HP Data Protector 8.00 Integration Guide for Microsoft Volume Shadow Copy Service

Contents

Publication history

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

Table 1 Edition history

Part number	Guide edition	Product
N/A	June 2013	Data Protector release 8.00
N/A	June 2013 (second edition)	Data Protector release 8.00

About this guide

This guide describes how to configure and use Data Protector with applications that use Volume Shadow Copy Service for backup.

Intended audience

This guide is intended for backup administrators responsible for planning, setting up, and maintaining network backups using the Data Protector Microsoft Volume Shadow Copy Service. It assumes you are familiar with:

- Basic Data Protector functionality
- Database administration

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

Documentation set

The Help and other guides provide related information.

NOTE: The documentation set available at the HP support website at http://support.openview.hp.com/selfsolve/manuals contains the latest updates and corrections.

Help

Data Protector provides Help topics and context-sensitive (F1) Help for Windows and UNIX platforms. Install the Help during the Data Protector setup procedure by selecting the installation component `English Documentation (Guides, Help)` (Windows systems) or `OB2-DOCS` (on UNIX systems). Once installed, the Help resides in the following directory:

Windows systems: `Data_Protector_home\help\enu`

UNIX systems: `/opt/omni/help/C/help_topics`

You can access the Help from the top-level directory of any installation DVD-ROM without installing Data Protector:

Windows systems: Open `DP_help.chm`.

UNIX systems: Unpack the zipped tar file `DP_help.tar.gz` and open `DP_help.htm`.

Guides

Data Protector guides are available in the electronic PDF format. Install the PDF files during the Data Protector setup procedure by selecting the installation component `English Documentation (Guides, Help)` (on Windows systems) or `OB2-DOCS` (on UNIX systems). Once installed, the guides reside in the following directory:

Windows systems: `Data_Protector_home\docs`

UNIX systems: `/opt/omni/doc/C`

You can also access the guides:

- From the **Help** menu of the Data Protector graphical user interface
- From the HP support website at http://support.openview.hp.com/selfsolve/manuals (where the most up-to-date guide versions are available)

Data Protector guides are:

- *HP Data Protector Getting Started Guide*

  This guide contains information to get you started with using Data Protector. It lists installation prerequisites, provides instructions on installing and configuring a basic backup environment and procedures for performing backup and restore. It also lists resources for further information.

- *HP Data Protector Concepts Guide*

  This guide describes Data Protector concepts and provides background information on how Data Protector works. It is intended to be used with the task-oriented Help.

- *HP Data Protector Installation and Licensing Guide*

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

- *HP Data Protector Troubleshooting Guide*

  This guide describes how to troubleshoot problems you may encounter when using Data Protector.

- *HP Data Protector Disaster Recovery Guide*

  This guide describes how to plan, prepare for, test, and perform a disaster recovery.

- *HP Data Protector Command Line Interface Reference*

  This guide describes the Data Protector command-line interface, command options, and their usage as well as provides some basic command-line examples. It is located in the following directory:

  Windows systems: `Data_Protector_home\docs\MAN`

  UNIX systems: `/opt/omni/doc/C/`

  On UNIX systems, you can use the `omniintro` man page to display a list of the available Data Protector commands. You can then execute the man `CommandName` command to retrieve information about each Data Protector command.

- *HP Data Protector Product Announcements, Software Notes, and References*

  This guide gives a description of new features of HP Data Protector 8.00. It also provides information on installation requirements, required patches, and limitations, as well as known issues and workarounds.

- *HP Data Protector Integration Guides*

  These guides describe how to configure and use Data Protector to back up and restore various databases and applications. They are intended for backup administrators and operators. There are six guides:

  - *HP Data Protector Integration Guide for Microsoft Applications: SQL Server, SharePoint Server, and Exchange Server*

    This guide describes the integrations of Data Protector with the following Microsoft applications: Microsoft SQL Server, Microsoft SharePoint Server, and Microsoft Exchange Server.

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

    This guide describes the integrations of Data Protector with Oracle Server, SAP R/3, and SAP MaxDB.

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

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

- ◦ *HP Data Protector Integration Guide for Sybase and Network Data Management Protocol Server*

  This guide describes the integrations of Data Protector with Sybase Server and Network Data Management Protocol Server.

- ◦ *HP Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*

  This guide describes the integration of Data Protector with the Microsoft Volume Shadow Copy Service. This guide also documents application writer specifics.

- ◦ *HP Data Protector Integration Guide for Virtualization Environments*

  This guide describes the integrations of Data Protector with virtualization environments: VMware Virtual Infrastructure, VMware vSphere, VMware vCloud Director, Microsoft Hyper-V, and Citrix XenServer.

- *HP Data Protector Zero Downtime Backup Concepts Guide*

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

- *HP Data Protector Zero Downtime Backup Administrator's Guide*

  This guide describes how to configure and use the integration of Data Protector with HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, HP 3PAR StoreServ Storage, and EMC Symmetrix Remote Data Facility and TimeFinder. It is intended for backup administrators or operators. It covers the zero downtime backup, instant recovery, and the restore of filesystems and disk images.

- *HP Data Protector Zero Downtime Backup Integration Guide*

  This guide describes how to configure and use Data Protector to perform zero downtime backup, instant recovery, and standard restore of Oracle Server, SAP R/3, Microsoft Exchange Server, and Microsoft SQL Server databases.

- *HP Data Protector Granular Recovery Extension User Guide for Microsoft Exchange Server*

  This guide describes how to configure and use the Data Protector Granular Recovery Extension for Microsoft Exchange Server. Graphical user interface of the Data Protector Granular Recovery Extension for Microsoft Exchange Server is integrated into the Microsoft Management Console. This guide is intended for Microsoft Exchange Server administrators and Data Protector backup administrators.

- *HP Data Protector Granular Recovery Extension User Guide for Microsoft SharePoint Server*

  This guide describes how to configure and use the Data Protector Granular Recovery Extension for Microsoft SharePoint Server. The Data Protector Granular Recovery Extension is integrated into Microsoft SharePoint Server Central Administration and enables you to recover individual items. This guide is intended for Microsoft SharePoint Server administrators and Data Protector backup administrators.

- *HP Data Protector Granular Recovery Extension User Guide for VMware vSphere*

  This guide describes how to configure and use the Data Protector Granular Recovery Extension for VMware vSphere. The Data Protector Granular Recovery Extension is integrated into VMware vCenter Server and enables you to recover individual items. This guide is intended for VMware vCenter Server users and Data Protector backup administrators.

- *HP Data Protector Deduplication*

  This technical white paper describes basic data deduplication concepts, principles of Data Protector integration with Backup to Disk devices and its use of deduplication. It also provides instructions how to configure and use deduplication in Data Protector backup environments.

- *HP Data Protector Integration with Autonomy IDOL Server*

  This technical white paper describes all aspects of integrating Data Protector with Autonomy IDOL Server: integration concepts, installation and configuration, Data Protector backup image indexing, full content search-based restore, and troubleshooting.

- *HP Data Protector Integration with Autonomy LiveVault*

  This technical white paper all aspects of integrating Data Protector with Autonomy LiveVault: integration concepts, installation and configuration, backup policy management, cloud backup, cloud restore, and troubleshooting.

Documentation map

Abbreviations

Abbreviations in the documentation map that follows are explained below. The documentation item titles are all preceded by the words "HP Data Protector".

Abbreviation	Documentation item
CLI	Command Line Interface Reference
Concepts	Concepts Guide
DR	Disaster Recovery Guide
GS	Getting Started Guide
GRE Exchange	Granular Recovery Extension User Guide for Microsoft Exchange Server
GRE SPS	Granular Recovery Extension User Guide for Microsoft SharePoint Server
GRE VMware	Granular Recovery Extension User Guide for VMware vSphere
Help	Help
Install	Installation and Licensing Guide
IG IBM	Integration Guide for IBM Applications: Informix, DB2, and Lotus Notes/Domino
IG MS	Integration Guide for Microsoft Applications: SQL Server, SharePoint Server, and Exchange Server
IG VSS	Integration Guide for Microsoft Volume Shadow Copy Service
IG O/S	Integration Guide for Oracle and SAP
IG Var	Integration Guide for Sybase and Network Data Management Protocol Server
IG VirtEnv	Integration Guide for Virtualization Environments
IG IDOL	Integration with Autonomy IDOL Server
IG LV	Integration with Autonomy LiveVault

Abbreviation	Documentation item
PA	Product Announcements, Software Notes, and References
Trouble	Troubleshooting Guide
ZDB Admin	ZDB Administrator's Guide
ZDB Concepts	ZDB Concepts Guide
ZDB IG	ZDB Integration Guide

Map

The following table shows where to find information of different kinds. Shaded squares are a good place to look first.

	Help	GS	Concepts	Install	Trouble	DR	CLI	PA	Integr. guides						ZDB			GRE		
									MS	O/S	IBM	Var	VSS	VirtEnv	Concepts	Admin	IG	Exchange	SPS	VMware
Backup	X	X	X						X	X	X	X	X	X	X	X	X			
CLI								X												
Concepts, techniques	X		X						X	X	X	X	X	X	X	X	X	X	X	X
Disaster recovery	X		X			X														
Installation, upgrade	X	X		X				X												
Instant recovery	X		X												X	X	X			
Licensing	X			X				X												
Limitations	X				X			X	X	X	X	X	X	X			X			
New features	X							X												
Planning strategy	X		X												X					
Procedures, tasks	X			X	X	X			X	X	X	X	X	X		X	X	X	X	X
Recommendations		X						X							X					
Requirements				X				X	X	X	X	X	X	X						
Restore	X	X	X						X	X	X	X	X	X		X	X	X	X	X
Supported configurations															X					
Troubleshooting	X			X	X				X	X	X	X	X	X		X	X	X	X	X

Integrations

Look in these guides for details of the integrations with the following software applications:

Software application	Guides
Autonomy IDOL Server	IG IDOL
Autonomy LiveVault	IG LV
IBM DB2 UDB	IG IBM
Informix Server	IG IBM
Lotus Notes/Domino Server	IG IBM
Microsoft Exchange Server	IG MS, ZDB IG, GRE Exchange

Software application	Guides
Microsoft Hyper-V	IG VirtEnv
Microsoft SharePoint Server	IG MS, ZDB IG, GRE SPS
Microsoft SQL Server	IG MS, ZDB IG
Microsoft Volume Shadow Copy Service (VSS)	IG VSS
Network Data Management Protocol (NDMP) Server	IG Var
Oracle Server	IG O/S, ZDB IG
SAP MaxDB	IG O/S
SAP R/3	IG O/S, ZDB IG
Sybase Server	IG Var
VMware vCloud Director	IG VirtEnv
VMware vSphere	IG VirtEnv, GRE VMware

Look in these guides for details of the integrations with the following families of disk array systems:

Disk array family	Guides
EMC Symmetrix	all ZDB
HP P4000 SAN Solutions	ZDB Concepts, ZDB Admin, IG VSS
HP P6000 EVA Disk Array Family	all ZDB, IG VSS
HP P9000 XP Disk Array Family	all ZDB, IG VSS
HP 3PAR StoreServ Storage	ZDB Concepts, ZDB Admin, IG VSS

Document conventions and symbols

Table 2 Document conventions

Convention	Element
Blue text: "Document conventions" (page 13)	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	Website addresses
Bold text	- Keys that are pressed - Text typed into a GUI element, such as a box - GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes
Italic text	Text emphasis
`Monospace` text	- File and directory names - System output - Code - Commands, their arguments, and argument values
`Monospace, italic` text	- Code variables - Command variables
`Monospace, bold` text	Emphasized monospace text

△ **CAUTION:** Indicates that failure to follow directions could result in damage to equipment or data.

① **IMPORTANT:** Provides clarifying information or specific instructions.

NOTE: Provides additional information.

:ּ॑ः **TIP:** Provides helpful hints and shortcuts.

Data Protector graphical user interface

Data Protector provides a graphical user interface for Microsoft Windows operating systems. For information about it, see the *HP Data Protector Help*.

Figure 1 Data Protector graphical user interface
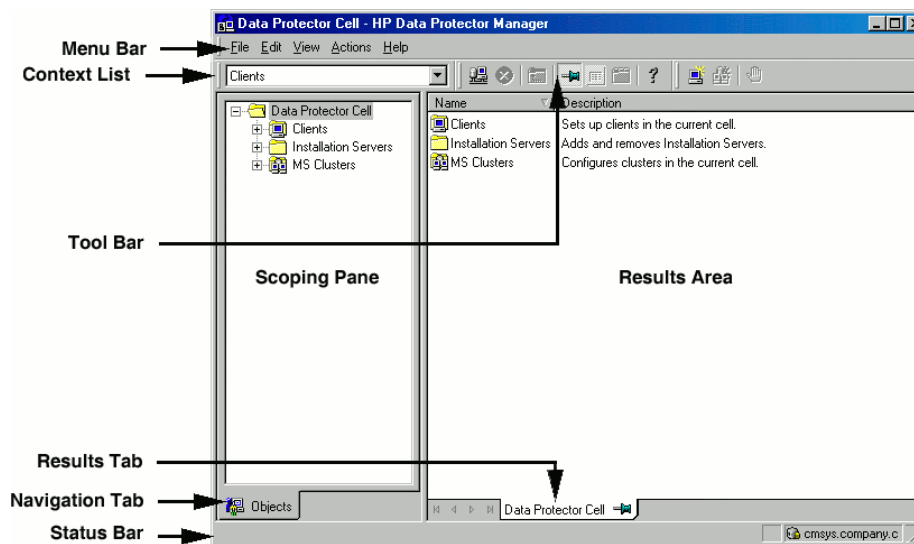
General information

General information about Data Protector can be found at http://www.hp.com/go/dataprotector.

HP technical support

For worldwide technical support information, see the HP support website:

http://www.hp.com/support

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

http://www.hp.com/go/e-updates

After registering, you will receive e-mail notification of product enhancements, new driver versions, firmware updates, and other product resources.

HP websites

For additional information, see the following HP websites:

- http://www.hp.com
- http://www.hp.com/go/software
- http://support.openview.hp.com/selfsolve/manuals
- http://www.hp.com/support/downloads

Documentation feedback

HP welcomes your feedback.

To make comments and suggestions about product documentation, please send a message with the subject line `Feedback on Data Protector documentation` to **AutonomyTPFeedback@hp.com**. All submissions become the property of HP.

1 Introduction

A traditional backup process is based on the direct communication between the backup application and the application whose data is backed up. This backup method requires from the backup application an individual interface for each application it backs up. To address the growing number of applications and their specific interfaces, a new coordinator between the actors of the backup and restore process was introduced on Microsoft Windows systems — the *Microsoft Volume Shadow Copy Service*.

Data Protector integrates with the Volume Shadow Copy service through the Data Protector Microsoft Volume Shadow Copy Service integration (**VSS integration**).

The Data Protector integration with the Microsoft Volume Shadow Copy Service can serve in two functions:

- *It provides support for certified VSS writers*. For a complete list of supported VSS writers and providers, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.

- *It forms the base for other integrations*, such as the Data Protector Microsoft Exchange Server 2010 integration.

This guide describes how to use the generic Data Protector VSS integration for backing up and restoring writer's data. For descriptions of application integrations, based on the VSS technology, see the appropriate guides.

Benefits of using the VSS integration

Advantages of using the Data Protector VSS integration are the following:

- A unified backup interface is provided for all applications that provide a writer.

- Data integrity is provided on application level, because it is provided by the writers. No interference is needed from the backup application.

- Due to the generic nature of the programming and user interfaces, support for new applications, application versions, or disk arrays (for example through certifications and so on) can be added much easier, even after the initial Data Protector release.

  Application integrations, such as the Data Protector Exchange Server 2010 integration, on the other hand offer more functionality and are better adapted to the application specifics than the generic VSS integration. However, due to additional implementation specifics, it may take longer to develop such integrations as compared to adding support through the generic Data Protector VSS integration.

2 Integration concepts

The Data Protector VSS integration uses the Microsoft Volume Shadow Copy Service and Virtual Disk Service interfaces and the VSSBAR agent to perform backup and restore of application data.

Volume Shadow Copy Service

Volume Shadow Copy Service (VSS) is a software service available on Microsoft Windows operating systems.

During backup, the service collaborates with a backup application (for example, Data Protector), applications to be backed up (usually database applications), shadow copy providers, and the operating system kernel to create a consistent, point-in-time shadow copy set. These shadow copy sets can be backed up to backup media or can be kept on a disk array for instant recovery or data mining.

During restore, the service collaborates with the backup application and database application to prepare for the restore operation and to perform the recovery of the restored data.

The Data Protector Volume Shadow Copy integration provides a unified communication interface that can coordinate backup and restore of an application regardless of their specific features. With this approach, a backup application does not need to handle each application to be backed up specifically. However, the production application as well as the backup application must conform to the VSS specification.

"Actors of the traditional backup model" (page 17) and "Actors of the Data Protector VSS integration backup model" (page 18) show the differences between the traditional backup model and the model with the Data Protector Microsoft Volume Shadow Copy Service integration.

Figure 2 Actors of the traditional backup model
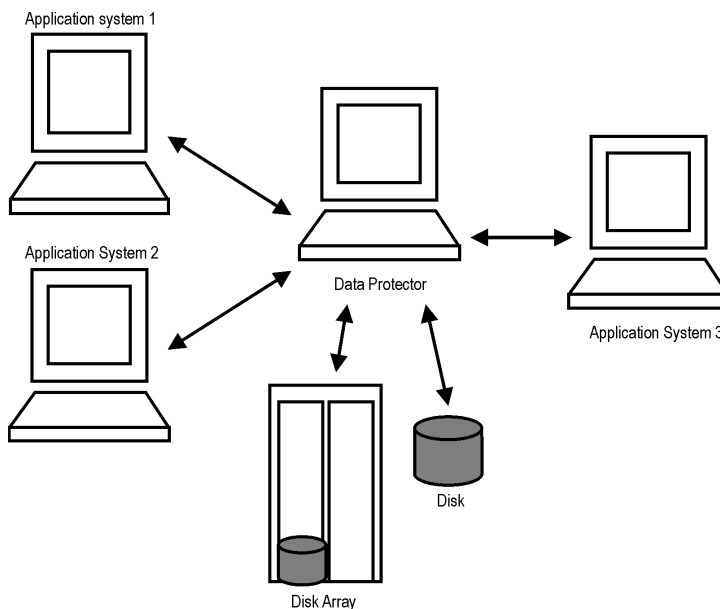
Figure 3 Actors of the Data Protector VSS integration backup model
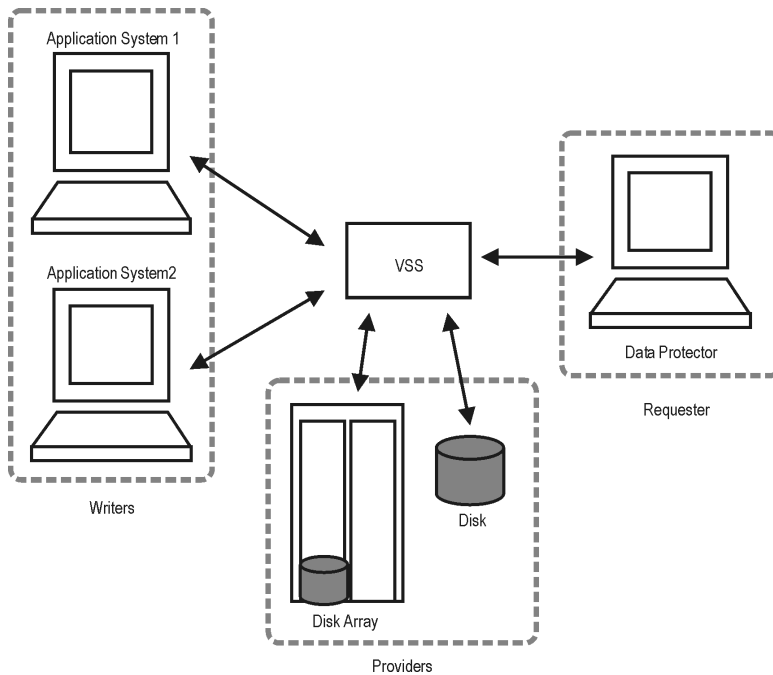
Without using the Volume Shadow Copy Service, Data Protector has to communicate with each application to be backed up individually. The Data Protector VSS integration introduces a unified backup and restore interface and provides the coordination among the participants of the backup and restore process. The main actors of a VSS based backup are the:

- **Requestor** – the backup application which requests a shadow copy creation.

- **Providers**, which create the shadow copies. Two types of *VSS* providers exist:

  ○ software providers

  ○ hardware providers, supplied by the hardware (disk array) vendors

- **Writers** – application-specific software modules which ensure application's data consistency when a shadow copy is created.

Virtual Disk Service

Virtual Disk Service (VDS) is a Microsoft Windows service that provides a common interface for managing storage hardware and disks and for creating volumes on these disks. Similar to VSS, the hardware capabilities are abstracted by **VDS providers**.

Two types of VDS providers exist:

- built-in software providers (at the operation system level)

- hardware providers, supplied by the hardware (disk array) vendors.

You can use the VDS provider to configure and manage disks, volumes, partitions, and hardware RAID subsystems. The disks can either be physical disks or virtual disks (LUNs). If your hardware providers support this, you can manage individual physical disks in a LUN. The software provider does not resolve LUNs to individual physical disks.

For some disk arrays, the Data Protector VSS integration can use the VDS hardware providers (if present) to gather additional information needed for instant recovery and to perform the instant recovery.

VSSBAR agent

The central part of the integration is the **VSSBAR agent**, which links Data Protector with the Microsoft Volume Shadow Copy Service. Data Protector Microsoft Volume Shadow Copy Service integration uses the VSSBAR agent for automatic browsing of VSS-aware writers, coordinating backup and restore. The VSSBAR agent is responsible for the following actions:

- detecting VSS writers
- examining and analyzing Writer Metadata Document (WMD)

  Writer Metadata Document (WMD) is metadata provided by each writer. Writers identify themselves by the metadata and instruct the backup application what to back up and how to restore the data. Thus, Data Protector follows the requirements provided by the writer when selecting the volumes to be backed up and the restore method.

- requesting shadow copy creation
- backing up writers' data to media
- coordinating restore session start
- restoring the Writer Metadata Document
- restoring writer's data from media
- requesting instant recovery

During the Data Protector VSS integration backup, Data Protector does not interact directly with each writer, but through the VSS interface. It uses the VSSBAR agent to coordinate the backup process. The consistency of data is a responsibility of the VSS writer and not dependent on Data Protector functionality. The shadow copy is created by the VSS provider.

Backup types

The available backup types depend on the type of the VSS provider used and its functionality:

- With the *software provider*, you can back up the data to tape (**standard Data Protector backup to tape**). This can be done only on the same system where the shadow copy is created.
- With the *hardware provider*, you can create target volumes, which can be used for backing up to tape or for instant recovery. Such a backup is a Data Protector **zero downtime backup (ZDB)**, where the hardware providers replace the disk array agent, which usually replicates the volumes. Thus, to create a ZDB backup specification when working with the Data Protector VSS integration, you must always select the hardware providers. For a general description of ZDB (split mirror or snapshot backup) and instant recovery concepts, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

  The backup system to which the replicas are presented and from where they are backed up to tape can either be:

  ◦ *A dedicated system* (Data Protector ZDB terminology: a dual-host configuration).

    In VSS terminology, a backup performed in such an environment is a **transportable backup**. It is available only in combination with a hardware provider.

  ◦ *The same system as the application system* (Data Protector ZDB terminology: a single-host configuration).

    In VSS terminology, a backup performed in such an environment is a **local or network backup**.

Figure 4 Local VSS backup
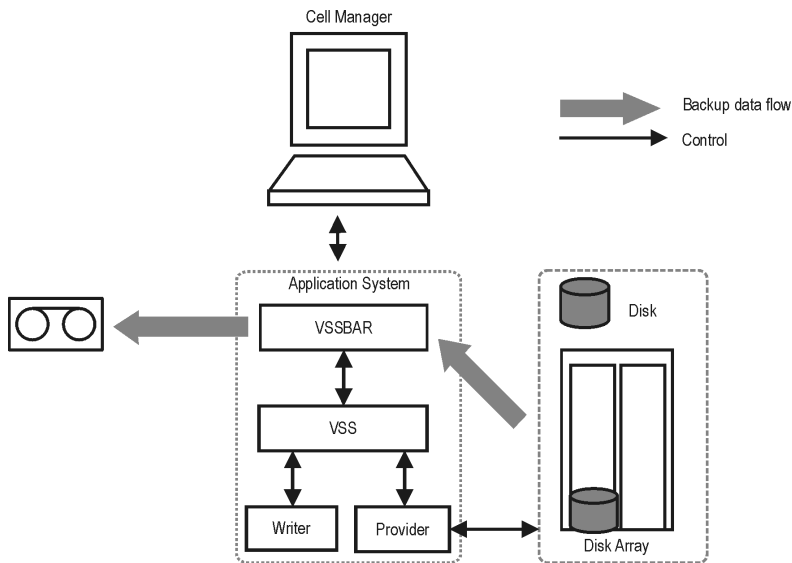
Figure 5 Transportable VSS backup
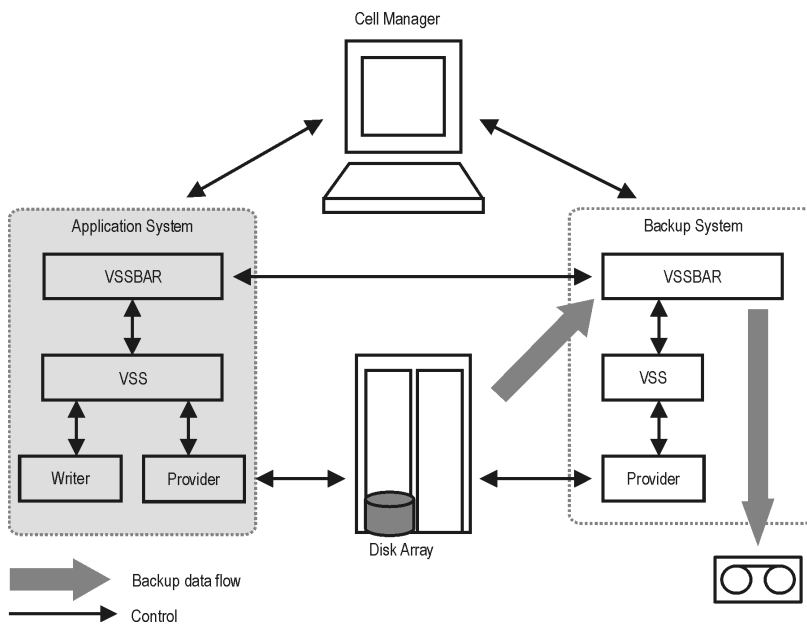
VSS database

The VSS database (VSSDB) is an extension to the Data Protector Internal Database (IDB) on the Cell Manager. It holds the VSS integration-specific information about the VSS backup sessions. The stored information can be used for the following purposes:

- Instant recovery

- Data mining (you can save backup components and writer metadata documents)

- Listing the backup sessions and viewing details about them

- Removing the backup sessions

- Enabling (presenting and mounting) and disabling (dismounting and unpresenting) the ZDB-to-disk and ZDB-to-disk+tape sessions' replicas

The VSSDB holds the VSS sessions' metadata. This metadata is stored in two parts of the VSSDB, persistent part and the non-persistent part. The **persistent** part holds information about all backup

sessions, while the **non-persistent** part holds information only on instant recovery–enabled sessions (ZDB-to-disk and ZDB-to-disk+tape sessions). Once a replica from the instant recovery-enabled session is rotated out of the replica set, the information about the session is deleted from the non-persistent part of the VSSDB and the session's target volumes are removed from the disk array. The information about the session is still available in the persistent part of the VSSDB and can be deleted only manually using the omnidbvss command. Note that ZDB-to-tape sessions and sessions created using the software provider are always recorded only to the persistent part of the VSSDB.

