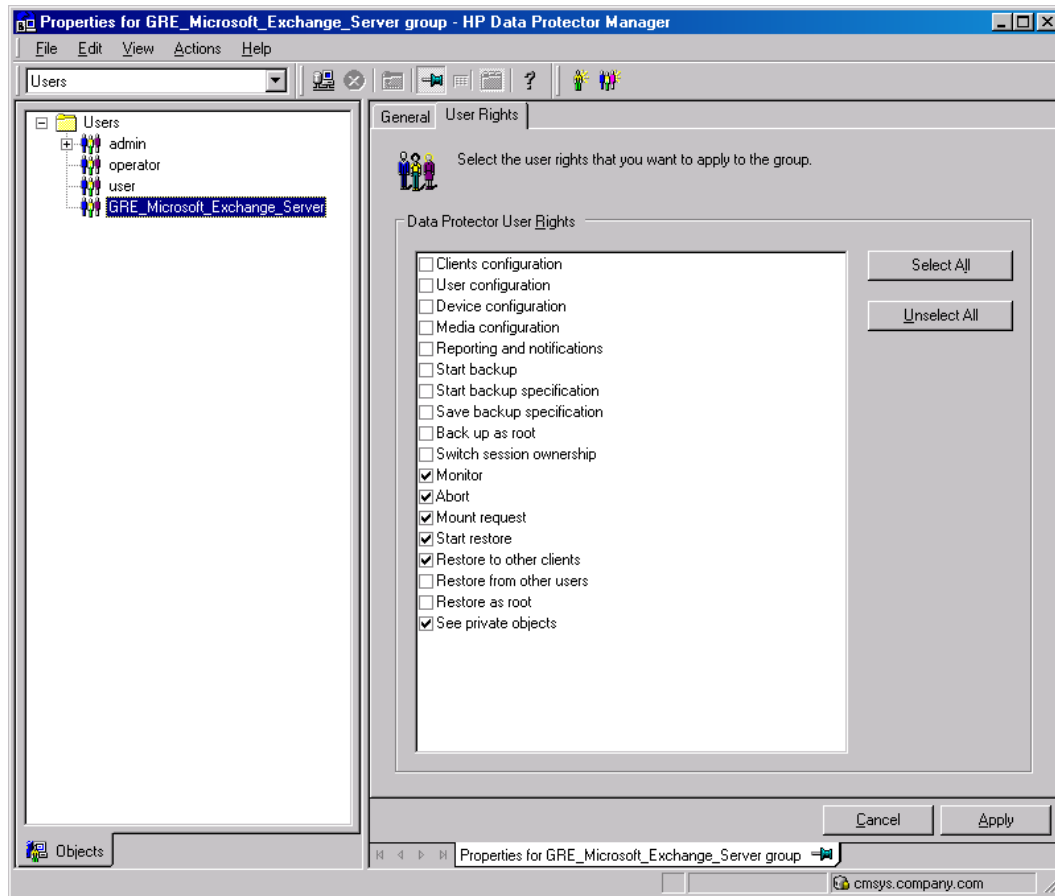
1. Open the Data Protector GUI (**Data Protector Manager**).
2. Create a new user group to be used by the extension, for example, GRE_Microsoft_Exchange_Server.

For details on adding user groups, see the *HP Data Protector Help* index: “adding, user groups”.

3. Assign the following Data Protector user rights to the GRE_Microsoft_Exchange_Server user group:
Monitor , Abort, Mount request, Start restore, Restore to other clients, and See private objects.

For details on assigning Data Protector user rights, see the *HP Data Protector Help* index: “changing, user rights”.

Figure 3: Assigning Data Protector user rights to the GRE_Microsoft_Exchange_Server user group



4. Each time you perform a new installation of the Granular Recovery Extension, add a Data Protector user to the GRE MS Exchange user group. Specify the following General User properties:

Name: SYSTEM, **Domain or UNIX Group:** NT AUTHORITY, **Client system:** *ComputerName*
(Specify the computer name which contains the node with the Granular Recovery Extension installation).

For the detailed procedure, see the *HP Data Protector Help* index: “adding, users”.

Other necessary privileges

Assign the following permissions to your GRE user account:

- creating Windows Registry keys
- setting the Windows registry key values

Privileges for executing Exchange Management cmdlet operations

To create a remote runspace for executing the Exchange Management cmdlet operations remotely, configure user credentials with specific Exchange Management roles. These operations are executed as part of Microsoft Exchange Server backup and restore operations and are necessary for successful operation of the extension.

Configure your GRE user account with the following Exchange privileges:

- a member of the specific Exchange Management built-in role groups with certain built-in management roles assigned:
 - the Organization Management role group
 - the Discovery Management role group
 - the Mailbox Import Export management role

As a member of the Organization Management role group, you are not assigned this management role by default. Assign the role to your GRE user account to be able to recover Exchange mailbox items to the original location in the original mailbox or to a Personal Folders file (.pst).

Note: Recovery to a .pst file requires you to create a network shared folder with read/write permissions granted to the Exchange Trusted Subsystem group.

- a member of the administrators group of the Microsoft Exchange Server system on which the extension is installed

The user credentials specified in the Remote Powershell Configuration dialog box are stored locally in the Windows Registry under the HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\OmniBackII\Plugins\exchgre directory, on the Microsoft Exchange Server system on which the extension is installed.

Important: Only user credentials (a username, password and domain name) of a single user can be stored. Each time you enter new user credentials the existing ones are overwritten. The encrypted password is stored on the Microsoft Exchange Server system on which the extension is installed.

For more information on the Exchange Management cmdlet operations and on how to assign the Exchange Management built-in management roles, see the Microsoft Exchange Server documentation.

Chapter 4: Backup

The Granular Recovery Extension for Microsoft Exchange Server does not provide you any backup solution. Back up your Microsoft Exchange Server databases using the HP Data Protector Microsoft Exchange Server 2010 integration.

For more information on the HP Data Protector Microsoft Exchange Server backup solution, see the *HP Data Protector Integration Guide for Microsoft Applications: SQL Server, SharePoint Server, and Exchange Server*.

The following backup types are supported: Full, Incremental, Copy, and Differential.

The following disk-based backups are not supported:

- Microsoft Volume Shadow Copy Service (VSS) zero downtime backup (ZDB) to Disk
- VSS ZDB to Disk + Tape
- VSS ZDB instant recovery transportable

For detailed information on backup types, see the *HP Data Protector Help* index: “backup types”.

Note: The recovery procedure does not differ by Data Protector backup type.

Chapter 5: Restore and recovery

Limitations

Data Protector Granular Recovery Extension limitations

- Databases or mailboxes encrypted by third party applications cannot be recovered.
- The **View** button in the action pane is meant for the future use. Thus, you cannot customize columns displayed in the results pane in the Cache Management and Status pages yet.
- The recovery of Public Folder mailboxes, team/site mailboxes, and linked mailboxes are not supported.
- The mount-point folder path that points to a mount volume location without a drive letter is not a supported restore location path.

Microsoft Exchange Server limitations

- Microsoft Exchange Server `Restore-Mailbox` cmdlet does not support recovery to special folders. Special folders: Inbox, Drafts, Sent Items, Deleted Items, Junk E-mail, Outbox, RSS Feeds, Sync Issues, Conversation History, Tasks, Calendars, and Contacts. As a result, items stored in the Exchange Server special folders cannot be browsed in the Granular Recovery — Recovery Settings page:
 - using the **Recover to an existing folder** option that sets the target of the original mailbox.
 - using the **Recover into different Mailbox** option that sets the target of the different mailbox.

You can recover items where the target is a special folder only when using the option **Recovery to original location**.

- The restore of Public Folder mailbox databases is not supported.
- In the Microsoft Exchange Server 2013 environment, deleted team mailboxes cannot be identified, hence you cannot filter them out.
- In the Microsoft Exchange Server 2013 environment, a single mailbox items recovery is not supported. Therefore, you need to use a combination of `New-MailboxRestoreRequest`, `Search-Mailbox`, and `New-SearchMailbox` cmdlet operations.

Considerations

Data Protector Granular Recovery Extension considerations

- Two or more instances of the Granular Recovery Extension GUI or the GUI and the command-line interface (CLI) cannot be used at the same time.

- Any Microsoft Exchange Server operations, such as creating or deleting the Recovery Database, that are performed outside the Granular Recovery Extension (GRE) using management tools such as the Exchange Management Console, while the GRE user interface is open, are not reflected in the GRE user interface.
- Multiple restore requests for the same backup object (mailbox database) from the same backup version are processed once.
- The **Export List ...** button is displayed in the action pane, if you select the HP Data Protector Granular Recovery Extension node. The button creates a list of all the contents displayed in the console tree: Cache Management, Status (Import Sessions / Recovery Sessions) and Settings Microsoft Exchange. The list can be exported in the following formats: text, Unicode text, comma-separated value (CSV), and Unicode CSV. This functionality is provided by default by the Microsoft Management Console (MMC).
- You can perform multiple recovery requests for the same restored database.
- In the GRE wizard, you can specify search criteria to narrow the list of items which you can select for recovery. After entering some values in the Mailbox Search Criteria page, and selecting one or more items for recovery in a specific mailbox folder, any items in this folder's subfolders will also be recovered if they meet the same search criteria.
- The Search Results page displays only three levels of folders, which does not affect your restore process. If you select an item on the third level, its child items will also be restored.
- Make sure the destination folder is empty before performing restore.

Tip: You can specify a new restore folder in the Restore Settings page of the Import From Backup wizard, and the extension creates a new folder.

- In the Mailbox Selection page of the Import from Backup wizard, the Mailbox user names starting with non-ASCII characters are grouped under --.
- Non-ASCII characters are not supported in paths. When typing the restore path, avoid non-ASCII characters. Otherwise the restore may fail.
- If a recovery database already exists in a target location on a mailbox server, the Granular Recovery Cache keeps all the versions without deleting them.
- You can perform multiple restore requests to the same target location.
- Multiple recovery databases can be created on a mailbox server.
- The number of recovery databases is limited by the disk space available in the temporary restore location.
- The Granular Recovery Cache keeps only one recovery database (RDB) mounted per Microsoft Exchange mailbox server, even though the restored database files are still on the disk once the recovery session finishes.

On a Microsoft Exchange mailbox server, there is only one recovery database stored.

For example, in a DAG environment, each mailbox node can contain one recovery database, but only on one server there can be a mounted recovery database.

Data Protector considerations

- The Granular Recovery Extension cannot restore or recover items from backup images created with the Data Protector Microsoft Volume Shadow Copy Service (VSS) integration.
- The Granular Recovery Extension does not support Instant Recovery (IR).

Microsoft Exchange Server considerations

- The option of recovering data into a PST file is only available for Microsoft Exchange Server 2010 environments using Microsoft Exchange Server 2010 SP1 or a later service pack.
- Moved or deleted mailbox databases cannot be searched (the Microsoft Exchange Server known issue). However, you can recover a moved mailbox once it was moved.

Deleted mailboxes are displayed in the Import From Backup wizard, on the Import from Backup — Mailbox Selection only if the retention period is not expired. Once the retention period is reached and the deleted mailbox is no longer available in the Import From Backup wizard, re-importing the mailbox database that contains the needed mailbox is necessary.

For details, see ["Importing mailbox databases" on page 26](#).

- In the Microsoft Exchange Server 2013 environment, a combination of `New-MailboxRestoreRequest`, `Search-Mailbox`, and `New-SearchMailbox` cmdlet operations induces storage space and performance factors which cannot be avoided due to Microsoft Exchange Server limitations.
- Make sure that active databases in a standalone or DAG configuration (the GRE supported Microsoft Exchange Server configuration) have two times of a source mailbox size storage space for creating temporary mailboxes during recovery.
- Make sure that target mailboxes to which recovery items are to be stored have sufficient storage space for storing the recovered items.

Restore and recovery flow

To restore and then recover the Microsoft Exchange Server items, follow the basic steps:

1. Import

a. Restore

The HP Data Protector Granular Recovery Extension uses the Data Protector Microsoft Exchange Server 2010 integration to restore the Microsoft Exchange Server databases.

Temporary restore location

First the Microsoft Exchange Server database files are saved to a temporary restore location. Restore the database files (.edb), checkpoint files (.chk), reserve transaction log files (.jrs), and transaction log files (.log) to the specified temporary restore location on a Microsoft Exchange Server system.

A recovery database (RDB) is created in the temporary location.

The restored files are located on the Microsoft Exchange Server system that is chosen as the restore target system. The default location is C:\Restore, but you can specify a different restore location.

For details, see ["Changing settings" on page 38](#).

After successful restore of mailbox databases from their backup images, the restored databases are available in the granular recovery cache.

b. Mount

Mount the restored database in the granular recovery cache to the Microsoft Exchange Server. Before browsing and recovering items you have to mount the restored database files.

2. Recover

Browse and recover the Microsoft Exchange Server items from the recovery database to the original mailbox database, or to any different location.

3. Dismount

Only one recovery database can remain mounted, for this reason when you no longer need to recover items from a recovery database:

- Dismount the recovery database from the Microsoft Exchange Server.

Once you dismount the recovery databases, the recovery databases are still in the Granular Recovery — Cache Management but their status are dismounted. At this point, you can still remount them if needed for another recovery session.

4. Removal

The recovery databases remain in the cache for 30 days (default value) or for as long as it is set in the retention time.

After the retention time expires, the database is dismounted and removed from the Granular Recovery Extension cache automatically, but the restored database files still exist in temporary restore location.

- Optionally, change the retention period. For detailed procedure, see ["Changing the retention period" on page 39](#).
- Optionally, remove the no longer needed recovery database from the cache manually before the retention period. For detailed procedure, see ["Removing databases" on page 37](#).

- The restored database files still remain in temporary restore location, you can completely remove them, by deleting the files from the temporary restore location manually.

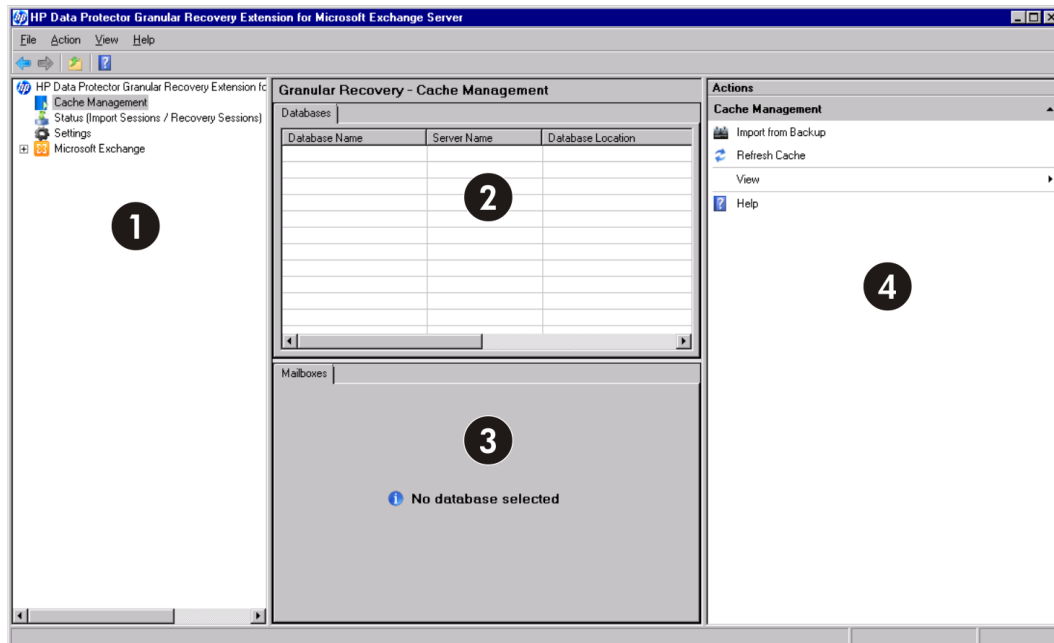
Opening the HP Data Protector Granular Recovery Extension GUI

To open the extension:

1. Log in to the Microsoft Exchange Server system where the Granular Recovery Extension is installed.
2. Click the Start button and then click the HP Data Protector Exchange GRE icon to open the extension's Graphical User Interface (GUI).

The HP Data Protector Granular Recovery Extension for Microsoft Exchange Server is started.

Figure 4: Main window of the Granular Recovery Extension Graphical User Interface (GUI)



Item	Description
1	Console tree
2	Result pane
3	Work pane
4	Action pane

Note: In the Microsoft Exchange Server 2010 environment, the Exchange Management tasks are managed using the Exchange Management Console (EMC) accessible by clicking the Microsoft Exchange icon in the console tree of the Granular Recovery Extension Graphical User Interface (GUI).

In the Microsoft Exchange Server 2013 environment, the Exchange Management tasks are managed by using the Exchange Administration Center (EAC). The EAC has its own web based graphical user interface and cannot be accessed through the extension's Microsoft Management Console (MMC) snap-in GUI.

3. In the console tree, click the Cache Management icon.

The empty **Granular Recovery - Cache Management** is displayed.

The remote powershell has to be configured before performing any granular recovery operations. See "[Remote powershell configuration](#)" below.

To import Microsoft Exchange Server databases, follow the procedure for "[Importing mailbox databases](#)" on the next page.

Remote powershell configuration

Configure a user account for executing the Exchange Management cmdlet operations remotely.

For details, see "[Privileges for executing Exchange Management cmdlet operations](#)" on page 17.

If no valid user credentials are specified for executing the Exchange Management cmdlet operations remotely, the Remote Powershell Configuration dialog box is displayed. Enter the required user credentials and click **OK**.

Figure 5: Remote powershell configuration

Remote Powershell Configuration

These credentials are used for remote powershell configuration and to execute exchange powershell cmdlets as part of granular recovery operations.

User Name: Administrator

Password: ●●●●●●●●●●●●

Domain: DOMAIN

Ok Cancel

Importing mailbox databases

Prerequisites

- Make sure sufficient disk space is available on the target Exchange Mailbox Server.

Considerations

- The HP Data Protector Granular Recovery Extension cannot restore or recover items from backup images created with the Data Protector Microsoft Volume Shadow Copy Service (VSS) integration.
- The HP Data Protector Granular Recovery Extension does not support Instant Recovery (IR).

Limitations

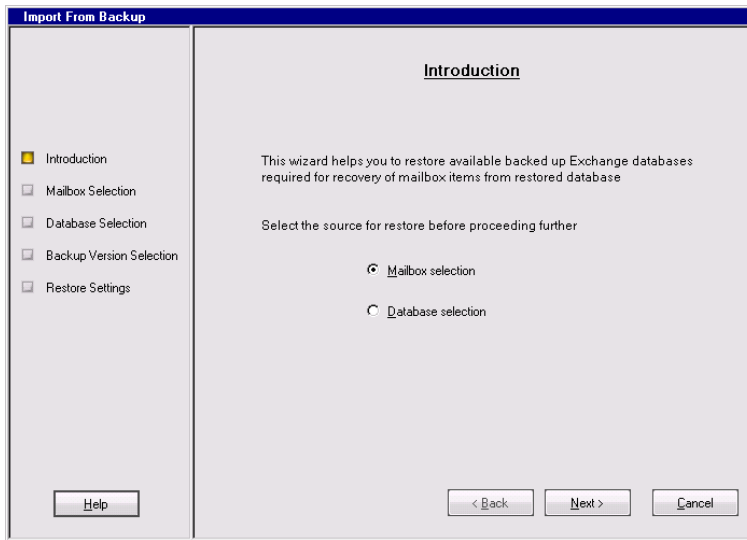
- In the Mailbox Selection page of the Import from Backup wizard, the deleted mailbox database names can be identified by the suffix @@De1eted in the mailbox display name. However, if there are more mailbox databases with the same display name, only one entry is displayed. To be able to import other mailboxes with the same name, contained in the database, use the Database selection page instead of Mailbox selection page.
- Deleted mailboxes are no longer displayed after the expiration of the retention period.
- If a mailbox is moved from a standalone environment to new standalone environment, an import from backup can be performed in the original environment. The backup version displayed is the one that still contains the mailbox in the original standalone environment.

Procedure

Before browsing and recovering items, you have to import databases:

1. In the console tree, click the **Cache Management** icon. The **Granular Recovery Cache Management** page is displayed in the results pane.
2. In the action pane under the Cache Management node, click **Import from Backup**. The Import From Backup wizard is displayed.

Figure 6: Selecting a restore object

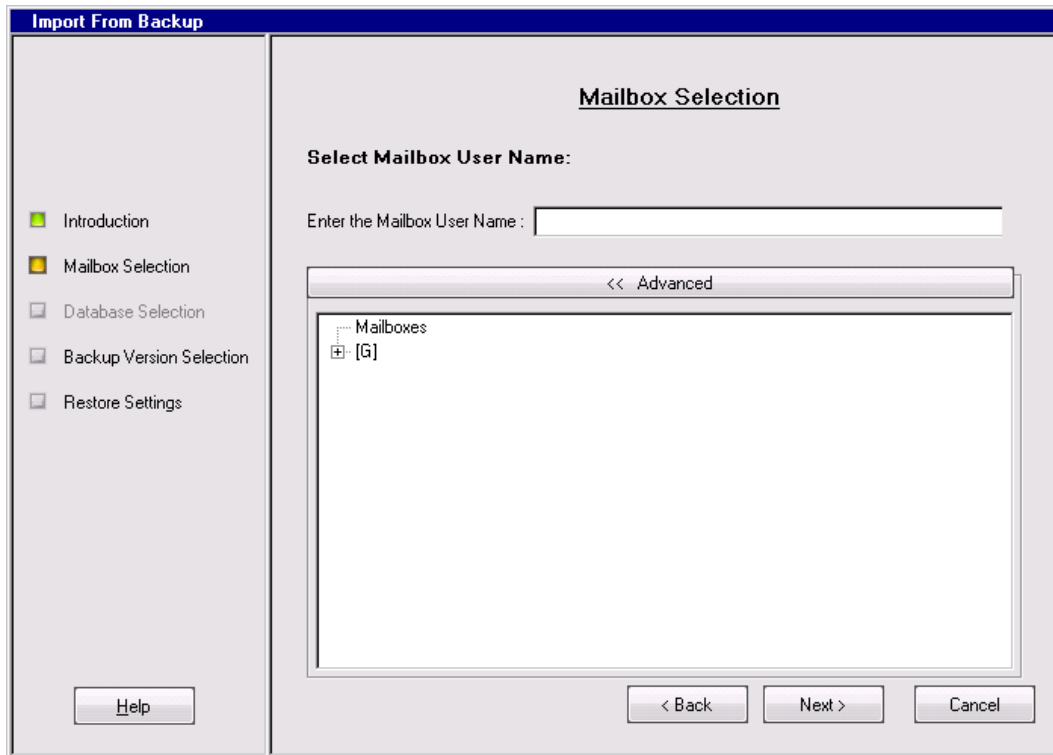


3. In the Introduction page, select a backup source:

- To import only a specific mailbox, select **Mailbox selection** and click **Next**. The Mailbox Selection page is displayed.

Tip: **Mailbox selection** is especially useful if the mailbox database is unknown.

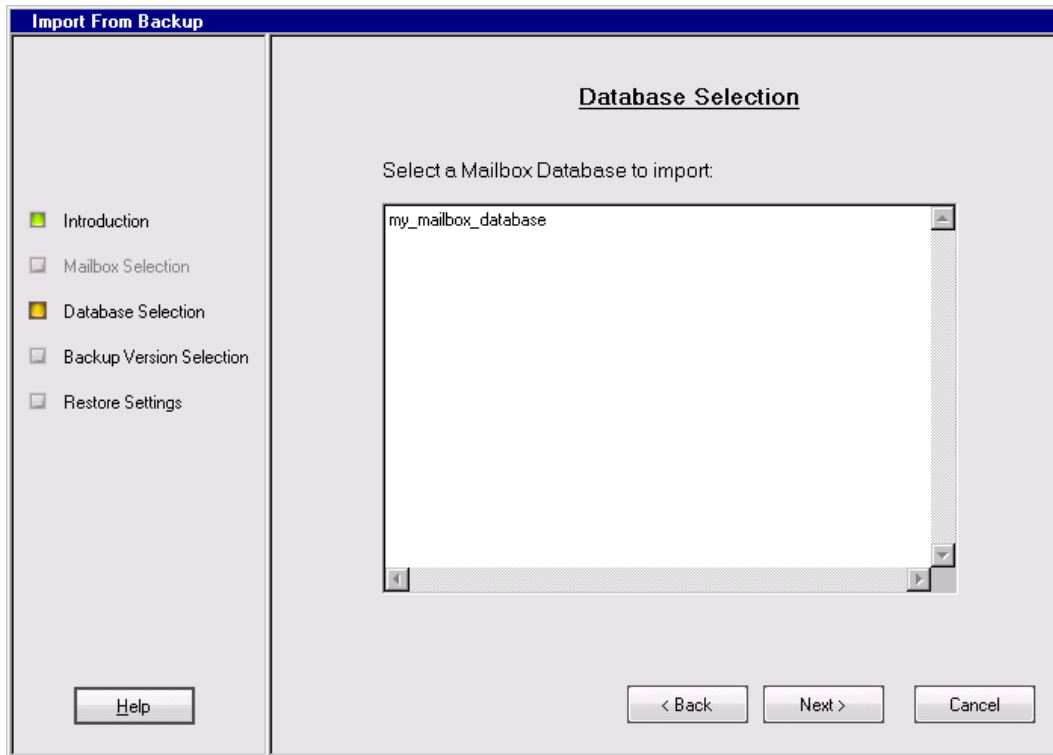
Figure 7: Selecting a desired mailbox



- i. Specify the mailbox user name, and click **Next**.
 - ii. The Backup Version Selection page is displayed. Select the backup data you want to restore and click **Next**.
- To import a complete database and all the mailboxes contained in the database, select **Database selection** and click **Next**.

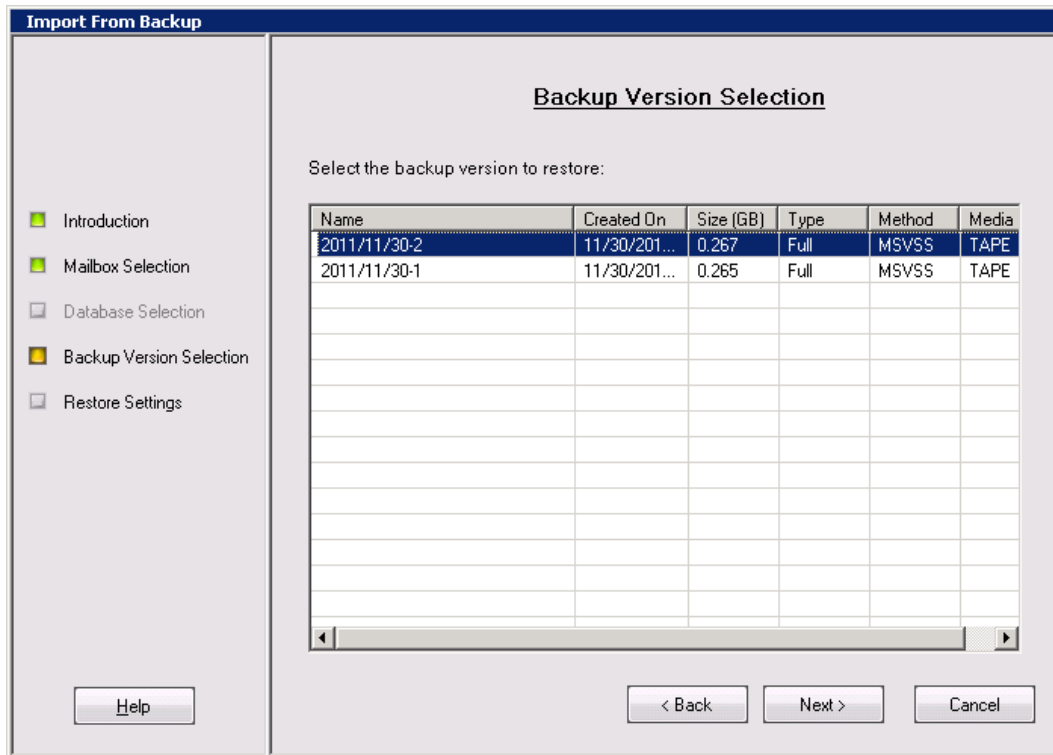
Note: The database selection for restore can be useful if, for any reason, the mailbox you want to recover is not visible in the Mailbox Selection page.

Figure 8: Selecting a mailbox database



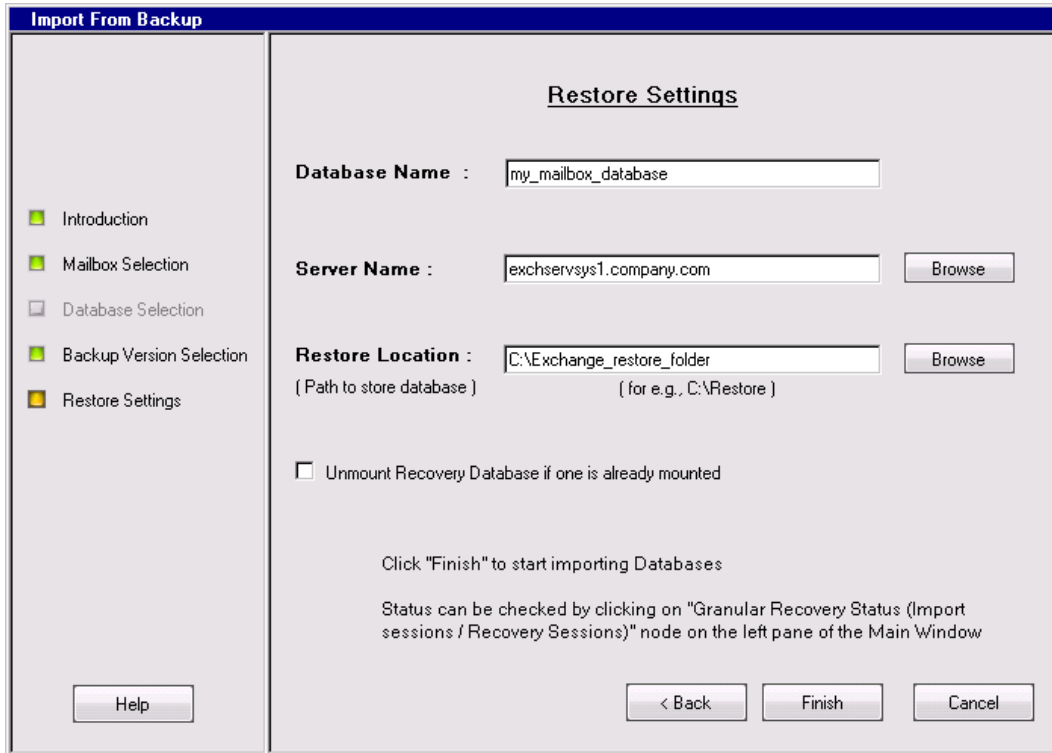
- i. The Database Selection page is displayed. Select the database you want to restore and click **Next**.
- ii. The Backup Version Selection page is displayed. Select the backup version you want to restore and click **Next**.

Figure 9: Selecting a backup version



- The Restore Settings page is displayed. Confirm or adjust the values of the Database Name, Server Name, and the Restore Location.

Figure 10: Adjusting the restore settings

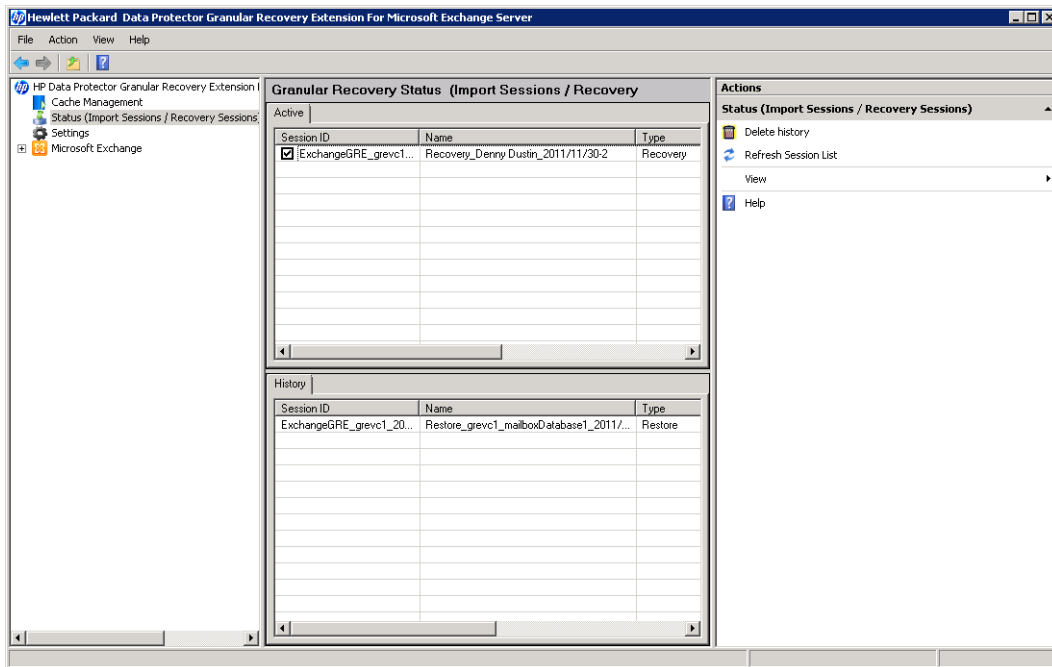


Make sure that the destination folder is empty before performing restore.

Tip: You can specify a new restore folder in the Restore Settings page of the Import From Backup wizard and the extension creates a new folder.

5. Optionally, you can dismount a recovery database if one already exists in the Granular Recovery Cache Management on the server specified in the Restore Settings page. Select the option and click **Finish**.
6. To monitor the restore session, in the console tree click **Status (Import Sessions / Recovery Sessions)**. The Granular Recovery Status (Import Sessions / Recovery Sessions) page is displayed.
7. To stop the restore session, click **Abort Sessions**.

Figure 11: Imported database as displayed in the GUI



Mounting databases

Prerequisites

- Make sure the mailbox database files are restored (imported) in the Granular Recovery — Cache Management.

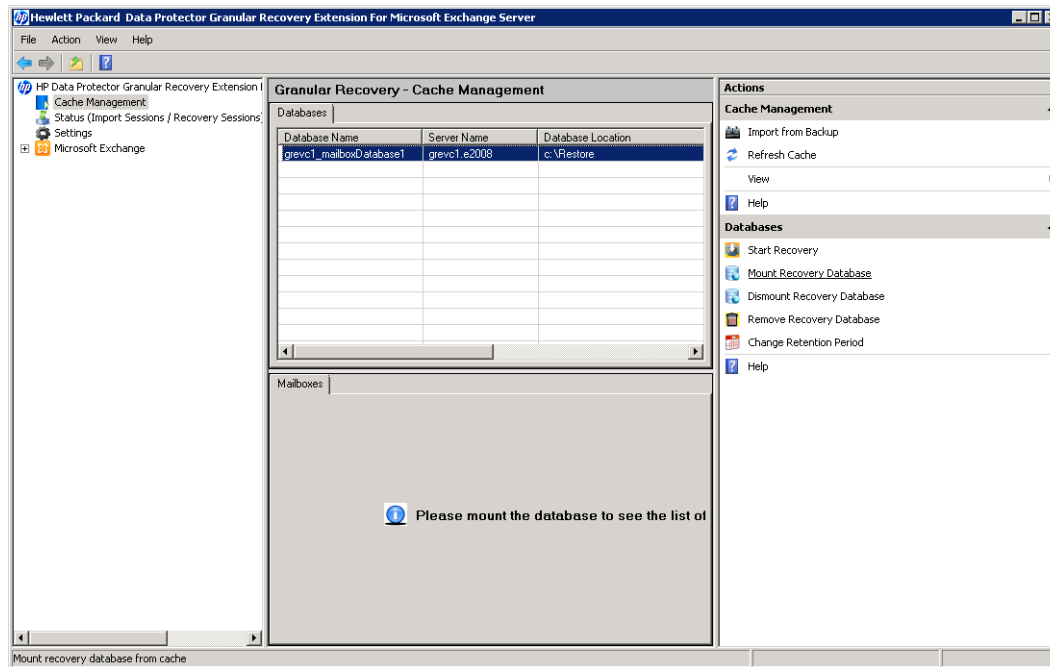
Procedure

After importing the mailbox database (once the restore process is finished), mount the database manually.

To mount the mailbox database manually:

1. In the console tree, select the Cache Management node.
2. In the results pane, select the database you want to mount.

Figure 12: Selecting a database



3. In the action pane under the Databases node, click **Mount Recovery Database**.

Once the database is mounted, in the work pane, the mailbox display name, size, and the last logged on username are displayed.

Starting recovery

Prerequisites

- Make sure the recovery database from which you want to recover Microsoft Exchange Server single items is mounted.

Limitations

- If you select contact items and mail items in the Search Results page, and recover them in one session, the contact items will not be recovered. You must select and recover the contact items separately.
- If several mailboxes with the same display name were disabled at the time of the backup and no other connected mailboxes in the database had the same display name, two deleted mailboxes are displayed in the Mailbox selection page after import, one with the suffix `@@Deleted` and one without.

If the Granular Recovery Extension does not find the `@@Deleted` suffix and the mailbox was deleted, the **Mailbox Personal Folder** option is disabled in the Mailbox Search Criteria page.

The `@@Deleted` suffix is not appended to the name of the deleted mailbox if there is no other mailbox with the same display name.

Procedure

To recover Microsoft Exchange Server single items:

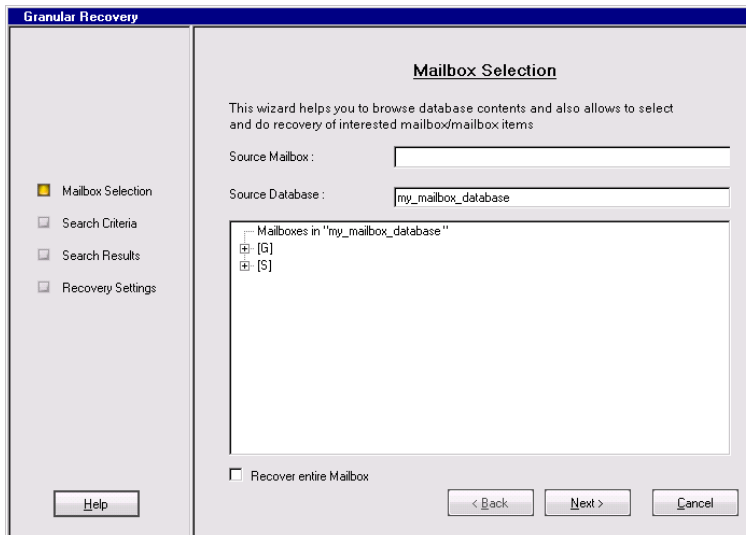
1. In the console tree, select the Cache Management node. In the results pane, select the recovery database from which you want to recover.
2. In the action pane under Databases, click **Start Recovery**.

Tip: A shortcut access to any action button, for example **Start Recovery**, is to right-click the database in the Granular Recovery—Cache Management.

3. In the Mailbox Selection page, select the mailbox for recovery. Optionally, to restore a complete mailbox folder, select the **Restore entire Mailbox** option.

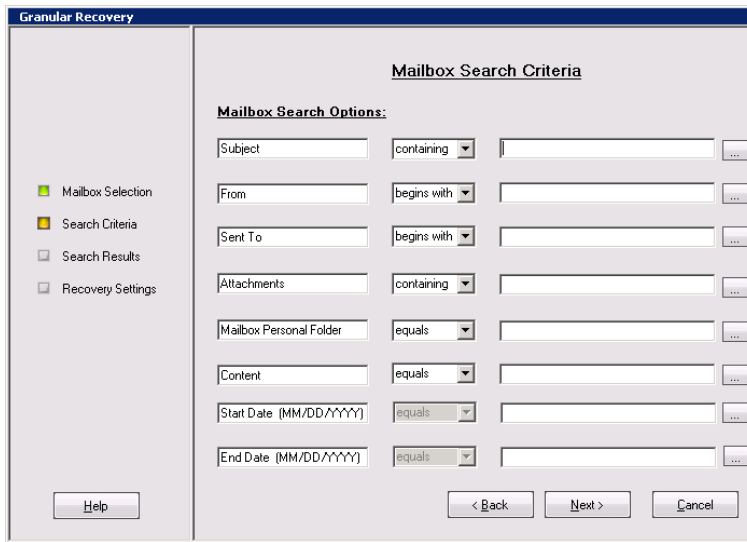
Click **Next**. The Mailbox Search Criteria page is displayed.

Figure 13: Selecting a mailbox



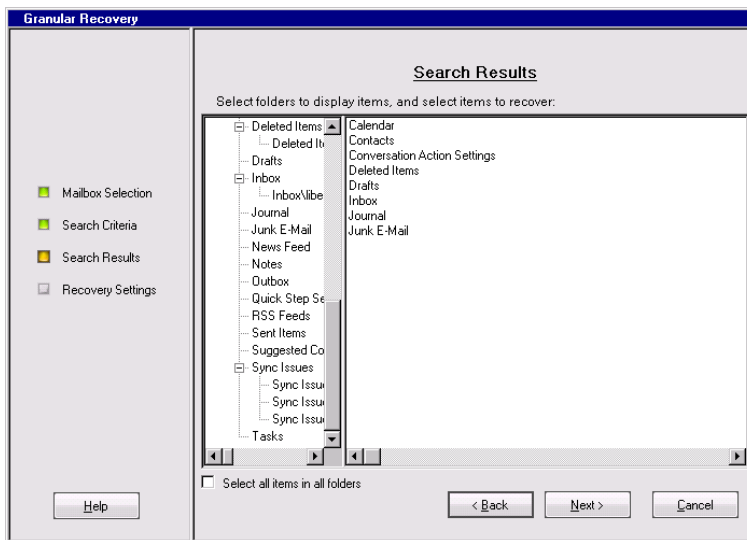
4. Filter e-mails by subject, author, recipient, attachment term, personal folder, or e-mail content, and click **Next**.