You can query and manage the items of the VSSDB using the omnidbvss command. For information on the command syntax, description, and examples, see the *HP Data Protector Command Line Interface Reference*.

The following tasks can be performed using the omnidbvss command:

- "Querying the VSSDB" (page 21)
- "Deleting backup sessions" (page 21)
- "Enabling and disabling replicas" (page 21)

Querying the VSSDB

Using the omnidbvss command, you can list:

- All backup sessions, as well as their details.
- All backup sessions based on a specific backup specification, as well as their details.
- All backup sessions, which were created before a specified date, as well as their details.
- Details on a specific backup session, identified by the session ID.

Using the omnidbvss command, you can save the documents about the backup components and writer metadata to a specified directory.

See the omnidbvss man page for the command syntax and examples.

Deleting backup sessions

Using the omnidbvss command, you can delete:

- A specific instant recovery-enabled backup session (a replica version), identified by the session ID, from the non-persistent part of the VSSDB and disk array, or only from the VSSDB.
- All instant recovery-enabled backup sessions based on a specific backup specification (a replica set) from the non-persistent part of the VSSDB and disk array, or only from the VSSDB.
- A specific backup session, identified by the session ID, from the persistent part of the VSSDB.
- All backup sessions based on a specific backup specification or session that were created before a specified date from the VSSDB.

See the omnidbvss man page for the command syntax and examples.

Enabling and disabling replicas

Using the omnidbvss command, you can present and mount (enable) on a backup system and dismount and unpresent (disable) from a backup system the replicas from:

- All instant recovery-enabled sessions.
- A specific instant recovery-enabled session, identified by the session ID.
- All instant recovery-enabled sessions based on a specific backup specification.

Changes in the environment

During VSS instant recovery sessions or VSSDB maintenance sessions using the omnidbvss command (remove, disable, enable), the state of target volumes in a replica is checked prior to

restore, remove, disable, and enable operations. Before restore, the state of the source volumes on the application system is also checked. If the check finds any changes that have not been tracked in the VSSDB, the session is aborted, notifying you about the specific changes. This check can be disabled by setting the `OB2VSS_IGNORE_BACKUP_DISK_CHANGES` and `OB2VSS_IGNORE_SOURCE_DISK_CHANGES` omnirc options to `1`.

During a VSS backup session, a replica is created and left on a disk array until the specified number of replicas rotated is reached. After that, the next replica to be created replaces the oldest replica in the set. If the oldest replica in the set cannot be replaced because you have manually (not using Data Protector) unpresented the replica from the backup system and presented it on some other system, the backup session continues and creates a new replica and leaves the old one on the array. Note that in this case, if there is no space on the disk array for creating a new replica, the backup session also fails.

If you set the `OB2VSS_IGNORE_BACKUP_DISK_CHANGES` omnirc option to `1`, the backup session continues by removing the old replica and creating a new one in its place.

3 Configuration

Prerequisites and limitations

This is a list of prerequisites and limitations for the Data Protector Microsoft Volume Shadow Copy Service integration. Additional limitations and recommendations that are not directly connected to the integration (such as operating system and GUI limitations) and disk array limitations are listed in the *HP Data Protector Product Announcements, Software Notes, and References*.

Prerequisites

- Before you begin, ensure that you have correctly installed and configured Data Protector, writers. If you want to perform ZDB, install and configure also the VSS and VDS hardware providers. See the:

  ○ Latest support matrices on the web for an up-to-date list of supported versions, platforms, devices, disk arrays, limitations, and other information.

    http://support.openview.hp.com/selfsolve/manuals

  ○ *HP Data Protector Installation and Licensing Guide* for information of how to install Data Protector on various architectures and how to install the Data Protector Microsoft Volume Shadow Copy Service integration and ZDB agents.

  ○ Writers and shadow copy providers documentation for instructions of how to install and configure writers and providers on your system.

- Install any needed Microsoft Windows patches or hotfixes. For a detailed list, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.

Transportable backup prerequisites

- The backup system must be configured to accept connections from the application system and the other way round.

- The following Data Protector and HP storage components must be installed and configured on both *the application and backup system*:

  ○ `MS Volume Shadow Copy Integration`

  ○ The appropriate Data Protector disk array agent

  ○ VSS hardware provider

  Depending on the operating system version and disk array, the following HP Storage components may be required:

  ○ VDS hardware provider

    – For ZDB, the VDS hardware provider is optional.

    – For instant recovery, certain limitations may apply, depending on the disk array family and version of the operating system. In some cases, the VDS hardware provider is required for a successful instant recovery. See "Instant recovery" (page 49).

Limitations

For a list of general Data Protector limitations, see the *HP Data Protector Product Announcements, Software Notes, and References*.

- To run a VSS integration backup, the writer's data must be on an NTFS filesystem. For hardware providers, this is not required.
- The VSS integration backup of writers which store their data on network shared volumes is not supported.
- The Data Protector Microsoft VSS integration does not provide any restore method for writers requesting a custom restore method. These writers are by default not presented by Data Protector.

  If a writer specifies a custom restore method, it is only possible to restore the writer's data as plain files using the Data Protector functionality. You can perform the custom restore manually. For additional information on the restore methods, see the writer's documentation.

- Preview is not possible for VSS backup and restore sessions.
- Backup preview is only available for VSS filesystem backup sessions.
- Dynamic disks (Logical Disk Manager partitions) can only be backed up with software providers.

ZDB and instant recovery limitations

- Not all VSS writers supported by Data Protector can be used with the instant recovery functionality. For a list of the VSS writers that support the instant recovery functionality, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.
- VSS backups with hardware provider are only supported for volumes on disk arrays. This is also valid for system disks. Instant recovery of system disks backups is not possible.
- The application system and the backup system that you use must have the same operating system version installed.
- During a local backup, with the option `Track the replica for instant recovery` selected, and with the option `Mount the replica` not selected, note the following:

  ◦ On Windows Server 2003, the created snapshots are unpresented on the application system after the backup.

  ◦ On Windows Server 2008, the created snapshots are only put offline after the backup, while they remain presented on the application system.

HP P6000 EVA Disk Array Family hardware provider limitations

- With the HP P6000 EVA Disk Array Family provider, if there are too many, for example more than four, source volumes selected for backup, the creation of snapshots may take too long and VSS may fail the snapshots creation process. It is strongly recommended not to specify too many objects in a backup specification, since they might reside on too many source volumes. For example, a Microsoft Exchange Server storage group should typically not reside on more than four source volumes, and only one storage group should be selected for backup if the storage group occupies four source volumes. The same limitation also applies to Microsoft SQL Server databases. If your backup session still fails, reduce the number of volumes further, for example to two.
- On Windows Server 2003 systems if a disk array from the HP P6000 EVA Disk Array Family is used, instant recovery is possible only if the VDS hardware providers is installed.

  On Windows Server 2008 systems, the VDS providers are optional. However, if the VDS provider was installed the time of the backup session, it must be installed at the time of the instant recovery session.

HP P9000 XP Disk Array Family hardware provider limitations

- If a disk array from the HP P9000 XP Disk Array Family is used and the P9000 XP Array VSS provide is in resync mode, the VDS providers are optional for instant recovery. However, if the VDS provider was installed the time of the backup session, it must be installed at the time of the instant recovery session.

  If the P9000 XP Array provider is in VSS compliant mode, the VDS providers are always required for instant recovery.

- You cannot switch between the VSS compliant mode and resync mode during execution of a backup session.

- If the P9000 XP Array provider is in resync mode, a maximum of three replicas (S-VOLs) can be created for a source volume (P-VOL).

- If the HP P9000 XP Disk Array Family VSS hardware provider is in the VSS compliant mode, and there is a high read-write load on the disk array, the replica creation may take a long time. Consequently, the backup session may fail.

  By default, Data Protector waits for a maximum of 45 minutes (2700 seconds) for mirrors to become synchronized. If synchronization is not completed when the waiting period expires, Data Protector aborts the backup session.

  You can increase the waiting period duration by setting the `OB2VSS_WAIT_TIMEOUT` omnirc option. The option value determines the period (in seconds) for which the Data Protector agent waits for mirrors that are created in one session to become fully synchronized. If the `OB2VSS_WAIT_TIMEOUT` option is set to more than 2 hours (7200 seconds), you should additionally increase the value of the global option `SmIdleTimeout` to be greater that the value of `OB2VSS_WAIT_TIMEOUT`. The value of `SmIdleTimeout` should be specified in minutes, rather than seconds.

  If backup sessions are lengthy and there is a possibility that the VSS backup window will not be met, it is highly recommended that you first switch the HP P9000 XP Disk Array Family VSS hardware provider to resync mode. Afterwards, consecutively run backup sessions until the backup session count reaches the current value of the **Number of replicas rotated** option. This will start creating disk mirrors on P9000 XP Array that will be used in scheduled backups.

- Second-level mirrors are not supported.

For supported backup and connectivity topologies, see the *HP Data Protector Zero Downtime Backup Administrator's Guide.*

Configuring the application system for instant recovery-enabled backup sessions

To be able to perform instant recovery-enabled backups and then instant recovery, you must resolve the application system using the `omnidbvss` command. During the resolve operation, Data Protector contacts VDS to get information about the storage IDs of the source volumes presented on the application system. The information about the storage IDs is required during instant recovery, in case such information cannot be gathered at restore time.

If resolving is not performed before a backup session, one of the following happens:

- The application system is resolved automatically during the backup session, if the `OB2VSS_DISABLE_AUTO_RESOLVE` option in the `omnirc` file is set to `0` (default). In this case, the backup time for creating a replica is prolonged.

- A ZDB-to-disk session fails notifying you to resolve the application system before the backup is run, if the `OB2VSS_DISABLE_AUTO_RESOLVE` option is set to `1`.

- A ZDB-to-disk+tape session completes with the warning that only backup to tape was performed not leaving the replica on the disk array and thus disabling instant recovery, if the `OB2VSS_DISABLE_AUTO_RESOLVE` option is set to `1`.

Always perform the resolve operation after:

- Installing or upgrading Data Protector.
- Your source volumes' configuration on the application system has changed (for example, you have modify the existing source volumes or you have presented new source volumes).
- You have added a new storage object (for example, a Microsoft Exchange Server storage group).

To resolve the application system, execute the following on any VSS client in the Data Protector cell:

```
omnidbvss –resolve {-apphost ApplicationSystem | –all}
```

where *ApplicationSystem* is the name of the application system you want to resolve. In a cluster environment (for example, Microsoft Exchange Server CCR), provide the virtual server name for the *ApplicationSystem* parameter to enable the resolve operation on both cluster nodes.

To resolve all application systems in a cell simultaneously, use the option –all.

Alternatively, you can set the omnirc option OB2VSS_ALWAYS_RESOLVE_SOURCES to 1. In this case, the source volumes on the application system are resolved automatically during every instant recovery-enabled backup session. However, the replica creation time is considerably prolonged. Do not set this option to 1 in a CCR Microsoft Exchange environment if you plan to back up only the database copy disks and leave the possibility for eventual instant recovery from these disk to the production database disks. In this case, the production database disks are never resolved during backup and this can cause such an instant recovery to fail.

Configuring HP P4000 SAN Solutions

This section describes how to configure the Data Protector HP P4000 SAN Solutions integration for use with the VSS integration.

Prerequisites

- Obtain or install **HP storage licenses and components**:
  - HP P4000 SAN Solutions VSS provider.
  - HP P4000 SAN Solutions DSM (Device Specific Module) for MPIO.

  For details of how to install these components, see the HP P4000 Windows Solution Pack documentation.

  For information on supported product versions, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.

- Obtain or install **Data Protector licenses and components**:
  - An instant recovery license.
  - HP P4000 Agent.

    For installation and licensing information, see the *HP Data Protector Installation and Licensing Guide*.

- Make sure the same operating system (and its version) is installed on the application and backup systems.
- Connect an HP P4000 SAN Solutions storage system to the application and backup systems through the SAN. Backup system must be connected to the same SAN as the HP P4000 SAN Solutions storage system.
- For VSS transportable backups, the source volumes must be presented to both the application and backup systems.

For supported backup and connectivity topologies, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

For general Data Protector and integration-specific limitations, see the *HP Data Protector Product Announcements, Software Notes, and References*.

Configuring the integration

Before you start configuration, make sure you met the prerequisites described in "Prerequisites" (page 31). Then set the login information for the SMI-S HP P4000 SAN Solutions provider running on a management system.

To set, delete, list, or check the login information, use the `omnidbp4000 --ompasswd` command. For command syntax and examples, see the *HP P4000 SAN Solutions* part of the *HP Data Protector Zero Downtime Backup Administrator's Guide* and the `omnidp4000` reference page in *HP Data Protector Command Line Interface Reference*.

The information you provide is kept in the HP P4000 SAN Solutions part of the ZDB database for SMI-S based integrations (SMISDB).

For each management system you configure, the following information is stored:

- Hostname as recognized in the IP network.

- Port number through which HP P4000 Agent communicates (default - 5988). If the SMI-S HP P4000 SAN Solutions provider accepts the SSL-based connection, the default port number is 5989.

- Management group name, and management group username and encrypted password of an HP P4000 SAN Solutions provider login account.

The HP P4000 SAN Solutions part of the SMISDB resides on the Cell Manager in the directory *Data_Protector_program_data*\server\db80\smisdb\p4000 (Windows systems) or /var/opt/omni/server/db80/smisdb/p4000/login (UNIX systems).

> ⓘ **IMPORTANT:** It is recommended to use the default port number (5988 for the non-SSL and 5989 for the SSL-based SMI-S P4000 SAN Solutions provider connection settings).
>
> It is also recommended to execute `omnidbp4000 --ompasswd --check [--host ClientName]` before backup or instant recovery to verify the configuration of SMI-S HP P4000 SAN Solutions provider.

Configuring the VSS hardware provider

Set the management group credentials, on both the application and the backup system. The credentials required for using the HP P4000 SAN Solutions VSS provider are:

- Management group name

- Management group user name and password

Depending on the operating system, use either the Authentication Console or edit the Windows registry. See the HP P4000 SAN Solutions VSS provider documentation for details.

After you complete this procedure, the HP P4000 SAN Solutions VSS provider is ready to be used for zero downtime backup and instant recovery.

Configuring HP P6000 EVA Disk Array Family

This section describes how to configure the Data Protector HP P6000 EVA Disk Array Family (P6000 EVA Array) integration for use with the VSS integration.

Prerequisites

- Obtain or install **HP storage licenses and components**:

  ◦ HP Virtual Controller Software (VCS or XCS) and Command View (CV) EVA.

  For installation instructions, see the VCS and CV EVA documentation. For information on supported product versions, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.

  ◦ HP Business Copy (BC) P6000 EVA microcode and license.

  ◦ HP MPIO Full Featured DSM (Device Specific Module) for HP P6000 EVA Disk Array Family.

  ◦ A P6000 EVA Array license for, at least, basic and snapshot operation.

- Obtain or install **Data Protector licenses and components**:

  ◦ An instant recovery license.

  ◦ HP P6000 / HP 3PAR SMI-S Agent.

  If this agent is not installed, only "Restore using Microsoft Virtual Disk Service" (switch of disks) is possible. For installation and licensing information, see the *HP Data Protector Installation and Licensing Guide*.

- Make sure the same operating system (and its version) is installed on the application and backup systems.

- Connect P6000 EVA Array to the application and backup systems through the SAN. Backup system must be connected to the same SAN as the P6000 EVA Array.

For supported backup and connectivity topologies, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

For general Data Protector and integration-specific limitations, see the *HP Data Protector Product Announcements, Software Notes, and References*.

Configuring the integration

Before you start configuration, make sure you met the prerequisites described in "Prerequisites" (page 28). Then set the login information for SMI-S P6000 EVA Array provider running on a management system.

To set, delete, list, or check the login information, use the `omnidbsmis -ompasswd` CLI command. For command syntax and examples, see the *HP Data Protector Command Line Interface Reference*.

The information you provide is kept in the ZDB database for the HP P6000 EVA Disk Array Family integration (SMISDB).

For each management system you configure, the following information is stored:

- Hostname as recognized in the IP network.

- Port number through which P6000 EVA SMI-S Agent communicates (default - 5988). If the SMI-S P6000 EVA Array provider accepts the SSL-based connection, the default port number is 5989.

- User name and encoded password for the SMI-S P6000 EVA Array provider login.

SMISDB resides on the Cell Manager in the directory `Data_Protector_program_data\server\db80\smisdb` (Windows systems) or `/var/opt/omni/server/db80/smisdb` (UNIX systems).

If a failover from the active to the standby management system happens, proceed as follows:

- If standby and failed management systems have the same hostname, no action is needed.

- If standby and failed management systems have different hostnames, remove the failed system from the Data Protector configuration, and then add the new management system.

ⓘ **IMPORTANT:** It is recommended to use the default port number (5988 for non-SSL and 5989 for the SSL-based SMI-S P6000 EVA Array provider connection settings).

It is also recommended to execute `omnidbsmis -ompasswd -check [-host ClientName]` before backup or instant recovery to verify the configuration of SMI-S P6000 EVA Array provider.

Configuring the VSS hardware provider

Open the HP EVA VSS Configuration Utility and follow the steps below:
1. Ensure that the correct management appliance IP address is specified.
2. Click **Login** and in the Login dialog box, enter the username and password.
3. Select **SnapClone** as the Snapshot Type.
4. Click **Select Disk Group** and select the disk group with the disks that you intend to back up for instant recovery.

Open the HP EVA VDS Configuration Utility and follow the steps below:
1. Ensure that the correct Appliance IP address is specified.
2. Click **Login** and in the Login dialog box, enter the username and password.
3. Click **Select Disk Group** and select the following disk groups:

- The disk group with disks created by the VSS provider.

- The disk group with the source disks.

After you complete this procedure, the VSS and VDS providers are ready to be used for zero downtime backup and instant recovery.

Configuring HP P9000 XP Disk Array Family

This section describes how to configure the Data Protector HP P9000 XP Disk Array Family integration for use with the Data Protector VSS integration.

Prerequisites

- Obtain or install *HP P9000 XP Disk Array Family licenses and components*:

  ◦ RAID Manager Library on the application and backup systems. For installation instructions, see the RAID Manager Library documentation.

    RAID Manager Library is firmware-dependent. Consult the HP sales representative for information on which version of RAID Manager Library to use.

  ◦ HP Business Copy (BC) P9000 XP microcode and a license for at least basic and HP Business Copy operation.

  ◦ P9000 XP Array VSS and VDS provider.

  ◦ HP MPIO Full Featured DSM (Device Specific Module) for HP P9000 XP Disk Array Family

For information on supported product versions, see the latest support matrices at http:// support.openview.hp.com/selfsolve/manuals.

- Obtain or install **Data Protector licenses and components**:

    ○ An instant recovery license.

    ○ The HP P9000 XP Agent.

      If this agent is not installed, it is not possible to restore data from the P9000 XP Array. For installation and licensing information, see the *HP Data Protector Installation and Licensing Guide*.

- Make sure the same operating system and version is installed on both application and backup systems.

- Connect the application and backup systems to the same P9000 XP Array.

- Assign LUNs to the respective ports.

- Pre-configure the LDEVs for the VSS snapshot creation and put them in the S-VOL host group. If you perform restore using VDS, you need to put new LDEVs in the S-VOL host group after the restore, since VDS restore switches the disks and the replica becomes the source volume. Add as many new LDEVs as were used for restore.

Configuring the VSS hardware provider

Before running ZDB sessions, open the HP XP VSS Hardware Provider Configuration Utility:

- Choose the configuration mode for your backups: VSS Compliant Mode or Resync Mode.

  NOTE: Changing the modes between backup sessions may have an impact on restore. For details, see "HP P9000 XP Array considerations" (page 54).

- Enable the use of snapshots (if supported by the P9000 XP Array and the VSS provider), select the Snap Type snapshot.

  Data Protector allows both, backups in resync mode and with snapshots for the same source volumes.

For details, see the HP P9000 XP Disk Array Family VSS hardware provider documentation.

Configuring the user authentication data

Before running zero downtime backup (ZDB) and instant recovery (IR) sessions on a disk array which is operating in the user authentication mode, add appropriate user credentials to the ZDB database (XPDB). For more information on the user authentication mode and on how to configure disk array user authentication data, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

Configuring HP 3PAR StoreServ Storage

This section describes how to configure the Data Protector HP 3PAR StoreServ Storage integration for use with the Data Protector Microsoft Volume Shadow Copy Service integration. Before you start with the configuration, ensure the prerequisites discussed in the following section are fulfilled.

Prerequisites

- Obtain or install **HP storage licenses and components**:

    ◦ HP 3PAR VSS Provider Software.

    For installation instructions, see the HP 3PAR StoreServ Storage documentation. For information on supported product versions, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.

- Obtain or install **Data Protector licenses and components**:

    ◦ Appropriate zero downtime backup extension and instant recovery extension licenses-to-use (LTU).

    ◦ HP 3PAR VSS Agent installed on both the application system and the backup system

    For licensing and installation information, see the *HP Data Protector Installation and Licensing Guide*.

- Make sure the same operating system (and its version) is installed on the application and backup systems.

- Connect a storage system of the HP 3PAR StoreServ Storage family to the application and backup systems through the SAN. The backup system must be connected to the same SAN as the storage system of the HP 3PAR StoreServ Storage family.

- Source volumes must have *snapshot space* (*copy space*) in a storage system's Common Provisioning Group (CPG) associated with.

For supported backup and connectivity topologies, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

For general Data Protector and integration-specific limitations, see the *HP Data Protector Product Announcements, Software Notes, and References* and the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

Configuring the Data Protector HP 3PAR StoreServ Storage integration

For configuration instructions, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*, part *HP 3PAR StoreServ Storage*, chapter *Configuration*.

Configuring the HP 3PAR VSS hardware provider

The only step you need to perform is to configure a VSS hardware provider user account on both the application system and the backup system.

The provider user account must have the same privileges as the user account that you added while configuring the Data Protector HP 3PAR StoreServ Storage integration: the *Edit* privilege level on the application system and the source volumes. Such a user account ensures the VSS hardware provider proper access to the 3PAR StoreServ Storage. The configured user accounts on the application and backup systems can be different. For instructions, see the HP 3PAR StoreServ Storage documentation.

After you complete this process, the HP 3PAR StoreServ Storage VSS provider is ready to be used in Data Protector zero downtime backup and instant recovery sessions.

4 Backup

Backup types

VSS backup types

The following VSS backup types are available with the Data Protector VSS integration:

- Local backup
- Transportable backup

See "Backup types" (page 19) for more details on VSS backup types.

Application specific backup types

For supported application specific backup types, such as `Full`, `Differential`, and so on, see the appropriate writer specific sections in "Writer specifics" (page 61).

Backup flow

1. The coordinator (the VSS service) identifies all writers that support the VSS feature and passes the list of available writers and their characteristics (Writer Metadata Document) back to Data Protector.
2. Data Protector examines Writer Metadata, identifies the volumes that contain the data to be backed up, and prepares a list of volumes (shadow copy set) that must be put into a consistent state and passes this information back to the coordinator, which informs available writers.
3. The VSSBAR agent notifies the writers about the shadow copy creation. The VSS mechanism ensures that there are no writes on the volume while the shadow copy is being created.

   The VSSBAR agent then passes the shadow copy creation requests to VSS.
4. After a shadow copy is created, the VSS service returns the related information to Data Protector. If an instant recovery enabled ZDB is performed, the VSSBAR agent coordinates the VDS agent to gather information on the created replica that is needed for instant recovery.
5. Data Protector backs up the data from the shadow copy to media and then notifies the VSS service that the shadow copy can be released. The shadow copy provider destroys the shadow copy that has been already backed up.

   "Local VSS backup" (page 20) shows the relations between the actors of a local VSS backup.
6. *Zero downtime backup:*

   The backup scenarios, which are not mutually exclusive, depend on the:

   - backup options selected in the backup specification
   - the ZDB type selected in the backup specification or at the start of the backup
   - the VSS backup type (local or transportable).

   The replica can be:

   - Kept on a disk array for rotation and instant recovery purposes.
   - Kept on a disk array only for data mining.
   - Moved to a backup medium from the application system (local backup) or backup system (transportable backup).
   - Presented and mounted to the backup system specified in the backup specification. It can be mounted in read–only or read/write mode.

- If not kept on the array, destroyed after the backup session completes (after data in the replica is moved to a backup medium).
- Manipulated using the Data Protector `omnidbvss` command.

7. The related information about the completed backup session is recorded in the VSSDB.

Zero downtime backup with the Data Protector VSS integration

The following ZDB types are available with the Data Protector VSS integration:

- ZDB to tape
- ZDB to disk
- ZDB to disk+tape

For more details on ZDB types, see the *HP Data Protector Zero Downtime Backup Concepts Guide*.

HP P4000 SAN Solutions

There is one replica type available when backing up using the HP P4000 SAN Solutions hardware providers:

- Snapshot

  The snapshot is dependent on previous snapshots and the original volume—when a snapshot is created, instead of writing to an original volume, new writes are redirected to a new location ("redirect on write"). When a resync (instant recovery) of a snapshot is performed, all dependent snapshots are automatically deleted.

HP P6000 EVA Disk Array Family

There are three replica types available when backing up using the P6000 EVA Array hardware providers:

- Snapshot *with* pre-allocation of disk space (*standard snapshot*). Snapshots can be created from original volumes or mirrorclones.
- Snapshot *without* pre-allocation of disk space (*vsnap* or virtually capacity-free snapshot). Vsnaps can be created from original volumes or mirrorclones.
- A full copy of the source volume (original virtual disk), independent of the original virtual disk (*snapclone*).

In the Data Protector GUI, you can select between snapshots or vsnaps (**Differential (Snapshot)**) and snapclones (**Plex (Clone/Mirror)**). To create vsnaps instead of standard snapshots, you must first configure the hardware provider to create vsnaps.

If you select **Keep the replica after the backup** or enable instant recovery (by selecting the **Track the replica for instant recovery**), you can further select the subtypes:

- Differential (Snapshot):
  - **Snapshot**
  - **MirrorClone snapshot**
- Plex (Clone/Mirror):
  - **Snapclone**

HP P9000 XP Disk Array Family

There are two replica types available:

- *Mirrors* (**Plex (Clone/Mirror)**)
- *Snapshots* (**Differential (Snapshots)**)

  Snapshots must be supported by the P9000 XP Array and the P9000 XP Array VSS hardware provider.

Additionally, there are two configuration modes available for the backup with the P9000 XP Array hardware providers:

- VSS compliant mode

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

  Note that if you select Differential (Snapshots) as the replica type, the VSS compliant mode is ignored and the resync mode is always used.

- Resync mode

  When the P9000 XP Array provider is in the resync mode, the source volume (P-VOL) and its replica (S-VOL) are in a suspended mirror relationship after backup. The maximum number of replicas (S-VOLs per P-VOL) rotated is three, provided that the 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.

HP 3PAR StoreServ Storage

There is one replica type available when backing up using the HP 3PAR StoreServ Storage hardware providers:

- Snapshot

  The snapshot is independent of previous snapshots—after it is created, new writes are not redirected to a new location but continue to be made on the original volume after its current contents are copied over ("copy on write").

Considerations

Microsoft Cluster environments and transportable backups

When using a VSS hardware provider in a Microsoft Cluster environment, a transportable shadow copy (replica) must be transported outside of the cluster if the original volume is mounted within the cluster.

As a result, Data Protector supports only the following configurations in a cluster:

- Local or network backup using a VSS software provider.
- Transportable backup, where the backup system is not part of the cluster (not a cluster node).

HP P4000 SAN Solutions

When a new snapshot is created with HP P4000 SAN Solutions, new disk space is allocated on the array where the new data is written (this method of creating a snapshot is known also as *redirect on write*). This means that newer snapshots depend on older ones and that you need *all snapshots in a chain* for the last snapshot to be valid.