Figure 14: Specifying the search criteria



5. In the Search Results page, select a folder from the mailbox. The items are displayed in a table.

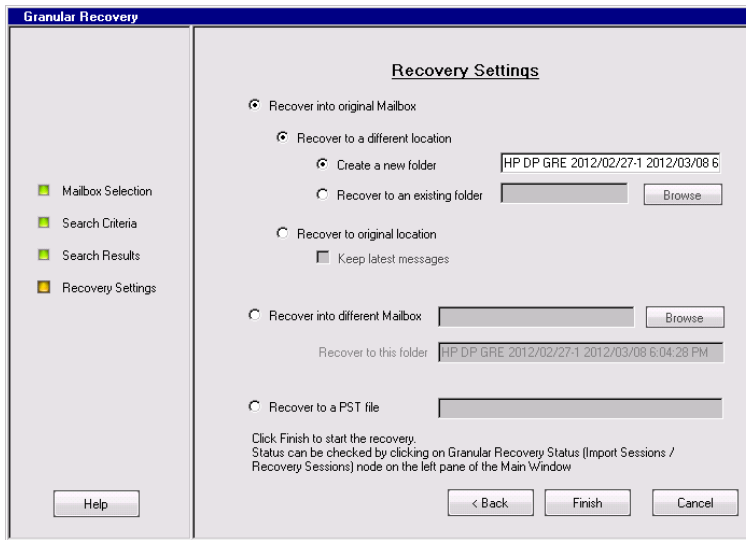
Figure 15: Selecting an item for recovery



Select the item you want to recover and click **Next**. The Recovery Settings page is displayed.

Tip: You can select multiple items by holding **Ctrl** or **Shift** button.

Figure 16: Specifying the target location



Tip: If **Keep latest messages** is not selected, the items are overwritten.

6. Specify the target recovery location: an existing Exchange Server mailbox or the name of the PST file to be created during recovery.

Tip: When you recover items to an existing folder of the mailbox, only the folders you created are displayed. Special folders, such as Inbox, Drafts, Sent Items, Deleted Items, Junk E-mail, Outbox, RSS Feeds, Sync Issues, Conversation History, Tasks, Calendars, and Contacts, cannot be set as a target, therefore are not displayed. Clicking the **Browse** button of **Recover to an existing folder** or, the **Browse** button of the **Recover into different Mailbox** does not display special folders. Only **Recover to original location** can recover items to special folders.

The PST file folder must be accessible from the Exchange Server system.

To use a remote system or current server, enter the path to the network shared folder in the following (UNC) format:

```
\\SystemName\FolderShareName\Filename.pst
```

Note: Make sure the local system user account performing the recovery has read and write permissions set on the network share folder in order to create a PST file. If the folder is located on a remote server system, the Data Protector MS Exchange Granular Recovery Extension component does not need to be installed on the remote system.

Click **Finish** to start the recovery operation and close the wizard.

Dismounting databases

Prerequisites

- Make sure the recovery database you want to dismount is mounted.

Considerations

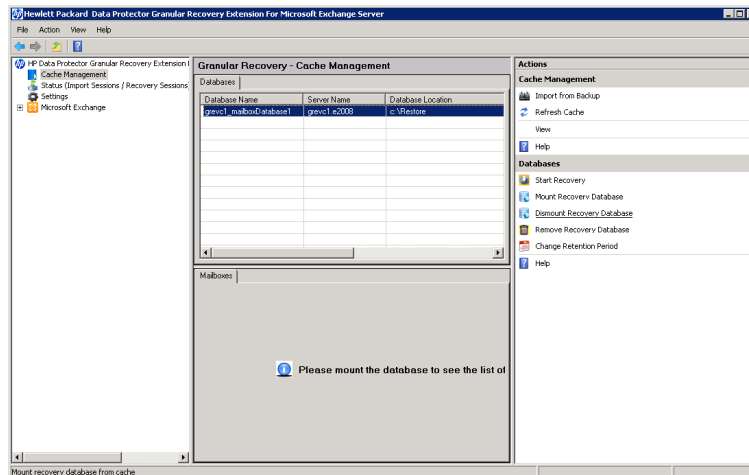
- Only one recovery database can stay mounted. For this reason, when you want to recover single items, or complete mailboxes from a different backup version or from a different mailbox database, you have to first dismount the mailbox database already presented to the Microsoft Exchange Server, the one displayed in the Granular Recovery Cache Management.

After the restore process is finished, the selected database is mounted in the Granular Recovery Cache Management. The mounted database is displayed in the results pane.

To dismount no longer needed databases from the Granular Recovery Cache Management:

1. In the console tree, select the Cache Management node. The Granular Recovery Cache Management page is displayed.
2. In the results pane, select the database. In the action pane under the Databases node, click **Dismount Recovery Database**.

Figure 17: Clicking Dismount



3. The confirmation dialog box is displayed. Click **Yes**. The database is dismounted, the mailbox information is no longer displayed in the work pane.

Removing databases

Prerequisites

- Make sure the recovery database you want to remove from the disk is restored.

Considerations

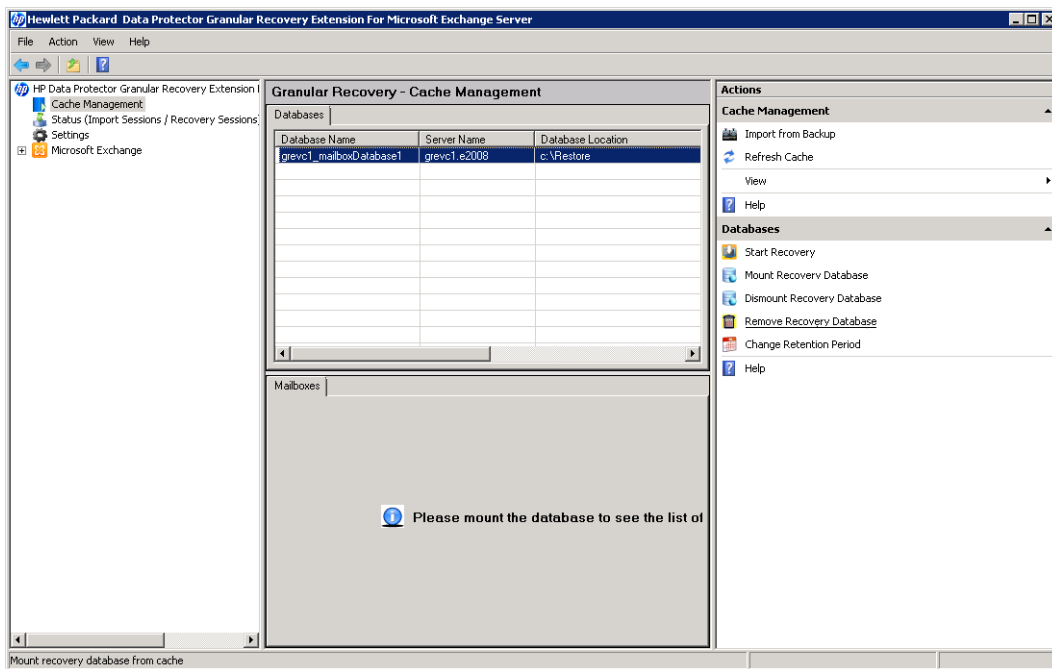
- The restored mailbox database is available in the Granular Recovery — Cache Management for 30 days (default value). After the retention period the database is deleted automatically.

Procedure

To remove no longer needed databases manually from the Granular Recovery — Cache Management and from the temporary restore location on the disk:

1. In the console tree, select the Cache Management node. The Granular Recovery Cache Management page is displayed.
2. In the results pane, select the database. In the action pane under the Databases node, click **Remove Recovery Database**.

Figure 18: Removing a database



3. The confirmation dialog box is displayed. Click **Yes**. The database is removed from the temporary restore location.

Changing settings

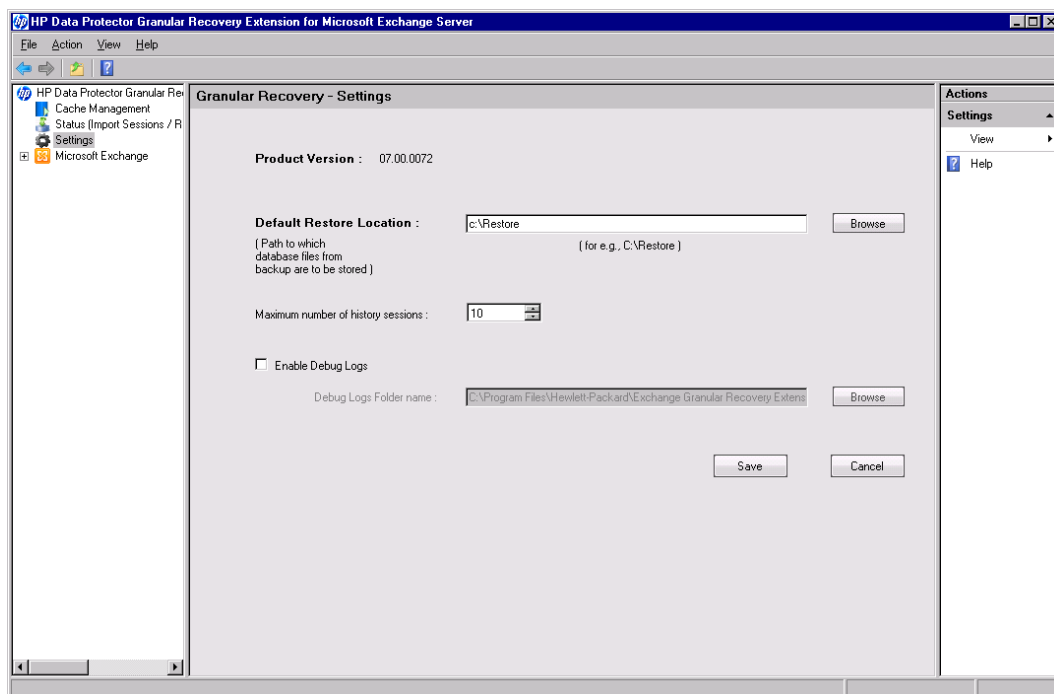
After the restore session is completed, the database files (.edb), checkpoint files (.chk), reserve transaction log files (.jrs), and transaction log files (.log) are copied to the temporary restore location c:\restore.

Procedure

To change the default settings of the extension:

1. In the console tree, click the **Settings** icon. The **Granular Recovery Settings** page is displayed in the results pane.
2. The temporary restore location is set to `c:\restore` (default). To change the temporary restore location, specify the new path by typing or browsing the new directory.
3. To set how many completed sessions (restore and recovery sessions) are displayed in the Granular Recovery Status page, specify the maximum number of history sessions.
4. To enable debugging of the Granular Recovery extension, select **Enable debug logs**. To change the default location of the debug files: type the new location, or specify a new location by clicking **Browse**, and click **Save**.

Figure 19: Changing default settings



Changing the retention period

Prerequisites

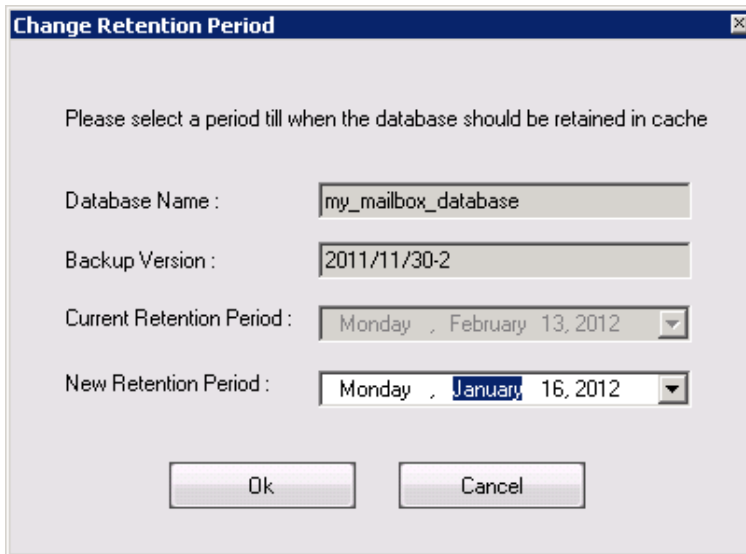
- Make sure the mailbox database is available in the granular recovery cache.

Procedure

The retention period of the mailbox databases is 30 days. After the expiration date the databases are removed from the cache automatically. To change the default value:

1. In the results pane, select the database.
2. In the action pane under the Databases node, click **Change Retention Period**. The Change Retention Period dialog box is displayed.
3. In the New Retention Period drop-down list, select the new date in the calendar.

Figure 20: Changing retention period



Chapter 6: Command line reference

The Data Protector Granular Recovery Extension for Microsoft Exchange Server offers a command line interface which you can use instead of the GUI.

Prerequisites

Configure a user account for executing the Exchange Management cmdlet operations remotely by using the command `--Config`, providing Username, Password, and Domain. The remote powershell has to be configured before performing any granular recovery operations using the CLI commands.

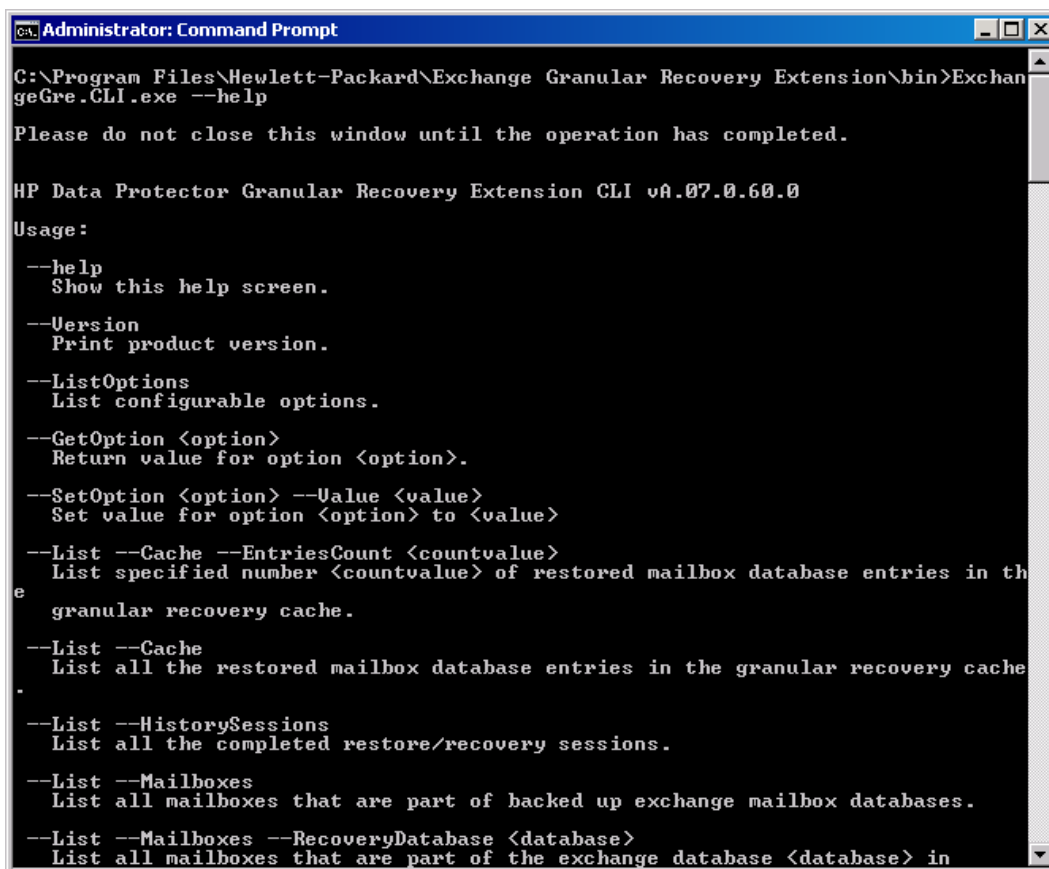
For details, see ["Privileges for executing Exchange Management cmdlet operations"](#) on page 17.

Note: If no valid user credentials are specified for remotely executing the Exchange Management cmdlet operations, the command displays an error message.

The command `ExchangeGre.CLI.exe` is located in the installation directory of the extension:

`C:\Program Files\Hewlett-Packard\Exchange Granular Recovery Extension\bin`

Figure 21: Retrieving the command line help



```
Administrator: Command Prompt
C:\Program Files\Hewlett-Packard\Exchange Granular Recovery Extension\bin>ExchangeGre.CLI.exe --help
Please do not close this window until the operation has completed.

HP Data Protector Granular Recovery Extension CLI vA.07.0.60.0
Usage:
--help
  Show this help screen.
--Version
  Print product version.
--ListOptions
  List configurable options.
--GetOption <option>
  Return value for option <option>.
--SetOption <option> --Value <value>
  Set value for option <option> to <value>
--List --Cache --EntriesCount <countvalue>
  List specified number <countvalue> of restored mailbox database entries in the granular recovery cache.
--List --Cache
  List all the restored mailbox database entries in the granular recovery cache.
--List --HistorySessions
  List all the completed restore/recovery sessions.
--List --Mailboxes
  List all mailboxes that are part of backed up exchange mailbox databases.
--List --Mailboxes --RecoveryDatabase <database>
  List all mailboxes that are part of the exchange database <database> in
```

SYNOPSIS

ExchangeGre.CLI.exe--Version | --Help

ExchangeGre.CLI.exe--List {--Cache [--EntriesCountNumber | --Verbose]| --
HistorySessions | --Mailboxes [--RecoveryDatabaseRecoveryDatabaseName] | --
BackupVersions {MailboxDBDatabaseName | MailboxMailboxName} | --AllBackupDatabases}

ExchangeGre.CLI.exe--Remove {--SessionsSessionID [SessionID...] | --AllSessions | --
RecoveryDatabaseDatabaseName--ServerComputerName}

ExchangeGre.CLI.exe {--MountDB | --DismountDB} --RecoveryDatabaseRecoveryDatabase --
MailboxDB Database --Server ComputerName

ExchangeGre.CLI.exe--SetRP--RecoveryDatabaseRecoveryDatabaseName--Server
ComputerName--PeriodNewDate

ExchangeGre.CLI.exe--Details--SessionSessionID

ExchangeGre.CLI.exe--Search {--MailboxMailboxName | --Cache--MailboxMailboxName--
MailboxDBDatabaseName}

ExchangeGre.CLI.exe--ListOptions | --GetOptionOptionName | --SetOptionOptionName--
ValueValue

ExchangeGre.CLI.exe--StartSession--Restore--DismountRDB [true | false]--
MailboxDBDatabaseName--BackupIDBackupVersion--ServerComputerName--
TargetLocationTargetFolderPath

ExchangeGre.CLI.exe--StartSession--Recovery--SrcMailboxSourceMailboxName--
RecoveryDatabaseRecoveryDatabaseName--MailboxDBDatabaseName {--RecoverWholeMailbox |
--Filter [FILTER_OPTIONS]} --RecoveryTargetType [ORG_MAILBOX | DIFF_MAILBOX | PST]

ORG_MAILBOX

OrgLocation--KeepLatestMsg [true | false] [RECOVERY_OPTIONS]

DIFF_MAILBOX

DiffMailbox [RECOVERY_OPTIONS] --TargetMailboxMailboxName

PST

pst--PSTFileNamePSTFileNameWithPath

RECOVERY_OPTIONS

--DiffLocation {--CreateNewFolder {NewFolderName | Default} | --
ExistingFolderFolderName}

FILTER_OPTIONS

```
Subject="Term" | Contents="Term" | Attachments="Term" | Senders="SenderName" |  
Recipients="RecipientName" | Folders="FolderName" | StartDate="Date" | EndDate="Date"
```

DESCRIPTION

ExchangeGre.CLI.exe is the command line interface of the Granular Recovery Extension for Microsoft Exchange Server. You can use it to perform queries, to restore, mount, recover, and dismount databases, to recover single items, and set different recovery options.

For a detailed description of available options see ["OPTIONS" below](#).

Note: The command line interface and the graphical user interface cannot be used at the same time.

OPTIONS

--Version	Displays the version of the extension.
--Help	Displays the usage synopsis for the ExchangeGre.CLI.exe command.

<pre>--List {--Cache -- HistorySessions -- Mailboxes[-- RecoveryDatabase RecoveryDatabaseName] - -BackupVersions[-- MailboxDB MailboxDatabaseID -- MailboxMailboxName] -- AllBackupDatabases}</pre>	<p>The <code>--List</code> option lists various information about the Granular Recovery Cache, mailboxes, sessions, and so on.</p> <p>The <code>--Cache</code> option lists the information of the restored mailbox databases in the Granular Recovery Cache Management:</p> <p>The database name, the Recovery Database ID with the granular recovery time stamp, for example <code>RDB_ExchangeGRE_dpexchange5_2011-11-15_1</code>, the server name, the status of the mailbox database (mounted or dismounted), and the backup version and the size of the database. With the <code>--Verbose</code> option you can see more details such as the database location (the directory to where the database files are restored), the date the database was restored, and the date of the expiration when the retention period is over the database is dismounted from the Granular Recovery Cache Management.</p> <p>The <code>--HistorySessions</code> option displays 10 restore and recovery sessions by default, with the following information: session ID, Name, Type Start Time, End time, and status (completed or failed).</p> <p>The <code>--Mailboxes</code> option displays all the mailbox names from the backed up Microsoft Exchange mailbox databases.</p> <p>The <code>--BackupVersions</code> option displays all backup versions available for a recovery database with information about the name, data of backup creation, the size of the backup, type of the backup, the method used to perform the backup, and the media type; as well as all backup versions of a recovery database of a specific mailbox.</p> <p>The <code>--AllBackupDatabases</code> option displays all the backup databases names.</p>
<pre>--Details-- SessionSessionID</pre>	<p>Displays details of the specified session.</p>
<pre>--ListOptions --GetOption OptionName --SetOption OptionName --Value Number</pre>	<p>The <code>--ListOptions</code> option displays the available options.</p> <p>The <code>--GetOptionOptionName</code> option retrieves the value of a specified option.</p> <p>The <code>--SetOptionOptionName--ValueNumber</code> option applies a value to a specified option.</p> <p>The following options are available:</p> <ul style="list-style-type: none"> • The <code>DebugFolder</code> option displays the path of the debug files when the <code>EnableDebugLogging</code> option is set to 1. • The <code>RestoreLocation</code> option displays the directory where the restored mailbox database files are located. • The <code>MaxHistorySessionsNumber</code> option displays how many latest restore and recovery sessions are displayed under the History Tab of the Status page the extension.

<pre>--Remove {-- SessionsSessionID [SessionID...] -- AllSessions -- RecoveryDatabase DatabaseName-- ServerComputerName}</pre>	<p>The --Remove option removes the following items:</p> <ul style="list-style-type: none"> • The --SessionsSessionID [SessionID...] option removes the specified sessions. • The --AllSessions option removes all the recovery sessions from the Granular Recovery Cache. • The --RecoveryDatabaseDatabaseName--ServerComputerName option removes a specified recovery database, or a mailbox database from the Granular Recovery Cache but not the disk.
<pre>{--Mount --Dismount} -- RecoveryDatabase RecoveryDatabaseName</pre>	<p>Mounts or dismounts the recovery database specified by --RecoveryDataBase RecoveryDatabaseName.</p>
<pre>--Search {-- MailboxMailboxName -- Cache-- MailboxMailboxName-- MailboxDBDatabaseName}</pre>	<p>Lists the mailbox user name and the database name in a backup or in the Granular Recovery Cache Management.</p>
<pre>--StartSession--Restore - -DismountRDB [true false]--MailboxDB DatabaseName--BackupID BackupVersion--Server ComputerName-- TargetLocation TargetFolderPath</pre>	<p>Restores mailbox databases, making them available for mount and then recovery.</p> <p>The --DismountRDB option controls whether the restore session dismounts any already mounted database in the Granular Recovery Cache Management. Available values are true or false.</p>

<pre>--StartSession--Recovery- -SrcMailbox SourceMailboxName-- RecoveryDatabase RecoveryDatabaseName-- MailboxDB DatabaseName {- -RecoverWholeMailbox -- Filter [FILTER_OPTIONS]}- -RecoveryTargetType [ORG_ MAILBOX DIFF_MAILBOX PST]</pre>	<p>Recovers complete mailboxes or single items.</p> <p>The <code>--SrcMailbox</code> option defines the mailbox from which the data is recovered. The <code>--RecoveryDatabase</code> option defines the Recovery Database named with the granular recovery time stamp, for example <code>RDB_ExchangeGRE_dpexchange5_2011-11-15_1</code>. The <code>--MailboxDB</code> option defines the mailbox database from which the recovery is performed. You can choose to recover a complete mailbox database by using the <code>--RecoverWholeMailbox</code> option, or filter the items using the <code>--Filter</code> option.</p> <p>Different recovery target locations are selected with the <code>--RecoveryTargetType</code> option:</p> <p>ORG_MAILBOX</p> <p><code>OrgLocation--KeepLatestMsg [true false] [RECOVERY_OPTIONS]</code></p> <p>If the <code>OrgLocation</code> option is specified, the recovery is performed to the original mailbox. If the <code>--KeepLatestMsg</code> option is set to true, the recovery session does not overwrite new messages.</p> <p>DIFF_MAILBOX</p> <p><code>DiffMailbox [RECOVERY_OPTIONS] -- TargetMailboxMailboxName</code></p> <p>If the <code>DiffMailbox</code> option is specified, the recovery is performed to another location, not the original location of items. If the <code>--TargetMailboxMailboxName</code> option is specified the recovery is performed to the new location specified by the <code>MailboxName</code>.</p> <p>PST</p> <p>If the <code>pst</code> option is specified, the items are recovered to a <code>.pst</code> file, specified with <code>--PSTFileNamePSTFileNameWithPath</code>.</p> <p>RECOVERY_OPTIONS</p> <p>The following options can be used to modify the recovery session:</p> <p>The <code>--DiffLocation</code> option sets the new location for the recovered items. The <code>--CreateNewFolderNewFolderName</code> option recovers items to a new location in a different folder than the original item's location (specified by <code>Default</code>).</p> <p>The <code>--ExistingFolderFolderName</code> uses an existing folder in the mailbox as a target for the recovered items.</p> <p>FILTER_OPTIONS</p> <p>You can filter your mailbox items by using the following filtering options:</p>
--	--

	<ul style="list-style-type: none">• The <code>Subject="Term"</code> options filters the subject of the e-mails. You can list more than one term by separating them with colons (:), for example <code>Subject="Report:Proposal"</code>, where the search terms are "Report" and "Proposal".• The <code>Contents="Term"</code> option searches the body of the e-mails for a specific term.• The <code>Attachments="Term"</code> option searches the attachments for a specific term.• The <code>Senders="SenderName"</code> option searches your mailbox for an author, or group list of the e-mail.• The <code>Recipients="RecipientName"</code> option searches the receiver of the e-mail.• The <code>Folders="FolderName"</code> option searches for a name of a folder in the mailbox.• The <code>StartDate="Date"</code> and the <code>EndDate="Date"</code> options filters the dates when the messages were sent and received.
--	--

Examples

Note: In the examples below, `ExchangeGre.CLI.exe` is omitted for simplicity.

Changing Granular Recovery Extension settings

To list available options for the granular recovery extension, specify:

```
--ListOptions
```

To get the value of the option "EnableDebugLogging", specify:

```
--GetOption EnableDebugLogging
```

To set the value for the option "EnableDebugLogging" to "1", specify:

```
--SetOption EnableDebugLogging --Value 1
```

Restoring a mailbox database from Data Protector backup

To list the backup databases, specify:

```
--List --AllBackupDatabases
```

To import a mailbox database from the Data Protector backup session with the ID 2011/09/08-5 to the temporary restore location `c:\restore`, specify:

Microsoft Exchange 2010 Server :

```
--StartSession --Restore --DismountRDB false --MailboxDB DatabaseName --BackupID  
2011/09/08-5 --Server computer.company.com --TargetLocation "c:\Restore"
```

Microsoft Exchange 2007 Server :

```
--StartSession --Restore --RemoveRSG false --MailboxDB DatabaseName --BackupID  
2011/09/08-5 --Server computer.company.com --TargetLocation "c:\Restore"
```

To import the mailbox database "Mailbox Database 0474359329" from the Data Protector backup session with the ID "2011/11/09-2" to the temporary restore location "c:\test\restore", specify:

```
--StartSession --Restore --DismountRDB true --MailboxDB "Mailbox Database 0474359329"  
--BackupID 2011/11/09-2 --Server dpexchange5.company.com --TargetLocation  
"C:\\Test\\restore1"
```

Listing mailbox database information

To list mailbox database information, such as database name, recovery database ID, server name, mount status, backup version and its size for all database entries in the granular recovery cache, specify:

```
--List --Cache
```

To list more specific details about the recovery sessions such as mailbox database name, server name, location of the restored database files, database ID, mount status, backup version, size of the database, retention period, specify:

```
--List --Cache --Verbose
```

To list mailbox database information for 20 database entries in the granular recovery cache, specify:

```
--List --Cache --EntriesCount 20
```

To list details about completed recovery sessions, such as sessions IDs, session name, type, date and time when the session started and ended, mount status, specify:

```
--List --HistorySessions
```

To list all mailboxes that are part of backed up Exchange mailbox databases, specify:

```
--List --Mailboxes
```

List all mailboxes that are part of the Exchange database "RDB_ExchangeGRE_dpexchange5_2011-11-15_1" granular recovery cache, specify:

```
--List --Mailboxes --RecoveryDatabase RDB_ExchangeGRE_dpexchange5_2011-11-15_1
```

To list all backup versions of Exchange mailbox database "Mailbox Database 0474359329", specify:

```
--List --BackupVersions --mailboxDB "Mailbox Database 0474359329"
```

To list all backup versions of backup Exchange mailbox databases that contain the mailbox "Administrator", specify:

```
--List --BackupVersions --Mailbox Administrator
```

To list backed up Exchange mailbox databases, specify:

```
--List --AllBackupDatabases
```


To show details about the completed restore session "ExchangeGRE_dpexchange5_2011-11-09_4", specify:

```
--Details --Session ExchangeGRE_dpexchange5_2011-11-09_4
```

Changing the retention period

To change the retention period of the recovery database "ExchangeGRE_dpexchange5_2011-11-09_4" on the server "dpexchange5.company.com" to end on January 15, 2012, specify:

```
--SetRP --RecoveryDatabase RDB_ExchangeGRE_dpexchange5_2011-11-15_1 --Server  
dpexchange5.company.com --Period 2012/01/15
```

Mounting a mailbox database

To mount a database, specify:

```
--MountDB --RecoveryDatabase RecoveryDatabaseName --MailboxDB DatabaseName --Server  
computer.company.com
```

To mount the restored mailbox "Mailbox Database 0474359329" database on to recovery database "RDB_ExchangeGRE_dpexchange5_2011-11-15_1" on the server "dpexchange5.company.com", specify:

```
--MountDB --RecoveryDatabase RDB_ExchangeGRE_dpexchange5_2011-11-15_1 --MailboxDB  
"Mailbox Database 0474359329" --Server dpexchange5.company.com
```

Microsoft Exchange 2007 Server :

recoverydatabase is used for simplicity reasons. In this version the recovery database is called recovery storage group.

Dismounting a mailbox database

To dismount a mailbox database from the Granular Recovery Cache and still keep the files in the temporary restore location:

```
--DismountDB --RecoveryDatabase RecoveryDatabaseName --MailboxDB DatabaseName --  
Server computer.company.com
```

To dismount the mounted Exchange database "Mailbox Database 0474359329" from the recovery database

"RDB_ExchangeGRE_dpexchange5_2011-11-15_1" on the Exchange server "dpexchange5.company.com", specify:

```
--DismountDB --RecoveryDatabase RDB_ExchangeGRE_dpexchange5_2011-11-15_1 --MailboxDB  
"Mailbox Database 0474359329" --Server dpexchange5.company.com
```

Searching a mailbox

To search a mailbox "john" in a backup, specify:

```
--Search --Mailbox john
```

To search a mailbox “john” in the mounted restored database “Mailbox Database 0474359329” in the Granular Recovery Cache, specify:

```
--Search --cache --Mailbox john --MailboxDB "Mailbox Database 0474359329"
```

Recovering items to the original location

In the following recovery examples, the recovery will be performed from the mailbox database "Mailbox Database 0474359329" mounted to the recovery database “RDB_ExchangeGRE_dpexchange5_2011-11-15_1”.

To restore mailbox database to a temporary restore location:

```
--StartSession --Restore --DismountRDB true --MailboxDB DataBaseName --BackupID ID --  
Server computer.company.com --TargetLocation C:/restore
```

To recover a complete mailbox to the original location replacing the old emails by the latest version, specify:

```
--StartSession --Recovery --SrcMailbox mailbox_name --RecoveryDatabase DBName --  
MailboxDB mailbox_name --RecoveryWholemailbox --RecoveryTargetType orgMailbox --  
OrgLocation --KeepLatestMsg true
```

To recover the folder “inbox” from the user mailbox “Administrator” to the original location, without overriding new messages, specify:

```
--StartSession --Recovery --SrcMailbox Administrator --RecoveryDatabase RDB_  
ExchangeGRE_dpexchange5_2011-11-15_1 --MailboxDB "Mailbox Database 0474359329" --  
Filter Folders=inbox --RecoveryTargetType orgMailbox --OrgLocation --KeepLatestMsg  
true
```

To recover only e-mails with the subject “market analysis” from the folder “inbox” from the user mailbox “john” to the original location, without overriding new messages, specify:

```
--StartSession --Recovery --SrcMailbox john --RecoveryDatabase RDB_ExchangeGRE_  
dpexchange5_2011-11-15_1 --MailboxDB "Mailbox Database 0474359329" --Filter  
subject="market analysis"|folders="inbox" --RecoveryTargetType orgMailbox --  
OrgLocation --KeepLatestMsg true
```

Recovering items to another location

To recover a complete mailbox to a new location for example a new mailbox, specify:

```
--StartSession --Recovery --SrcMailbox mailbox_name --RecoveryDatabase  
RecoveryDatabaseName --MailboxDB mailbox_name --RecoveryWholemailbox --  
RecoveryTargetType
```

To recover the complete user mailbox “Administrator” to the mailbox “john” to a new folder “recovered mailbox”, specify:

```
--StartSession --Recovery --SrcMailbox Administrator --RecoveryDatabase RDB_  
ExchangeGRE_dpexchange5_2011-11-15_1 --MailboxDB "Mailbox Database 0474359329" --  
Filter Folders=inbox --RecoveryTargetType diffMailbox --DiffLocation --  
CreateNewFolder "recovered mailbox" --TargetMailbox john
```

To recover only e-mails with the subject “market analysis” from the folder “inbox” from the user mailbox “john” to a different default location, without overriding new messages, specify:

```
--StartSession --Recovery --SrcMailbox john --RecoveryDatabase RDB_ExchangeGRE_
dpexchange5_2011-11-15_1 --MailboxDB "Mailbox Database 0474359329" --Filter
subject="market analysis"|folders="inbox" --RecoveryTargetType orgMailbox --
OrgLocation --KeepLatestMsg true --DiffLocation --CreateNewFolder default
```

To recover only e-mails with the subject “market analysis” from the folder “inbox” from the user mailbox “john” to a different (existing) folder “recovered data items”, without overriding new messages, specify:

```
--StartSession --Recovery --SrcMailbox john --RecoveryDatabase RDB_ExchangeGRE_
dpexchange5_2011-11-15_1 --MailboxDB "Mailbox Database 0474359329" --Filter
subject="market analysis"|folders="inbox" --RecoveryTargetType orgMailbox --
OrgLocation --KeepLatestMsg true --DiffLocation --ExistingFolder "recovered data
items"
```

To recover a complete mailbox to a .pst file, specify:

```
--StartSession --Recovery --SrcMailbox mailbox_name --RecoveryDatabase
RecoveryDatabaseName --RecoveryDatabase MailboxDB mailbox_name --RecoveryWholemailbox
--RecoveryTargetType pst
```

To recover the complete user mailbox “john” to the file “C:\recovered\john.pst”, specify:

```
--StartSession --Recovery --SrcMailbox john --RecoveryDatabase RDB_ExchangeGRE_
dpexchange5_2011-11-15_1 --MailboxDB "Mailbox Database 0474359329" --
RecoverWholeMailbox --RecoveryTargetType pst --PSTfilename "C:\\recovered\\john.pst"
```

Removing sessions

To remove a completed recovery session with the ID 2011/09/08-5 from the Granular Recovery Cache, specify:

```
--Remove --Session 2011/09/08-5
```

To remove all the recovery sessions from the Granular Recovery Cache, specify:

```
--Remove --AllSessions
```

Removing recovery databases

To remove the mailbox database “RDB_ExchangeGRE_dpexchange5_2011-11-15_1” from disk on the server “dpexchange5.company.com”, specify:

```
--Remove --RecoveryDatabase RDB_ExchangeGRE_dpexchange5_2011-11-15_1 --Server
dpexchange5.company.com
```


Chapter 7: Troubleshooting

This chapter lists general checks and verifications, plus problems you might encounter when using the Data Protector Granular Recovery Extension for Microsoft Exchange Server.

- For Microsoft Exchange Server troubleshooting information, see the troubleshooting sections of the Microsoft Exchange Server parts of the *HP Data Protector Integration Guide for Microsoft Applications: SQL Server, SharePoint Server, and Exchange Server*.
- For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

Before you begin

- To enable debugging, see "[Enabling debugging option](#)" below.
- Make sure that the latest official Data Protector patches are installed. See the *HP Data Protector Help* index: "patches" on how to verify this.
- For general Data Protector limitations, as well as recognized issues and workarounds, see the *HP Data Protector Product Announcements, Software Notes, and References*
- For an up-to-date list of supported versions, platforms, and other information, see <http://support.openview.hp.com/selfsolve/manuals>.

Debugging

Enabling debugging option

1. To enable the debugging option, click **Settings** in the console tree. The Granular Recovery Settings page is displayed.
2. Select the **Enable Debug Logs** option. The Debug Logs folder name field is activated. Specify a new location of the folder, and click **Save**.

For a list of known issues and workarounds, see "[Known issues and workarounds](#)" on the next page.

Known issues and workarounds

Search Criteria Results page remains empty after at least one search keyword is entered

Problem

In the Granular Recovery wizard, in the Mailbox Search Criteria page, after you enter a term to search for and then click **Next**, no results appear in the Search Results page even if the search criteria are met for some items.

Action

To display a list of search results, proceed as follows:

1. Dismount the recovery database. For detailed procedure see ["Dismounting databases" on page 37](#).
2. Using the Cache Management page of the Granular Recovery Extension GUI, identify the folder where the database has been restore into. From this folder, remove the subfolder that has the string CatalogData in its name.
3. In the Exchange Management Shell, run the following command:

```
set-mailboxDatabase DatabaseName -indexenabled $false.
```
4. Mount the recovery database again. For detailed procedure, see ["Mounting databases" on page 32](#).
5. Start the recovery once again and follow the Granular Recovery wizard.
6. In the Mailbox Search Criteria page, re-enter the search keywords and click **Next**.

Manual removal of temporary mailboxes created by the extension

Problem

Normally, temporary mailboxes created by the extension during the recovery process, are also deleted automatically after the process completes. However, in exceptional cases (for example, if the extension stops working), the temporary mailboxes created by the extension, are not deleted automatically after the process completes and have to be removed manually.

Action

You can identify such mailboxes by the prefix DP_Recovery or DP_SEARCH.

In the Microsoft Exchange Server 2010 environment, use the Exchange Management Console (EMC) to remove the redundant mailboxes manually.

In the Microsoft Exchange Server 2013 environment, use the Exchange Administration Center (EAC) or the Exchange Management Shell (EMS) to remove the redundant mailboxes manually.

Search for mailbox items fails and reports an error

Problem

In the Microsoft Exchange Server 2010 environment, in the Granular Recovery wizard, in the Mailbox Search Criteria page, after you trigger a search using the specified criteria, the search fails with the following error:

```
ERROR debugs - Powershell error : Restore-Mailbox : Error was found for DP_Recovery_ExchangeGRE_2011-10-03_38 (DP_Recovery_ExchangeGRE_2011-10-03_38@mail.hp-dp.com) because: Error occurred in the step: Opening source mailbox. Failed to open mailbox by GUID with error: The operation failed. error code: -1056749260 + CategoryInfo: InvalidOperation: (0:Int32) [Restore-Mailbox], RecipientTaskException + FullyQualifiedErrorId: DD312EA7,Microsoft.Exchange.Management.RecipientTasks.RestoreMailbox
```

Action

Install the Update Rollup 2 for Exchange Server 2010 Service Pack 1 (KB2425179) or a subsequent update. You can obtain the Update Rollup 2 from the website <http://www.microsoft.com/download/en/details.aspx?id=12938>.

Mailboxes are missing from the list in the Import from Backup wizard

Problem

In the Import from Backup wizard, in the Mailbox Selection page, when you browse for user mailboxes after clicking **Advanced**, some mailboxes are missing from the Mailboxes tree, although they exist in the backup image of the Exchange Server database.

There can be different reasons for such a problem, including time synchronization problems within the Data Protector cell and Exchange Server problems in retrieving the mailbox metadata.

Action

If you know to which mailbox database the desired mailbox belongs, you can start the process anew by selecting the appropriate database. Follow the steps:

1. Check **Back** to return to the Introduction page of the wizard.
2. Select **Database selection**, click **Next**, and follow the wizard to complete the importing process.

Mounting a restored database fails

Problem

In the Microsoft Exchange Server 2013 environment, mounting a restored database might end up with an error.

The problem occurs when the restored database is in the "dirty-shutdown state". The database should be in the "clean-shutdown state" to be mounted successfully.

Action

1. Bring the restored database to the "clean-shutdown state" by executing the Microsoft Exchange Server `eseutil.exe recovery` command.

For more information on the `eseutil.exe` utility, see the Microsoft Exchange Server documentation.

2. Retry to mount the restored database.

Interprocess communication error being reported by the GUI

Problem

After attempting to trigger an operation from the Granular Recovery Extension graphical user interface (GUI), the following error message is displayed in a dialog box:

The communication object, `System.ServiceModel.Channels.ServiceChannel`, cannot be used for communication because it is in the `Faulted` state.

All subsequent user actions in the GUI fail with the same error.