If several backup specifications contain the same source volume, instant recovery from a ZDB session of one backup specification may cause that the newer ZDB sessions based on other backup specifications are not available for instant recovery.

For example, the backup specification BSpec1 contains volumes C: and G: and the backup specification BSpec2 volumes G: and M:.

An instant recovery from a ZDB session of BSpec1 will remove all subsequent snapshots of the volumes C: and G:. This will also prevent further instant recovery sessions from all ZDB sessions of BSpec2 that were created after the ZDB session for BSpec1 that is being used for instant recovery, since the snapshots for the volume G: are no longer available.

As a result, you must consider the following:

- It is recommended to limit the number of volumes that are included in the backup specification to a minimum needed.
- Perform instant recovery for sessions and not for particular volumes.
- Any newer replica (snapshot or smartclone) created outside of Data Protector, will prevent instant recovery from a session that contains the volume which was replicated outside of Data Protector.

HP P6000 EVA Disk Array Family

- When the cloning process for a source volume is in progress, another replica of the same source volume cannot be created.
- Only one type of target volumes per source volume can exist on a P6000 EVA Array at the same time. For example, creation of a snapclone may fail if a snapshot or vsnap for the same source volume exists on the array. You should first delete such replicas using the Data Protector command omnidbvss, or using Command View EVA if they were not created by Data Protector.

  Note that in case of a MirrorClone Snapshot, the omnidbvss command can delete only the snapshots. The MirrorClone used as the source for the snapshot cannot be deleted using omnidbvss and must be deleted manually using EVA Command View, by first detaching the MirrorClone and then deleting the VDisk.

- On Windows Server 2003 systems, the VDS hardware provider is required for instant recovery.
- On Windows Server 2003 systems, when the VDS hardware provider is not used, the backup system is left enabled (in read/write mode) even if the option **Enable backup system in read/write mode** was not selected. As a result, you must ensure that you do not delete or break these shadow copies when working with other (third-party) tools.

Replica creation and reuse

A new replica is created and added to the replica set when the specified **Number of replicas rotated** is not reached.

The oldest replica in the set is deleted and the new one is created when the specified **Number of replicas rotated** is reached.

HP P9000 XP Disk Array Family

- For backups (*mirrors*) created in the *resync mode*, the value of the **Number of replicas rotated** option is limited to a maximum of three replicas (S-VOLs) per source volume (P-VOL), provided that the `MU# range` option in the P9000 XP Array VSS hardware provider configuration is set to `0-2` or `0, 1, 2`. When the same P-VOL is used in more than one backup specification, the sum of values specified in the **Number of replicas rotated** option in all backup specifications should not exceed 3.

  This limitation also applies for the following backup scenario in a Microsoft Exchange Server CCR environment with the same backup system selected: Suppose you create three backups of a source volume on the production database system with one backup specification, and then create three backups of the source volume on the database copy system with a different specification. Together, this would create six replicas. However, Data Protector limits the number and supports only three replicas together in such a case. Therefore, you can choose to back up source volumes only on the production database system or only on the database copy system. Otherwise, after three replicas are created with one backup specification, the next backup with the second backup specification will fail.

  Snapshots are not limited to 3 replicas per source volume. The `MU#` number for snapshots can be configured in the VSS provider configuration.

- You can change the P9000 XP Array VSS hardware provider mode between different backup sessions using the same backup specification, but this is not recommended when you use replica set rotation, because restore of such backup data may fail.

  For details, see "HP P9000 XP Array considerations" (page 54).

- On Windows Server 2003 systems, when the VDS hardware provider is not used, the backup system is left enabled (in read/write mode) even if the option **Enable backup system in read/write mode** was not selected. As a result, you must ensure that you do not delete or break these shadow copies when working with other (third-party) tools.

HP 3PAR StoreServ Storage

- After a new snapshot is created with HP 3PAR StoreServ Storage, the data on the original volume which is about to be overwritten by new data is first copied to the snapshot space associated with the volume (this method of creating a snapshot is known also as *copy on write*). This means that newer snapshots are independent of older ones.

- HP recommends not to manually set retention periods on target volumes created by Data Protector. Volume retention period, when set, prevents a volume from being deleted from the 3PAR StoreServ Storage when Data Protector rotates it out of the replica set. In this case, after replica set rotation, the volume becomes an orphaned volume that unnecessarily occupies storage space. Due to a problem in the HP 3PAR VSS hardware provider version 2.1.0.11, failed deletions of such volumes are reported as successful.

Configuration check

With instant recovery, it is possible to *selectively* restore separate writer's components if they reside on separate source volumes. Instant recovery of separate components also requires that the volumes with components data should not contain any other data. The configuration check detects, whether there is more than one component on the volume and whether there is any other data besides the component's data.

At backup time, Data Protector can check whether the individual components can be selectively restored using the instant recovery functionality. At instant recovery time, Data Protector can check whether the data required for the components restore is available.

In case of the Microsoft Exchange Server Writer, it is checked, whether the whole storage group should be restored, or whether the separate database stores can be restored individually. In case of the MSDE Writer, it is checked, whether the user, system, and log files are on separate volumes.

If the check fails for a component during the backup, you will not be able to select this component for the instant recovery, you will need to recover all writer components. If the check fails during the instant recovery, the instant recovery session will fail.

Configuration check modes

You can select between three configuration check modes:

- **Strict**

  If any file or folder on the volume does not belong to the component, the ZDB to disk, ZDB to disk+tape, or instant recovery fails.

- **Non-strict**

  If any folder on the volume does not belong to the component, the ZDB to disk, ZDB to disk+tape, or instant recovery fails.

- **Disabled**

  The check detects whether there is more than one component on the volume or there is any other data besides the component's data on the volume, but the session does not fail in any case.

> ① **IMPORTANT:** If the configuration check is disabled for an instant recovery, you will lose the data that does not belong to a component, but resides on the same volume as the component. Disable configuration check only if instant recovery cannot be performed with an enabled configuration check and only after you make sure that this will not result in a loss of data. In case of a data loss, the data that does not belong to a component, but resides on the same volume, will be lost.

The configuration check should not be disabled except for the cases when instant recovery cannot be performed with an enabled configuration check. Due to the specific behavior, some writers (not applicable to the MSDE Writer and Microsoft Exchange Server writers) may create temporary files on components volumes during backup and instant recovery causing the check failure. In such cases, instant recovery is not possible without disabling the check.

Creating backup specifications

The procedure below shows how to back up Microsoft VSS objects using the Data Protector GUI. Some writers have specific limitations. For writers specific prerequisites, limitations, and additional steps see the appropriate sections in "Writer specifics" (page 61).

To create a new backup specification for the VSS integration, proceed as follows:

1. In the **HP Data Protector Manager**, switch to the **Backup** context.
2. In the Scoping Pane, expand **Backup Specifications**.
3. Right-click **MS Volume Shadow Copy Writers** and click **Add Backup**.

4. In the **Create New Backup** dialog box, select the backup type. You can choose between the following types:

    *   **Local or network backup**

        This type is used for single-host VSS backup. To perform a zero downtime backup (ZDB), you need a hardware provider. Otherwise, no hardware provider is required for this type of backup.

    *   **VSS transportable backup**

        Use this option to create shadow copies on the application system and present them to the backup system, which can perform the backup to tape.

        A hardware provider is required for this type of backup.

    Figure 6 Selecting VSS transportable backup
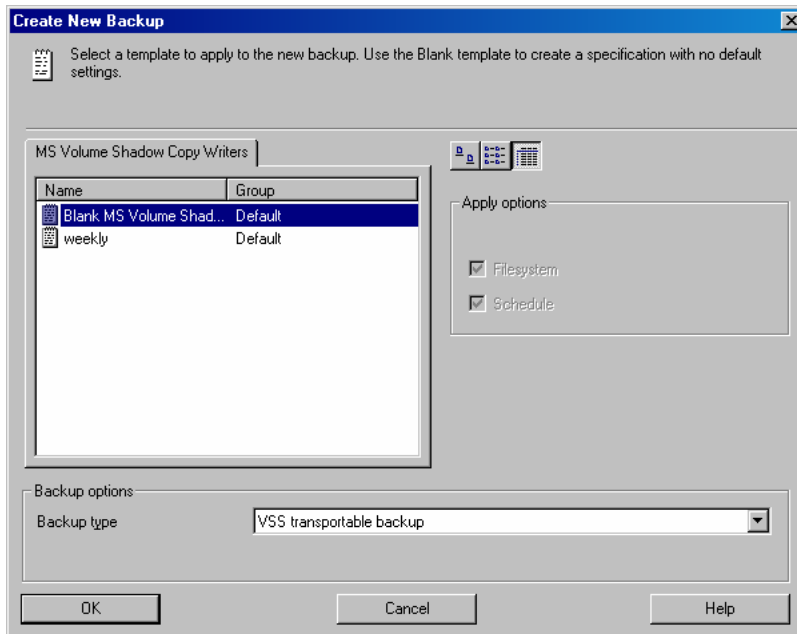
5. In Application system, specify the name of the client that has the VSSBAR agent installed.

    When backing up cluster-aware writers (such as SQL Server via the MSDE Writer, or Exchange Server in LCR or CCR environment), specify the virtual server name given in the particular writer resource group.

    Zero downtime backup:

    Specify the following options:

    *   *For local or network backup*, select **Use hardware provider** to enable ZDB and specify other backup options.

    *   *For a transportable backup*, the name of the backup system from where the shadow copy can be backed up to tape or where you want your shadow copies to be presented and mounted after the backup. The hardware provider is used automatically.

    *   To keep the replica after backup, select **Keep the replica after backup**.

    *   To enable instant recovery, select **Track the replica for instant recovery**. **Keep the replica after backup** is selected automatically.

    *   If you select **Keep the replica after the backup**, you can specify the replica type:

Plex (Clone/Mirror)

A shadow copy, independent from its source volume is created:

- ○ This replica type is not supported by the HP P4000 SAN Solutions and HP 3PAR StoreServ Storage hardware providers.

- ○ With the HP P6000 EVA Disk Array Family hardware provider, a snapclone is created.

- ○ With the HP P9000 XP Disk Array Family hardware provider, the replica created depends on the provider mode (clone in the case of VSS compliant mode, or mirror in the case of resync mode). The selected subtype is ignored.

Differential (Snapshot)

A shadow copy, dependent on its source volume is created.

- ○ With the HP P4000 SAN Solutions or HP 3PAR StoreServ Storage hardware provider, a snapshot is created.

- ○ With the HP P6000 EVA Disk Array Family hardware provider, you can select between a snapshot (depending on the provider configuration, this is a standard snapshot or vsnap) or a MirrorClone snapshot. See "P6000 EVA specific options" (page 42).

- ○ With the HP P9000 XP Disk Array Family hardware provider, a snapshot is created.

Default

The replica type is determined by the VSS hardware provider configuration.

Not available for instant recovery.

A provider may support one or both types. If you select an unsupported replica type, the backup will fail.

For more information on supported replica types, see "Zero downtime backup with the Data Protector VSS integration" (page 33).

- To mount the replicas on the backup system, select **Mount the replica on the backup system**.

Figure 7 VSS transportable backup options
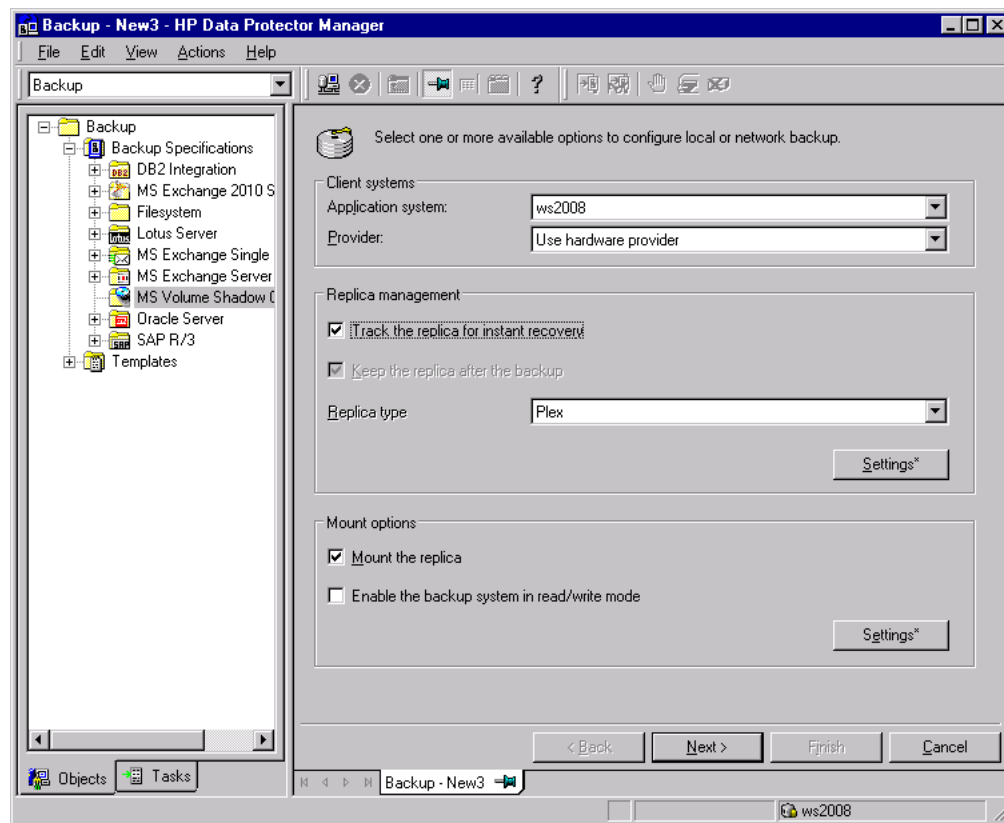
- To change the replica management and mount options, click on **Settings** to open the Replica management or the Mount dialog box.

  NOTE:

  Whenever you change an option on the page that affects these additional settings, the buttons are marked with an asterisk (*****). After you open the settings window and review or modify the options, the asterisk is removed.

  For descriptions of replica options, see Table 4 (page 42) and Table 5 (page 42) or press **F1**.

  For descriptions of mount options, see Table 6 (page 42) or press **F1**.

  Click **Next**.

6. In this page, the selection of your VSS client is displayed.

   On Windows Server 2008, you can specify the **User and group/domain** options. For information on these options, press **F1**.

   Click **Next**.

7. Select the backup objects you want to back up. Make sure that in case of instant-recovery enabled sessions, you select all objects (for example all storage groups, all virtual machines …) residing on a specific source volume that will be backed up.

   You can specify a *full client backup* by selecting the top-level item (the name of the client), a single writer backup, or a writer's component backup by selecting a lower-level item.

   If a full client backup is selected, Data Protector checks which writers exist on the client and backs up all of them at backup time.

   If a writer requires all of its components to be backed up, lower-level items are automatically selected. If you select such a writer for backup, all its components will be backed up. If a

writer has no components to be backed up, it is not displayed in the list of writers, and is not backed up when the full client is selected.

> **NOTE:** On a disk array of the HP P6000 EVA Disk Array Family, to perform an instant recovery, do not select too many objects because they might reside on more than four source volumes. If the objects reside on more than four source volumes, create several backup specifications where objects do not reside on more than four source volumes.

The **Filesystem** item displays all mounted disks. If another disk is mounted to a directory on a disk, the parent disk name is displayed twice. The first name represents the parent disk name (for example `c:`), while the second name represents the container for the mountpoint (for example `c:\mnt\1`). To select the mounted disk, select the container for the mountpoint.

8. For backup to tape, select the devices you want to use for the backup to tape. Click **Properties** to set the device concurrency, media pool, and preallocation policy. For more information on these options, click **Help**. If you do not select a device, only backup to disk will be available.

> ⊙ **IMPORTANT:** If you do not configure devices and do not select the option **Track the replica for instant recovery**, you will not be able to restore the data with Data Protector.

In case of backup to tape, you can also specify whether you want to create additional copies (mirrors) of the backup during the backup session. Specify the desired number of mirrors by clicking the **Add mirror** and **Remove mirror** buttons. Select separate devices for the backup and for each mirror.

For detailed information on the object mirror functionality, see the *HP Data Protector Help*.

9. Following the wizard, select the backup options and schedule your backup.

> 💡 **TIP:** If you are not sure about selecting the backup options, keep the default values.

For details about the options common to all Data Protector backup specifications, see the *HP Data Protector Help*.

10. Once you have defined all backup options and the schedule, you need to name and save the newly-created backup specification. You have now completed the creation of a Microsoft Volume Shadow Copy Writers backup specification.

11. You can review the newly-created and saved backup specification in the **Backup** context, under the specified group of backup specifications.

12. You can run backup using one of the following methods:

    - Schedule the backup of an existing Microsoft Volume Shadow Copy Writers backup specification using the Data Protector Scheduler.

    - Start an interactive backup of an existing Microsoft Volume Shadow Copy Writers backup specification.

Backup options

Table 3 VSS-specific backup options

Option	Description
Pre-exec	Specify a command that will be started by `vssbar.exe` on the application system directly before replica creation. Do not use double quotes. Type only the name of the command, not the pathname. The command must reside in *Data_Protector_home*\bin.
Post-exec	Specify a command that will be started by `vssbar.exe` on the application system directly after the replica creation. Do not use double quotes. Type only the name of the command, not the pathname. The command must reside in *Data_Protector_home*\bin.

ZDB options

Table 4 Replica management options

Option	Description
Configuration check mode	If you have selected **Track the replica for instant recovery**, specify the **Configuration check mode**. Configuration check applies to the disk backups that are to be used for instant recovery and does not apply to the tape backup. For information on the options, see "Configuration check" (page 36).
Number of replicas rotated	During ZDB sessions, Data Protector creates a new replica and leaves it on the array until the specified **Number of replicas rotated** is reached (specify if you selected **Keep the replica after the backup**). After that, the oldest replica is deleted and a new one created. Default: 3 The maximum number for vsnaps and standard snapshots is limited by the P6000 EVA storage system. Data Protector does not limit the number of replicas rotated, but the session fails if the limit is exceeded.

Table 5 P6000 EVA specific options

Option	Description

Replica sub type	Available if **Keep the replica after the backup** is selected.
	
	Sets the replica subtype. Depending on the **Replica type** you selected, the following subtypes can be selected:
	
	• Plex (Clone/Mirror):
	
	  ○ **SnapClone**
	
	• Differential (Snapshot):
	
	  ○ **Snapshot**
	
	  ○ **MirrorClone Snapshot**
	
	Disabled if **Replica type** is set to **Default**.

Wait for data copy to complete	Available if Replica subtype is set to **SnapClone**. During VSS backup session, a target storage allocation for the snapclone replica is created first. At that point, the replica is only virtual but immediately available for use. However, in the background, copy process runs for some time to copy all data from the source to the target storage. To avoid the slowdown of the copy process caused by further actions, for example Exchange integrity check process, use this option. The backup resumes after the time specified in **Wait up to n minutes** (default: 60) is reached even if the copy process is not finished.
Attempt to minimize wait	This option resumes the backup process after the copy process has finished if it is finished before the specified time. Selected by default if Wait for data copy to complete is selected.

NOTE: During the creation of the backup specification, the disk array type is not determined yet by Data Protector and the P6000 EVA options are available regardless of the disk array used. If a different disk array is used, the options are ignored when the backup is performed.

Table 6 Mount options

Option	Description
Mount the replica	Select this option to mount the replica on backup system. The mount path depends on the **Root of the mount point on the backup system** and **Add directories to the mount path** options.
Enable backup system in read/write mode	Select this option to enable read/write access to the mounted replica's disks on the backup system. Note that instant recovery of disks that have been activated in read/write mode is not recommended, since setting the write permission allows the disks' data to be changed after

Table 6 Mount options *(continued)*

Option	Description
	the backup time. Default: not selected if **Track the replica for instant recovery** is on, or selected if only **Keep the replica after the backup** is on.
Root of the mount point on the backup system	Specifies the root directory under which the filesystems from the replica will be mounted. Where exactly the filesystems are mounted depends on how you specify **Add directories to the mount path**. The default mount point is `c:\mnt`.
Add directories to the mount path	This option allows discrete control over the created mount points. When Session ID is used in path creation, this guarantees unique mount points. The options define which subdirectories will be created in the directory specified with the **Root of the mount path on the backup** system option. Example: Root directory: `c:\mnt` Application system: `app.comp.com` Backup session ID: `2008-02-22-4` Mount path on the application system: `E:\disk1` If **Hostname** is selected: `c:\mnt\app.comp.com\E\disk1` If **Hostname + session ID** is selected (default): `c:\mnt\app.comp.com\2008-02-22-4\E\disk1` If **Session ID** is selected: `c:\mnt\2008-02-22-4\E\disk1` If **Session ID + hostname** is selected: `c:\mnt\2008-02-22-4\app.comp.com\E\disk1`
Automatically dismount the filesystems at destination mount points	If the mount points are in use (for example, a previous session may be mounted) and this option is selected, Data Protector will try to dismount the mounted filesystems. If the option is not selected and the mount points are in use, or if the option is selected and the dismounting fails, the session will fail. Default: not selected.

Scheduling backup sessions

For more detailed information on scheduling, see the *HP Data Protector Help* index: "scheduled backups".

To schedule a Microsoft Volume Shadow Copy Writers backup specification, perform the following steps in the Data Protector GUI:

1. In the **HP Data Protector Manager**, switch to the **Backup** context.
2. In the Scoping Pane, expand **Backup**, then **Backup Specifications**. Click **MS Volume Shadow Copy Writers**.

    A list of available backup specifications is displayed in the Results Area.
3. Double-click the backup specification you want to schedule and click the **Schedule** tab to open the **Schedule** property page.
4. In the **Schedule** property page, select a date in the calendar and click **Add** to open the **Schedule Backup** dialog box.
5. Specify **Recurring**, **Time options**, **Recurring options**, and **Session options**.

    Note that the backup type for ZDB sessions that are to be used for instant recovery is set to **Full**.

    In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the **Split mirror/snapshot backup** option.

Figure 8 Scheduling a backup

Figure 9 Scheduling a ZDB backup

6. Click **OK** to return to the **Schedule** property page.
7. Click **Apply** to save the changes.

Starting backup sessions

An interactive backup can be started using the Data Protector GUI by following these steps:

1. In the **HP Data Protector Manager**, switch to the **Backup** context.

2. In the Scoping Pane, expand **Backup**; then expand the **Backup Specifications** and the **MS Volume Shadow Copy Writers** items.

3. Right-click the backup specification you want to use, and then select **Start Backup** from the pop-up menu.

   The **Start Backup** dialog box appears.

   Select the backup type and the network load (**High**, **Medium**, or **Low**).

   Note that the backup type for ZDB sessions that are to be used for instant recovery is set to **Full**.

   In the case of a ZDB-to-disk or a ZDB-to-disk+tape session, specify the **Split mirror/snapshot backup** option.

   For a description of network load, see the *HP Data Protector Help*.

4. Click **OK**. Upon successful completion of the backup session, a `Session Completed Successfully` message appears.

5 Restore

You can restore the Data Protector Microsoft Volume Shadow Copy Service integration backup objects using the Data Protector GUI.

Data Protector offers two methods for restoring the writers:

- From backup media to the application system on LAN (standard restore). See "Standard restore" (page 46).

- Using the instant recovery functionality. See "Instant recovery" (page 49).

Standard restore

Limitation for custom restore methods

- Data Protector Microsoft VSS integration does not automatically provide any restore method for writers requesting custom restore. If a writer specifies custom restore method, it is only possible to restore the writer's data as plain files using the Data Protector file restore functionality. You can use the `Restore Into` option to specify an alternate restore path for these plain files. You can then perform the custom restore from these plain files manually. For information on writer's custom restore, see the writer's documentation.

Restore modes

Data Protector offers two restore modes:

- **component restore** using the VSS service

- **file restore** using the Data Mover Agent (DMA) instead of VSS.

By default, Data Protector restores writer components using the VSS service.

Component restore

When restoring writer components, Data Protector works with the VSS service to ensure data consistency:

1. Data Protector first restores the metadata, which was collected during the backup. Then it uses the metadata to identify the backup components and determine the restore method.
2. Data Protector restores the data from the backup media to the locations specified in the backup metadata, following the writers' instructions regarding any additional checking or processing.
3. After a successful restore from the backup media, Data Protector notifies the coordinator and the writers can now access the newly-restored data and start the internal processing, for example database recovery.

File restore

For a successful restore of a writer component, all files comprising this component must be restored. If a restore of a single file fails, the restore of the a whole component fails. Data Protector offers an additional restore mode for restoring single files that does not use the Microsoft Volume Shadow Copy Service, thus solving this problem. This mode can also be used for restoring to systems that do not support VSS or do not have a VSS writer installed.

When restoring files or a group of files, DMA is started and the files are restored using the standard Data Protector filesystem restore procedure.

ⓘ **IMPORTANT:** As the file restore mode does not utilize VSS services, additional tasks that are performed after a component restore – such as database recovery – are not performed and your application data may be left in an inconsistent state, requiring additional manual procedures before the application is recovered.

Restoring using the GUI

The procedure below shows how to restore Microsoft VSS components using the Data Protector GUI. Some writers require custom restore procedures and/or have specific limitations. See also the appropriate sections in "Writer specifics" (page 61).

To restore Microsoft VSS objects using the Data Protector GUI, proceed as follows:

1. In the Data Protector GUI, switch to the **Restore** context.
2. Expand **Restore Objects**, expand **MS Volume Shadow Copy Writers**, expand the client from which you want to restore the data, and then click **MS Volume Shadow Copy Writers**. In the Results Area, a list of writers, which were backed up on this client, is displayed.
3. Select the Restore mode:

    - **Restore components**

      With this option selected, whole components are restored using the Volume Shadow Copy Service. Individual files cannot be selected for restore.

    - **Restore files to temporary location**

      With this option, you can select individual files or a group of files that were backed up using the selected writer. The files are restored using the Data Mover Agent and not the Volume Shadow Copy Service.

4. In the Results Area, select the writers or writers' components (for component restore) or files or a group of files (for file restore mode).

Figure 10 Restore objects
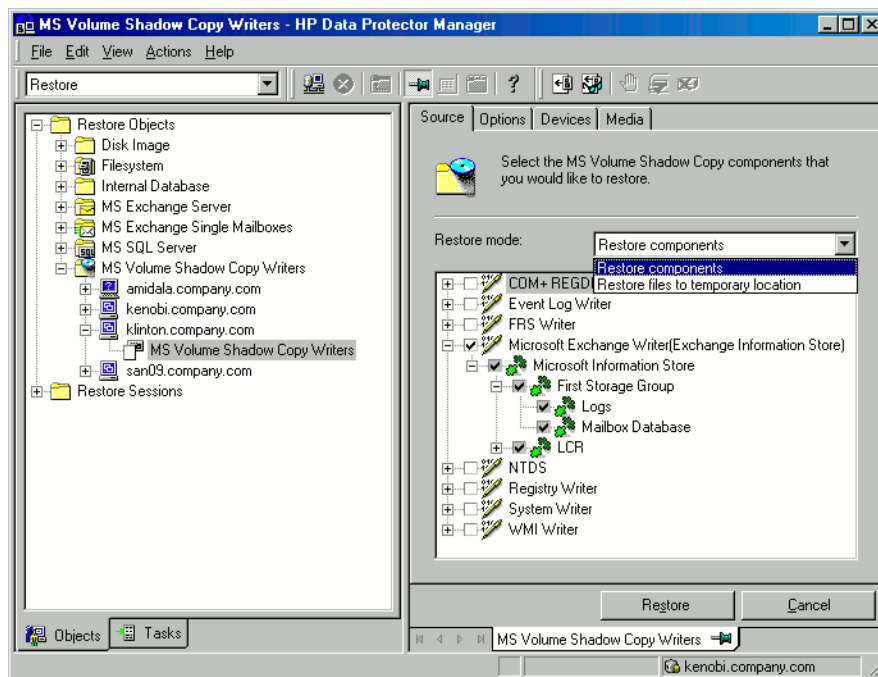
You can select the top-level item (full writer restore) or only specific components. If you select a full writer restore, but some components of this writer were not backed up in the same session, the unavailable components are shaded and cannot be selected. To select the version (the date of a backup), right-click the object name and click **Properties**. The last available backup version is selected by default, however, you can select a different version from the drop-down list.