This error is reported when the Granular Recovery Extension GUI is open for a period of time that is longer than the Internet Information Services (IIS) recycling time period. The IIS unloads the Exchange GRE Web Service among other web services, resulting in a communication failure between the service and the GUI.

Action

Close and then re-launch the Granular Recovery Extension GUI.

An Exchange GRE recovery or restore operation fails due to insufficient permission

Problem

If an Granular Recovery Extension recovery or restore operation fails due to insufficient permissions even when the user who executes the Granular Recovery Extension (either GUI or CLI) has sufficient permissions, the issue may arise due to insufficient permission for the local SYSTEM account.

This account needs appropriate permissions due to the following:

- The Granular Recovery Extension Web Service runs with Local SYSTEM privileges.
- The Granular Recovery Extension Web Service allows the Data Protector Inet service, which by default runs under the Windows local System user account, to execute or launch the Data Protector Exchange Integration agent. The restore session is performed using the same user account.

Action

Grant the local SYSTEM user account appropriate permissions to restore Microsoft Exchange Server databases and create recovery databases:

1. Close the currently running Exchange Granular Recovery Extension client (GUI or CLI).
2. Stop the Granular Recovery Extension Web Service.
3. Provide appropriate permissions for the local SYSTEM account.
4. Start the Granular Recovery Extension and continue the Granular Recovery Extension operations.

When a request comes from the GUI or CLI, the Internet Information Services (IIS) start the Granular Recovery Extension Web Service automatically.

The message Adding snap-in to console... is displayed for a long time

Problem

When you open the Granular Recovery Extension graphical user interface (GUI), the message MMC cannot initialize the snap-in is displayed, followed by the Adding snap-in to console... which is displayed for a long time.

When you open the GUI and locate the console tree, the Cache Management, Status and Settings nodes are not loaded.

The issue may appear if Internet Information Services (IIS) or if its associated services are not running and the Granular Recovery Extension GUI cannot communicate with the Exchange GREWeb Service.

Action

Make sure that IIS and its associated services are up and running:

1. Open the Internet Information Services (IIS) Manager, locate and select the ExchangeGre node in the console tree under the Default Web Site. The Exchange GRE Web Service home page is displayed.
2. Right-click the ExchangeGre node, click Manage Application, and then click Browse.

The snap-in is added to the MMC. When you open the GUI, the Cache Management, Status and Settings nodes are displayed. The message is not displayed anymore.

If during the step 1, the Exchange GRE Web Service home page is not displayed:

1. When the Adding snap-in to console... is displayed, click Cancel to close the message.
2. Open the Server Manager, select the Roles node. Under the Roles Summary locate the WebServer (IIS) and select it.
3. The Web Server (IIS) is displayed. Under the System Services, verify if all the services are up and running, including the Application Host Helper Service, IIS Admin Service, and the World Wide Web Publishing Service.
4. Re-open the Granular Recovery Extension GUI.

The About HP Data Protector Granular Recovery Extension for Microsoft Exchange Server does not display the product build number

Problem

You select Help **About HP Data Protector Granular Recovery Extension for Microsoft Exchange Server**, but the product version is not displayed.

The issue is caused by a Microsoft Management Console (MMC) known issue, due to a string caching mechanism in the MUI cache. The registry is not updated automatically after the upgrade. The MUI cache does not get cleared and the product version is not displayed in the MMC snap-in.

Action

To manually delete strings in the MUI cache and re-install the snap-in:

1. Click the Start button, click Run, type REGEDIT, and click OK.
2. In the Registry Editor, locate one of the following hierarchy (if they exist):

```
HKEY_USERS\S-1-5-21-61196776-1057610366-2591919248-500_Classes\Local  
Settings\Software\Microsoft\Windows\Shell\MuiCache
```

```
HKEY_USERS\S-1-5-21-2765349584-3720068851-1520285658-500_Classes\Local  
Settings\MuiCache\96\52C64B7E
```

3. Delete the following registry keys of the Granular Recovery Extension for Microsoft Exchange Server:

```
@C:\Program Files\Hewlett-Packard\Exchange Granular Recovery  
Extension\bin\GreSnapInResource.dll, -114
```

```
@C:\Program Files\Hewlett-Packard\Exchange Granular Recovery  
Extension\bin\GreSnapInResource.dll, -115
```

```
@C:\Program Files\Hewlett-Packard\Exchange Granular Recovery  
Extension\bin\GreSnapInResource.dll, -116
```

```
@C:\Program Files\Hewlett-Packard\Exchange Granular Recovery  
Extension\bin\GreSnapInResource.dll, -117
```

Close the Registry Editor.

4. Run the following command:

```
%windir%\Microsoft.NET\Framework64\v2.0.50727\InstallUtil.exe "C:\Program  
Files\Hewlett-Packard\Exchange Granular Recovery  
Extension\bin\ExchangeGre.MmcGui.dll" /install
```

5. Click the Help menu and then click the **About HP Data Protector Granular Recovery Extension for Microsoft Exchange Server** again, the product build number is displayed.

Glossary

A

access rights

See user rights.

ACSL (StorageTek specific term)

The Automated Cartridge System Library Server (ACSL) software that manages the Automated Cartridge System (ACS).

Active Directory (Windows specific term)

The directory service in a Windows network. It contains information about resources on the network and makes them accessible to users and applications. The directory services provide a consistent way to name, describe, locate, access, and manage resources regardless of the physical system they reside on.

AES 256-bit encryption

The Data Protector software encryption, based on the AES-CTR (Advanced Encryption Standard in Counter Mode) encryption algorithm that uses random keys of 256-bit length. The same key is used for both encryption and decryption. With AES 256-bit encryption, data is encrypted before it is transferred over a network and before it is written to media.

AML (ADIC/GRAU specific term)

Automated Mixed-Media library.

AMU (ADIC/GRAU specific term)

Archive Management Unit.

application agent

A component needed on a client system to back up or restore online database integrations. See also Disk Agent.

application system (ZDB specific term)

A system the application or database runs on. The application or database data is located on source volumes. See also backup system and source volume.

archive logging (Lotus Domino Server specific term)

Lotus Domino Server database mode where transaction log files are overwritten only after they have been backed up.

archived log files (Data Protector specific term)

Files that keep track of changes made to the Data Protector Internal Database (IDB). They are used for online or offline IDB restore and recovery where the IDB needs to be recreated either in its latest possible state, beyond the time of the most recent IDB backup session, or in a state between the times of two consecutive IDB backup sessions.

archived redo log (Oracle specific term)

Also called offline redo log. If the Oracle database operates in the ARCHIVELOG mode, as each online redo log is filled, it is copied to an archived log destination. This copy is the archived redo log. The presence or absence of an archived redo log is determined by the mode the database is using: ARCHIVELOG - The filled online redo log files are archived before they are reused. The database can be recovered if an instance or a disk fails. A "hot" backup can be performed only when the database is running in this mode. NOARCHIVELOG - The filled online redo log files are not archived. See also online redo log.

ASR set

A collection of files stored on several diskettes, required for proper reconfiguration of the replacement disk (disk partitioning and logical volume configuration) and automatic recovery of

the original system configuration and user data that was backed up during the full client backup. These files are stored as an ASR archive file on the Cell Manager in the directory <Data_Protector_program_data>\Config\Server\dr\asr (Windows systems) or /etc/opt/omni/server/dr/asr (UNIX systems), as well as on the backup medium. After a disaster occurs, the ASR archive file is extracted to diskettes which you need to perform ASR.

audit logs

Data files to which auditing information is stored.

audit report

User-readable output of auditing information created from data stored in audit log files.

auditing information

Data about every backup session that was performed over an extended, user-defined period for the whole Data Protector cell.

autochanger

See library.

autoloader

See library.

Automatic Storage Management (ASM) (Oracle specific term)

A filesystem and volume manager integrated into Oracle which manages Oracle database files. It eliminates complexity associated with data and disk management and optimizes performance by providing striping and mirroring capabilities.

auxiliary disk

A bootable disk that has a minimal operating system with networking and

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.

B

BACKINT (SAP R/3 specific term)

A Data Protector interface program that lets the SAP R/3 backup programs communicate with the Data Protector software via calls to an open interface. For backup and restore, SAP R/3 programs issue commands through the Data Protector backint interface.

backup API (Oracle specific term)

The Oracle interface between the Oracle backup/restore utility and the backup/restore media management layer. The interface defines a set of routines to allow reading and writing of data to media and also to allow creation, searching, and removing of backup files.

backup chain

See restore chain.

backup device

A device configured for use with Data Protector that can write data to and read data from storage media. This can be, for example, a standalone DDS/DAT drive or a library.

backup generation

One backup generation includes one full backup and all incremental backups until the next full backup.

backup ID

An identifier of an integration object that equals the session ID of the backup of this object. The backup ID is preserved when an object is copied, exported, or imported.

backup object

A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database/application entity or a disk image (rawdisk). A backup object is defined by: client name (hostname of the Data Protector client where the backup object resides), mount point (for filesystem objects - the access point in a directory structure on the client where the backup object is located (drive on Windows systems and mount point on UNIX systems), for integration objects - backup stream identification, indicating the backed up database/application items), description (for filesystem objects - uniquely defines objects with identical client name and mount point, for integration objects - displays the integration type), and type (for filesystem objects - filesystem type, for integration objects - "Bar").

backup owner

Each backup object in the IDB has an owner. The default owner of a backup is the user who starts the backup session.

backup session

A process that creates a copy of data on storage media. The activities are specified in a backup specification or an interactive session. The result of a backup session is a set of media, which was written to, also called the backup or media set. See also backup specification, incremental backup, and full backup.

backup set

A complete set of integration objects associated with a backup.

backup set (Oracle specific term)

A logical grouping of backed up files that are created using the RMAN backup command. A backup set is a complete set of files associated with a backup. The files can be multiplexed to increase performance. A backup set contains either datafiles or archived logs, but not both together.

backup specification

A list of objects to be backed up, together with a set of devices or drives to be used; backup options for all objects in the specification; and days and time that you want backups to be performed. The objects are entire disks/volumes or parts of them such as files, directories, or even the Windows Registry (for example). File selection lists, such as include-lists and exclude-lists, can be specified. All clients configured in one backup specification are backed up at the same time in one backup session using the same backup type (full or incremental).

backup system (ZDB specific term)

A system connected to target volumes of one or multiple application systems. The backup system is typically connected to a backup device to perform the backup of the data in a replica. See also application system, target volume, and replica.

backup types

See incremental backup, differential backup, transaction backup, full backup, and delta backup.

backup view

Data Protector provides different views of your backup specifications: By Type - according to the type of data available for backups/templates. Default view. By Group - according to the group to which backup specifications/templates belong. By Name - according to the name of

backup specifications/templates. By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.

BC (EMC Symmetrix specific term)

Business Continuances are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices.

BC Process (EMC Symmetrix specific term)

A protected storage environment solution, which uses specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices. See also BCV.

BCV (EMC Symmetrix specific term)

Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splittable mirrors of the primary EMC Symmetrix SLDs that need to be protected. See also BC and BC Process.

Boolean operators

The Boolean operators for the full text search functionality of the Data Protector Help system are AND, OR, NOT, and NEAR. Used when searching, they enable you to define your query precisely by creating a relationship between search terms. If no operator is specified in a multi-word search, AND is used by default. For example, the query manual disaster recovery is equivalent to manual AND disaster AND recovery.

boot volume/disk/partition

A volume/disk/partition containing files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.

BRARCHIVE (SAP R/3 specific term)

An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process. See also BRBACKUP and BRRESTORE.

BRBACKUP (SAP R/3 specific term)

An SAP R/3 backup tool that allows an online or offline backup of the control file, of individual data files, or of all tablespaces and, if necessary, of the online redo log files. See also BRARCHIVE and BRRESTORE.

BRRESTORE (SAP R/3 specific term)

An SAP R/3 tool that can be used to restore files of the following type: Database data files, control files, and online redo log files saved with BRBACKUP; Redo log files archived with BRARCHIVE; Non-database files saved with BRBACKUP. You can specify files, tablespaces, complete backups, log sequence numbers of redo log files, or the session ID of the backup. See also BRARCHIVE and BRBACKUP.

BSM

The Data Protector Backup Session Manager, which controls the backup session. This process always runs on the Cell Manager system.

C

CAP (StorageTek specific term)

The Cartridge Access Port built into the door panel of a library. The purpose is to enter or eject media.

Catalog Database (CDB)

A part of the Data Protector Internal Database (IDB) that contains information about backup, restore, object copy, object consolidation, object verification, and media management sessions. This part of the IDB is always local to the cell. It is stored in the embedded database. See also MMDB.

catalog protection

Defines how long information about backed up data (such as filenames and file attributes) is kept in the IDB. See also data protection.

CDB

See Catalog Database (CDB).

CDF file (UNIX systems specific term)

A Context Dependent File is a file consisting of several files grouped under the same pathname. The system ordinarily selects one of the files using the context of the process. This mechanism allows machine dependent executables, system data, and device files to work correctly from all hosts in a cluster while using the same pathname.

cell

A set of systems that are under the control of a Cell Manager. The cell typically represents the systems on a site or an organizational entity, which are connected to the same LAN. Central control is available to administer the backup and restore policies and tasks.

Cell Manager

The main system in the cell where the essential Data Protector software is installed and from which all backup and restore activities are managed. The GUI used for management tasks can be located on a different system. Each cell has one Cell Manager system.

centralized licensing

Data Protector allows you to configure centralized licensing for the whole enterprise environment consisting of several cells. All Data Protector licenses are installed and kept on the Enterprise Cell Manager system. You can then allocate licenses to specific cells to suit your needs. See also MoM.

Centralized Media Management Database (CMMDB)

See CMMDB.

Certificate Server

A Windows certificate server can be installed and configured to provide certificates for clients. It provides customizable services for issuing and managing certificates for the enterprise. These services issue, revoke, and manage certificates employed in public key-based cryptography technologies.

Change Journal (Windows specific term)

A Windows filesystem feature that logs a record of each change as it occurs to the files and directories on a local NTFS volume.

Change Log Provider

A module that can be queried to determine which objects on a filesystem have been created, modified, or deleted.

channel (Oracle specific term)

An Oracle Recovery Manager resource allocation. Every allocated channel starts a new Oracle process, which performs backup, restore, and recovery actions. The type of channel allocated determines the type of media used: type 'disk' or type 'sbt_tape'. If the specified channel is of type 'sbt_tape' and Oracle is integrated with Data Protector, the server process will attempt to read backups from or write data files to Data Protector.

circular logging (Microsoft Exchange Server and Lotus Domino Server specific term)

Circular logging is a Microsoft Exchange Server database and Lotus Domino Server database mode, in which the transaction log file contents are periodically overwritten after the corresponding data is committed to the database. Circular logging reduces disk storage space requirements.

client backup

A backup of all volumes (filesystems) mounted on a Data Protector client. What is actually backed up depends on how you select objects in a backup specification. If you select the check box next to the client system name, a single backup object of the Client System type is created. As a result, at the time of the backup, Data Protector first detects all volumes that are mounted on the selected client and then backs them up. On Windows clients, CONFIGURATION is also backed up. If you individually select all volumes that are mounted on the client system, a separate backup object of the Filesystem type is created for each volume. As a result, at the time of the backup, only the selected volumes are backed up. Volumes that have been potentially mounted on the client after the

backup specification was created are not backed up.

client or client system

Any system configured with any Data Protector functionality and configured in a cell.

cluster continuous replication (Microsoft Exchange Server specific term)

Cluster continuous replication (CCR) is a high availability solution that uses cluster management and failover options to create and maintain an exact copy (CCR copy) of a storage group. A storage group is replicated to a separate server. CCR removes any single point of failure in your Exchange back-end servers. You can perform backups using VSS on your passive Exchange Server node where a CCR copy is located and thus reducing the load on the active node. A CCR copy is used for disaster recovery since you can switch to the CCR copy in a few seconds. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and can be backed up (using VSS) like an ordinary storage group. See also Exchange Replication Service and local continuous replication.

cluster-aware application

It is an application that supports the cluster Application Programming Interface. Each cluster-aware application declares its own critical resources (disk volumes (on Microsoft Cluster Server), volume groups (on HP Serviceguard), application services, IP names and addresses ...).

CMD script for Informix Server (Informix Server specific term)

A Windows CMD script that is created in INFORMIXDIR when an Informix Server database is configured. The CMD script

is a set of system commands that export environment variables for Informix Server.

CMMDB

The Data Protector Centralized Media Management Database (CMMDB) is the result of merging MMDBs from several cells in the MoM environment. It allows you to share high-end devices and media across multiple cells in a MoM environment. One cell can control the robotics, serving the devices that are connected to systems in other cells. The CMMDB must reside on the Manager-of-Managers. A reliable network connection between the MoM cell and the other Data Protector cells is highly recommended. See also MoM.

COM+ Class Registration Database (Windows specific term)

The COM+ Class Registration Database and the Windows Registry store COM+ application attributes, class attributes, and computer-level attributes. This guaranties consistency among these attributes and provides common operation on top of these attributes.

command device (HP P9000 XP Disk Array Family specific term)

A dedicated volume in the disk array which acts as the interface between a management application and the disk array's storage system. It cannot be used for data storage and only accepts requests for operations that are then executed by the disk array.

command-line interface (CLI)

A set commands that you can use in shell scripts to perform Data Protector configuration, backup, restore, and management tasks.

concurrency

See Disk Agent concurrency.

container (HP P6000 EVA Disk Array Family specific term)

Space on a disk array, which is pre-allocated for later use as a standard snapshot, vsnap, or snapclone.

control file (Oracle and SAP R/3 specific term)

A data file that contains entries specifying the physical structure of the database. It provides database consistency information used for recovery.

copy set (HP P6000 EVA Disk Array Family specific term)

A pair that consists of the source volumes on a local P6000 EVA and their replica on a remote P6000 EVA. See also source volume, replica, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

CRS

The Data Protector Cell Request Server process (service), which runs on the Cell Manager, starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager. The CRS runs under the account root on UNIX systems. On Windows systems it runs under the account of the user, specified at installation time.

CSM

The Data Protector Copy and Consolidation Session Manager process controls the object copy and object consolidation sessions and runs on the Cell Manager system.

D**data file (Oracle and SAP R/3 specific term)**

A physical file created by Oracle that contains data structures such as tables and indexes. A data file can only belong to one Oracle database.

data protection

Defines how long the backed up data on media remains protected, that is, Data Protector will not overwrite it. When the protection expires, Data Protector will be able to reuse the media in one of the next backup sessions. See also catalog protection.

Data Protector user account

You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

data replication (DR) group (HP P6000 EVA Disk Array Family specific term)

A logical grouping of HP P6000 EVA Disk Array Family virtual disks. It can contain up to eight copy sets provided they have common characteristics and share a common HP CA P6000 EVA log. See also copy set.

data stream

Sequence of data transferred over the communication channel.

Data_Protector_home

A reference to the directory containing Data Protector program files (on Windows Vista, Windows 7, Windows Server 2008, and Windows Server 2012) or the directory containing Data Protector program files and data files (on other Windows operating systems). Its default path is %ProgramFiles%\OmniBack, but the path can be changed in the Data Protector Setup Wizard at installation time. See also Data_Protector_program_data.

Data_Protector_program_data

A reference to the directory containing Data Protector data files on Windows Vista, Windows 7, Windows Server 2008, and Windows Server 2012. Its default path is %ProgramData%\OmniBack, but the path can be changed in the Data Protector Setup Wizard at installation time. See also Data_Protector_home.

database library

A Data Protector set of routines that enables data transfer between Data Protector and a server of an online database integration, for example, Oracle Server.

database parallelism

More than one database is backed up at a time if the number of available devices allows you to perform backups in parallel.

database server

A computer with a large database stored on it, such as the SAP R/3 or Microsoft SQL database. A server has a database that can be accessed by clients.

Dbobject (Informix Server specific term)

An Informix Server physical database object. It can be a blob space, db space, or logical log file.

DC directory

A directory that contains DC binary files, one for each configured Data Protector backup medium. DC directories constitute the Detail Catalog Binary Files part of the Data Protector Internal Database. See also Detail Catalog Binary Files (DCBF) and Internal Database (IDB).

DCBF

See Detail Catalog Binary Files (DCBF).

delta backup

A delta backup is a backup containing all the changes made to the database from the last backup of any type. See also backup types.

Detail Catalog Binary Files (DCBF)

A part of the Data Protector Internal Database that stores names, versions, and metadata of the backed up items. It consists of DC directories with DC binary files. See also DC directory and Internal Database (IDB).

device

See backup device.

device chain

A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.

device group (EMC Symmetrix specific term)

A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to: -Identify and work with a subset of the available

EMC Symmetrix devices. -Get configuration, status, and performance statistics by device group. -Issue control operations that apply to all devices in the device group.

device streaming

A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes writing to the tape, and so on. If data is written to the tape slower than it is delivered to the device then the device is streaming. Streaming significantly improves the use of space and the performance of the device.