For application-specific options, see "Writer specifics" (page 61).

5. In the **Options** property page, select the MS Volume Shadow Copy specific restore options. See "Restore options" (page 48).

6. In the **Devices** and **Media** property pages, the devices and media for restore are automatically selected.

    Note that you can change the device used for the restore. Therefore, you have the possibility of using a different device for a restore than the one that was used for the backup. See the *HP Data Protector Help* index: "selecting devices for restore".

7. Click **Restore**. Review your selection, and then click **Finish** to start a restore session.

    The restore session messages are displayed in the Results Area.

8. If you are restoring a VSS writer that requires a custom restore, continue manually, using the writer-specific methods, if it is provided by a writer. See the writer's documentation.

Restore options

The following restore options are specific to the Data Protector Microsoft Volume Shadow Copy Service integration.

NOTE: Do not use these options for Microsoft Exchange Server 2007 writer since there are other Exchange Server 2007 specific options provided for restore to another location. See the Microsoft Exchange Server 2007 writer specifics.

Restore to another client

By default, the components or files are restored to the client from which the application data was backed up. However, you may restore the data to another VSS client if you specify the **Restore to another client** option. The new target Microsoft VSS client must be a part of the Data Protector cell. For component restore, it must also run on the same platform and have the MS Volume Shadow Copy Integration software component installed. For file restore, the MS Volume Shadow Copy Integration software component is not required.

Restore into the following directory

By default, you restore the data to the same directory from which it was backed up (it can be on the original client or on some other client which you selected).

However, if you specify the **Restore into the following directory** option, your data will be restored to another directory. When defining the restore location, you can specify the path to the directory where you want to restore your data.

Restoring using the CLI

Use the `omnidbvss -get session` *SessionKey* command to get details about the session that you want to use for restore.

For example:

```
omnidbvss -get session_persistent 2011/01/26-2:computer1
===========================================================
Session ID:              2011/01/26-2
Barlist Name:            SQL_NEW_DB_SIMPLE
Bar Hostname:            computer1.company.com
Backup Type:             FULL
Instant Restore:         FALSE
Disk-Only:               FALSE

        Component Name
===========================================================
[0]     /SqlServerWriter(SQL Server 2008:SQLWriter)/COMPUTER1/New_DB
```

The objects that were backed up in this session are listed under Component Name. See the `omnidbvss` reference page in the *HP Data Protector Command Line Interface Reference* for details.

Execute the `omnir` command:

```
omnir -vss —barhost ClientName -session BackupID -tree TreeName1 [-tree
TreeName2...] -conf_check {strict|non-strict|disabled}
```

See the `omnir` reference page in the *HP Data Protector Command Line Interface Reference* for details.

> **NOTE:** *BackupID* is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup session no longer exists, but the media set created in an object copy session still exists. To restore the objects, you must specify the session ID of the original *backup* session (that is, the backup ID) and not the session ID of the *object copy* session.

The `omnir` syntax does not let you specify from which object copy to restore if several copies of the same object exist. This is only possible using the Data Protector GUI by setting the media allocation priority list.

To recover the components from the above example, execute:

```
omnir -vss —barhost server1.company.com -session 2010/1/12-09 -tree
/Microsoft Exchange Writer(Exchange Information Store)/Microsoft
Information Store/Store1/Logs —tree /Microsoft Exchange Writer(Exchange
Information Store)/Microsoft Information Store/Store1/File -conf_check
disabled
```

Instant recovery

After ZDB to disk+tape or ZDB to disk, restore can be performed using instant recovery. After ZDB to disk, instant recovery is the only possible restore method. Backup session information is saved in the IDB, and the array-specific information required for instant recovery is saved in the VSS database (VSSDB).

Instant recovery consists of two phases:

- An instant recovery session is started and the relevant data files are restored back to the system.
- A writer performs a recovery of the application data, which is based on the restored data. For some writers, database recovery is performed automatically by the database engine. See the writer specifics for details.

> **IMPORTANT:** During instant recovery, the whole replica is restored. This means not only are the originally selected backup objects restored, but the complete content of all the volume groups that contained them. After the instant recovery, the content is returned to the state when the replica was created.

For detailed information on instant recovery concepts, see the *HP Data Protector Zero Downtime Backup Concepts Guide* and *HP Data Protector Zero Downtime Backup Administrator's Guide*.

You can *selectively* restore separate writer's components only if they reside on separate target volumes. Instant recovery of separate components also requires that the target volumes with components data should not contain any other data. If you do not select all of the writer's components residing on the target volumes to be restored, the restore fails.

Before starting an instant recovery, Data Protector checks:

- Whether the data that is required for the components restore is available. If the check fails, the instant recovery session also fails, ensuring that no data that does not belong to the components is lost.
- If any changes that are not controlled by Data Protector happened in a disk array environment (for example, the target volumes to be restored are presented to some other system as it is

recorded in the VSSDB, or the source volumes to be replaced are not presented on the application system). If Data Protector detects such changes (by comparing the state of the environment with data in the VSSDB), the instant recovery session aborts. However, you can force the session to continue by setting the `OB2VSS_IGNORE_BACKUP_DISK_CHANGES` and `OB2VSS_IGNORE_SOURCE_DISK_CHANGES omnirc` options to `1`.

You can perform an instant recovery using either the Volume Shadow Copy or Virtual Disk services or the Data Protector disk array integration agents:

- **Using Microsoft Virtual Disk Service (VDS)**

  The instant recovery is initiated by Data Protector and performed by the VDS provider. Information about the replica is deleted from the VSSDB, therefore it is not possible to perform another instant recovery using that replica.

- **Using Microsoft Volume Shadow Copy Service LUN resync**

  The instant recovery is performed by the VSS hardware provider. The actual instant recovery method depends on the disk array and VSS hardware provider settings. The VSS LUN resync functionality must be supported by the operating system (Microsoft Windows Server 2008 R2).

  For a list of hardware providers that support this type of recovery, see the latest support matrices on the web http://support.openview.hp.com/selfsolve/manuals.

- **Using a Data Protector disk array integration agent** (HP P4000 Agent, HP P6000 / HP 3PAR SMI-S Agent, HP P9000 XP Agent, or HP 3PAR VSS Agent)

  The instant recovery is performed by the appropriate Data Protector disk array agent.

The available instant recovery methods depend on the disk array, the way the instant recovery is performed (either by using one of the hardware providers or Data Protector disk array agents), and on the type of the backup.

Instant recovery methods

Disk array technologies support several methods of instant recovery:

- Switch of disks
- Copy of replica data
- Restore of snapshot data
- Resync of volumes

Some of the methods also enable you to keep the source volumes while with others, the source volumes are not retained. For details, see the description of each method. For a complete overview of all available instant recovery methods for a specific disk array, see "Supported instant recovery methods for HP P4000 SAN Solutions" (page 53), "Supported instant recovery methods with HP P6000 EVA Disk Array Family" (page 54), "Supported instant recovery methods for HP P9000 XP Disk Array Family" (page 54), or "Supported instant recovery methods for HP 3PAR StoreServ Storage" (page 56).

Switch of disks

With this method, the source volumes are unpresented and the target volumes (replica) are presented in the place of the source volume. You can select to retain the old source volumes. However, you cannot retain the replica and perform another instant recovery from this replica.

If the restore session fails or is aborted, the VSS integration reverts the application and backup environment back to the state in which they were before the instant recovery.

Advantages:

- Instant recovery is very fast.

- The original source volume can be retained after instant recovery.

- By selecting the **Restore using Microsoft Virtual Disk Service** option, you can restore your data without using HP P6000 / HP 3PAR SMI-S Agent or HP P9000 XP Agent.

Disadvantages:

- It is not possible to perform another instant recovery using the same replica.

- This method of instant recovery changes the physical location of the application data, since, after instant recovery, a replica becomes the source volume. The application starts running on the physical disks that were used for backup. If the selected replica and its source volume belong to separate disk groups, the two disk groups are also switched during instant recovery.

To use this method of instant recovery, in the Data Protector GUI, select the **Restore using Microsoft Virtual Disk Service** option.

Copy of replica data

Copy of replica data with the source volume retained

With this method, a new copy from a replica is created in the same disk group in which the source volumes reside. When the copy is finished, a switch of disks between the source and the new copy is performed. The old source volumes are retained.

If the restore session fails or is aborted, the VSS integration reverts the application and backup environment back to the state in which they were before the instant recovery.

The replica is automatically retained thus you can perform another restore from this backup.

To delete the replica, you can choose to manually remove it using the `omnidbvss` command or wait until it is deleted automatically due to replica rotation.

Advantages:

- The original source volumes are retained after restore.

- Disk group of the source volume is not changed after instant recovery.

- Another restore from the same backup is possible.

Disadvantages:

- Restore is not as fast as with the "switch of disks" method.

- Physical location of the application production data changes.

To perform this type of restore, select the **Copy replica data to the source location** option and leave selected the **Retain source for forensics** option in the Data Protector GUI.

Copy of replica data with the source volume not retained

With this method, the source volumes are directly overwritten by the replica by copying data from the target volumes back to the source volumes. The old source volumes are not retained and if the restore session fails after the copy process has already started, the original application data residing on the source volumes is lost. If you try to abort the session at this point, the abort operation is rejected and the session continues.

The replica is automatically retained thus you can perform another restore from this backup.

To delete the replica, you can choose to manually remove it using the `omnidbvss` command or wait until it is deleted automatically due replica rotation.

Advantages:

- Disk group of the source volume is not changed after instant recovery.
- Physical location of the application production data remains the same.
- Another restore from the same backup is possible.

Disadvantages:

- Restore is not as fast as with the "switch of disks" method.
- The original source volumes are lost during restore.

To perform this type of restore, select the **Copy replica data to the source location** option and clear the **Retain source for forensics** option in the Data Protector GUI.

Restore snapshot data to the source volume

Restore snapshot data to the source volume with the source volume retained

With this method, the data from the snapshots is copied back to the original volumes.

If the restore session fails or is aborted, the VSS integration reverts the application and backup environment back to the state in which they were before the instant recovery.

The replica is automatically retained thus you can perform another restore from this backup.

To delete the replica, you can choose to manually remove it using the `omnidbvss` command or wait until it is deleted automatically due replica rotation.

Advantages:

- The original source volumes are retained after restore.
- Disk group of the source volume is not changed after instant recovery.
- Another restore from the same backup is possible.
- Because a snapshot contains only differential data, the amount of copied data is usually lower compared to copying a replica that is a full copy of its source (for example a snapclone).

Disadvantages:

- Restore is not as fast as with the "switch of disks" method.

To perform this type of restore, select the **Copy replica data to the source location** option and leave selected the **Retain source for forensics** option in the Data Protector GUI.

Restore snapshot data to the source location with the source volume not retained

With this method, a snapshot of the source volume is made first and then the data from the snapshots is copied back to the original volumes.

Advantages:

- Disk group of the source volume is not changed after instant recovery.
- Physical location of the application production data remains the same.
- Another restore from the same backup is possible.
- Because a snapshot contains only differential data, the amount of copied data is usually lower compared to copying a replica that is a full copy of its source (for example a snapclone).

Disadvantages:

- Restore is not as fast as with the "switch of disks" method.
- The original source volumes are lost during restore.

To perform this type of restore, select the **Restore snapshot data to the source location** option and clear the **Retain source for forensics** option in the Data Protector GUI.

Resync of volumes

With this method, the target volumes are resynchronized with the original source volumes.

With this type of instant recovery, HP P9000 XP Agent (SSEA) synchronizes a source volume (P-VOL) with its target volume (S-VOL) and then splits the pair during the instant recovery session.

Depending on the selected resync type, SSEA waits for re-synchronization or copy process to complete (normal restore, selected with **Resync replica data to the source volume**), before making the source volume available. Alternatively, the source volume can be made immediately available while the re-synchronization or copy process is running in the background (quick restore, selected with **Quick-resync replica data to the source volume**).

If you try to abort the session after the copy process has started, the abort operation is rejected and the session continues.

> ⓘ **IMPORTANT:** The original source volumes are not retained after instant recovery. If the session fails after the copy process has already started, the application data residing on the original source volumes is lost.

Limitations

- Instant recovery cannot be performed after a point-in-time recovery. To be able to perform instant recovery, you first need to run backup with the instant recovery options set.

- The number of replicas available for instant recovery is limited by **Number of replicas rotated**, which sets the size of the replica set. You can view these replicas in the GUI, in the **Instant Recovery** context by expanding **Restore Sessions**. Replicas are identified by the backup specification name and the session ID. Other information, such as time when a particular replica was created, is also provided. Alternately, you can use the omnidbvss command to list sessions. For information about omnidbvss, see the *HP Data Protector Command Line Interface Reference*.

- Concurrent instant recovery sessions utilizing the same backup system or the same application system are not supported. Restores utilizing the same backup system or the same application system must be performed serially.

- For instant recovery of Exchange Server 2007 data to a different location, the Microsoft Virtual Disk Service, HP P4000 SAN Solutions, HP P6000 EVA Disk Array Family, HP P9000 XP Disk Array Family, and HP 3PAR StoreServ Storage restore options are not available.

For additional writer specific limitations, see "Writer specifics" (page 61).

HP P4000 SAN Solutions considerations

- The following table lists instant recovery methods available with the HP P4000 SAN Solutions

Table 7 Supported instant recovery methods for HP P4000 SAN Solutions

Instant recovery method		HP P4000 Agent
Restore snapshot data to the source volume	with the source volume retained	NO
	with source volume not retained	YES

- HP P4000 SAN Solutions replicas depend on previous replicas in a chain. This means that in order to restore a particular replica, newer replicas are deleted. When performing instant recovery, try to use the newest replica possible and then gradually restore other replicas to the desired point in time, to avoid accidental replica loss.

- Any newer replicas created outside of Data Protector will prevent instant recovery from a session that contains the volume which was replicated outside of Data Protector.

HP P6000 EVA Array considerations

- The following table lists instant recovery methods available with the HP P6000 EVA Disk Array Family:

Table 8 Supported instant recovery methods with HP P6000 EVA Disk Array Family

Instant recovery method		VDS hardware provider	HP P6000 / HP 3PAR SMI-S Agent
Switch of disks		YES [1]	YES [1]
Copy of replica data	with the source retained	NO	YES
	with the source not retained	NO	YES
Restore snapshot data to the source volume	with the source volume retained	NO	YES
	with source volume not retained	NO	YES

[1] Available if the backup was performed using Snapclone as subtype.

- If you restore using P6000 EVA Array, before instant recovery, it is recommended to check if the SMI-S CIMOMs were configured properly in the Data Protector cell by executing the `omnidbsmis -ompasswd -check [-host ClientName]` command.

  This command performs a health check of you environment, which may help identify such potential problems as wrong user name or password provided, a broken network connection, a DNS resolution problem, and so on.

- After instant recovery, restored filesystems are mounted to the same mount points or drive letters as they were at the backup time. If these mount points or drive letters have other filesystems mounted, these filesystems are automatically dismounted before instant recovery, and the restored filesystems are mounted afterwards.

- You can restore snapshots or MirrorClone snapshots only to the source volume.

- On Windows Server 2008 systems, if the VDS P6000 EVA Array hardware provider was installed at the time of the backup, the VDS P6000 EVA Array hardware provider is required for instant recovery. On Windows Server 2003 systems, the VDS P6000 EVA Array hardware provider is always required for instant recovery.

- If the VDS P6000 EVA Array hardware provider is not used, Data Protector can perform fewer checks during instant recovery which may decrease the ability to revert changes in case of failure.

- If using Command View 9.2.1, instant recovery of a snapshot or MirrorClone snapshot can only be performed after the cache on the Command View server has been refreshed (triggered by default every 30 minutes).

HP P9000 XP Array considerations

- The following table lists instant recovery methods available with the HP P9000 XP Array:

Table 9 Supported instant recovery methods for HP P9000 XP Disk Array Family

Instant recovery method	VDS hardware provider	HP P9000 XP Agent
Switch of disks	YES [1]	NO
Resync of volumes	NO	YES [2]

Instant recovery method		VDS hardware provider	HP P9000 XP Agent
Restore snapshot data to the source volume	with the source volume retained	NO	NO
	with source volume not retained	NO	YES [2]

[1] hardware provider in VSS compliant mode

[2] hardware provider in resync mode

- Instant recovery depends on the P9000 XP Array VSS hardware provider mode which was used during backup. If the VSS compliant mode was used, you can restore your data only using VDS. If the resync mode was used, you can restore your data only using SSEA.

- Changing the selected P9000 XP Array VSS hardware provider mode between different backup sessions using the same backup specification is not recommended. If the mode is changed while the replica set rotation count is set to more than 1, instant recovery will fail in the following situations:

  - If you perform a backup in one mode and then the same backup in the other mode, restore of the backup performed in the VSS compliant mode (the "switch of disks" restore) will fail because the relationship between the S-VOL and its P-VOL will be detected from the other backup performed in the resync mode. This pair relationship should not be removed during restore and thus a switch of the source volume with the replica cannot be performed.

    To restore such a backup, perform one of the following prior to restore:

    - Using the omnidbvss command, manually remove S-VOLs that were created during backups in resync mode.

    - Set the option OB2VSS_FORCE_INSTANT_RECOVERY in the omnirc file and enable the restore option **Retain source for forensics** in the GUI. When this option is selected, Data Protector preserves the pair relationship between the S-VOL and its P-VOL created in the backup with resync mode.

  - If you perform restore using re-synchronization (the "copy of replica data" restore) and the production source volume (P-VOL) is not presented on the system, the restore session is aborted. This may happen after at least two backup sessions are run using the same backup specification in different modes and you run a "switch of disks" restore before a "copy of replica data" restore.

    Under such circumstances, perform one of the following prior to restore:

    - On the application system, manually present the P-VOL.

    - Set the option OB2VSS_FORCE_INSTANT_RECOVERY in the omnirc file.

- In VSS compliant mode, the VDS P9000 XP Array hardware provider is required for instant recovery, even if the backup was performed without it. In such a case, install the VDS P9000 XP Array hardware provider and execute the omnidbvss -resolve command before you attempt an instant recovery:

  - to resolve the source disks: omnidbvss -resolve -apphost *ApplicationClient*

  - to resolve the target disks: omnidbvss -resolve session *SessionID*

The following restore considerations apply to restore using the HP P9000 XP Agent only:

- You cannot start another instant recovery session using the same disk on the application system at the same time. A session can be started only after the preceding session completes synchronization.
- After instant recovery, the restored filesystems are mounted to the drive letters as were used at the time when backup was run. If these drive letters have other filesystems mounted, these filesystems are automatically dismounted before instant recovery, and the restored filesystems are mounted afterwards.

HP 3PAR StoreServ Storage considerations

- The following table lists instant recovery methods available with the HP 3PAR StoreServ Storage:

Table 10 Supported instant recovery methods for HP 3PAR StoreServ Storage

Instant recovery method		HP 3PAR VSS Agent
Restore snapshot data to the source volume	with the source volume retained	NO
	with source volume not retained	YES

- HP 3PAR StoreServ Storage replicas do not depend on previous replicas in a chain.

Instant recovery procedure

Limitations

- Instant recovery of Cluster Shared Volumes is not supported.
- Snapshots can be restored only to the original LUN (the LUN from where the replica was created) and to the same client.

Procedure

1. In the **HP Data Protector Manager**, switch to the **Instant Recovery** context.
2. In the Scoping Pane, expand **MS Volume Shadow Copy Writers** under **Restore Objects**, expand the backup specification from which you want to restore, and click the backup session from which you want to restore.

Figure 11 Selecting writers components for instant recovery
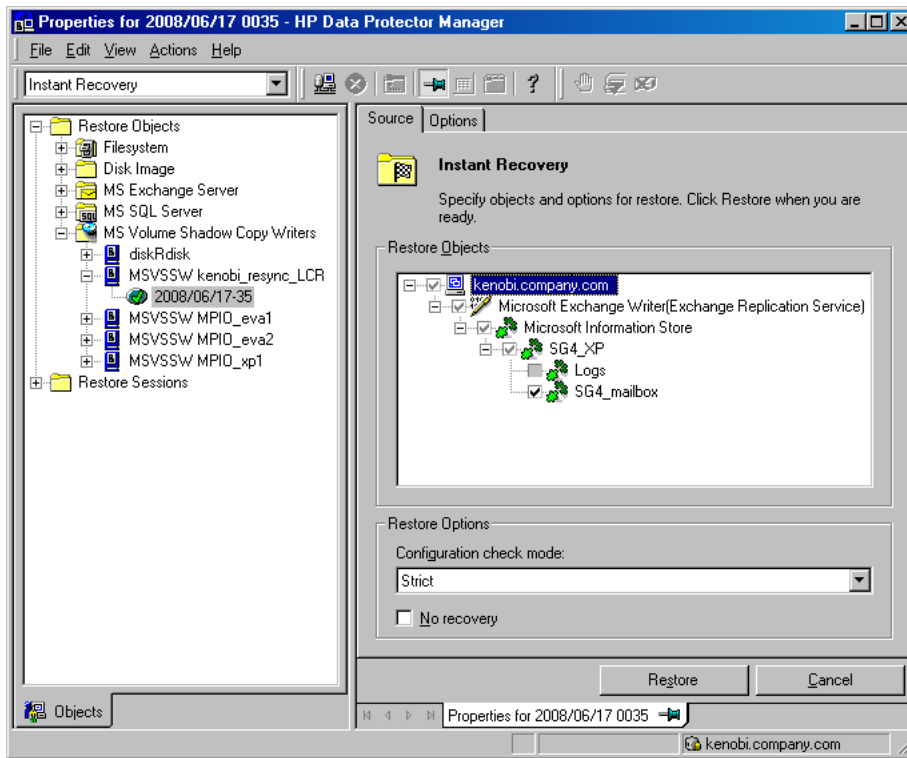
3. In the Source property page, specify writers and/or components for recovery.

   Select the configuration check mode. For more information, see "Configuration check" (page 36).

   To perform additional steps after the instant recovery before recovering the database, select the **No recovery** option. Data Protector will not perform a database recovery and you can finish the recovery steps (such as applying transaction logs) later on manually.

   If the option is not selected, all recovery steps are finished. In such a case, additional tasks cannot be performed. This option depends on the application writer that is being recovered.

   For any additional application-specific options, see "Writer specifics" (page 61).

4. Under the **Options** tab, select the instant recovery type:

   - **Restore using Microsoft Virtual Disk Service**

   - **Restore using Microsoft Volume Shadow Copy Service LUN resync**

   - Restore using the disk array agent

   See "Instant recovery methods" (page 50) for details on which type is supported with which disk array.

   For instant recovery using the VDS hardware provider or Disk Agent, to retain the source volume, select **Retain source for forensics**.

   When a recovery using a disk array agent is selected, additional options are available.

   - **Restore using HP P4000**

     The restore method: **Restore snapshot data to the source volume**.

     For details on these methods see "Instant recovery methods" (page 50).

   - **Restore using HP P6000 EVA SMI-S**

     The restore method: **Copy replica data to the source volume**, **Restore snapshot data to the source volume**, or **Switch the replica**. For details on these methods see "Instant recovery methods" (page 50).

If you select **Copy replica data to the source volume**, a full copy of the replica is created in the source storage (the copy is then switched with the source volume). A virtual copy is created and immediately available for use. However, in the background, the copy process runs for some time to copy all data from the replica to the source storage. To avoid the slowdown of the copy process caused by application use or the integrity check process (Exchange Server), use this option. The restore resumes after the time specified in **Wait up to n minutes** (default: 60) is reached (even if the copy process is not finished) or after the copy process has finished if it is finished before the specified time.

- **Restore using HP P9000 XP**

  The restore method: **Resync replica data to the source volume, Quick-Resync replica data to the source volume (internal swap)**, or **Restore snapshot data to the source volume**.

  For details on these methods see "Instant recovery methods" (page 50).

- **Restore using HP P10000 3PAR**

  The restore method: **Restore snapshot data to the source volume**.

  For details on these methods see "Instant recovery methods" (page 50).

For details, press **F1**.

Figure 12 Selecting instant recovery options (P4000 SAN Solutions integration)
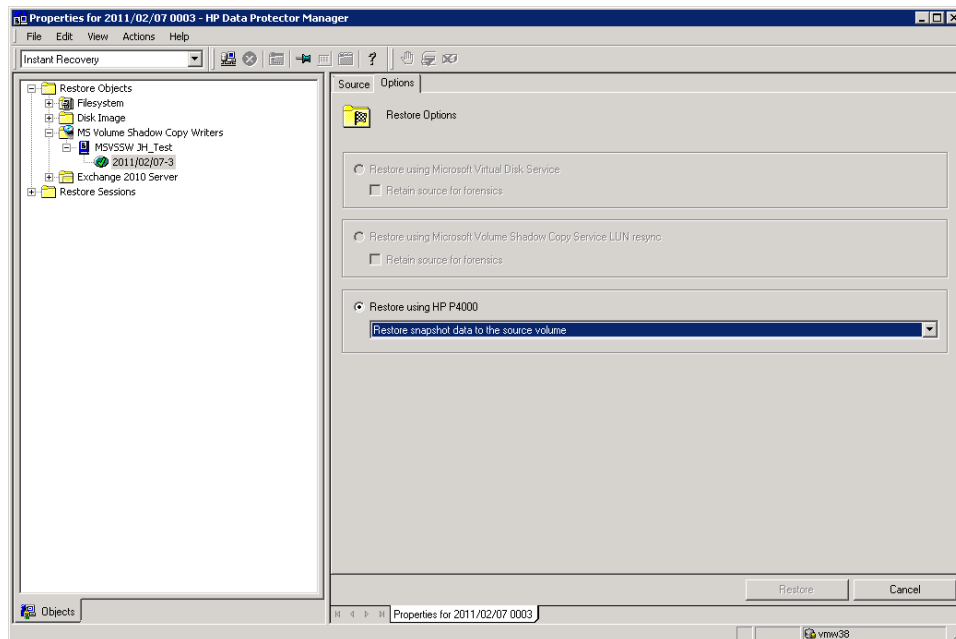
Figure 13 Selecting instant recovery options (P6000 EVA Array integration)
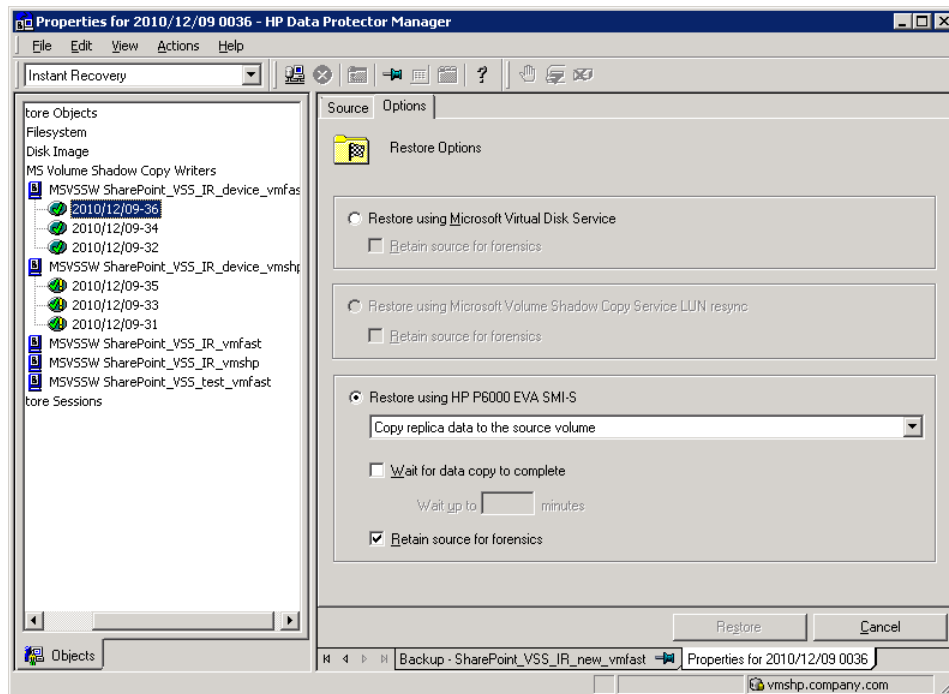
Figure 14 Selecting instant recovery options (P9000 XP Array integration, split mirror)
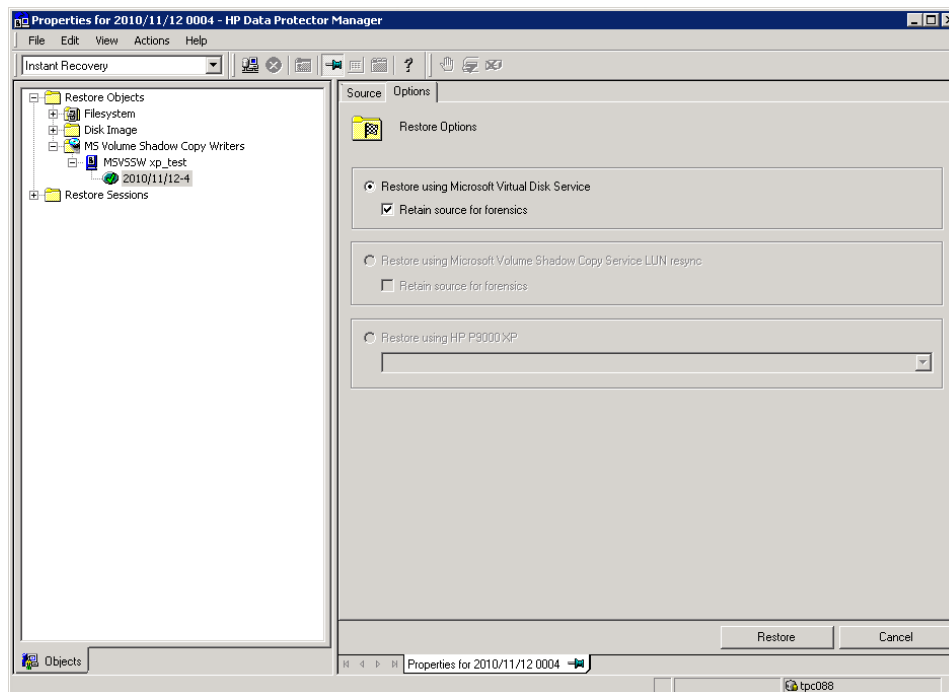
Figure 15 Selecting instant recovery options (P9000 XP Array integration, snapshot)
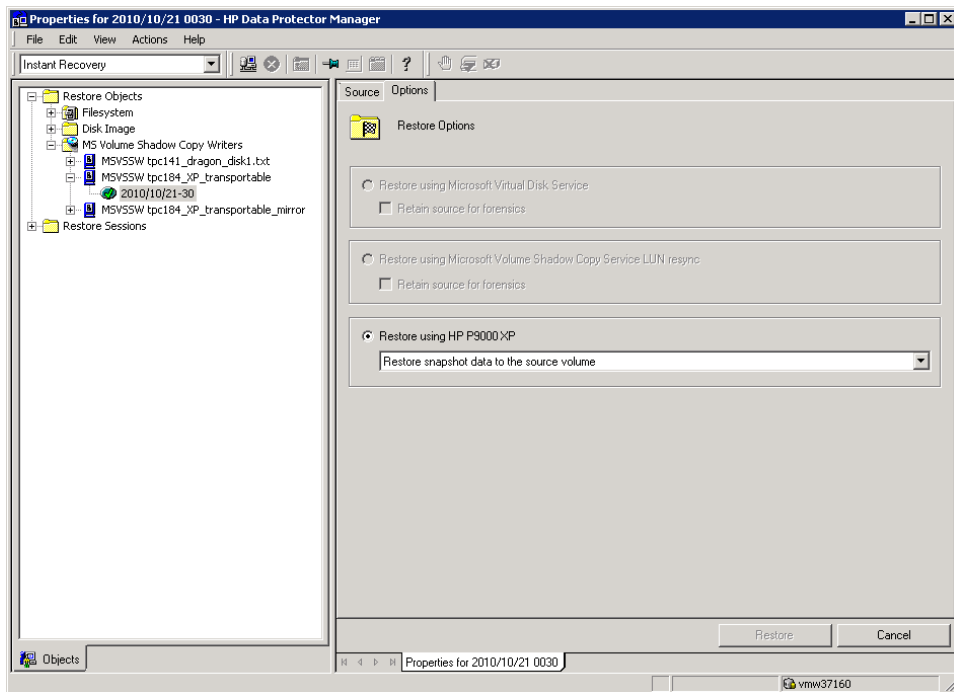
Figure 16 Selecting instant recovery options (3PAR StoreServ integration)
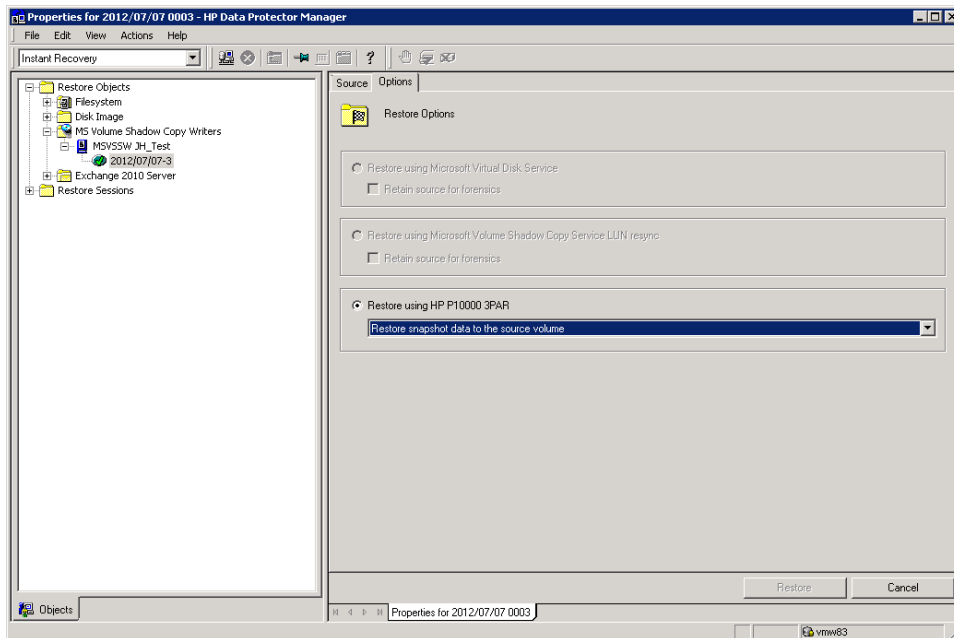
5. Click **Restore**.
6. Perform any additional writer specific steps. See "Writer specifics" (page 61).

Clusters

To perform an instant recovery in the cluster environment, use the above instant recovery procedure. When selecting the client name, use virtual client names instead of physical systems. For details on performing an instant recovery in cluster configurations, see the *HP Data Protector Zero Downtime Backup Administrator's Guide*.

6 Writer specifics

This chapter provides specific information about VSS writers, that you need to take into account before backing up or restoring the writers.

VSS writers either come with the Windows operating system or with applications. For a complete list of supported VSS writers and providers, see the latest support matrices at http://support.openview.hp.com/selfsolve/manuals.

The Data Protector Microsoft VSS integration does not provide any restore method for writers requesting a custom restore. If a writer specifies a custom restore method, it is only possible to restore the writer's data as plain files using the Data Protector functionality. You can perform the custom restore manually. For additional information on the restore methods, see the writer's documentation.

NOTE: Writers requiring custom restore methods are by default not shown by Data Protector. The `omnirc` option `OB2VSS_SHOWALLWRITERS` must be set to `1` for all writers to be displayed.

"Writer description" (page 61) provides a description of VSS writers.

Table 11 Writer description

Writer name	Description	Restore method
Certificate Authority Writer	This is a system writer, used to back up and restore Certificate Authority (CA) Service database. This service issues, revokes, and manages certificates employed in public key-based cryptography technologies.	Files are restored after a system restart.
Cluster Service Writer	This VSS writer uses a custom API, and is used to back up and restore Cluster Service on Microsoft Cluster Server (MSCS). The Cluster Service is a component on Windows Server systems used to control server cluster activities on cluster nodes. It is fundamental to the operation of the cluster.	Custom restore method
COM+ REGDB Writer	This VSS writer uses a custom API, and is used to back up and restore COM+ Database Service. This service provides automatic distribution of events to subscribing COM+ components.	Custom restore method
DHCP Jet Writer	This is a system writer, used to back up and restore DHCP Service database. DHCP Service provides dynamic IP address assignment and network configuration for Dynamic Host Configuration Protocol (DHCP) clients.	Files are restored after a system restart.
Event Log Writer	This is a system writer, used to back up and restore Event Logs. Event Logs are files where the Windows operating system saves information about events, such as starting and stopping services or the logging on and logging off of a user.	Files are restored after a system restart.
FRS Writer	This VSS writer uses a custom API, and is used to back up and restore File Replication Service data. File Replication Service is a multithreaded replication engine that replicates system policies and logon scripts stored in System Volume (`SYSVOL`). FRS can also replicate data for Distributed File System (Dfs), copy and maintain shared files and folders on multiple servers simultaneously.	Custom restore method
IIS Metabase Writer	This is a system writer, used to back up and restore Microsoft Internet Information Server (IIS). IIS 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).	Files are restored after a system restart.
MSDE Writer	This is a writer used to back up and restore Microsoft SQL Server 2000/2005. SQL Server is a database management	See "MSDE writer specifics" (page 100).

Table 11 Writer description (continued)

Writer name	Description	Restore method
	system that can respond to queries from client machines formatted in the SQL language.	
Microsoft SQL 2005 Writer	This is a writer used to back up and restore Microsoft SQL Server 2005. SQL Server is a database management system that can respond to queries from client machines formatted in the SQL language.	Standard VSS restore and instant recovery.
Microsoft Data Protection Manager Writer	This is a writer used to back up and restore Microsoft Data Protection Manager. Microsoft Data Protection Manager is a server that creates and stores replicas of clients and uses them for recovering the data on clients	See "Microsoft Data Protection Manager writer specifics" (page 63).
Microsoft Exchange Server Writer	This is a writer used to back up and restore Microsoft Exchange Server. Microsoft Exchange Server is a mail and groupware server.	See "Microsoft Exchange Server 2003 writer specifics" (page 66), "Microsoft Exchange Server 2007 writer specifics" (page 72), or "Microsoft Exchange Server 2010 writer specifics" (page 86).
Microsoft Virtual Server 2005 Writer	This is a writer used to back up and restore Microsoft Virtual Server 2005. Microsoft Virtual Server 2005 is a virtualization platform for Microsoft Windows Server systems. Data Protector supports live backup of individual virtual machines and the Virtual Server configuration, ensuring data consistency of the backup and restore. Hardware providers are not supported if a Virtual Server Machine is in online mode; use a software provider or put the Virtual Server Machine in offline mode. For details about Virtual Server online and offline mode, see the Microsoft Virtual Server documentation. Cluster configurations are not supported, only individual nodes can be backed up.	Standard VSS restore and instant recovery.
Microsoft Hyper-V Writer	This is a writer used to back up and restore Microsoft Virtual Server 2008 Hyper-V configuration and individual or all virtual machines running on the server. Software and hardware providers are supported during online and offline backup. Cluster-aware backups are not supported.	Standard VSS restore and instant recovery. For writer specifics, see "Microsoft Hyper-V writer specifics" (page 93).
SharePoint Services Writer	This is a reference writer used to back up and restore Microsoft Office SharePoint Server 2007 (MOSS). MOSS 2007 is an information portal used to connect and exchange expertise between people and teams. Only a single server configuration (farm) is supported. The OSearch VSS writer and SPSearch VSS writer are used by the reference writer to back up and restore the index files of user and help content. They must not be used for backup and restore.	Standard VSS restore. For writer specifics, see "Microsoft SharePoint Services writer specifics" (page 97).
NTDS Writer	This is a system writer used to back up and restore Microsoft Active Directory on Windows Server systems. Active Directory Service is a Windows server directory service that enables you to manage data structures distributed over a network. For example, Active Directory Service stores information about user accounts, passwords, phone numbers, profiles, and installed services. It provides methods for storing directory data and making this data available to network users and administrators.	To restore Active Directory, boot into Directory restore mode. Files will be restored if they can be overwritten.
Registry Writer	This VSS writer uses a custom API, and is used to back up and restore Windows Registry. Windows Registry is a database	Custom restore method

Table 11 Writer description *(continued)*

Writer name	Description	Restore method
	repository of information containing the Windows system configuration.	
Remote Storage Writer	This is a system writer used to back up and restore Remote Storage Service (RSS). RSS is used to automatically move infrequently accessed files from local to remote storage. Remote files are recalled automatically when the file is opened.	Files are restored after a system restart.
Removable Storage Manager Writer	This is a system writer used to back up and restore Removable Storage Manager Service. This service manages removable media, drives, and libraries.	Files are restored after a system restart.
System Writer	This is a system writer that backs up a specific set of Windows dynamic link libraries (DLL).	Files are restored after a system restart.
TermServLicencing Writer	This is a system writer that backs up Windows Terminal Services. These services provide a multi-session environment that allows client systems to access a virtual Windows desktop session and Windows-based programs running on the server.	Files are restored after a system restart.
WINS Jet Writer	This is a system writer, used to back up and restore Windows Internet Name Service (WINS). WINS is a dynamic replicated database service that can register and resolve NetBIOS names to IP addresses used on a TCP/IP network.	Files are restored after a system restart.
WMI Writer	This is a system writer, used to back up and restore Windows Management Instrumentation (WMI). WMI is a unified management infrastructure in Windows for monitoring system resources.	Files are restored after a system restart.

Microsoft Data Protection Manager writer specifics

Backup

Microsoft Data Protection Manager (DPM) is a server application that creates replicas of the clients, synchronizes them through LAN, and stores these replicas as snapshots.

The Data Protection Manager writer is used to back up:

- the Data Protection Manager database and the Data Protection Manager Report database
- the *latest version* of the DPM replicas.

IMPORTANT: To ensure data consistency, schedule a DPM replica synchronization before starting a backup.

The DPM uses DPM snapshots for restore. These snapshots are *not* backed up. To be able to recreate DPM snapshots you must manually schedule a backup of the replica each time after the DPM creates a new replica.

Two backup types are supported:

- *Full* (for the DPM databases and replicas)
- *Incremental* (replicas only).

If you select unsupported backup types (*Copy* or *Differential*) when scheduling the backup, Data Protector will abort the backup and display an error message.

Prerequisite

The MSDE writer (used for backing up the DPM databases) must be installed.

Limitations

- Hardware providers are not supported with DPM.
- If you start a backup while an incremental synchronization of the DPM replica is still in progress, the backup gets corrupted, although Data Protector does not report any errors. In case of synchronization with consistency check, Data Protector automatically aborts the backup session.

Figure 17 Selecting Microsoft Data Protection Manager database and replicas
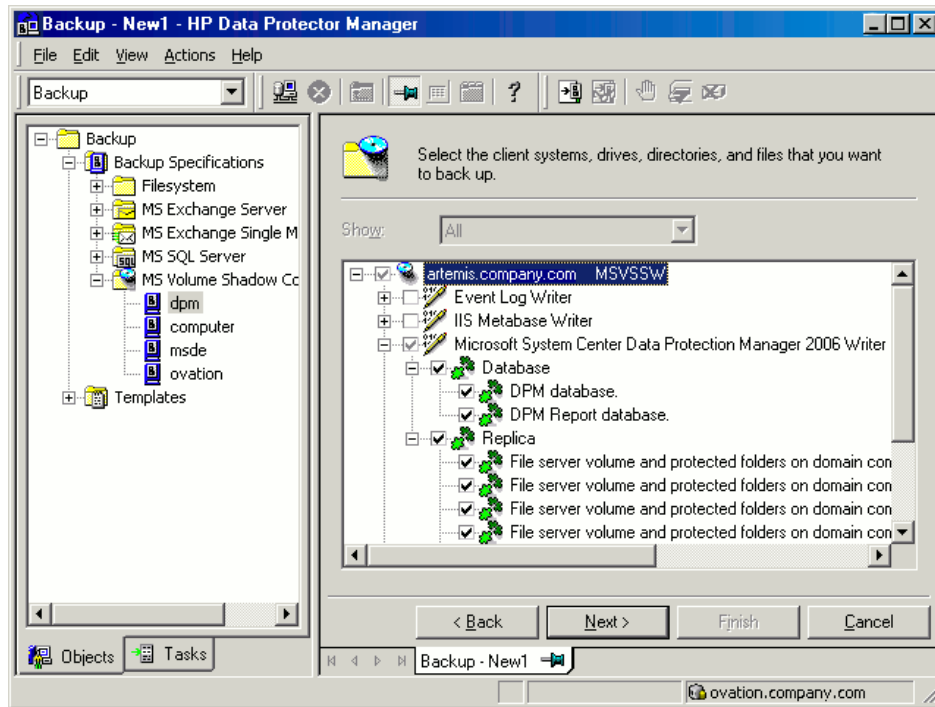
Restore

When restoring the DPM writer, you can:

- Restore the DPM *server* first and then use the DPM to restore clients.

  In case of a disaster, when the entire DPM server is lost, perform a standard disaster recovery procedure first and continue with restoring the DPM server. See "Restore the DPM server first" (page 65)

- Restore individual DPM *clients* directly, without using the DPM server (for example, if you cannot restore the DPM server or if you want to avoid the additional step of recreating the DPM snapshot). When restoring the DPM clients directly you can select between component restore and file restore modes. See "Restore the DPM clients directly" (page 65).

NOTE: Although the Data Protection Manager databases could also be restored using the MSDE writer, this method is not recommended, because DPM is *not* shut down automatically as with the DPM writer. If you really need to use this writer, shut down the DPM server manually.

Limitations

- Restore to another server is not supported by the Data Protection Manager writer.
- Parallel restore to different clients is not supported.

Restore the DPM server first

1. Start the DPM administrator console and add disks to the storage pool so that you have enough free space to restore the replicas.

   Ensure that the DPM writer (service) is started.

2. Switch to the Data Protector **Restore** context. Expand **Restore** and **Microsoft Volume Shadow Copy Writers** and select the client from which you want to restore the data.

3. In the Results Area, expand the DPM writer and select *only* the Data Protection Manager databases.

   Proceed as with general VSS writer restore. See "Standard restore" (page 46) for the general VSS writer restore procedure.

4. Execute the DPM command `DpmSync -Sync` to reallocate replicas.

5. Switch back to the Data Protector **Restore** context, and select and restore the necessary *replicas*.

Figure 18 Restoring the Microsoft Data Protection Manager client
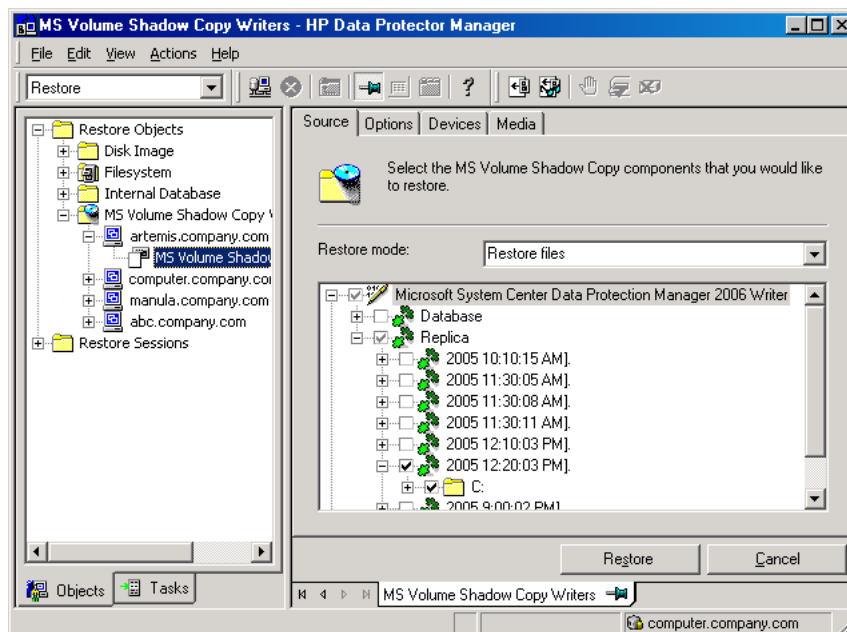
6. Use the DPM to restore individual clients.

> ① **IMPORTANT:** The DPM console does not automatically check for new or restored snapshots. Before you can start the restore of clients, you must use the Data Protection Manager to recreate a DPM snapshot.

   a. In the DPM console, open the **Recovery** context. Under the **Browse** tab, select the server, right click on the restored replica, and select **Create shadow copy now**.
   b. Select and restore the new snapshot to the client.

Restore the DPM clients directly

1. Switch to the **Restore** context. Expand **Restore** and **Microsoft Volume Shadow Copy Writers** and select the client from which you want to restore the data.

2. Select the restore modes:

   - **Restore Components**

     Use this mode *only* if the client to which you want to restore supports VSS, for example if you restore to Windows Server 2003 clients.

You can restore only entire replicas.

- **Restore Files**

   The client does not need to support VSS and you can restore individual folders or files.

3. When selecting the DPM writer for restore, select *only* the **Replica** components. Do not select the DPM database.
4. Click the **Options** tab, and under **Restore to another client** enter the name of the target client. Click **Next**.
5. Proceed as with general VSS writer restore. See for the general VSS writer restore procedure.

Microsoft Exchange Server 2003 writer specifics

Backup

Backup types

The Microsoft Exchange Server Writer supports the following Microsoft Exchange backup types:

- *Full* - backs up databases, transaction logs, and checkpoint files. The transaction logs are truncated.
- *Incremental* - backs up the transaction logs to record changes since the last full or incremental backup. The transaction logs are truncated. Not available with the hardware providers.
- *Differential* - similar to incremental backup, but the transaction logs are not truncated. Not available with the hardware providers.
- *Copy* - a Full backup, but the logs are not truncated. This type of backup is not intended for use in recovering failed systems.

Limitations

- A combination of VSS snapshot backups and non-VSS backups (for example, incremental stream backups) is not supported.
- You can back up only the whole server or full storage groups. Single stores cannot be backed up.
- Circular logging must be disabled.
- With Exchange Server 2003, only one VSS backup session can be running on the same application system at once. With Exchange Server 2007, only one VSS backup session backing up a specific storage group can be running on the same application system at once. Consequently, if you start such a backup session while another one is in progress, the latter waits until the first one finishes.
- You must create a separate incremental backup specification for each backup specification that will be used for instant recovery (has the **Track the replica for instant recovery** option enabled).

   For example, you *cannot* create the backup specification BSpec1 for the storage group SGroup1 with snapshots as the backup type and the backup specification BSpec2 for the storage group SGroup2 with snapclone as backup type and then use only one incremental backup specification containing both storage groups.

Recommendations

- With the HP SMI-S P6000 EVA Array provider, an Exchange storage group should not reside on more than four source volumes. However, the recommended configuration is that transaction logs reside on one source volume and the stores on another source volume. This configuration enables you to perform a rollforward recovery if only the stores are lost.

  Create *one backup specification for each storage group* and schedule them in a row.

Rollforward recovery requirements

To ensure that a rollforward recovery from your backup is possible, consider the following:

Standard restore: Transaction logs must be backed up to enable the rollforward operation.

Instant recovery: Rollforward of a storage group can be performed with incremental or differential ZDB to tape of the same objects as were backed up in instant recovery enabled sessions (ZDB to disk or ZDB to disk + tape). If the logs are still available on disk, the restore of last full backup will roll forward the store to the latest state (recorded in the logs).

Consistency check

A backup of the Microsoft Exchange Server database is considered as successful only if the consistency check of the replicated datafiles succeeds. The consistency check is enabled by default. To disable the consistency check, click on a created backup specification, right-click **Microsoft Exchange Writer** in the Source tab, and then click **Additional options**. In this page, you can also specify to throttle-back the consistency check for a second after the specified number of input/output operations.

The consistency check can also be run before instant recovery.

Restore

The following scenarios are possible:

- One or more databases are corrupted, but the log files are not damaged. In this case the database is restored and transaction logs are applied—a **rollforward recovery** from the loss of one or more databases.

- The log files are corrupted or missing. In this case all databases and log files need to be restored. A rollforward recovery of the database is not possible—a **point-in-time restore** after loss of a log file.

Rollforward recovery from the loss of one or more databases

For a rollforward recovery:

1. Dismount all stores from the storage group in which the target store resides using Microsoft Exchange System Manager.
2. In the Data Protector GUI, switch to the **Restore** context. Expand **Restore Objects** and **MS Volume Shadow Copy Writers** and select the client from which you want to restore the data.

   In the Results Area, expand Microsoft Exchange Writer and select the stores you want to recover. The **Logs** component is shaded and cannot be selected.

   NOTE: Individual stores (databases) are always restored from the latest backup. The **Properties** menu, which enables you to specify backup version, is not available at store level.

Figure 19 Selecting Microsoft Exchange Server stores for rollforward recovery

3. Proceed as with general VSS writer restore. See "Standard restore" (page 46) for the general VSS writer restore procedure.
4. Mount all stores from the storage group in which they reside using Exchange System Manager. Selected stores are recovered.

Point-in-time restore after loss of a log file

To perform a point-in-time restore:

1. Start Exchange System Manager and check if the storage group is already dismounted. If not, dismount the whole group.
2. Switch to the **Restore** context. Expand **Restore Objects** and **Microsoft Volume Shadow Copy Writers** and select the client from which you want to restore the data.

   In the Results Area, expand Microsoft Exchange Writer and select the whole storage group. Do not select individual stores.

Figure 20 Selecting Microsoft Exchange Server stores for point-in-time restore
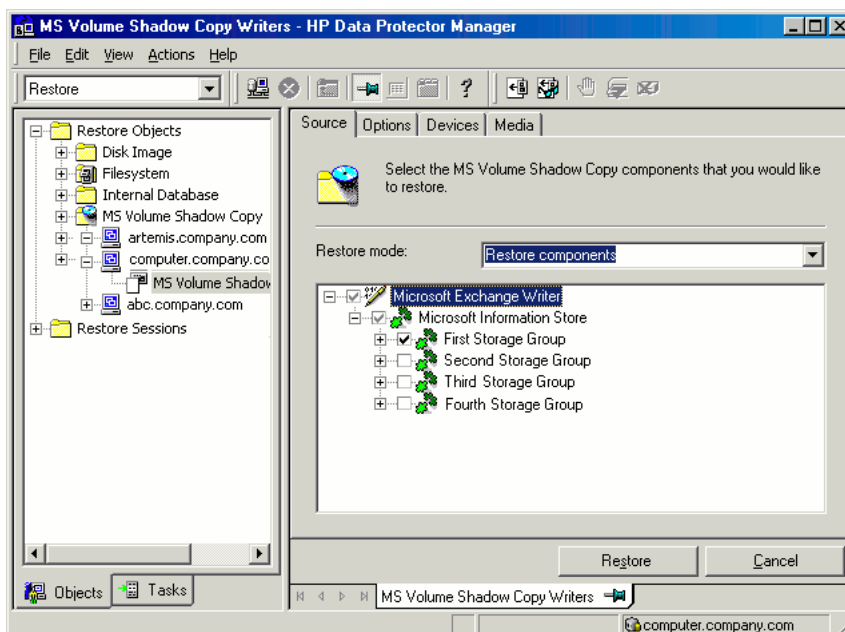
3. Proceed as with general VSS writer restore. See "Standard restore" (page 46) for the general VSS writer restore procedure.
4. Mount the stores from the storage group in which the target stores reside using Exchange System Manager. All stores are mounted and put in the state as they were at the last selected full, incremental, or differential backup.