DHCP server

A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic IP address assignment and network configuration for DHCP clients.

differential backup

An incremental backup that backs up changes made since the last full backup. To perform this type of backup, specify the Incr1 backup type. See also incremental backup.

differential backup (Microsoft SQL Server specific term)

A database backup that records only the data changes made to the database after the last full database backup. See also backup types.

differential database backup

A differential database backup records only those data changes made to the database after the last full database backup.

directory junction (Windows specific term)

Directory junctions use the reparse point concept of Windows. An NTFS 5 directory junction allows you to redirect a directory/file request to another location.

disaster recovery

A process to restore a client's main system disk to a state close to the time when a (full) backup was performed.

disaster recovery operating system

See DR OS.

Disk Agent

A component needed on a client to back it up and restore it. The Disk Agent controls reading from and writing to a disk. During a backup session, the Disk Agent reads data from a disk and sends it to the Media Agent, which then moves it to the device. During a restore session the Disk Agent receives data from the Media Agent and writes it to the disk. During an object verification session the Disk Agent receives data from the Media Agent and performs the verification process, but no data is written to disk.

Disk Agent concurrency

The number of Disk Agents that are allowed to send data to one Media Agent concurrently.

disk group (Veritas Volume Manager specific term)

The basic unit of data storage in VxVM system. A disk group can consist of one or more physical volumes. There can be more than one disk group on the system.

disk image (rawdisk) backup

A high-speed backup where Data Protector backs up files as bitmap images. A disk image (rawdisk) backup does not track the files and directory

structure stored on the disk, but stores a disk image structure on byte level. You can perform a disk image backup of either specific disk sections or a complete disk.

disk quota

A concept to manage disk space consumption for all or a subset of users on a computer system. This concept is used by several operating system platforms.

disk staging

The process of backing up data in several phases to improve the performance of backups and restores, reduce costs of storing the backed up data, and increase the data availability and accessibility for restore. The backup stages consist of backing up data to one media type first (for example disk) and later copying it to a different media type (for example tape).

distributed file media format

A media format, available with the file library, which supports a space efficient type of synthetic backup called virtual full backup. Using this format is a prerequisite for virtual full backup. See also virtual full backup.

Distributed File System (DFS)

A service that connects file shares into a single namespace. The file shares can reside on the same or on different computers. DFS provides client access to the resources in a location-transparent manner.

DMZ

The Demilitarized Zone (DMZ) is a network inserted as a "neutral zone" between a company's private network (intranet) and the outside public network (Internet). It prevents outside users from getting direct access to company servers in the intranet.

DNS server

In the DNS client-server model, this is the server containing information about a portion of the DNS database that makes computer names available to client resolvers querying for name resolution across the Internet.

DR image

Data required for temporary disaster recovery operating system (DR OS) installation and configuration.

DR OS

An operating system environment in which disaster recovery runs. It provides Data Protector with a basic runtime environment (disk, network, tape, and filesystem access). It has to be installed on disk or loaded into memory and configured before the Data Protector disaster recovery can be performed. DR OS can be either temporary or active. A temporary DR OS is used exclusively as a host environment for some other operating system restore along with the target operating system configuration data. It is deleted after the target system is restored to the original system configuration. An active DR OS not only hosts the Data Protector disaster recovery process but can also be a part of the restored system because it replaces its own configuration data with the original configuration data.

drive

A physical unit that receives data from a computer system and can write it onto a magnetic medium. It can also read the data from the medium and send it to the computer system.

drive index

A number that identifies the mechanical position of a drive inside a library device.

This number is used by the robotic control to access a drive.

drive-based encryption

The Data Protector drive-based encryption uses the encryption functionality of the drive. While performing the backup, the drive encrypts both the data and the metadata that is written to the medium.

E

EMC Symmetrix Agent

A Data Protector software module that prepares the EMC Symmetrix environment for backup and restore operations.

emergency boot file (Informix Server specific term)

The Informix Server configuration file `ixbar.<server_id>` that resides in the directory `<INFORMIXDIR>/etc` (on Windows systems) or `<INFORMIXDIR>\etc` (on UNIX systems). `<INFORMIXDIR>` is the Informix Server home directory and `<server_id>` is the value of the `SERVERNUM` configuration parameter. Each line of the emergency boot file corresponds to one backup object.

encrypted control communication

Data Protector secure communication between the clients in the Data Protector cell is based on a Secure Socket Layer (SSL) that uses export grade SSLv3 algorithms to encrypt control communication. Control communication in a Data Protector cell is all communication between Data Protector processes, except the data transfer from Disk Agent (and Integrations) to Media Agent, and the other way round.

encryption key

A 256-bit randomly generated number used by the Data Protector encryption algorithm to encode information during backups for which AES 256-bit software encryption or drive-based encryption has been specified. The same key is used for subsequent decryption of the information. Encryption keys for a Data Protector cell are stored in a central keystore on the Cell Manager.

encryption KeyID-StoreID

Combined identifier used by the Data Protector Key Management Server to identify and administer encryption keys used by Data Protector. KeyID identifies the key within the keystore. StoreID identifies the keystore on the Cell Manager. If Data Protector has been upgraded from an earlier version with encryption functionality, there may be several StoreIDs used on the same Cell Manager.

enhanced incremental backup

Conventional incremental backup backs up files that have changed since a previous backup, but has certain limitations in detection of changes. Unlike conventional incremental backup, enhanced incremental backup reliably detects and backs up also renamed and moved files, as well as files with changes in attributes.

enterprise backup environment

Several cells can be grouped together and managed from a central cell. The enterprise backup environment includes all clients located in several Data Protector cells which are managed and administered from a central cell using the Manager-of-Managers concept. See also MoM.

Event Log (Data Protector Event Log)

A central repository of all Data Protector related notifications. By default, all notifications are sent to the Event Log. The Event Log is accessible only to Data Protector users in the Admin group and to Data Protector users who are granted the Reporting and notifications user rights. You can view or delete all events in the Event Log.

Event Logs (Windows specific term)

Files in which Windows logs all events, such as the starting or stopping of services and the logging on and off of users. Data Protector can back up Windows Event Logs as part of the Windows configuration backup.

Exchange Replication Service (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that represents storage groups that were replicated using either local continuous replication (LCR) or cluster continuous replication (CCR) technology. See also cluster continuous replication and local continuous replication.

exchanger

See library.

exporting media

A process that removes all data about backup sessions, such as systems, objects, and filenames, which reside on the media from the IDB. Information about the media and their relation to a pool is also removed from the IDB. The data on the media remains unchanged. See also importing media.

**Extensible Storage Engine (ESE)
(Microsoft Exchange Server specific
term)**

A database technology used as a storage system for information exchange in Microsoft Exchange Server.

F**failover**

Transferring of the most important cluster data, called group (on Windows systems) or package (on UNIX systems) from one cluster node to another. A failover can occur mostly because of software or hardware failures or maintenance on the primary node.

**failover (HP P6000 EVA Disk Array
Family specific term)**

An operation that reverses the roles of source and destination in HP Continuous Access + Business Copy (CA+BC) P6000 EVA configurations. See also HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

FC bridge

See Fibre Channel bridge.

Fibre Channel

An ANSI standard for high-speed computer interconnection. Using either optical or copper cables, it allows the high speed bi-directional transmission of large data files and can be deployed between sites kilometers apart. Fibre Channel connects nodes using three physical topologies: point-to-point, loop, and switched.

Fibre Channel bridge

A Fibre Channel bridge or multiplexer provides the ability to migrate existing parallel SCSI devices, like RAID arrays, solid state disks (SSD), and tape libraries

to a Fibre Channel environment. On one side of the bridge or multiplexer there is a Fibre Channel interface while on the other side there are parallel SCSI ports. The bridge or multiplexer enables SCSI packets to be moved between the Fibre Channel and parallel SCSI devices.

file depot

A file containing the data from a backup to a file library device.

file jukebox device

A device residing on disk consisting of multiple slots used to store file media.

file library device

A device which resides on a disk emulating a library with several media, hence containing multiple files, referred to as file depots.

File Replication Service (FRS)

A Windows service that replicates the domain controller store logon scripts and group policies. FRS also enables replication of Distributed File System (DFS) shares between systems and allows any server to perform replication activity.

file tree walk

The process of traversing a filesystem to determine which objects have been created, modified, or deleted.

file version

The same file can be backed up multiple times in case of full backups and incremental backups (if the file changed). If the log level ALL is selected for backup, Data Protector retains one entry in the IDB for the filename itself and one for each version (date/time) of the file.

filesystem

The organization of files on a hard disk. A filesystem is backed up so that the file attributes and the file contents are stored on the backup media.

first-level mirror (HP P9000 XP Disk Array Family specific term)

A mirror of an internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which can be further mirrored itself, producing second-level mirrors. For Data Protector zero downtime backup and instant recovery purposes, only first-level mirrors can be used. See also primary volume and mirror unit (MU) number.

flash recovery area (Oracle specific term)

A directory, filesystem, or Automatic Storage Management (ASM) disk group managed by Oracle that serves as a centralized storage area for files related to backup, restore, and database recovery (recovery files). See also recovery files.

formatting

A process that erases any data contained on a medium and prepares it for use with Data Protector. Information about media (media ID, description, and location) is saved in the IDB as well as on the respective media (media header). Use the Force operation option to reformat Data Protector media with non-protected data. Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.

free pool

An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.

full backup

A backup in which all selected objects are backed up, whether or not they have been recently modified. See also backup types.

full database backup

A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.

full mailbox backup

A full mailbox backup is a backup of the entire mailbox content.

full ZDB

A ZDB-to-tape or ZDB-to-disk+tape session in which all selected objects are streamed to tape, even if there are no changes from the previous backup. See also incremental ZDB.

G

global options

A set of options that define behavior of the entire Data Protector cell. The options are stored in a plain text file on the Cell Manager.

group (Microsoft Cluster Server specific term)

A collection of resources (for example disk volumes, application services, IP names and addresses) that are needed to run a specific cluster-aware applications.

GUI

A graphical user interface provided by Data Protector for easy access to all configuration, administration, and operation tasks. It is available for Microsoft Windows operating systems.

H**hard recovery (Microsoft Exchange Server specific term)**

A Microsoft Exchange Server database recovery that is performed after a restore by the database engine, using transaction log files.

heartbeat

A cluster data set with a time stamp carrying information about the operational status of a particular cluster node. This data set or packet is distributed among all cluster nodes.

Hierarchical Storage Management (HSM)

A method for optimizing the use of expensive hard disk storage by migrating less frequently used data to less expensive optical platters. When needed, the data is migrated back to hard disk storage. This balances the need for fast retrieval from hard disk with the lower cost of optical platters.

Holidays file

A file that contains information about holidays. You can set different holidays by editing the Holidays file that resides on the Cell Manager at the following location: <Data_Protector_program_data>\Config\Server\holidays (Windows systems) and /etc/opt/omni/server/Holidays (UNIX systems).

hosting system

A working Data Protector client used for Disk Delivery Disaster Recovery with a Data Protector Disk Agent installed.

HP Business Copy (BC) P6000 EVA (HP P6000 EVA Disk Array Family specific term)

A local replication software solution that enables creation of point-in-time copies (replicas) of the source volumes using the snapshot and clone capabilities of the P6000 EVA firmware. See also replica, source volume, snapshot, and HP Continuous Access + Business Copy (CA+BC) P6000 EVA.

HP Business Copy (BC) P9000 XP (HP P9000 XP Disk Array Family specific term)

An HP P4000 SAN Solutions configuration that enables creation and maintenance of internal copies of LDEVs for various purposes, such as data duplication and backup. The copies (secondary volumes or S-VOLs) can be separated from the primary volumes (P-VOLs) and connected to a different system. For Data Protector zero downtime backup purposes, P-VOLs should be available to the application system, and one of the S-VOL sets should be available to the backup system. See also LDEV, HP Continuous Access (CA) P9000 XP, Main Control Unit (MCU), application system, and backup system.

HP Command View (CV) EVA (HP P6000 EVA Disk Array Family specific term)

The user interface that enables you to configure, manage, and monitor your P6000 EVA storage system. It is used to perform various storage management tasks, for example, creating virtual disk families, managing storage system hardware, and creating snapshots, snapclones, or mirrorclones of virtual disks. The HP Command View EVA software runs on the HP Storage Management Appliance, and is accessed

by a Web browser. See also HP P6000 / HP 3PAR SMI-S Agent and HP SMI-S P6000 EVA Array provider.

HP Continuous Access (CA) P9000 XP (HP P9000 XP Disk Array Family specific term)

An HP P9000 XP Disk Array Family configuration that enables creation and maintenance of remote copies of LDEVs for purposes such as data duplication, backup, and disaster recovery. HP CA P9000 XP operations involve main (primary) disk array units and remote (secondary) disk array units. The main disk array units are connected to the application system and contain primary volumes (P-VOLs), which store original data. The remote disk array units are connected to the backup system and contain secondary volumes (S-VOLs). See also HP BC P9000 XP (HP P9000 XP Disk Array Family specific term), Main Control Unit (MCU), and LDEV.

HP Continuous Access + Business Copy (CA+BC) P6000 EVA (HP P6000 EVA Disk Array Family specific term)

An HP P6000 EVA Disk Array Family configuration that enables creation and maintenance of copies (replicas) of the source volumes on a remote P6000 EVA, and later use of these copies as the source for local replication on this remote array. See also HP BC P6000 EVA, replica, and source volume.

HP P6000 / HP 3PAR SMI-S Agent

A Data Protector software module that executes all tasks required for the HP P6000 EVA Disk Array Family integration. With the P6000 / 3PAR SMI-S Agent, the control over the array is established through an appropriate SMI-S provider, which directs communication between incoming requests and the storage system's native interface. See

also HP Command View (CV) EVA and HP SMI-S P6000 EVA Array provider.

HP P9000 XP Agent

A Data Protector software component that executes all tasks needed by the Data Protector HP P9000 XP Disk Array Family integration. It communicates with the P9000 XP Array storage system via the RAID Manager Library.

HP SMI-S P6000 EVA Array provider

An interface used for controlling HP P6000 EVA Disk Array Family. SMI-S P6000 EVA Array provider runs as a separate service on the HP Storage Management Appliance system and acts as a gateway between incoming requests and HP Command View EVA. With the Data Protector HP P6000 EVA Disk Array Family integration, SMI-S P6000 EVA Array provider accepts standardized requests from the HP P6000 / HP 3PAR SMI-S Agent, communicates with HP Command View EVA for information or method invocation, and returns standardized responses. See also HP P6000 / HP 3PAR SMI-S Agent and HP Command View (CV) EVA.

I

ICDA (EMC Symmetrix specific term)

EMC's Symmetrix Integrated Cached Disk Arrays (ICDA) is a disk array device that combines a set of physical disks, a number of FWD SCSI channels, an internal cache memory, and control and diagnostic software commonly referred to as the microcode.

IDB

See Internal Database (IDB).

IDB recovery file

A file that maintains information about completed IDB backup sessions and the

backup media and backup devices used in them. If available, it significantly simplifies and speeds up offline recovery of the Internal Database in case of a Cell Manager disaster. Its filename is obdrindex.dat.

importing media

A process that re-reads all data about backup sessions which are on the medium back into the IDB. This then allows for fast and convenient access to the data on the media. See also exporting media.

incremental (re-)establish (EMC Symmetrix specific term)

A BCV or SRDF control operation. In BCV control operations, an incremental establish causes the BCV device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired. In SRDF control operations, an incremental establish causes the target (R2) device to be synchronized incrementally and to function as an EMC Symmetrix mirrored medium. The EMC Symmetrix devices must have been previously paired.

incremental backup

A backup that selects only files that have changed since a previous backup. Several levels of incremental backup are available, which enables detailed control of restore chain length. See also backup types.

incremental backup (Microsoft Exchange Server specific term)

A backup of the Microsoft Exchange Server data that has changed since the last full or incremental backup. With the incremental backup, only the transaction log files are backed up. See also backup types.

incremental mailbox backup

An incremental mailbox backup backs up all the changes made to the mailbox after the last backup of any type.

incremental restore (EMC Symmetrix specific term)

A BCV or SRDF control operation. In BCV control operations, an incremental restore reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data that was written to the BCV device during the time of the original pair split, and the data that was written to the standard device during the split is overwritten with data from the BCV mirror. In SRDF control operations, an incremental restore reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. However, the source (R1) devices are updated with only the data that was written to the target (R2) device during the time of the original pair split, and the data that was written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

incremental ZDB

A filesystem ZDB-to-tape or ZDB-to-disk+tape session in which only changes from the last protected full or incremental backup are streamed to tape. See also full ZDB.

Incremental1 Mailbox Backup

An incremental1 mailbox backup backs up all the changes made to the mailbox after the last full backup.

Inet

A process that runs on each UNIX system or service that runs on each Windows system in the Data Protector cell. It is responsible for communication

between systems in the cell and for starting other processes needed for backup and restore. The Inet service is started as soon as Data Protector is installed on a system. The Inet process is started by the inetd daemon.

Information Store (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that is responsible for storage management. Information Store in Microsoft Exchange Server manages two kinds of stores: mailbox stores and public folder stores. A mailbox store consists of mailboxes that belong to individual users. A public folder store contains public folders and messages that are shared among several users. See also Key Management Service and Site Replication Service.

Informix Server (Informix Server specific term)

Refers to Informix Dynamic Server.

initializing

See formatting.

Installation Server

A computer system that holds a repository of the Data Protector installation packages for a specific architecture. The Installation Server is used for remote installation of Data Protector clients. In mixed environments at least two Installation Servers are needed: one for UNIX systems and one for Windows systems.

instant recovery (ZDB specific term)

A process in which a replica, produced by a ZDB-to-disk or a ZDB-to-disk+tape sessions, is used to restore the contents of the source volumes to their states at the time at which the replica was created, avoiding the need to perform a restore

from tape. Depending on the application/database concerned, this may be all that is required, or other steps, such as the application of transaction log files, may be required for full recovery. See also replica, zero downtime backup (ZDB), ZDB to disk, and ZDB to disk+tape.

integration object

A backup object of a Data Protector integration, such as Oracle or SAP MaxDB.

Internal Database (IDB)

An entity in Data Protector that keeps information regarding which data was backed up, to which media it was backed up, how and when backup, restore, and other sessions were run, which devices, libraries, and disk arrays are configured, and so on. It is implemented with an embedded database and a collection of proprietary data files which reside on the Cell Manager. See also DC directory and Detail Catalog Binary Files (DCBF).

Internet Information Server (IIS) (Windows specific term)

Microsoft Internet Information Server is a network file and application server that supports multiple protocols. Primarily, IIS transmits information in Hypertext Markup Language (HTML) pages by using the Hypertext Transport Protocol (HTTP).

ISQL (Sybase specific term)

A Sybase utility used to perform system administration tasks on Sybase SQL Server.

J

jukebox

See library.

jukebox device

A device consisting of multiple slots used to store either optical or file media. When being used to store file media, the jukebox device is known as the 'file jukebox device'.

K

Key Management Service (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that provides encryption functionality for enhanced security. See also Information Store and Site Replication Service.

keychain

A tool that eliminates the supply of a passphrase manually when decrypting the private key. It needs to be installed and configured on the Installation Server if you perform remote installation using secure shell.

keystore

All encryption keys are centrally stored in the keystore on the Cell Manager and administered by the Key Management Server (KMS).

KMS

Key Management Server (KMS) is a centralized service that runs on the Cell Manager and provides key management for the Data Protector encryption functionality. The service is started as soon as Data Protector is installed on the Cell Manager.

L

LBO (Symmetric specific term)

A Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as

one entity and can only be restored as a whole.

LDEV (HP P9000 XP Disk Array Family specific term)

A logical partition of a physical disk of a disk array of the HP P9000 XP Disk Array Family. An LDEV is the entity that can be replicated using the split-mirror or snapshot functionality of such disk array. See also HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and replica.

library

Also called autochanger, jukebox, autoloader, or exchanger. A library contains media in repository slots. Each slot holds one medium (for example, DDS/DAT). Media are moved between slots and drives by a robotic mechanism, allowing random access to media. The library can contain multiple drives.

lights-out operation or unattended operation

A backup or restore operation that takes place outside of normal business hours without an operator. This implies that no operator personnel is present to work with the backup application or service mount requests, for example.