Limitations

- Shadow copies cannot be restored to alternate locations on the backup system.
- You cannot restore the shadow copy to the Recovery Storage Group.
- Rollforward recovery cannot be performed after a point-in-time restore.
- Restore to Microsoft Exchange Server 2007 from backup made with Microsoft Exchange Server 2003 is not supported.

Instant recovery

With Microsoft Exchange Server 2003, you can restore a whole storage group or individual stores only to the original location. With Microsoft Exchange Server 2007, you can restore a whole storage group or individual stores to the original location or to a different location. For details, see "Microsoft Exchange Server 2007 writer specifics" (page 72).

Note that restore to the original location will fail if not all objects residing on the target volume are selected for restore. For example, if you have four stores and transaction logs on the target volume, and you select only one store for restore, the restore session will fail. Thus, the following configuration scenarios are possible when restoring Microsoft Exchange Server data:

- Transaction logs and database stores are on the same target volume.

  You cannot select only database stores for instant recovery in the GUI or CLI. If transaction logs and/or database stores are lost, the whole storage group (or all objects residing on the target volume) need to be recovered.

  In this case, you can perform only *point-in-time* recovery. Transaction logs will be replaced with the backed up transaction logs.

- Transaction logs and database stores are on different target volumes.

  You can select only the database stores for instant recovery in the GUI or CLI. If a database store is lost, it can be recovered separately provided that it resides alone on the target volume to be recovered. Otherwise, all stores residing on the target volume must be selected for restore. If transaction logs are lost, the whole storage group should be recovered.

  In this case, you can perform either *point-in-time* or *rollforward* recovery.

  To perform a *point-in-time* recovery of Microsoft Exchange Server writer data, select the whole storage group. Transaction logs will be replaced with the backed up transaction logs.

  To perform a *rollforward* recovery, select only the database stores and original location. The existing transaction logs will be applied to restored databases. However, rollforward recovery will not be possible if *point-in-time* recovery of the same backup session was performed before.

Prerequisites

- Manually dismount the databases to be recovered.

Post-instant recovery steps

- Manually re-mount the database stores.

Database recovery

You can run a database recovery from the **Instant Recovery** context of the Data Protector GUI. This option is available, if you have created a separate backup specification for **Incremental/Differential** backup with the same object and description as you have in the backup specification for instant recovery. Such an **Incremental/Differential** backup is based on the **Full** backup with the selected instant recovery option. You can select an **Incremental/Differential** backup in the **Instant Recovery** context and start restore. Instant recovery will be performed and transaction logs will be automatically applied to the recovered storage group.

User scenario for Microsoft Exchange Server 2003 backup and restore

This section provides examples of backup and restore policies for Microsoft Exchange Server 2003. Three examples are provided, one for ZDB (transportable VSS snapshots) and two for different types of restore.

Example - VSS transportable backup

This example describes a backup scenario for Microsoft Exchange Server 2003 using VSS transportable backup together with a disk array of the HP P6000 EVA Disk Array Family. The data is to be backed up on tape once a day for the storage group containing critical mailboxes and every second day for all other storage groups.

The storage groups should be backed up separately.

The following example setup is possible:

- A Microsoft Exchange Server 2003 is running on the application system and is connected to a disk array of the HP P6000 EVA Disk Array Family. A tape drive is connected to the application system for restore purposes.

- A separate backup system is connected to a tape library.

- Two storage groups are configured on the Microsoft Exchange Server 2003, each containing two stores.

  The first storage group, named Lists_Group, contains company-wide mailing lists. The storage group contains two stores, Staff_Store for general mailing lists and Team_Store for team mailing lists.

  The second storage group, named Mailbox_Group, contains a larger number of individual mailboxes for employees, which are given higher priority. This way the mailboxes are still available even if mailing lists are down. The mailboxes are set up in two separate stores, one store for each department, and named Sales_Store and Support_Store.

- Every storage group has a separate backup specification to reduce the time needed for the shadow copy creation. See "ZDB related problems" (page 106) for details of how to avoid issues with too long copy creation.

Figure 21 Microsoft Exchange Server 2003 storage groups
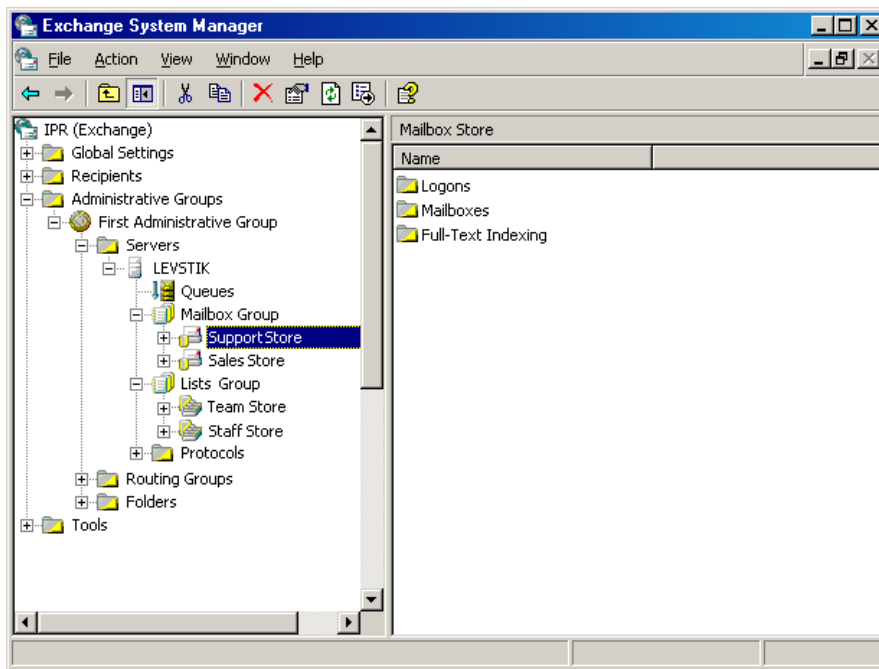
For the first storage group (Lists_Group), create the following backup specification:

1. Select **VSS transportable backup** as the backup type. Select the application and backup system and specify **Snapshot (Differential)** as the replica type.

2. Expand the Microsoft Exchange Server 2003 writer and select the first storage group (Lists_Group) for backup.

3. Using the Data Protector scheduler, schedule backup sessions for a VSS transportable backup at night. In the Recurring box, specify **Full** backup **Weekly** on Wednesdays and Fridays and **Incremental** backup **Weekly** on other days. Set the backup protection to 2 days.

For the second storage group (Mailbox_Group), create the backup specification that can be used for instant recovery:

1. Select **VSS transportable backup** as the backup type. Select the application and backup system and set the **Track the replica for instant recovery** option.

2. Expand the Microsoft Exchange Server 2003 writer and select the second storage group (Mailbox_Group) for backup.

3. Using the Data Protector scheduler you need to schedule backup sessions for a VSS transportable backup every day. In the Recurring box, select **Daily**, set **Recurring** options to 1 day and set the backup protection to 4 days.

Example restore scenario for Microsoft Exchange Server 2003

In this example the Microsoft Exchange Server 2003 is configured as in "Example - VSS transportable backup" (page 70) and the backup policy given in the example was implemented.

Example 1 Example - restoring a single store

The second store (Support_Store) in the second storage group (Mailbox_Group), which contains user mailboxes, is damaged, but the transaction logs and other stores are not damaged. Therefore, a rollforward recovery will be performed only for this store.

To perform a rollforward recovery, proceed as follows:

1. Start the Microsoft Exchange Manager and dismount all stores in Mailbox_Group, but do not dismount the stores from the first storage group.

2. Start Data Protector, go to the **Instant Recovery** context and expand the application system. Select the last backup session of the backup specification with enabled instant recovery options and expand Microsoft Exchange Server 2003 Writer. Select Support_Store under Mailbox_Group for instant recovery.

   Click **Restore**.

3. After the session is restored, start the Microsoft Exchange Manager. Mount all stores in the second storage group.

Example 2 Example - restoring a complete storage group after loss of transaction logs

In this example, the first store (Staff_Store) and transaction logs in the first storage group (Lists_Group) are damaged but the second store (Mailbox_Group) is intact. Since the transaction logs are damaged, a point-in-time restore of the whole storage group will be performed.

To restore the complete Lists_Group storage group, perform the following steps:

1. Start Exchange System Manager and check if the first storage group (Lists_Group) is already dismounted. If not, dismount the whole group.

2. Start Data Protector. In the Restore context expand the application system. Select the last backup session and then Microsoft Exchange Server 2003 Writer. Select the first storage group (Lists_Group).

   Start the restore.

3. Start Microsoft Exchange Manager and mount all stores in Lists_Group.

Microsoft Exchange Server 2007 writer specifics

Concepts

This section gives details of additional Microsoft Exchange Server 2007 features.

Continuous replication

Microsoft Exchange Server 2007 offers two models of replication for data protection that are supported by Data Protector.

- local continuous replication (**LCR**)

  With LCR, you can create and maintain an exact copy (LCR copy) of databases in a storage group. LCR copies are used in the event of data corruption, since you can switch the Exchange server to use LCR copies in only a few seconds. If an LCR copy is used for backup and is located on a different disk from the original data, the load on a production database is minimal.

- cluster continuous replication (**CCR**)

  CCR has the same characteristics as LCR. The only difference is that in the CCR environment, databases and transaction logs are replicated to separate servers. Therefore, CCR copies can be used for disaster recovery. You can perform VSS backups on the passive Exchange Server node where the CCR copy is located and thus reduce the load on the active node.

The replicated storage groups are represented as a new Exchange Server writer instance, Exchange Replication Service. They are backed up in the same way as original or production storage groups.

If using LCR or CCR with Standby Continuous Replication (SCR) configured, backups can only be performed on the source side of the SCR. Microsoft does not support backups on the SCR target side. For further information on SCR and supported SCR configurations, see the Microsoft website.

Restore to original or another location

With Microsoft Exchange Server 2007 writer, you can restore your data not only to the original location (from which the backup was performed) but also to a different location. You can restore:

- A whole storage group
- A single store

In both cases the respective LCR or CCR copies can also be restored.

You can restore data to:

- The original storage group
- A different storage group
- A non-Exchange location – With this restore method, after the restore is completed, the Recovery Storage Group (RSG) can be created automatically.
- A recovery server – This restore method restores data to different client and different storage group.

If you restore to a different storage group, you can access single mailboxes or individual e-mail messages at a different location without changing the original storage group content. Furthermore, if the whole server is destroyed, restoring to a different Exchange Server system (recovery server) will minimize the time window during which your mailboxes will be unavailable.

Backup

Backup types

The Microsoft Exchange Server Writer supports the following Microsoft Exchange backup types:

- *Full* - backs up databases, transaction logs, and checkpoint files. The transaction logs are truncated.
- *Incremental* - backs up the transaction logs to record changes since the last full or incremental backup. The transaction logs are truncated. Not available with the hardware providers.
- *Differential* - similar to incremental backup, but the transaction logs are not truncated. Not available with the hardware providers.
- *Copy* - a Full backup, but the logs are not truncated. This type of backup is not intended for use in recovering failed systems.

Limitations

- A combination of VSS snapshot backups and non-VSS backups (for example, incremental stream backups) is not supported.
- You can back up only the whole server or full storage groups. Single stores cannot be backed up.
- Circular logging must be disabled.
- With Exchange Server 2003, only one VSS backup session can be running on the same application system at once. With Exchange Server 2007, only one VSS backup session backing up a specific storage group can be running on the same application system at once.

Consequently, if you start such a backup session while another one is in progress, the latter waits until the first one finishes.

- You must create a separate incremental backup specification for each backup specification that will be used for instant recovery (has the **Track the replica for instant recovery** option enabled).

  For example, you *cannot* create the backup specification BSpec1 for the storage group SGroup1 with snapshots as the backup type and the backup specification BSpec2 for the storage group SGroup2 with snapclone as backup type and then use only one incremental backup specification containing both storage groups.

Recommendations

- With the HP SMI-S P6000 EVA Array provider, an Exchange storage group should not reside on more than four source volumes. However, the recommended configuration is that transaction logs reside on one source volume and the stores on another source volume. This configuration enables you to perform a rollforward recovery if only the stores are lost.

  Create *one backup specification for each storage group* and schedule them in a row.

Rollforward recovery requirements

To ensure that a rollforward recovery from your backup is possible, consider the following:

Standard restore: Transaction logs must be backed up to enable the rollforward operation.

Instant recovery: Rollforward of a storage group can be performed with incremental or differential ZDB to tape of the same objects as were backed up in instant recovery enabled sessions (ZDB to disk or ZDB to disk + tape). If the logs are still available on disk, the restore of last full backup will roll forward the store to the latest state (recorded in the logs).

Consistency check

A backup of the Microsoft Exchange Server database is considered as successful only if the consistency check of the replicated datafiles succeeds. The consistency check is enabled by default. To disable the consistency check, click on a created backup specification, right-click **Microsoft Exchange Writer** in the Source tab, and then click **Additional options**. In this page, you can also specify to throttle-back the consistency check for a second after the specified number of input/output operations.

The consistency check can also be run before instant recovery.

LCR and CCR environments

In LCR and CCR environments, the replicated storage groups are represented as a new instance of Exchange Server writer, **Exchange Replication Service**. The replicated storage groups are backed up in the same way as original (production) storage groups.

You can select any combination of storage groups for backup. However, you cannot select the original and its replicated storage group in the same backup specification. See Figure 22 (page 75).

Figure 22 Selecting a replicated Microsoft Exchange Server 2007 storage group
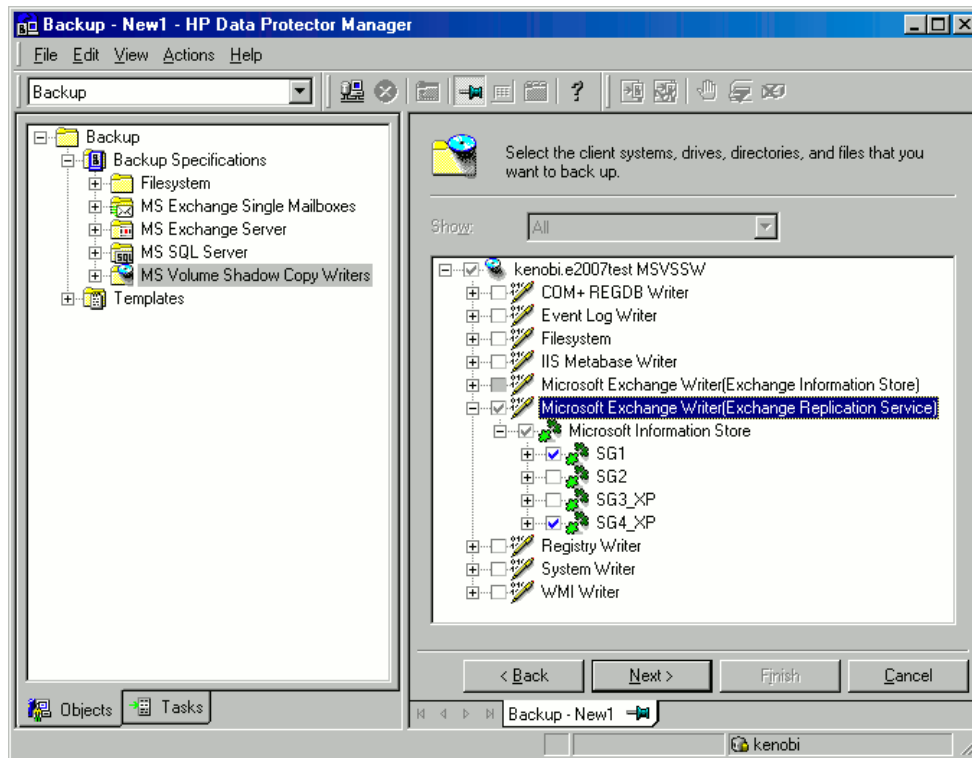
Cluster support in a CCR environment

In a CCR environment, a cluster node from which you want a backup to be performed can be selected, regardless of which instance (Information Store or Replication Service) resides on this node. If you select the cluster node, Data Protector backs up any available instance on this node ignoring the selection of the instance in the GUI even if, for example, you select the replicated storage group (Exchange Replication Service) as backup object.

To specify the cluster node from which you want to perform a backup of any instance residing on this node, right-click an Exchange writer and click **Additional options**. In the MS Exchange additional options dialog box, select the node under **Back up any available instance from node**. See .

Figure 23 Additional options for Microsoft Exchange Server 2007 in a CCR environment
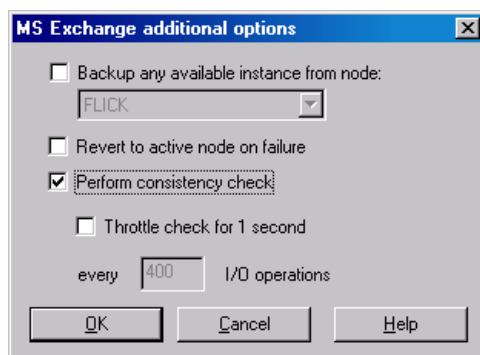
When backing up a Replication Service instance, backup may fail due to any of the following reasons:

• The selected node is not available.

• The status of the storage group to be backed up is not "Healthy".

- Data Protector is not running on the selected node.
- `Vssbar.exe` cannot be started on the selected node.

To avoid the session failing, select the option **Revert to active node on failure** in the same dialog box. The backup will be restarted on the original server (active cluster node) and the original storage group will be backed up. This option is ignored during backup of an Information Store instance.

Prerequisites

- Before creating a backup specification and before running a backup, ensure that the Exchange Server 2007 copy status is "Healthy". Otherwise, you cannot browse the objects to be backed up during creating backup specification, or the backup will fail. Note that the status "Initializing" is not acceptable.

Limitations

- In CCR environments, if a replicated storage group (Replication Service instance) is selected for backup, no other VSS writer (for example, filesystem writer or SQL Server writer) can be selected to be backed up in the same backup session. Because in this case, two writers would be located on different systems, but Data Protector limits VSS backup sessions to involve only one system.
- In LCR and CCR environments, each storage group can have only one store.
- CCR clusters can have only two nodes.

Considerations

- Transaction log files are truncated after each full or incremental Microsoft Exchange Server backup. The LCR and CCR clustering technologies, however, guarantee that logs that have not been replicated are not deleted. Thus, running backups in a mode that truncates logs may not actually free space. This may happen if replication of logs has not completed yet.
- Before using P9000 XP Array VSS hardware provider in the resync mode, consider all applicable limitations of this mode. The limitations are listed in "HP P9000 XP Disk Array Family" (page 36) and "Limitations" (page 84). Due to these limitations, it is recommended to rather use P9000 XP Array provider in the VSS compliant mode, or P6000 EVA Array VSS provider, or VSS software provider.

Restore

You can restore Microsoft Exchange Server data by performing a standard restore or instant recovery session:

- See "Standard restore" (page 77).
- See "Instant recovery" (page 83).

Limitations

- Individual databases (stores) can only be restored from the latest backup. To restore the databases to an earlier point in time, the complete storage group must be restored.

Standard restore

The following scenarios are possible:

- One or more databases are corrupted, but the log files are not damaged. In this case the database is restored and transaction logs are applied—a **rollforward recovery** from the loss of one or more databases.

- The log files are corrupted or missing. In this case all databases and log files need to be restored. A rollforward recovery of the database is not possible—a **point-in-time restore** after loss of a log file.

- Some data in a mailbox has been lost. For example, an e-mail has been deleted by mistake. Restore of the mailbox to an earlier point in time is needed. For details, see "Restoring individual mailboxes" (page 79).

Consistency check

Optionally, to specify options for the consistency check of a Microsoft Exchange writer, right-click the writer and click **Additional options**.

Rollforward recovery from the loss of one or more databases

For a rollforward recovery:

1. Dismount all stores from the storage group in which the target store resides using Microsoft Exchange System Manager.
2. In the Data Protector GUI, switch to the **Restore** context. Expand **Restore Objects** and **MS Volume Shadow Copy Writers** and select the client from which you want to restore the data.

   In the Results Area, expand Microsoft Exchange Writer and select the stores you want to recover. The **Logs** component is shaded and cannot be selected.

   NOTE: Individual stores (databases) are always restored from the latest backup. The **Properties** menu, which enables you to specify backup version, is not available at store level.

Figure 24 Selecting Microsoft Exchange Server stores for rollforward recovery

3. Proceed as with general VSS writer restore. See "Standard restore" (page 46) for the general VSS writer restore procedure.
4. Mount all stores from the storage group in which they reside using Exchange System Manager. Selected stores are recovered.

Point-in-time restore after loss of a log file

To perform a point-in-time restore:

1. Start Exchange System Manager and check if the storage group is already dismounted. If not, dismount the whole group.
2. Switch to the **Restore** context. Expand **Restore Objects** and **Microsoft Volume Shadow Copy Writers** and select the client from which you want to restore the data.

    In the Results Area, expand Microsoft Exchange Writer and select the whole storage group. Do not select individual stores.

Figure 25 Selecting Microsoft Exchange Server stores for point-in-time restore
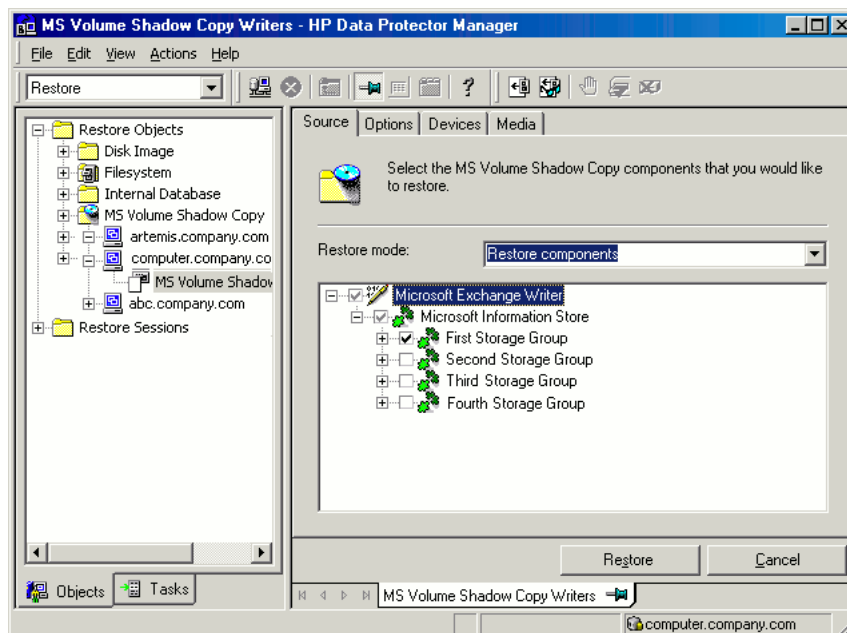
3. Proceed as with general VSS writer restore. See "Standard restore" (page 46) for the general VSS writer restore procedure.
4. Mount the stores from the storage group in which the target stores reside using Exchange System Manager. All stores are mounted and put in the state as they were at the last selected full, incremental, or differential backup.

Restoring individual mailboxes

To restore an individual mailbox:

1. Open the Data Protector GUI. In the Context List, click **Restore**. Expand **Restore Objects** and **MS Volume Shadow Copy Writers** and select the client from which you want to restore the data.

2. In the Results Area, expand the Microsoft Exchange Writer and select the storage group that contains this particular mailbox. Right-click the storage group, click **Properties**, and specify the desired point in time.

   Right-click the storage group and click **Restore as**. In the MS Exchange additional options dialog box, select the option **Restore to a non-Exchange location and create RSG**. In the **Original** drop-down list, select the database that contains the mailbox, click **Add**, and then click **OK**.

3. Proceed as described in the general VSS writer restore procedure.

4. After the database has been restored, mount the database: open the Exchange Management Shell or some other GUI tool like the Exchange Server Disaster Recovery Analyzer Tool.

   To list all available databases, execute:

   ```
   [PS] C:\>get-mailboxdatabase
   ```

   ```
   Name                 Server        StorageGroup          Recovery
   ----                 ------        ------------          --------
   LCR_store1           TPC181        LCR_sg1               False
   sg3_store1           TPC181        sg3_local             False
   store1               exchclu3      sg1                   False
   sg3_store2           TPC181        sg3_local             False
   sg1_store5           TPC181        First Storage Group   False
   sg4_store1           TPC181        sg4                   False
   sg3_store2           TPC181        DP RSG                True
   ```

   In the example above, the database that has been restored to the Recovery Storage Group `DP RSG` is named `sg3_store2`.

   To mount the database, execute:

   ```
   [PS] C:\>mount-database -identity "TPC181\DP RSG\sg3_store2"
   ```

5. Suppose the mailbox in question belongs to the person called John Doe. To extract the mailbox items from the restored database and move the items to the John Doe's mailbox to the folder RSG, execute:

```
[PS] C:\>restore-mailbox -RSGMailbox "John Doe" -RSGDatabase
"TPC181\DP R SG\sg3_store2" -identity "Doe" -TargetFolder RSG

Confirm
Are you sure you want to perform this action?
Recovering mailbox content from the mailbox 'John Doe' in the recovery
database 'TPC181\DP RSG\sg3_store2' into the mailbox for 'John Doe
(Doe@dp2.com)'. The operation can take a long time to complete.
[Y] Yes  [A] Yes to All  [N] No  [L] No to All  [S] Suspend  [?] Help
(default is "Y"):Y


Identity                         : dp2.com/Users/John Doe
DistinguishedName                : CN=John Doe,CN=Users,DC=dp2,DC=com
DisplayName                      : John Doe
Alias                            : Doe
LegacyExchangeDN                 : /o=DataProtector/ou=First Administrative Gro
                                   up/cn=Recipients/cn=Doe
PrimarySmtpAddress               : Doe@dp2.com
SourceServer                     : TPC181.dp2.com
SourceDatabase                   : TPC181\DP RSG\sg3_store2
SourceGlobalCatalog              : TPC136
SourceDomainController           :
TargetGlobalCatalog              : TPC136
TargetDomainController           :
TargetMailbox                    : dp2.com/Users/John Doe
TargetServer                     : TPC181.dp2.com
TargetDatabase                   : TPC181\sg3_local\sg3_store2
MailboxSize                      : 40553300B
IsResourceMailbox                : False
SIDUsedInMatch                   :
SMTPProxies                      :
SourceManager                    :
SourceDirectReports              :
SourcePublicDelegates            :
SourcePublicDelegatesBL          :
SourceAltRecipient               :
SourceAltRecipientBL             :
SourceDeliverAndRedirect         :
MatchedTargetNTAccountDN         :
IsMatchedNTAccountMailboxEnabled :
MatchedContactsDNList            :
TargetNTAccountDNToCreate        :
TargetManager                    :
TargetDirectReports              :
TargetPublicDelegates            :
TargetPublicDelegatesBL          :
TargetAltRecipient               :
TargetAltRecipientBL             :
TargetDeliverAndRedirect         :
Options                          : Default
SourceForestCredential           :
TargetForestCredential           :
TargetFolder                     : \RSG\Recovered Data - John Doe - 02.1
                                   0.2008 16:49:59
PSTFilePath                      :
RsgMailboxGuid                   : 0441be6c-46f6-4d8f-8562-ab615731ae89
RsgMailboxLegacyExchangeDN       : /O=DATAPROTECTOR/OU=FIRST ADMINISTRATIVE GRO
                                   UP/CN=RECIPIENTS/CN=DOE
RsgMailboxDisplayName            : John Doe
RsgDatabaseGuid                  : deb5029b-4737-4fea-8c2d-ece24007e75d
StandardMessagesDeleted          : 0
AssociatedMessagesDeleted        : 0
DumpsterMessagesDeleted          : 0
MoveType                         : Restore
MoveStage                        : Completed
StartTime                        : 02.10.2008 16:50:06
EndTime                          : 02.10.2008 16:50:12
```

```
StatusCode                    : 0
StatusMessage                 : This mailbox in the recovery storage group d
                                atabase has been restored to the target user
                                 mailbox.
ReportFile                    : C:\Program Files\Microsoft\Exchange Server\L
                                ogging\MigrationLogs\restore-Mailbox20081002
                                -164958-8141342.xml
```

Figure 26 Restored mailbox items
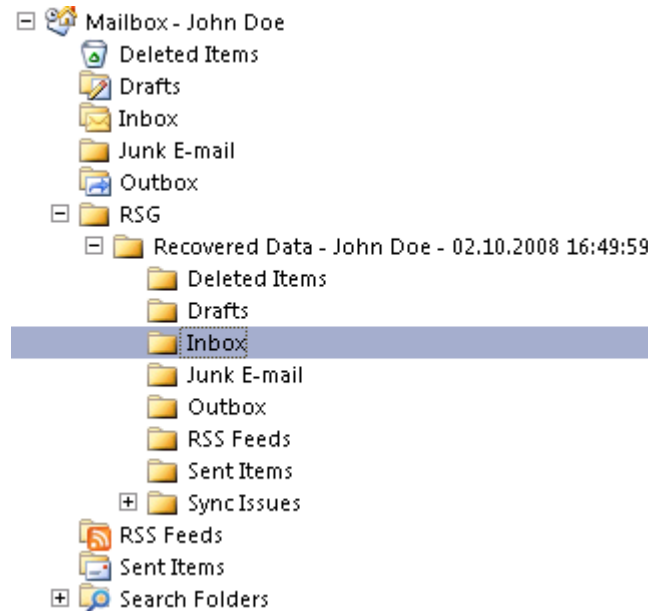
Restoring a LCR or CCR copy to the original location

If you restore an LCR or CCR copy to the original location, the restore will be performed to the original database (Exchange Information Store) and not to the database copy (Exchange Replication Service).

1. Right-click a storage group, a store, or logs and click **Restore as**.

2. In the MS Exchange additional options dialog box, select the target location for the components you want to restore: target server, target storage group, and target stores. The following options are available:

- **Restore to a different store**

  This option is selected by default.

  Select this option to select the target stores for each store to be restored (original stores). First, select the target system from the Target server name drop-down list, and then choose the desired pairs of stores by selecting entries in the Original and Target drop-down lists. Note that only stores cannot be restored, so logs are automatically selected in the restore session to different location.

- **Restore to a non-Exchange location**

  Select this option to restore your data to a non-Exchange location. In this case, the restored data will not be managed by Exchange Server and a Recovery Storage Group (RSG) will not be created. You can manually create an RSG after the restore session completes. First, select the target system from the Target server name drop-down list, and then choose the desired stores using the Original drop-down list.

- **Restore to a non-Exchange location and create RSG**

  Select this option to restore your data to a non-Exchange location. After restore, Data Protector will create a Recovery Storage Group called `DP RSG` on the target server. The selected stores and logs will be restored to this recovery group. First, select the target system from the Target server name drop-down list, and then choose the desired stores using the Original drop-down list.

By default, the storage group is restored in the `C:\Omni` directory. To select another location use the **Restore into location** option. Click **Browse** and select the desired location.

> (!) **IMPORTANT:**   The selected directory must be either empty or you can specify to create a new directory. If the directory is not empty, the restore session fails.

Note that you can only specify the restore location for the storage group and not for a specific store.

Figure 27 Restore to different location options (Exchange Server 2007 Writer)

Instant recovery

With Microsoft Exchange Server 2003, you can restore a whole storage group or individual stores only to the original location. With Microsoft Exchange Server 2007, you can restore a whole storage group or individual stores to the original location or to a different location. For details, see "Microsoft Exchange Server 2007 writer specifics" (page 72).

Note that restore to the original location will fail if not all objects residing on the target volume are selected for restore. For example, if you have four stores and transaction logs on the target volume, and you select only one store for restore, the restore session will fail. Thus, the following configuration scenarios are possible when restoring Microsoft Exchange Server data:

- Transaction logs and database stores are on the same target volume.

  You cannot select only database stores for instant recovery in the GUI or CLI. If transaction logs and/or database stores are lost, the whole storage group (or all objects residing on the target volume) need to be recovered.

  In this case, you can perform only *point-in-time* recovery. Transaction logs will be replaced with the backed up transaction logs.

- Transaction logs and database stores are on different target volumes.

  You can select only the database stores for instant recovery in the GUI or CLI. If a database store is lost, it can be recovered separately provided that it resides alone on the target volume to be recovered. Otherwise, all stores residing on the target volume must be selected for restore. If transaction logs are lost, the whole storage group should be recovered.

  In this case, you can perform either *point-in-time* or *rollforward* recovery.

  To perform a *point-in-time* recovery of Microsoft Exchange Server writer data, select the whole storage group. Transaction logs will be replaced with the backed up transaction logs.

To perform a *rollforward* recovery, select only the database stores and original location. The existing transaction logs will be applied to restored databases. However, rollforward recovery will not be possible if *point-in-time* recovery of the same backup session was performed before.

- Some data in a mailbox has been lost. For example, an e-mail has been deleted by mistake. Restore of the mailbox to an earlier point in time is needed. You need to perform a *point-in-time* recovery, which is available in both configurations, when the transaction logs and databases reside on the same storage volume and when they reside on different storage volumes. The procedure for restoring an individual mailbox is similar as when performing standard restore (see "Restoring individual mailboxes" (page 79)). However, you also need to consider instant recovery-specific information that is described in this section.

Prerequisites

- For restore to a different location, `Inet` must run under a domain account which needs to be a member of the following groups on the local system:

  - Administrators

  - Exchange Server Administrators

- While configuring restore to a different storage group in the Data Protector GUI, only mounted stores are displayed in the Target drop-down list. To enable all stores for target selection, mount the dismounted stores and click **Reload** in the **MS Exchange additional options**.

> ⊘ **IMPORTANT:** During restore to a non-Exchange location and creation of Recovery Storage Group (RSG), Data Protector deletes RSG that may already exists at the target location.

Limitations

- In CCR environments where original database resides on a different P6000 EVA Array than the database copy, restore of the database copy to the location where original database resides will fail, since restore to a different disk array is not supported. In this case, perform manual failover of Exchange Server before the restore.

- Restore to a different location can be only performed using Microsoft Virtual Disk Service.

  Therefore, with P9000 XP Array VSS hardware provider, restore to a different location is possible only after the backup session was run in the VSS compliant mode.

- With P9000 XP Array VSS hardware provider, if a backup of an LCR or CCR storage group copy is run in the resync mode, restore of data to the original location (original database or Exchange Information Store) fails.

  The reason for such a behavior is that source volumes, on which the database copy or Exchange Replication Service resides, are in the suspended mirror relationship with the ZDB replica, which should be restored to the original database.

  Under such circumstances, the restore process first unpresents the source volumes from the Exchange Server database copy, and presents them to the Exchange Server original database afterwards. These actions result in shortage of source disks presented to the database copy. After restore, you need to set up your LCR or CCR environment again.

  To enable such restore process, set the `omnirc` option `OB2VSS_FORCE_INSTANT_RECOVERY` to `1`, perform restore, and after it, manually set up the LCR or CCR environment.

- With Microsoft Exchange Server 2007 running in a CCR environment, whose database copy resides on the backup system, there are particular situations where you need to consider additional steps for performing instant recovery:

  - When P9000 XP Array VSS hardware provider is used, and the application system, where the production database resides, and the backup system, where the database copy resides, are connected to different disks arrays that are not configured in the same

SAN. In such a configuration, the application system does not see LUNs on the backup system and the other way round.

- ○ When P6000 EVA Array VSS hardware provider is used and the disk arrays of the application system and the backup system are not controlled by the same Command View.

- ○ When the P9000 XP Array VSS hardware provider is used in the resync mode, regardless of the number of disk arrays are used.

In the above situations, there is only one instant recovery scenario. For a successful instant recovery, follow the steps:

1. Fail over Exchange Server. This will cause your production database to reside on the former backup system, where the database copy resided before failover.
2. Perform instant recovery.
3. After the instant recovery session completes, fail back Exchange Server.

   This action makes all backup LUNs available on the same backup system where the database copy resides.

NOTE: It is recommended that both production and replication server systems are using disks of the same disk array.

For additional restore limitations, which apply also to instant recovery, see "Limitations" (page 84).

Restoring an LCR or CCR copy to the original location

See "Restoring a LCR or CCR copy to the original location" (page 81).

Post–instant recovery steps

1. Manually re-mount the database stores. In case of restore to Recovery Storage Group, re-mount the database stores in this RSG.
2. If you restored an LCR or CCR copy to the original database, perform additional steps:

   - In case of an LCR restore, it is recommended to *seed* the restored database to synchronize the original database with its copy.

     For more information, see the webpage http://technet.microsoft.com/en-us/library/aa995973.aspx.

   - In case of a CCR restore, you may need to perform additional steps.

△ **CAUTION:** Do not move the clustered mailbox server to the other node (using the `Move-ClusteredMailboxServer` cmdlet) without first completing the procedure below. If you move the server at this point, data loss may occur.

Perform the following steps to enable normal operation of the original and copy databases:

   a. On the application system, mount the restored store.
   b. On the passive node, where the copy of the database exists, delete transaction logs.
   c. On the passive node, use the `Update-StorageGroupCopy` cmdlet to seed the storage group copy or to re-synchronize the original storage group and its copy.
   d. On the passive node, use the `Resume-StorageGroupCopy` cmdlet to resume the storage group copy.

Database recovery

You can run a database recovery from the **Instant Recovery** context of the Data Protector GUI. This option is available, if you have created a separate backup specification for **Incremental**/**Differential** backup with the same object and description as you have in the backup specification for instant

recovery. Such an **Incremental/Differential** backup is based on the **Full** backup with the selected instant recovery option. You can select an **Incremental/Differential** backup in the **Instant Recovery** context and start restore. Instant recovery will be performed and transaction logs will be automatically applied to the recovered storage group.

Troubleshooting

For general VSS troubleshooting, see "Troubleshooting" (page 102).

Problem

Backup session waits 10 minutes to finish

By default, Data Protector waits 600 seconds for Microsoft Exchange Server 2007 writer to stabilize.

Action

Although this stabilization delay is recommended by Microsoft, you can change the waiting period by setting the OB2VSS_EXCHANGE_WRITER_STABILIZATION omnirc option. Specify the waiting period in seconds.

For details on how to set the option, see the *HP Data Protector Help* index: "omnirc options".

Microsoft Exchange Server 2010 writer specifics

Introduction

This section describes Microsoft Exchange Server 2010 writer specifics, when the Data Protector Microsoft Shadow Copy Service integration is used.

NOTE: It is recommended that you use the *Data Protector Microsoft Exchange Server 2010 integration* instead of the generic VSS integration.

The Data Protector Microsoft Exchange Server 2010 integration offers additional functionality and simplifies backup configuration, enabling you to back up multiple database copies (active and passive) from different system in DAG in one session, and so on. See the *HP Data Protector Integration Guide for Microsoft Applications: SQL Server, SharePoint Server, and Exchange Server* and the *HP Data Protector Zero Downtime Backup Integration Guide*.

Microsoft Exchange Server 2010 concepts

Microsoft Exchange Server 2010 introduces the Database Availability Group (DAG), which is an evolution of the Exchange Server 2007 SCR and CCR concept. Each DAG can consist of up to 16 systems, which can host multiple active and/or passive database copies. Passive copies are kept consistent with the active copy and the DAG heals itself automatically if data on a particular node gets corrupted. For example, if a passive copy gets corrupted, the passive copy is re-seeded. If an active copy can no longer be healed, one of the passive copies becomes the active copy.

Both, the active copy and the related passive copies represent the same database and can be exchanged during restore. This means that you can restore an active copy from a passive copy backup or the other way round.

The most common scenarios involving Data Protector would therefore include (but are not limited to):

- backing up passive copies thus avoiding additional loads on the active copy during backup

- restoring a lost or corrupt passive copy, avoiding large volumes of network traffic when passive copies are re-seeded

- restoring the database to a point in time (for example for investigation purposes)

See "Restore scenarios" (page 88) for additional restore scenarios.

Integrating Data Protector Microsoft Volume Shadow Copy Service integration with Exchange Server 2010

Using this integration, you can back up the Microsoft Exchange Server 2010 writers. Two writers are used for Exchange Server 2010 backup and restore:

- Microsoft Exchange Writer (used for active copies)
- Microsoft Exchange Replica Writer (used for passive copies)

Using the Exchange writers, only databases from a physical node (system) in a DAG can be backed up. The DAG—which is a virtual entity and not recognized by Data Protector—cannot be backed up.

Configuration

Prerequisites

- The `Data Protector MS Volume Shadow Copy Service Integration` component must be installed on all DAG nodes (systems where Microsoft Exchange Server 2010 is installed) which you want to back up.

Licensing

As you need to back up individual clients, you need one online-extension license for each system on which the Data Protector `MS Volume Shadow Copy Integration` component is installed.

Configuring

To configure the integration for ZDB and ensure that disk resolving works, it is recommended that you run the following command on the clients that you will back up from:

`omnidbvss –resolve –apphost ClientName`

where `ClientName` is the name of the Microsoft Exchange Server 2010 system, for example `server1.company.com`. For details, see the `omnidbvss` reference page.

If the command is not run in advance, Data Protector will automatically run this command during the first ZDB session which may slow down the backup session.

Backup

With the Data Protector Microsoft Volume Shadow Copy Service integration you can back up the Exchange writer or the Exchange Replica Writer on a physical client or individual Mailbox Database copies (either active or passive). You cannot select the database files or logs of a Mailbox Database individually.

The following Microsoft Exchange Server 2010 backup types are supported:

- Full
- Incremental (only with the software provider)
- Differential (only with the software provider)
- Copy

Limitations

- From the GUI, the additional options Perform consistency check and Throttle check for 1 second cannot be set for the Microsoft Exchange Replica Writer. The Perform consistency check option is enabled by default.
- For ZDB, the Strict and Non-strict configuration check modes are not supported. Select Disabled in the Configuration check mode drop-down list to disable the check. When this option is

disabled, Data Protector does not check if there are files on the source volume that were not selected for backup or instant recovery.

- You cannot combine a backup chain from full and incremental backups from different Microsoft Exchange servers. You can also not combine a restore chain from passive and active copies.

- You must create a separate incremental backup specification for each backup specification that will be used for instant recovery (has the **Track the replica for instant recovery** option enabled).

  For example, you *cannot* create the backup specification BSpec1 for the database MailboxDatabase1 with snapshots as the backup type and the backup specification BSpec2 for the storage group MailboxDatabase2 with snapclone as backup type and then use only one incremental backup specification containing both databases.

Creating a backup specification

To create a backup specification, follow the procedure described in .

When selecting the backup object (database), consider the following:

- To create a backup specification for an active copy, select the **Microsoft Exchange Writer**.

- To create a backup specification for a passive copy, select the **Microsoft Exchange Replica Writer**.

(!) **IMPORTANT:** After a failover, a passive copy becomes an active copy and an active copy becomes passive. In such a case, the database backup object will fail, as the backup specification is not automatically updated and Data Protector tries to back up an active copy using the Exchange Replica writer which is intended for passive copies. You must manually update the backup specification and select the correct writer.

Restore

Follow the standard restore procedure as described in . The restore scenarios in this section describe only Microsoft Exchange Server 2010 specifics.

Restore scenarios

Restoring a passive copy

You can restore a passive copy from a chain of active or passive copy backups.

Restore when both the database files and logs are lost

1. Suspend the database replication using the Exchange Management Console.
2. Right-click the database that you want to restore and open the Properties window. In the Maintenance page, select **This database can be overwritten by a restore**.

Figure 28 Enabling restore for a database
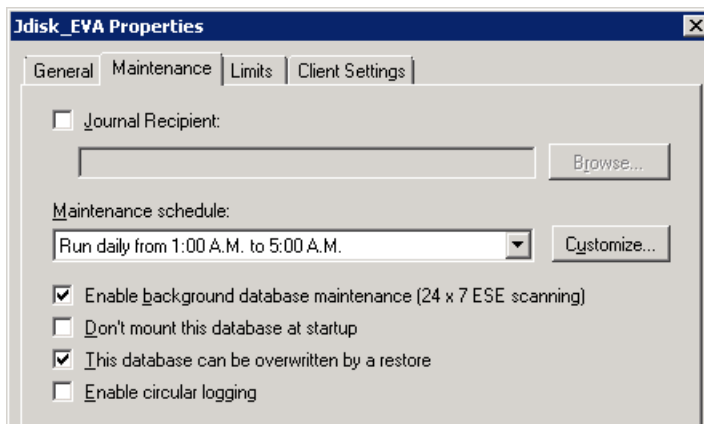
3. In the Data Protector GUI, select the **Restore** context. In the Scoping Pane click **MS Volume Shadow Copy** and select the client from which the database was backed up. You can select any client, not only the one to which you restore.
4. Depending on whether you backed up an active or passive copy, expand Microsoft Exchange Writer or Microsoft Exchange Replica Writer and select the backed up database copy.

Figure 29 Selecting a whole database
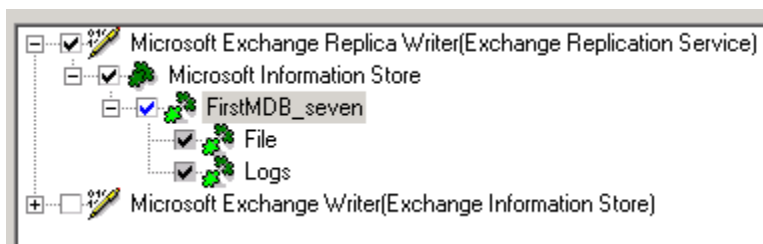
To perform a restore to a different client than the one from which the backup was performed, go to the Options page and select the client from the **Restore to another client** drop-down list.

Figure 30 Restore to another client

5. Specify the device and media options. For details, press **F1**.
6. Start the restore.
7. Resume the replication.

Restore when only database files are lost

If only database files are lost, but the logs are still available on the system, you can restore only the database files:

1. Suspend the database replication using the Exchange Management Console.

2. Right-click the database that you want to restore and open the Properties window. In the **Maintenance** tab, select **This database can be overwritten by a restore**.

3. In the Data Protector GUI, select the **Restore** context. In the scoping pane, expand **MS Volume Shadow Copy** and select the client from which the database was backed up. You can select any client, not only the one to which you restore.

4. Depending on whether you backed up an active or passive copy, expand **Microsoft Exchange Writer** or **Microsoft Exchange Replica Writer** and expand the database backup copy.

   Select **File**. See "Selecting the File component" (page 90).

Figure 31 Selecting the File component

To perform a restore to a different client than the one from which the backup was performed, go to the Options page and select the destination client from the **Restore to another client** drop-down list. See "Restore to another client" (page 89).

5. Specify the device and media options. For details, press **F1**.
6. Start the restore.
7. Resume the replication

Restoring an active copy

You can restore an active copy from a chain of full and incremental backups of an active or any passive copy.

⚠ **IMPORTANT:** As the passive copy is not always immediately updated, a point-in-time recovery (database and transaction logs files are restored) using such a backup image may not result in the last state of the active copy as it was at the time of the backup.

Restore when both the database files and logs are lost

1. Dismount the database using the Exchange Management Console.
2. Suspend the database replication. In the Exchange management console, right-click the database that you will restore and open the Properties window. In the Maintenance page, select **This database can be overwritten by a restore**. See "Enabling restore for a database" (page 89)
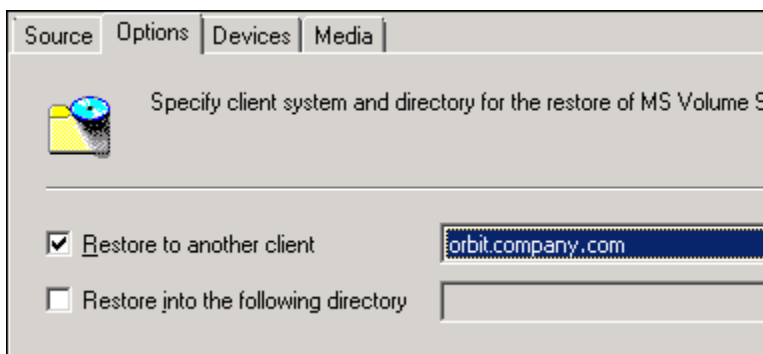3. In the Data Protector GUI, select the **Restore** context. In the scoping pane, expand MS Volume Shadow Copy and select the client from which the database was backed up. You can select any client, not only the one to which you restore.
4. Depending on whether you backed up an active or passive copy, expand the **Microsoft Exchange Writer or Microsoft Exchange Replica Writer** and select the database.

   If the system you backed up from currently hosts the active copy, the backed up copy will be automatically restored as an active copy.

   If the system you backed up from currently hosts a passive copy (for example, if a failover occurred), select the client which hosts the current active copy from the **Restore to another client** drop-down list. See "Restore to another client" (page 89).

5. Specify the device and media options. For details, press **F1**.
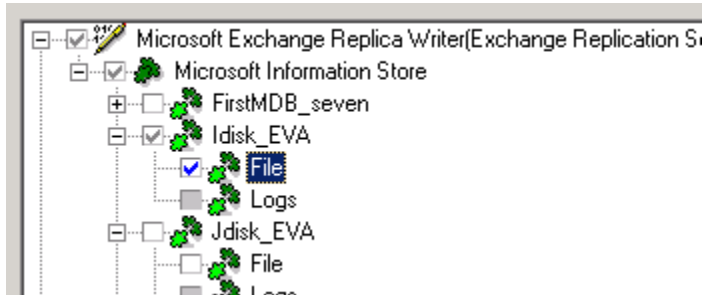
6. Start the restore.
7. Mount the database using the Exchange Management Console.
8. Resume the replication.

Restore when only database files are lost

1. Dismount the database using the Exchange Management Console.
2. Suspend the database replication. In the Exchange management console, right-click the database that you will restore and open the Properties window. In the Maintenance page, select **This database can be overwritten by a restore**. See Figure 28 (page 89).
3. In the Data Protector GUI, select the **Restore** context. In the scoping pane, expand MS Volume Shadow Copy and select the client from which the database was backed up. You can select any client, not only the one to which you restore.
4. Depending on whether you backed up an active or passive copy, expand Microsoft Exchange Writer or Microsoft Exchange Replica Writer and expand the backed up database.
5. Select **File**.

   If the system you backed up from currently hosts the active copy, the backed up copy will be automatically restored as an active copy.

   If the system you backed up from currently hosts a passive copy (for example, if a failover occurred), select the client which hosts the current active copy from the **Restore to another client** drop-down list. See "Restore to another client" (page 89).
6. Specify the device and media options. For details, press **F1**.
7. Start the restore.
8. Resume the replication.

Point-in-time restore

For a point-in-time restore, both active and the passive copies must be restored to the same point in time or you must manually perform a full re-seed for passive copies. To avoid a full re-seed of passive copies, you must restore the active and passive copies from the same backup copy.

1. Dismount the database using the Exchange Management Console.
2. Right-click the database that you want to restore and open the Properties window. In the Maintenance page, select **This database can be overwritten by a restore**. See "Enabling restore for a database" (page 89)
3. Suspend the database replication.
4. Restore the active copy. See "Restoring an active copy" (page 90).

   Then select the database backup copy to be used for restore, right-click the database and select Properties. Select the point in time to which to restore the active copy.

Figure 32 Selecting a point in time
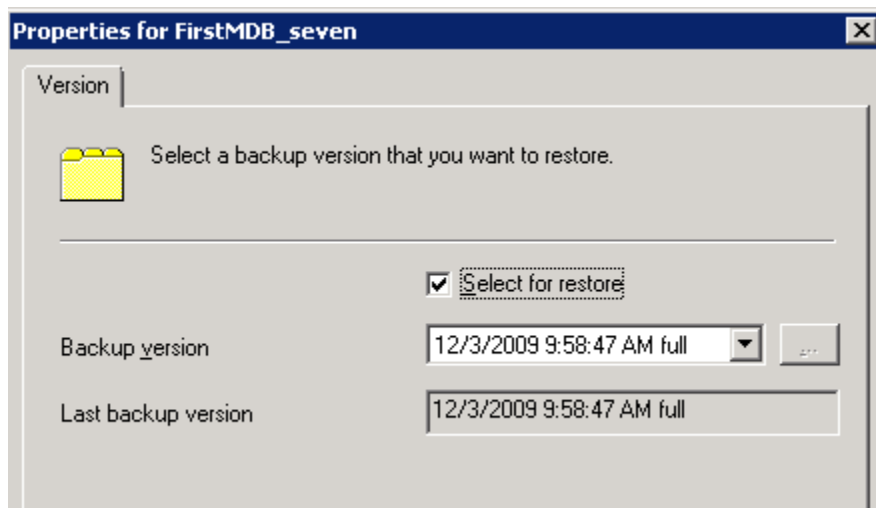
5. If you do not want to perform a full re-seed of the passive copies, restore the passive copies. See "Restoring a passive copy" (page 88).

   When selecting the database backup copy to be used for restore, select the same database backup copy that was used for the active copy, right-click the database, and select **Properties**. Select the same point in time as was selected for the active copy.

6. Mount the active database using the Exchange Management Console.

7. Resume the replication.

Instant recovery

Use the standard instant recovery procedure as described in "Backup" (page 32). For general prerequisites and conditions (for example dismounting the Microsoft Exchange Server databases, suspending the replication), see "Restore scenarios" (page 88). Additional instant recovery limitations are listed in this section.

Limitations

- You can perform instant recovery only on the client where the ZDB was created. The only exception is when performing an instant recovery of a P6000 EVA Array snapclone, which can be restored to a different client using the Data Protector CLI.

- Instant recovery of a P6000 EVA Array snapclone to a different client is not supported through the Data Protector GUI.

- With the HP P9000 XP Disk Array Family provider in the resync mode or with P6000 EVA storage systems, each connected to its own system, you cannot perform a point-in-time recovery from a single ZDB session. To be able to recover both an active copy and passive copies to the same point in time, you need multiple ZDB sessions that were performed at the same time, one for the active and one for each passive copy.

Instant recovery of a P6000 EVA Array snapclone to a different client using the CLI

1. Use the `omnidbvss` command to get details about the session that you want to use for instant recovery. For example:

```
omnidbvss -get session 2009/12/16-10
================================================================
Session ID:            2009/12/16-10
Barlist Name:          Xdisk_XPcompliant_trans
Bar Hostname:          server1.company.com
Backup Type:           FULL
Instant Restore:       TRUE
```

```
Disk-Only:              FALSE

Component Name
==================================================================
[0]     /Microsoft Exchange Writer(Exchange Information Store)/
Microsoft Information Store/Xdisk_XPbox/Logs
[1]     /Microsoft Exchange Writer(Exchange Information Store)/
Microsoft Information Store/Xdisk_XPbox/File
```

The objects that were backed up in this session are listed under `Component Name`.

See the `omndbvss` reference page in the *HP Data Protector Command Line Interface Reference* for details on the `omnidbvss` command.

2. Run the omnir command:

```
omnir -vss -instant_restore -barhost ClientName1 -session SessionID
-tree TreeName1 [-tree TreeName2...] -destination ClientName2
[VSS_INSTANT_RECOVERY_OPTIONS]
```

See the `omnir` reference page in the *HP Data Protector Command Line Interface Reference* for details on the `omnir` command.

To recover the components from the above example, run:

```
omnir -vss -instant_restore -barhost server1.company.com -session
2009/12/16-10 -tree /Microsoft Exchange Writer(Exchange Information
Store)/Microsoft Information Store/Xdisk_XPbox/Logs -tree /Microsoft
Exchange Writer(Exchange Information Store)/Microsoft Information
Store/Xdisk_XPbox/File -destination server2.company.com -conf_check
disabled
```

Troubleshooting

For general VSS troubleshooting, see "Troubleshooting" (page 102).