LISTENER.ORA (Oracle specific term)

An Oracle configuration file that describes one or more Transparent Network Substrate (TNS) listeners on a server.

load balancing

By default, Data Protector automatically balances the usage of devices selected for backup, so that they are used evenly. Load balancing optimizes the device usage by balancing the number of objects written to each device. Since load balancing is done automatically during

backup time, you do not need to manage how the data is actually backed up. You just specify the devices to be used. If you do not want to use load balancing, you can select which device will be used for each object in the backup specification. Data Protector will access the devices in the specified order.

local and remote recovery

Remote recovery is performed if all Media Agent hosts specified in the SRD file are accessible. If any of them fails, the disaster recovery process fails over to the local mode. This means that the target system is searched for locally attached devices. If only one device is found, it is automatically used. Otherwise, Data Protector prompts you to select the device, which will be used for restore.

local continuous replication (Microsoft Exchange Server specific term)

Local continuous replication (LCR) is a single-server solution that creates and maintains an exact copy (LCR copy) of a storage group. An LCR copy is located on the same server as the original storage group. When an LCR copy is created, it is kept up to date through change propagation (log replay) technology. The replication feature in LCR guarantees that logs that have not been replicated are not deleted. The implication of this behavior is that running backups in a mode that deletes logs may not actually free space if replication is sufficiently far behind in its log copying. An LCR copy is used for disaster recovery because you can switch to the LCR copy in a few seconds. If an LCR copy is used for backup and if it is located on a different disk than the original data, then the I/O load on a production database is minimal. A replicated storage group is represented as a new instance of Exchange writer called Exchange Replication Service and

can be backed up (using VSS) as a normal storage group. See also cluster continuous replication and Exchange Replication Service.

lock name

You can configure the same physical device several times with different characteristics, by using different device names. The lock name is a user specified string that is used for locking all such device configurations to prevent collision if several such devices (device names) are used concurrently. Use an identical lock name for all device definitions which use the same physical device.

log_full shell script (Informix Server UNIX systems specific term)

A script provided by ON-Bar that you can use to start backing up logical log files when Informix Server issues a logfull event alarm. The Informix Server ALARMPROGRAM configuration parameter defaults to the <INFORMIXDIR>/etc/log_full.sh, where <INFORMIXDIR> is the Informix Server home directory. If you do not want logical logs to be backed up continuously, set the ALARMPROGRAM configuration parameter to <INFORMIXDIR>/etc/no_log.sh.

logging level

An option that determines the amount of details on files and directories written to the IDB during backup, object copying, or object consolidation. You can always restore your data, regardless of the logging level used during backup. Data Protector provides four logging levels: Log All, Log Directories, Log Files, and No Log. The different logging level settings mainly influence the IDB growth and the convenience of browsing data for restore.

logical-log files

This applies to online database backup. Logical-log files are files in which modified data is first stored before being flushed to disk. In the event of a failure, these logical-log files are used to roll forward all transactions that have been committed as well as roll back any transactions that have not been committed.

login ID (Microsoft SQL Server specific term)

The name a user needs to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.

login information to the Oracle Target Database (Oracle and SAP R/3 specific term)

The format of the login information is <user_name>/<password>@<service>, where: <user_name> is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have Oracle SYSDBA or SYSOPER rights. <password> must be the same as the password specified in the Oracle password file (orapwd), which is used for authentication of users performing database administration. <service> is the name used to identify an SQL*Net server process for the target database.

login information to the Recovery Catalog Database (Oracle specific term)

The format of the login information to the Recovery (Oracle) Catalog Database is <user_name>/<password>@<service>, where the description of the user name, password, and service name is the same as in the Oracle SQL*Net V2 login information to the Oracle target database.

In this case, <service> is the name of the service to the Recovery Catalog Database, not the Oracle target database. Note that the Oracle user specified here must be the owner of the Oracle Recovery Catalog.

Lotus C API (Lotus Domino Server specific term)

An interface for the exchange of backup and recovery information between Lotus Domino Server and a backup solution, like Data Protector.

LVM

A Logical Volume Manager is a subsystem for structuring and mapping physical disk space to logical volumes on UNIX systems. An LVM system consists of several volume groups, where each volume group has several volumes.

M

Magic Packet

See Wake ONLAN.

mailbox (Microsoft Exchange Server specific term)

The location to which e-mail is delivered, which is set up by the administrator for each user. If a set of personal folders is designated as the e-mail delivery location, e-mail is routed from the mailbox to this location.

mailbox store (Microsoft Exchange Server specific term)

A part of the Information Store that maintains information in user mailboxes. A mailbox store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

Main Control Unit (MCU) (HP P9000 XP Disk Array Family specific term)

An HP P9000 XP Disk Array Family unit that contains primary volumes (P-VOLs) for the HP CA P9000 XP Array or HP CA+BC P9000 XP Array configuration and acts as a master device. See also HP Business Copy (BC) P9000 XP, HP Continuous Access (CA) P9000 XP, and LDEV.

maintenance mode

An operating mode you can initiate on the Cell Manager to prevent changes in the Internal Database. It enables you to perform various maintenance tasks, such as upgrading and patching the installation.

make_net_recovery

make_net_recovery is an Ignite-UX command, which allows you to create a recovery archive over the network onto the Ignite-UX server or any other specified system. The target system can be recovered across subnets after booting either from a bootable tape created by the Ignite-UX make_boot_tape command or the system boots directly from the Ignite-UX server. Booting directly from the Ignite-UX server can be automated with the Ignite-UX bootsys command or interactively specified on the boot console.

make_tape_recovery

make_tape_recovery is a command on Ignite-UX which creates a bootable recovery (installation) tape, customized for your system and enables you unattended disaster recovery by connecting the backup device directly to the target system and booting the target system from the bootable recovery tape. The backup device has to be locally connected to the client during the creation of the archive and recovery of the client.

Manager-of-Managers

See MoM.

MAPI (Microsoft Exchange specific term)

The MAPI (Messaging Application Programming Interface) is the programming interface that lets applications and messaging clients interact with messaging and information systems.

MCU

See Main Control Unit (MCU).

Media Agent

A process that controls reading from and writing to a device, which reads from or writes to a medium (typically a tape). During a backup session, a Media Agent receives data from the Disk Agent and sends it to the device for writing it to the medium. During a restore or object verification session, a Media Agent locates data on the backup medium and sends it to the Disk Agent for processing. For a restore session, the Disk Agent then writes the data to the disk. A Media Agent also manages the robotics control of a library.

media allocation policy

Determines in which sequence media are used for backup. The Strict allocation policy directs Data Protector to prompt for a specific medium. The Loose policy directs Data Protector to prompt for any suitable medium. The Formatted First policy directs Data Protector to give preference to unknown media, even if unprotected media are available in the library.

media condition

The quality of a medium as derived from the media condition factors. Heavy usage and age result in an increased number of

read and write errors with tape media. Media need to be replaced when they are marked as POOR.

media condition factors

The user-assigned age threshold and overwrite threshold used to assert the condition of a medium.

media label

A user-defined identifier used to describe a medium.

media location

A user-defined physical location of a medium, such as "building 4" or "off-site storage".

media management session

A session performing some action on a medium, such as initializing, scanning the content, verifying data on a medium, or copying a medium.

media pool

A set of media of the same type (such as DDS) used and tracked as a group. Media are formatted and assigned to a media pool.

media set

The result of a backup session is data backed up on a group of media called media set. Depending on the media usage policy, several sessions can share the same media.

media type

The physical type of media, such as DDS or DLT.

media usage policy

The media usage policy controls how new backups are added to the already used media. It can be Appendable, Non-Appendable, or Appendable for incrementals only.

medium ID

A unique identifier assigned to a medium by Data Protector.

merging

This defines one mode to resolve file conflicts during restore. If the file to be restored already exists at the destination, the one with the more recent modification date is kept. Files not present on the disk are always restored. See also overwrite.

Microsoft Exchange Server

A "client-server" messaging and a workgroup system that offers a transparent connection to many different communication systems. It provides users with an electronic mail system, individual and group scheduling, online forms, and workflow automation tools. It provides a developer with a platform on which to build custom information-sharing and messaging-service applications.

**Microsoft Management Console (MMC)
(Windows specific term)**

An administration model for Windows-based environments. It provides a simple, consistent, and integrated administration user interface allowing management of many applications through the same GUI, provided that the applications adhere to the MMC model.

Microsoft SQL Server

A database management system designed to meet the requirements of distributed "client-server" computing.

Microsoft Volume Shadow Copy Service (VSS)

A software service that provides a unified communication interface to coordinate backup and restore of a VSS-aware application regardless of its specific features. This service collaborates with the backup application, writers, shadow

copy providers, and the operating system kernel to implement the management of volume shadow copies and shadow copy sets. See also shadow copy, shadow copy provider, replica, and writer.

mirror (EMC Symmetrix and HP P9000 XP Disk Array Family specific term)

See target volume.

mirror rotation (HP P9000 XP Disk Array Family specific term)

See replica set rotation.

mirror unit (MU) number (HP P9000 XP Disk Array Family specific term)

A non-negative integer number that determines a secondary volume (S-VOL) of an internal disk (LDEV) located on a disk array of the HP P9000 XP Disk Array Family. See also first-level mirror.

mirrorclone (HP P6000 EVA Disk Array Family specific term)

A dynamic replica of a storage volume, which is kept updated with changes made to the original storage volume via a local replication link. Replication between the original storage volume and its mirrorclone can be suspended. For each storage volume, a single mirrorclone can be created on the disk array.

MMD

The Media Management Daemon process (service) (MMD) runs on the Data Protector Cell Manager and controls media management and device operations. The process is started when Data Protector is installed on the Cell Manager.

MMDB

The Media Management Database (MMDB) is a part of the IDB that contains information about media, media pools, devices, libraries, library drives, and slots

configured in the cell, as well as the Data Protector media used for backup. In an enterprise backup environment, this part of the database can be common to all cells. See also CMMDB and Catalog Database (CDB).

MoM

Several cells can be grouped together and managed from a central cell. The management system of the central cell is the Manager-of-Managers (MoM). The cells are called MoM clients. The MoM enables you to configure and manage multiple cells from a central point.

mount point

The access point in a directory structure for a disk or logical volume, for example /opt or d:. On UNIX systems, the mount points are displayed using the bdf or df command.

mount request

A screen prompt that tells you to insert a specific medium into a device. Once you respond to the mount request by providing the required medium and confirm the mount request, the session continues.

MSM

The Data Protector Media Session Manager, which runs on the Cell Manager and controls media sessions, such as copying media.

multisnapping (HP P6000 EVA Disk Array Family specific term)

Simultaneous creation of target volumes so that the backup data is consistent not only on each individual target volume, but also across all the volumes that constitute a snapshot. See also snapshot.



OBDR capable device

A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.

obdrindex.dat

See IDB recovery file.

object

See backup object.

object consolidation

The process of merging a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object. The process is a part of the synthetic backup procedure. The result is a synthetic full backup of the specified backup object.

object consolidation session

A process that merges a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object.

object copy

A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.

object copy session

A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.

object copying

The process of copying selected object versions to a specific media set. You can

select object versions from one or several backup sessions to be copied.

object ID (Windows specific term)

The object IDs (OIDs) enable access to NTFS 5 files no matter where in the system the files reside. Data Protector treats the OIDs as alternate streams of the files.

object mirror

A copy of a backup object created using object mirroring. Object mirrors are often referred to as object copies.

object mirroring

The process of writing the same data to several media sets during a backup session. Data Protector enables you to mirror all or some backup objects to one or more media sets.

object verification

The process of verifying the data integrity of backup objects, from the Data Protector point of view, and the ability of Data Protector to deliver them to the required destination. The process can be used to provide a level of confidence in the ability to restore object versions created by backup, object copy, or object consolidation sessions.

object verification session

A process that verifies the data integrity of specified backup objects or object versions and the ability of selected Data Protector network components to deliver them to a specified host. Object verification sessions can be run interactively, or as specified in automated post-backup, or scheduled specifications.

offline backup

A backup during which an application database cannot be used by the

application. For simple backup methods (non ZDB), the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the whole backup period (~minutes/hours). For instance, for backup to tape, until streaming of data to the tape is finished. For ZDB methods, the database is also put into the quiescent state, but for the period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process. See also zero downtime backup (ZDB) and online backup.

offline recovery

Offline recovery is performed if the Cell Manager is not accessible, for example, due to network problems. Only standalone and SCSI library devices can be used for offline recovery. Cell Manager can only be recovered offline.

offline redo log

See archived redo log.

ON-Bar (Informix Server specific term)

A backup and restore system for Informix Server. ON-Bar enables you to create a copy of your Informix Server data and later restore the data. The ON-Bar backup and restore system involves the following components: the onbar command, Data Protector as the backup solution, the XBSA interface, and ON-Bar catalog tables, which are used to back up dobjects and track instances of dobjects through multiple backups.

ONCONFIG (Informix Server specific term)

An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the onconfig file located in

the directory <INFORMIXDIR>\etc (on Windows systems) or <INFORMIXDIR>/etc/ (on UNIX systems).

online backup

A backup performed while a database application remains available for use. The database is placed into a special backup mode of operation for the time period that the backup application requires access to the original data objects. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly. For simple backup methods (non ZDB), backup mode is required for the whole backup period (~minutes/hours). For instance, for backup to tape, until streaming of data to tape is finished. For ZDB methods, backup mode is required for the short period of the data replication process only (~seconds). Normal database operation can then be resumed for the rest of the backup process. In some cases, transaction logs may also have to be backed up to allow a consistent database to be restored. See also zero downtime backup (ZDB) and offline backup.

online recovery

A type of the Internal Database recovery that you can use when the Cell Manager is accessible. In this case, the Cell Manager runs sessions, the sessions are logged into the IDB, and you can monitor the session progress using the GUI.

online redo log (Oracle specific term)

Redo logs that have not been archived, but are either available to the instance for recording database activity or are filled and waiting to be archived or reused. See also archived redo log.

OpenSSH

A set of network connectivity tools used to access remote machines securely, by using a variety of authentication and encryption methods. It needs to be installed and configured on the Installation Server and the client if you perform remote installation using secure shell.

Oracle Data Guard (Oracle specific term)

Oracle Data Guard is Oracle's primary disaster recovery solution. Oracle Data Guard is able to maintain up to nine standby databases, each of which is a real-time copy of the production (primary) database, to protect against corruptions, data failures, human errors, and disasters. If a failure occurs on the production (primary) database, then a failover to one of the standby databases which becomes the new primary database is possible. In addition, planned downtime for maintenance can be reduced because the production processing can be moved from the current primary database to a standby database and back quickly.

Oracle instance (Oracle specific term)

Each installation of an Oracle database on one or more systems. One computer system can have several instances of a database running.

ORACLE_SID (Oracle specific term)

A unique name for an Oracle Server instance. To switch among Oracle Servers, specify the desired <ORACLE_SID>. The <ORACLE_SID> is included in the CONNECT DATA parts of the connect descriptor in a TNSNAMES.ORA file and in the definition of the TNS listener in the LISTENER.ORA file.

original system

The system configuration backed up by Data Protector before a computer disaster hits the system.

overwrite

An option that defines one mode to resolve file conflicts during restore. All files are restored from a backup even if they are older than existing files. See also merging.

ownership

The ownership of a backup determines who can restore from the backup. The user who starts an interactive backup is the session owner. If a user starts an existing backup specification without modifying it, the session is not considered interactive. In that case, if the backup owner has been defined in the backup specification, they remain the session owner. Otherwise, the session owner becomes the user who started the backup in question. For the scheduled backups, by default, the session owner is for the UNIX Cell Manager: root.sys@<Cell Manager>, and for the Windows Cell Manager, the user that was specified during the installation of the Cell Manager. It is possible to modify the ownership, so that the specific user becomes the session owner. When copying or consolidating backup objects, the owner of the resulting objects is the user who started the original backup session.

P

P1S file

P1S file contains information on how to format and partition all disks installed in the system during Enhanced Automated Disaster Recovery (EADR). It is created during full backup and is saved on backup medium and on Cell Manager into the

directory <Data_Protector_program_data>\Config\Server\dr\p1s (Windows systems) or /etc/opt/omni/dr/p1s (UNIX systems) with the filename recovery.p1s.

package (HP ServiceGuard and Veritas Cluster Specific Term)

A collection of resources (for example volume groups, application services, IP names, and addresses) that are needed to run a specific cluster-aware application.

pair status (HP P9000 XP Disk Array Family specific term)

The status of a disk pair (secondary volume and its corresponding primary volume) of a disk array of the HP P9000 XP Disk Array Family. Depending on the circumstances, the paired disks can be in various states. The following states are particularly important for the operation of the Data Protector HP P9000 XP Agent:

PAIR - The secondary volume is prepared for zero downtime backup. If it is a mirror, it is completely synchronized, and if it is a volume to be used for snapshot storage, it is empty.

SUSPENDED - The link between the disks is suspended. However, the pair relationship is still maintained, and the secondary disk can be prepared for zero downtime backup again at a later time.

COPY - The disk pair is currently busy and making a transition into the PAIR state. If the secondary volume is a mirror, it is re-synchronizing with the primary volume, and if it is a volume to be used for snapshot storage, its contents are getting cleared.

parallel restore

Restoring backed up data to multiple disks at the same time (that is, in parallel) by running multiple Disk Agents, that receive data from one Media Agent. For the parallel restore to work, select data that is located on different disks or logical

volumes and during backup, the data from the different objects must have been sent to the same device using a concurrency of 2 or more. During a parallel restore, the data for multiple objects selected for restore is read from media at the same time, thereby improving performance.

parallelism

The concept of reading multiple data streams from an online database.

phase 0 of disaster recovery

Preparation for disaster recovery - the prerequisite condition for a successful disaster recovery.

phase 1 of disaster recovery

Installation and configuration of DR OS, establishing previous storage structure.

phase 2 of disaster recovery

Restoration of operating system (with all the configuration information that defines the environment) and Data Protector.

phase 3 of disaster recovery

Restoration of user and application data.

physical device

A physical unit that contains either a drive or a more complex unit such as a library.

post-exec

A backup option that executes a command or script after the backup of an object or after the entire session completes. Post-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows systems and as shell scripts on UNIX systems. See also pre-exec.

pre- and post-exec commands

Pre- and post-exec commands are used to perform additional action before and after a backup or restore session. They are not supplied by Data Protector. You need to create your own commands. They can be written as executables or batch files on Windows systems and as shell scripts on UNIX systems.

prealloc list

A subset of media in a media pool that specifies the order in which media are used for backup.

pre-exec

A backup option that executes a command or script before the backup of an object or before the entire session is started. Pre-exec commands are not supplied by Data Protector. You need to create your own. They can be written as executables or batch files on Windows systems and as shell scripts on UNIX systems. See also post-exec.

primary volume (P-VOL) (HP P9000 XP Disk Array Family specific term)

An internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family for which a secondary volume (S-VOL), either its mirror or a volume to be used for its snapshot storage, exists. In the HP CA P9000 XP and HP CA+BC P9000 XP configurations, primary volumes are located in the Main Control Unit (MCU). See also secondary volume (S-VOL) and Main Control Unit (MCU).

protection

See data protection and catalog protection.

public folder store (Microsoft Exchange Server specific term)

The part of the Information Store that maintains information in public folders. A

public folder store consists of a binary rich-text .edb file and a streaming native internet content .stm file.

public/private backed up data

When configuring a backup, you can select whether the backed up data will be: public, that is, visible (and accessible for restore) to all Data Protector users private, that is, visible (and accessible for restore) only to the owner of the backup and administrators

R

RAID

Redundant Array of Independent Disks.

RAID Manager Library (HP P9000 XP Disk Array Family specific term)

A software library that is used for accessing the configuration, status, and performance measurement data of a P9000 XP Array storage system, and for invoking operations on the disk array. It translates function calls into sequences of low-level SCSI commands. See also HP P9000 XP Agent.

RAID Manager P9000 XP (HP P9000 XP Disk Array Family specific term)

A software application that provides a command-line interface to disk arrays of the HP P9000 XP Disk Array Family. It offers an extensive set of commands for reporting and controlling the status of a P9000 XP Array storage system, and for performing various operations on the disk array.

rawdisk backup

See disk image backup.

RCU

See Remote Control Unit (RCU).

RDBMS

Relational Database Management System.

RDF1/RDF2 (EMC Symmetrix specific term)

A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.

Recovery Catalog (Oracle specific term)

A set of Oracle tables and views that are used by Recovery Manager to store information about Oracle databases. This information is used by Recovery Manager to manage the backup, restore, and recovery of Oracle databases. The recovery catalog contains information about the physical schema of the Oracle target database, data file and archived log backup sets, data file copies, archived redo logs, and stored scripts.

Recovery Catalog Database (Oracle specific term)

An Oracle database that contains a recovery catalog schema. You should not store the recovery catalog in your target database.

recovery files (Oracle specific term)

Recovery files are Oracle specific files that reside in the flash recovery area: the current control file, online redo logs, archived redo logs, flashback logs, control file autobackups, datafile copies, and backup pieces. See also flash recovery area.

Recovery Manager (RMAN) (Oracle specific term)

An Oracle command-line interface that directs an Oracle Server process to back up, restore, or recover the database it is connected to. RMAN uses either the

recovery catalog or the control file to store information about backups. This information can be used later in restore sessions.

RecoveryInfo

When backing up Windows configuration files, Data Protector collects the information about the current system configuration (information on disk layout, volume, and network configuration). This information is needed for disaster recovery.

recycle or unprotect

A process that removes the data protection from all backed up data on a medium, allowing Data Protector to overwrite it during one of the next backups. Data that belongs to the same session(s) but resides on other media is also unprotected. Recycling does not actually alter the data on the medium.

redo log (Oracle specific term)

Every Oracle database has a set of two or more redo log files. The set of redo log files for the database is known as the database's redo log. Oracle uses the redo log to record all changes made to data.

Remote Control Unit (RCU) (HP P9000 XP Disk Array Family specific term)

An HP P9000 XP Disk Array Family unit that acts as a slave device to the Main Control Unit (MCU) in the HP CA P9000 XP or HP CA+BC P9000 XP configuration. In bidirectional configurations, the RCU can also act as an MCU.

Removable Storage Management Database (Windows specific term)

A Windows service used for managing removable media (such as tapes and disks) and storage devices (libraries). Removable Storage allows applications

to access and share the same media resources.

reparse point (Windows specific term)

A system-controlled attribute that can be associated with any directory or file. The value of a reparse attribute can have user-controlled data. The format of the data is understood by the application that stored the data and a filesystem filter that was installed to interpret the data and process such files. Whenever the filesystem encounters a file with a reparse point, it attempts to find the filesystem filter associated with the data format.

replica (ZDB specific term)

An image, at a particular point in time, of the data in source volumes that contain user-specified backup objects. Depending on the hardware or software with which it is created, the image may be an independent exact duplicate (clone) of the storage blocks at the physical disk level (for example, a split mirror), or a virtual copy (for example, a snapshot). From perspective of a basic operating system, the complete physical disk containing a backup object is replicated. However, if a volume manager is used on a UNIX system, the whole volume/disk group containing a backup object is replicated. If partitions are used on a Windows system, the whole physical volume containing the selected partition is replicated. See also snapshot, snapshot creation, split mirror, and split mirror creation.

replica set (ZDB specific term)

A group of replicas, all created using the same backup specification. See also replica and replica set rotation.

replica set rotation (ZDB specific term)

The use of a replica set for regular backup production: Each time the same backup

specification requiring the use of a replica set is run, a new replica is created and added to the set, until the maximum number for the set is reached. After that, the oldest replica in the set is replaced, maintaining the maximum number of replicas in the set. See also replica and replica set.

restore chain

Backup images that are needed to restore a backed up object to the state it was in at the selected point in time. In general, a restore chain of an object consists of its full backup image and one or more related incremental backup images.

restore session

A process that copies data from backup media to a client.

resync mode (HP P9000 XP Disk Array Family VSS provider specific term)

One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the resync mode, the source volume (P-VOL) and its replica (S-VOL) are in the suspended mirror relationship after a backup. The maximum number of replicas (S-VOLs per a P-VOL) rotated is three provided that MU range is 0-2 or 0, 1, 2. Restore from a backup in such a configuration is possible only by re-synchronization of an S-VOL with its P-VOL. See also VSS compliant mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), MU number, and replica set rotation.

RMAN (Oracle specific term)

See Recovery Manager.

RSM

The Data Protector Restore Session Manager controls restore and object

verification sessions. This process always runs on the Cell Manager system.

RSM (Windows specific term)

Removable Storage Manager (RSM) includes a media management service that facilitates communication among applications, robotic changers, and media libraries. It enables multiple applications to share local robotic media libraries and tape or disk drives and to manage removable media.

S

scanning

A function which identifies the media in a device. This synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library). It is useful to perform a scan and check the actual media in the device if someone has manually manipulated media without using Data Protector to eject or enter, for example.

Scheduler

A function that controls when and how often automatic backups occur. By setting up a schedule, you automate the start of backups.

secondary volume (S-VOL) (HP P9000 XP Disk Array Family specific term)

An internal disk (LDEV) of a disk array of the HP P9000 XP Disk Array Family which is paired with another LDEV: a primary volume (P-VOL). It can act as a mirror of the P-VOL or as a volume to be used for the P-VOL's snapshot storage. An S-VOL is assigned a SCSI address different from the one used for the P-VOL. In an HP CA P9000 XP configuration, the S-VOLs acting as mirrors can be used as failover devices in a MetroCluster configuration. See also primary volume (P-VOL) and Main Control Unit (MCU).

session

See backup session, media management session, and restore session.

session ID

An identifier of a backup, restore, object copy, object consolidation, object verification, or media management session, consisting of the date when the session ran and a unique number.

session key

This environment variable for the pre- and post-exec script is a Data Protector unique identification of any session, including preview sessions. The session key is not recorded in the database, and it is used for specifying options for the omnimnt, omnistat, and omniabort CLI commands.

shadow copy (Microsoft VSS specific term)

A volume that represents a duplicate of the original volume at a certain point in time. The data is then backed up from the shadow copy and not from the original volume. The original volume continues to change as the backup process continues, but the shadow copy of the volume remains constant. See also Microsoft Volume Shadow Copy Service and replica.

shadow copy provider (Microsoft VSS specific term)

An entity that performs the work on creating and representing the volume shadow copies. Providers own the shadow copy data and expose the shadow copies. Providers can be software (for example, system providers) or hardware (local disks, disk arrays). See also shadow copy.

shadow copy set (Microsoft VSS specific term)

A collection of shadow copies created at the same point in time. See also shadow copy and replica set.

shared disks

A Windows disk on another system that has been made available to other users on the network. Systems with shared disks can be backed up without a Data Protector Disk Agent installed.

Site Replication Service (Microsoft Exchange Server specific term)

The Microsoft Exchange Server service that permits compatibility with Microsoft Exchange Server 5.5 by emulating the Exchange Server 5.5 directory service. See also Information Store and Key Management Service.

slot

A mechanical position in a library. Each slot can hold a medium, such as a DLT tape. Data Protector references each slot with a number. To read a medium, a robotic mechanism moves the medium from a slot into the drive.

SMB

See split mirror backup.

SMBF

The Session Messages Binary Files (SMBF) part of the IDB stores session messages generated during backup, restore, object copy, object consolidation, object verification, and media management sessions. One binary file is created per session. The files are grouped by year and month.

SMI-S Agent (SMISA)

See HP P6000 / HP 3PAR SMI-S Agent.

snapshot (HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP 3PAR StoreServ Storage specific term)

A type of target volumes created using a specific replication technology.

Depending on the disk array model and the chosen replication technique, a range of snapshot types with different characteristics is available. Basically, each snapshot may be either a virtual copy, still reliant upon the contents of the source volume, or an independent duplicate (clone) of the source volume. See also replica and snapshot creation.

snapshot backup

See ZDB to tape, ZDB to disk, and ZDB to disk+tape.

snapshot creation (HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP 3PAR StoreServ Storage specific term)

A replica creation process in which copies of the selected source volumes are created using storage virtualization technology. Such a replica is considered to be created at a particular point in time, and is immediately available for use. However, with certain snapshot types, a background data copying process continues to run on the disk array after the moment of the replica creation. See also snapshot.

source (R1) device (EMC Symmetrix specific term)

An EMC Symmetrix device that participates in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote EMC Symmetrix unit. An R1 device must be assigned to an RDF1 group type. See also target (R2) device.

source volume (ZDB specific term)

A storage volume containing data to be replicated.

sparse file

A file that contains data with portions of empty blocks. Examples are: -A matrix in which some or much of the data contains zeros -files from image applications -high-speed databases If sparse file processing is not enabled during restore, it might be impossible to restore this file.

split mirror (EMC Symmetrix and HP P9000 XP Disk Array Family specific term)

A type of target volumes created using a specific replication technology. A split-mirror replica provides independent duplicates (clones) of the source volumes. See also replica and split mirror creation.

split mirror backup (EMC Symmetrix specific term)

See ZDB to tape.

split mirror backup (HP P9000 XP Disk Array Family specific term)

See ZDB to tape, ZDB to disk, and ZDB to disk+tape.

split mirror creation (EMC Symmetrix and HP P9000 XP Disk Array Family specific term)

A replica creation technique, in which a pre-configured set of target volumes (a mirror) is kept synchronized with a set of source volumes until the time at which a replica of the contents of the source volumes is required. Then, the synchronization is stopped (the mirror is split) and a split mirror replica of the source volumes at the time of the split remains in the target volumes. See also split mirror.

split mirror restore (EMC Symmetrix and HP P9000 XP Disk Array Family specific term)

A process in which data backed up in a ZDB-to-tape or a ZDB-to-disk+tape session is first copied from the backup media to a replica, and from the replica to the source volumes afterwards. Individual backup objects or complete sessions can be restored using this method. See also ZDB to tape, ZDB to disk+tape, and replica.

sqlhosts file or registry (Informix Server specific term)

An Informix Server connectivity information file (on UNIX systems) or registry (on Windows systems) that contains the names of each of the database servers and any aliases to which the clients on a host computer can connect.

SRD file

A text file in the Unicode (UTF-16) format, generated during CONFIGURATION backup of a Windows system and stored on the Cell Manager. It contains system information required for installing and configuring the operating system on the target system in the case of a disaster. See also target system.

SRDF (EMC Symmetrix specific term)

The EMC Symmetrix Remote Data Facility is a business continuation process that enables effective, real-time data replication of SLDs between dislocated processing environments. These environments could be situated within the same root computer environment or separated by long distances.

SSE Agent (SSEA)

See HP P9000 XP Agent.

sst.conf file

The file `/usr/kernel/drv/sst.conf` is required on each Data Protector Solaris client to which a multi-drive library device is connected. It must contain an entry for the SCSI address of the robotic mechanism of each library device connected to the client.

st.conf file

The file `/kernel/drv/st.conf` is required on each Data Protector Solaris client with a backup device connected. It must contain device information and a SCSI address for each backup drive connected to that client. A single SCSI entry is required for a single-drive device and multiple SCSI entries are required for a multi-drive library device.

stackers

Devices with multiple slots for media storage usually with only one drive. A stacker selects media from the stack sequentially. In contrast, a library can randomly select media from its repository.

standalone file device

A file device is a file in a specified directory to which you back up data.

Storage Group (Microsoft Exchange Server specific term)

A collection of mailbox stores and public folder stores that share a set of transaction log files. Exchange Server manages each storage group with a separate server process.

storage volume (ZDB specific term)

A storage volume represents an object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, file systems, or other objects may exist.

Typically, these can be created or exist within a storage system such as a disk array.

StorageTek ACS library (StorageTek specific term)

Automated Cartridge System is a library system (also known as Silo) consisting of one Library Management Unit (LMU) and one to 24 Library Storage Modules (LSM) connected to the unit.

switchover

See failover.

Sybase Backup Server API (Sybase specific term)

An industry-standard interface developed for the exchange of backup and recovery information between a Sybase SQL Server and a backup solution like Data Protector.

Sybase SQL Server (Sybase specific term)

The server in the Sybase "client-server" architecture. The Sybase SQL Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.

SYMA

See EMC Symmetrix Agent.

synthetic backup

A backup solution that produces a synthetic full backup, an equivalent to a conventional full backup in terms of data, without putting stress on the production servers or the network. A synthetic full backup is created from a previous full backup and any number of incremental backups.

synthetic full backup

The result of an object consolidation operation, where a restore chain of a backup objects is merged into a new, synthetic full version of this object. A synthetic full backup is equivalent to a conventional full backup in terms of restore speed.

System Backup to Tape (SBT) (Oracle specific term)

An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.

system databases (Sybase specific term)

The four system databases on a newly installed Sybase SQL Server are the: - master database (master) -temporary database (tempdb) -system procedure database (sybssystemprocs) -model database (model).

System Recovery Data file

See SRD file.

System State (Windows specific term)

The System State data comprises the Registry, COM+ Class Registration database, system startup files, and the Certificate Services database (if the server is a certificate server). If the server is a domain controller, Active Directory services and the SysVol directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.

system volume/disk/partition

A volume/disk/partition containing operating system files. Microsoft

terminology defines the system volume/disk/partition as the volume/disk/partition containing files required for the initial step of the boot process.

SysVol (Windows specific term)

A shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.

T

tablespace

A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.

tapeless backup (ZDB specific term)

See ZDB to disk.

target (R2) device (EMC Symmetrix specific term)

An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type. See also source (R1) device.

target database (Oracle specific term)

In RMAN, the target database is the database that you are backing up or restoring.

target system (disaster recovery specific term)

A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore this system to the original system configuration. The difference between a faulty system and a target system is that a target system has all faulty hardware replaced.

target volume (ZDB specific term)

A storage volume to which data is replicated.

Terminal Services (Windows specific term)

Windows Terminal Services provide a multi-session environment that allows clients to access a virtual Windows desktop session and Windows-based programs running on the server.

thread (Microsoft SQL Server specific term)

An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.

TimeFinder (EMC Symmetrix specific term)

A business continuation process that creates an instant copy of single or multiple EMC Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system (s).

TLU

See Tape Library Unit.

TNSNAMES.ORA (Oracle and SAP R/3 specific term)

A network configuration file that contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.

transaction

A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes

transaction backup

Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred.

transaction backup (Sybase and SQL specific term)

A backup of the transaction log providing a record of changes made since the last full or transaction backup.

transaction log backup

Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.

transaction log files

Files that record transactions of the database modifications, and provide fault tolerance in case of a database disaster.

transaction log table (Sybase specific term)

A system table in which all changes to the database are automatically recorded.

transportable snapshot (Microsoft VSS specific term)

A shadow copy that is created on the application system and can be presented to the backup system which performs the backup. See also Microsoft Volume Shadow Copy Service (VSS).

U

unattended operation

See lights-out operation.

user account (Data Protector user account)

You can use Data Protector only if you have a Data Protector user account, which restricts unauthorized access to Data Protector and to backed up data. Data Protector administrators create this account specifying a user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

User Account Control (UAC)

A security component in Windows Vista and Windows Server 2008 that limits application software to standard user privileges until an administrator authorizes an increase in privilege level.

user disk quotas

NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. Data Protector backs up user disk quotas on the whole system and for all configured users at a time.

user group

Each Data Protector user is member of a User Group. Each User Group has a set

of user rights that are given to every user in that User Group. The number of User Groups with their associated user rights can be defined as desired. Data Protector provides three default user groups: admin, operator, and user.

user profile (Windows specific term)

Configuration information retained on a user basis. This information includes desktop settings, screen colors, network connections, and so on. When the user logs on, the user profile is loaded and the Windows environment is set accordingly.

user rights

User rights or access rights are the permissions needed to perform specific Data Protector tasks. Configuring a backup, starting a backup session, or starting a restore session are typical user rights. Users have the access rights of the user group to which they belong.

user_restrictions file

A file that restricts specific user actions, which are available to Data Protector user groups according to the user rights assigned to them, to be performed only on specific systems of the Data Protector cell. Such restrictions apply only to Data Protector user groups other than admin and operator.

V

vaulting media

The process of storing media to a safe and remote place. The media are brought back to the data center when they are needed for restore or are ready for reuse in subsequent backups. The vaulting procedure depends on your company's backup strategy and policies for data protection/reliability.

verify

A function that lets you check whether the Data Protector data on a specified medium is readable. Additionally, consistency within each block can be checked if the backup was performed with the cyclic redundancy check (CRC) option ON.

Virtual Controller Software (VCS) (HP P6000 EVA Disk Array Family specific term)

The firmware that manages all aspects of storage system operation, including communication with HP Command View EVA through the HSV controllers. See also HP Command View (CV) EVA.

Virtual Device Interface (Microsoft SQL Server specific term)

This is a Microsoft SQL Server programming interface that allows fast backup and restore of large databases.

virtual disk (HP P6000 EVA Disk Array Family specific term)

A unit of storage allocated from a storage pool of a disk array of the HP P6000 EVA Disk Array Family. A virtual disk is the entity that can be replicated using the snapshot functionality of such disk array. See also source volume and target volume.

virtual full backup

An efficient type of synthetic backup where data is consolidated using pointers instead of being copied. It is performed if all the backups (the full backup, incremental backups, and the resulting virtual full backup) are written to a single file library that uses distributed file medium format.

Virtual Library System (VLS)

A disk-based data storage device hosting one or more virtual tape libraries (VTLs).

virtual server

A virtual machine in a cluster environment defined in a domain by a network IP name and address. Its address is cached by the cluster software and mapped to the cluster node that is currently running the virtual server resources. This way all requests for a particular virtual server are cached by a specific cluster node.

virtual tape

An archival storage technology that backs up data to disk drives in the same way as if it were being stored on tape. Benefits of virtual tape systems include improved backup and recovery speed and lower operating costs. See also Virtual Library System (VLS) and virtual tape library (VTL).

virtual tape library (VTL)

Data storage virtualization software that is able to emulate tape devices and libraries thus providing the functionality of traditional tape-based storage. See also Virtual Library System (VLS).

volser

A VOLume SERial number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/GRAU and StorageTek devices.

volume group

A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.

volume mountpoint (Windows specific term)

An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to

the target volume. Provided that the volume is mounted, users and applications can refer to the data on the mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

Volume Shadow Copy Service

See Microsoft Volume Shadow Copy Service (VSS).

VSS

See Microsoft Volume Shadow Copy Service (VSS).

VSS compliant mode (HP P9000 XP Disk Array Family VSS provider specific term)

One of two P9000 XP Array VSS hardware provider operation modes. When the P9000 XP Array provider is in the VSS compliant mode, the source volume (P-VOL) and its replica (S-VOL) are in simplex, unpaired state after a backup. Therefore the number of replicas (S-VOLs per a P-VOL) rotated is not limited. Restore from a backup in such a configuration is possible only by switching of the disks. See also resync mode, source volume, primary volume (P-VOL), replica, secondary volume (S-VOL), and replica set rotation.

VxFS

Veritas Journal Filesystem.

VxVM (Veritas Volume Manager)

A Veritas Volume Manager is a system for managing disk space on Solaris platforms. A VxVM system consists of an arbitrary group of one or more physical volumes organized into logical disk groups.

W

Wake ONLAN

Remote power-up support for systems running in power-save mode from some other system on the same LAN.

Web reporting

The Data Protector functionality that allows you to view reports on backup, object copy, and object consolidation status and Data Protector configuration using the Web interface.

wildcard character

A keyboard character that can be used to represent one or many characters. The asterisk (*), for example, typically represents one or more characters, and the question mark (?) typically represents a single character. Wildcard characters are often used in operating systems as a means of specifying more than one file by name.

Windows configuration backup

Data Protector allows you to back up Windows CONFIGURATION, including Windows Registry, user profiles, Event Logs, and WINS and DHCP server data (if configured on a system) in one step.

Windows Registry

A centralized database used by Windows to store configuration information for the operating system and the installed applications.

WINS server

A system running Windows Internet Name Service software that resolves Windows networking computer names to IP addresses. Data Protector can back up WINS server data as part of the Windows configuration.

writer (Microsoft VSS specific term)

A process that initiates change of data on the original volume. Writers are typically applications or system services that write persistent information on a volume. Writers also participate in the shadow copy synchronization process by assuring data consistency.

X

XBSA Interface (Informix Server specific term)

ON-Bar and Data Protector communicate with each other through the X/Open Backup Services Application Programmer's Interface (XBSA).

Z

ZDB

See zero downtime backup.

ZDB database (ZDB specific term)

A part of the IDB, storing ZDB related information such as source volumes, replicas and security information. The ZDB database is used for ZDB, instant recovery, and split mirror restore. See also zero downtime backup (ZDB).

ZDB to disk (ZDB specific term)

A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time. Multiple replicas, produced using the same backup specification at different times, can be kept in a replica set. A replica from ZDB to disk can be restored using the instant recovery process. See also zero downtime backup (ZDB), ZDB to tape, ZDB to disk+tape, instant recovery, and replica set rotation.

ZDB to disk+tape (ZDB specific term)

A form of zero downtime backup where the replica produced is kept on the disk array as a backup of the source volumes at a specific point in time, in the same way as with ZDB to disk. However, data in the replica is also streamed to a backup medium, as with ZDB to tape. If this backup method is used, data backed up in the same session can be restored using the instant recovery process, the standard Data Protector restore from tape, or with specific disk array families, split mirror restore. See also zero downtime backup (ZDB), ZDB to disk, ZDB to tape, instant recovery, replica, and replica set rotation.

ZDB to tape (ZDB specific term)

A form of zero downtime backup where data in the replica produced is streamed to a backup medium, typically tape. Instant recovery is not possible from such a backup, so the replica need not be retained on the disk array after backup completion. The backed up data can be restored using standard Data Protector restore from tape. With specific disk array families, split mirror restore can also be used. See also zero downtime backup (ZDB), ZDB to disk, ZDB to disk+tape, instant recovery, and replica.

zero downtime backup (ZDB)

A backup approach in which data replication techniques, provided by a disk array, are used to minimize the impact of backup operations on an application system. A replica of the data to be backed up is created first. All subsequent backup operations are performed on the replicated data rather than the original data, while the application system can be returned to normal operation. See also ZDB to disk, ZDB to tape, ZDB to disk+tape, and instant recovery.

Index

B

- backup 19
- backup solutions 8

C

- changing
 - settings 38
- CLI 41, 47, 49, 51
- command line reference 41
- configuring 15
 - remote powershell 25

D

- debugging
 - before you begin 53
 - enabling 53
- dismounting
 - database 49
 - databases 37

G

- Granular Recovery Extension
 - installing 11
- granularity 8

H

- HP Data Protector Granular Recovery Extension
 - environment 12

I

- importing
 - databases 26
- installation 11

- installing

 - Granular Recovery Extension 11

- installing extension

 - system 13

K

- known issues
 - workarounds 54

L

- listing
 - content database 49

M

- mailbox
 - CLI 49
- mounting
 - database 49
 - databases 32
- multiple
 - mailboxes 9

O

- opening
 - Granular Recovery Extension GUI 24

P

- prerequisites
 - before installing 11

R

- recover granularity capability 8
- recovering
 - another location 50
 - original location 50
- recovery
 - different locations 9

- granularity 9

- restore 20

- restore flow 22

- removing

 - database 51

 - database from disk 51

 - databases 37

- removing extension

 - system 14

- restore

 - requests 9

- restoring

 - database 47

S

- search 9

- searching 49

- starting

 - recovery 33

- system 12

T

- troubleshooting 53

 - before you begin 53

We appreciate your 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 Granular Recovery Extension User Guide for Microsoft Exchange Server (Data Protector 8.10)

Just 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.