Problem

After a failover in a DAG occurs, Data Protector reports that some Microsoft Exchange writer components could not be found:

```
[Major] From: OB2BAR_VSSBAR@server5.company.com "MSVSSW" Time: 12/7/2009
1:16:40 PM
```

```
Failed to find component that would match tree:
```

```
'/Microsoft Exchange Replica Writer(Exchange Replication Service)'.
```

Action

You can solve the issue in two different ways:

- Update the backup specifications to use the correct writers.
- Perform a failover so that the database returns to the original state (active or passive), the issue will no longer be reported.

Microsoft Hyper-V writer specifics

Concepts

Microsoft Hyper-V writer is the successor to the Microsoft Virtual Server 2005 writer and is supported on Microsoft Windows Server 2008. It is a VSS writer with similar set of functionality as Virtual Server. With both, it is possible to perform backups and restores of virtual machines. With Hyper-V, it is possible to perform online backups using hardware providers.

For backup and restore specifics of Hyper-V writer, see "Backup" (page 94) and "Restore" (page 96).

Prerequisites

- Windows Server 2008 Service Pack 2 must be installed. For more information, see http://support.microsoft.com/kb/948465.

Backup

With the Hyper-V VSS writer, it is possible to back up:

- The Hyper-V configuration
- Virtual machines

Figure 33 Selecting Microsoft Hyper–V VSS writer backup objects
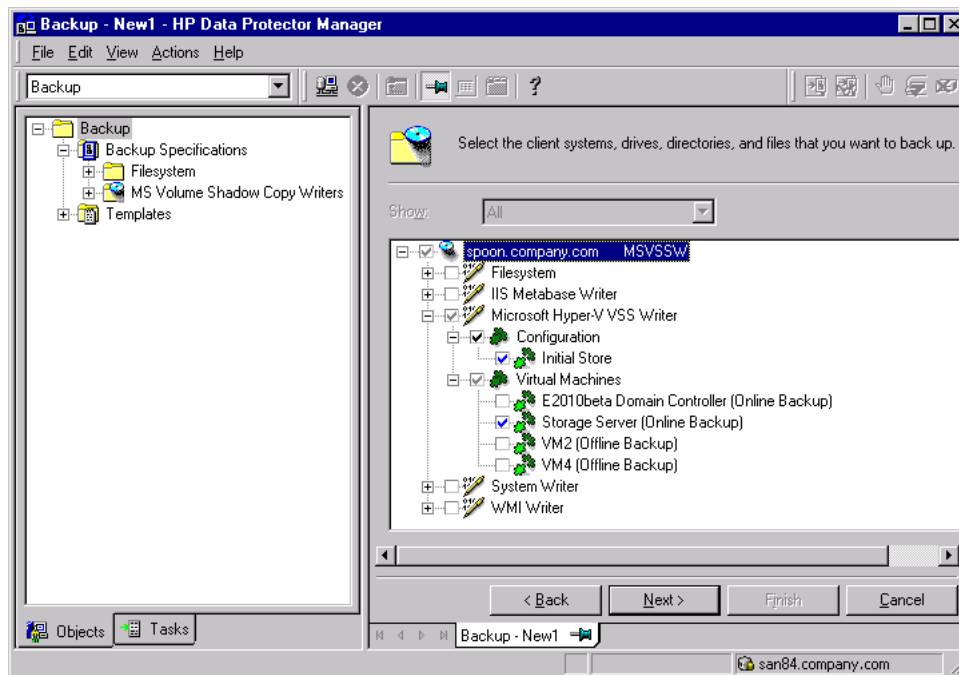
The following two types of backups are supported by the Data Protector VSS integration:

- *Online backup*

  Online backup of Hyper-V writer data is possible using a software provider or hardware provider.

  If using a hardware provider for an online backup process, Hyper-V writer creates a shadow copy, which is replicated together with the *vhd* file via the VSS hardware provider. Afterwards, the shadow copy is presented to the Hypervisor System. In the case of transportable backup, the shadow copy is afterwards unpresented from the Hypervisor System and presented to the backup system, specified in the backup specification.

  After restore of an online backup session, the virtual machine is always turned off, regardless of the state in which it was before the restore.

- *Offline backup*

Offline backup is performed in the following cases:

- The guest operating system – any other non-Microsoft guest operating system that is supported by Hyper-V – is not VSS enabled.

- The guest operating system does not have Hyper-V VSS integration services installed.

- The virtual machine to be backed up is turned off.

Before offline backup, the virtual machine is automatically suspended (if not already) and resumed after the backup.

After restore of an offline backup session, the virtual machine is suspended, regardless of the state it was in before the restore.

The advantage of offline backup is that virtual machines are restored to the state they were in at backup time, including the state of applications running at the time of backup. This is possible because a virtual machine is in a suspended state during the backup.

Prerequisites

- For online backup:

  - The guest operating system must have the Hyper-V VSS integration services installed and should not use dynamics disks.

  - The snapshot file (avhd file) needs to be configured on the same volume as the virtual disk file (vhd file).

  - The virtual machine to be backed up must be online.

  - Automatic mounting of new volumes must be enabled on the host (hypervisor) system. To enable automatic mounting, run MOUNTVOL /E on the hypervisor system.

- For offline backup, if the prerequisites for online backup are met, a virtual machine to be backed up must be put in the offline or suspended state manually.

Limitations

- Cluster-aware backups are not supported.

- ZDB to disk, ZDB to disk+tape, and instant recovery of a cluster node is not supported.

- Only whole virtual machines can be backed up or restored. However, with ZDB to disk and ZDB to disk + tape only a whole source volume can be backed up, regardless of the number of virtual machines residing on the volume.

- Backup of virtual machines configured to use physical disks is not supported.

- Backup of dismounted disks is not supported. A possible workaround is to configure mount points for these volumes as described in: support.microsoft.com/kb/947021.

- ZDB to disk, ZDB to disk+tape, and instant recovery of online virtual machines are not supported for HP P4000 SAN Solutions.

Backup from a physical cluster node

When backing up from a cluster node consider the following:

- Offline backups trigger a failover.

  During offline backups the virtual machine to be backed up is suspended for a moment. The cluster server recognizes this as a failure and initiates the cluster failover. To avoid such a failover, perform one of the following:

  ○ Run only online backups.

  ○ Before running an offline backup, manually put the virtual machine into the "Saved" state using Failover Cluster Administrator.

- Whenever a failover happens, you need to use another backup specification.

  After a failover, the hostname where the virtual machine is running changes. Since this change is not reflected in the original backup specification, you need to create a new backup specification to specify the new hostname as the application system name.

Restore

Restore of Hyper-V writer data is possible to the original or to a different location. During a restore to a different location, the Hyper-V writer checks whether a virtual machine with the same identity already exists on the system. If it does, the Hyper-V writer removes the virtual machine from the system before restore and imports the restored virtual machine. If such a virtual machine does not exist on the system, it means that a restore session to this system is in progress or that the virtual machine was already removed. It is possible to perform restore or instant recovery of Hyper-V virtual machines to any Hyper-V system, which has a Hyper-V writer.

Restore from a physical cluster node

ZDB to tape sessions from a physical cluster node can be restored to any cluster node using a standard restore procedure. To successfully restore such a session to a cluster node, perform the following steps:

Backup sessions from a physical cluster node can be restored to any cluster node. To successfully restore such a session to a cluster node, perform the following steps:

1. Delete the cluster group of the virtual machine to be restored using Failover Cluster Management.
2. Perform standard restore of the virtual machine on the cluster node where the virtual machine disk is active. For the standard restore procedure, see "Standard restore" (page 46).
3. Re-create the virtual machine cluster group using Failover Cluster Management.

Troubleshooting

Problem

Backup session of a Microsoft Hyper-V virtual machine ends unexpectedly

When backing up a Microsoft Hyper-V virtual machine, the session ends unexpectedly with an error similar to the following:

```
[Major] From: OB2BAR_VSSBAR@computer.company.com "MSVSSW"
Time: 2/1/2011 11:29:03 AM [145:575]
Writer 'Microsoft Hyper-V VSS Writer' failed to prepare files
for backup:
        Reported state:        VSS_WS_FAILED_AT_POST_SNAPSHOT
        Expected state:        VSS_WS_WAITING_FOR_BACKUP_COMPLETE
        Failure code:          VSS_E_WRITERERROR_NONRETRYABLE
```

The following are possible causes:

- Automatic mounting is disabled on the host (hypervisor) system. For details, see http://support.microsoft.com/kb/2004712.
- There are issues inside the virtual machine, such as not enough free disk space for shadow copies, use of a non-NTFS filesystem, and so on.

Action

Check if automatic mounting is enabled on the host hypervisor system. For example:

```
diskpart.exe
Microsoft DiskPart version 6.1.7600
Copyright (C) 1999-2008 Microsoft Corporation.
On computer: TPC021

DISKPART> automount
Automatic mounting of new volumes enabled.
```

If automatic mounting is disabled, enable it by running the following command:

```
MOUNTVOL /E
```

If the issue persists even though automatic mounting is enabled, check the application logs inside the virtual machine to determine the cause.

Microsoft SharePoint Services writer specifics

Concepts

The Microsoft SharePoint Services writer is a *reference* writer that integrates with the Windows VSS framework, allowing backup applications to back up and restore Microsoft SharePoint data. This writer has dependencies on:

- Search writers:
  - OSearch VSS writer
  - SPSearch VSS writer
- SQL writers:
  - MSDE writer for Microsoft SQL Server 2000
  - Microsoft SQL 2005 writer for Microsoft SQL Server 2005

With the SharePoint Services writer you can back up and restore:

- the configuration database
- the central administration content database
- other content databases
- shared services provider databases
- search databases
- index files.

Backup types

The SharePoint Services writer supports the following Microsoft Office SharePoint Server 2007 backup type:

- Full (for databases and index files)

Limitations

- Multi-server SharePoint configurations (farms) are not supported.

Prerequisites

- The SharePoint Services writer and SQL writer are not started by default. Ensure that both writers are properly installed and registered. The SharePoint Services writer must be registered using the SharePoint command line administration tool:

  ```
  stsadm -o  registerwsswriter
  ```

Backup

Limitations

- If you back up individual writer components, to ensure that data is synchronized, the following combinations of items must be backed up in the same session:

  - The configuration database and central administration content database

  - Search databases and corresponding index files

  This means that, for example, you should not back up just a search database. Instead, always *select also* the corresponding index file.

- Do not use the *standalone* (not part of the reference) OSearch and SPSearch writers for backing up the index files. If you use them, the search index will have to be reindexed.

See "Selecting Microsoft SharePoint Services writer and the corresponding search writers" (page 99) for an example of how to select writer components.

Figure 34 Selecting Microsoft SharePoint Services writer and the corresponding search writers
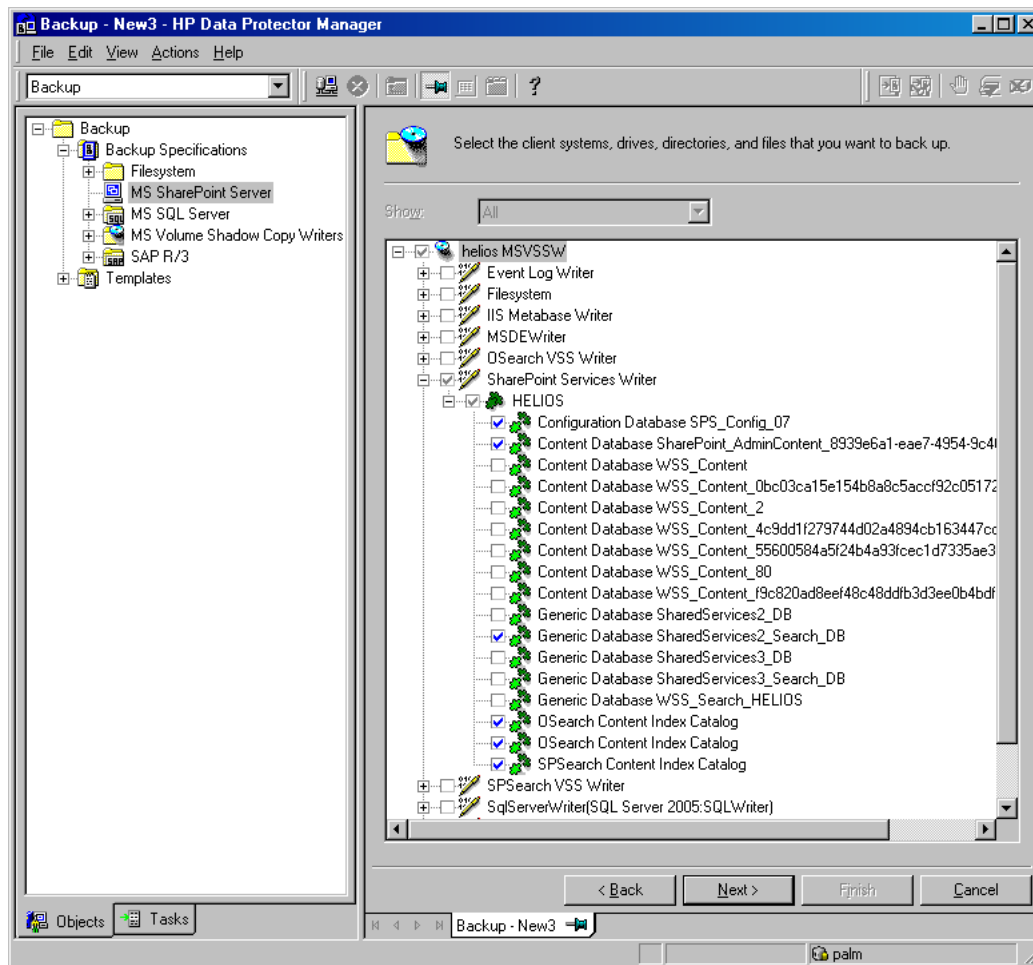
Restore

To ensure that data is synchronized the following items must be restored in the same session:

- Configuration database and central administration content database
- Search databases and corresponding index files

This means that, for example, you should not restore just a search database. Instead, always *select also* the corresponding index file.

See "Selecting Microsoft SharePoint Services writer and the corresponding search writers for restore" (page 100).

Because the configuration database and the central administration content database contain system-specific information, you can restore them only to an environment that you configure to be precisely the same, including all software updates, server names, and number of servers.

Prerequisites

- Before performing a restore, stop the following services:
  - Windows SharePoint Services Administration
  - Windows SharePoint Services Search

- ◦ Windows SharePoint Services Timer
- ◦ Office SharePoint Server Search
- If you restore the whole farm, you must shut down the Internet Information Server (IIS).

Limitations

- The VSS restore mode **Restore files to temporary location** is not supported.
- The VSS restore option **Restore to another client** is not supported.

Figure 35 Selecting Microsoft SharePoint Services writer and the corresponding search writers for restore
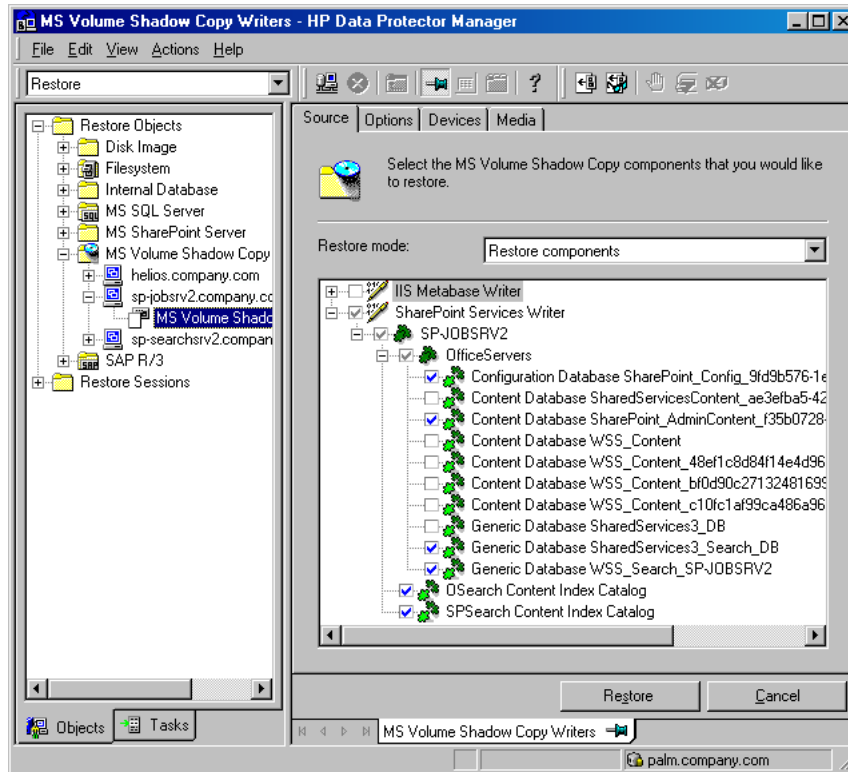
MSDE writer specifics

Restore

MSDE writer is used to back up and restore Microsoft SQL database.

ⓘ **IMPORTANT:** Before restoring the SQL system databases (`master`, `model`, `msdb` and `pub`), you have to stop the SQL service.

Figure 36 MSDE writer
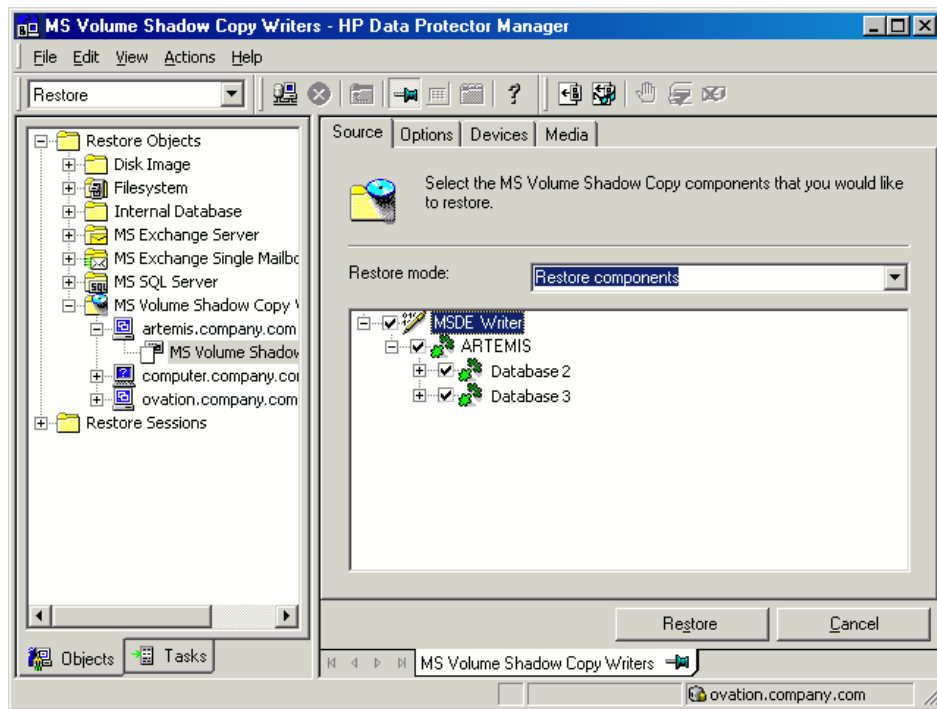
When you expand the MSDE Writer item in the Results Area, all Microsoft SQL Server instances are displayed. Each instance contains all databases it includes. System databases (`master`, `model`, `msdb` and `pub`) are always listed there.

ⓘ **IMPORTANT:** If system databases are restored, the whole internal database structure will be changed.

NOTE: Only point-in-time restore is possible. Rollforward restore is not supported.

User databases will be restored only if it is possible to overwrite the files. MSDE writer will take the user databases offline before the restore, while SQL service will have to be stopped manually in order to restore the system databases.

7 Troubleshooting

This chapter lists problems you might encounter when using the Data Protector Microsoft Volume Shadow Copy Service integration.

For general Data Protector troubleshooting information, see the *HP Data Protector Troubleshooting Guide*.

Before you begin

- Ensure that the latest official Data Protector patches are installed. See the *HP Data Protector Help* index: "patches" for information of how to verify this.
- See the *HP Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See http://support.openview.hp.com/selfsolve/manuals for an up-to-date list of supported versions, platforms, and other information.

Checks and verifications

- On the application and backup systems, examine system errors reported in the `debug.log` file.
- With the P9000 XP Array integration, ensure that RAID Manager Library is correctly installed on both the application and backup systems. In addition, check if the `libsvrrm.dll` file exists in the *RMLIB_home* directory.
- Names of all clients in a Data Protector cell must match DNS entries and the Data Protector `cell_server` file entries. If the names do not match, a warning message is displayed. In this case, do not use the VSS integration as it can behave unpredictably. First reset network settings on the client and then re-import the client.

Problems

Problem

Microsoft Exchange Server 2003 writer backup failed

When you start a backup of a Microsoft Exchange Server 2003 Writer, the following error is displayed:

```
[MAJOR]: Writer 'Microsoft Exchange Writer' failed to prepare files
for backup.
```

The reason may be the failure of the previous backup and, due to a Microsoft Exchange Server 2003 issue, the Exchange Server Writer could not perform a proper cleanup.

Action

Restart the Information Store.

For more information, see the webpage http://support.microsoft.com/kb/945424/en-us.

Problem

Microsoft Exchange Server 2003 aborts the backup if the shadow copy creation takes over 20 seconds

If an Exchange Server 2003 Writer is being backed up, the session can fail with VSSBAR reporting:

```
Snapshot could not be created.
```

In the application event log on the application system, the following event is recorded:

```
Event Type: Error
Event Source: ESE Event
Category: (16)
Event ID: 2004
Information Store (4916) Shadow copy 3 time-out (20000 ms).
```

Action

The following can help to solve the problem:

- Limit the number of users that are accessing the management system.

- Reduce the number of volumes in a snapshot set. Create a backup specification dedicated to each storage group instead of one specification for the whole server. See "Backup" (page 66).

Problem

Backup of a Microsoft Exchange Server 2007 CCR database copy fails

During a backup session of a database copy in a Microsoft Exchange Server CCR environment, Data Protector reports a major error notifying that the backup session has failed.

This problem may occur in CCR environments, due to a Microsoft Exchange Server 2007 issue with incorrect display of database copy states in Exchange Management Console. In such a case, a database copy in a "Failed" state may be displayed as "Healthy".

Action

To make the Exchange Management Console show the real status of a database copy, perform either of the following actions:

- Update Microsoft Exchange Server 2007 with Service Pack 1.
- Perform re-seeding procedure as follows:
  1. On the passive node, suspend the replication by using `Suspend-StorageGroupCopy` cmdlet.
  2. Delete all log files from the `Logs` directory of the database copy.
  3. Seed the database copy or re-synchronize the original database and its copy by using the `Update-StorageGroupCopy` cmdlet.
  4. Resume the database copy by using the `Resume-StorageGroupCopy` cmdlet.
  5. On the passive node, check the status of the Exchange Replication Service by using `vssadmin list writers` command. If the status is not `stable`, restart Microsoft Exchange Replication Service.

Problem

Exchange Replication Service writer instance in LCR environment is not displayed in the Data Protector GUI

In the Data Protector GUI, while creating a backup specification, the `Microsoft Exchange Writer(Exchange Replication Service)` object is not displayed on the pane for selection of backup objects.

This problem may occur in LCR environments, due to a Microsoft Exchange Server 2007 issue with incorrect display of database copy states in Exchange Management Console. In such a case, a database copy in a "Failed" state may be displayed as "Healthy".

Action

To make the Exchange Management Console show the real status of a database copy and to enable selection of the backup object, restart Microsoft Exchange Replication Service.

Problem

After the restore of system writers was aborted, the Windows operating system is corrupted when you restart it

If the restore of some system writers (for example, System Writer) is aborted for any reason (hardware or software failure, manually aborted, and so on), the Windows operating system may be corrupted after the restart (for example, the GUI or some system services cannot be started, and so on).

Action

Depending on the nature of the corruption, repair or re-install the operating system from the Windows installation CD-ROM.

Problem

Data loss during a point-in-time restore of a Microsoft Exchange Server 2003/2007 database

Though a point-in-time restore session finishes successfully, some of the data is not restored, because existing logs may interfere with the mount procedure.

Action

Manually delete the log files before running the point-in-time restore.

Problem

Creation of an RSG fails when an RSG for the same storage group already exists

When you start a standard restore or instant recovery of the Microsoft Exchange Server 2007 writer with the option `Restore to a non-Exchange location and create RSG`, the following error is displayed:

```
[Major]
Application specific function 'PostRestoreEndExt' failed with error:
'The mailbox database that you specified is already associated with
a recovery mailbox database.
```

This error appears if an RSG for a storage group or mailbox database you restore is already created on some other system.

Action

Since only one RSG can exist for the same storage group and since Data Protector can only delete an RSG that exists on the target restore location, you need to manually delete the RSG on the other system and restart the restore.

Problem

VSS writers end up in Failed state after backup

After performing a backup session, VSS application writers may consistently end up in `Failed` state.

For example, during a backup of the Microsoft SQL Server, using the SQL 2005 writer, the writer fails. After restarting the writer, the writer is in good state. However, during the next backup session, the writer again fails.

The issue has been noticed for the SQL 2005 writer but may apply to other writers as well.

Use the `vssadmin list writer` command to check the state of the writers. For example, when you run the command before the backup session:

```
vssadmin list writers
vssadmin 1.1 - Volume Shadow Copy Service administrative command-line
(C) Copyright 2001-2005 Microsoft Corp.

Writer name: 'SqlServerWriter'
```

```
   Writer Id: {a65faa63-5ea8-4ebc-9dbd-a0c4db26912a}
   Writer Instance Id: {c249df3c-f9d5-4f4a-8797-03724a60771c}
   State: [1] Stable
   Last error: No error
```

After the backup session:

```
C:\Program Files>vssadmin list writers
vssadmin 1.1 - Volume Shadow Copy Service administrative command-line
(C) Copyright 2001-2005 Microsoft Corp.

Writer name: 'SqlServerWriter'
   Writer Id: {a65faa63-5ea8-4ebc-9dbd-a0c4db26912a}
   Writer Instance Id: {e3d115c5-9c1a-47f6-8404-bfdba4c37890}
   State: [8] Failed
   Last error: Non-retryable error
```

Action

The reasons for such behavior often include incorrectly configured or expired user accounts, expired passwords, and other connectivity issues between the application and the Volume Shadow Copy Service. Ensure that your accounts are set up properly.

For Microsoft SQL Server, see the KB article http://support.microsoft.com/kb/919023 for information of how to resolve the Microsoft SQL Server and VSS connectivity issues.

Problem

Backup or restore fail when Remote BLOB Storage (RBS) with the FILESTREAM provider is used for SQL Server 2008

A backup session ends unexpectedly with the following error message:

```
[Major] From: OB2BAR_VSSBAR@computer.company.com
<mailto:OB2BAR_VSSBAR@computer.company.com>
"MSVSSW"  Time: 2/3/2011 3:42:06 PM
[145:575]   Writer 'SqlServerWriter' failed to prepare files for
backup:
        Reported state:     VSS_WS_FAILED_AT_PREPARE_SNAPSHOT
        Expected state:     VSS_WS_WAITING_FOR_BACKUP_COMPLETE
        Failure code:       VSS_E_WRITERERROR_NONRETRYABLE

[Major] From: OB2BAR_VSSBAR@computer.company.com
<mailto:OB2BAR_VSSBAR@computer.company.com>
"MSVSSW"  Time: 2/3/2011 3:42:06 PM
   Cannot perform backup of:
   '/SqlServerWriter(SQL Server 2008:SQLWriter)/BELMAVM20/SHAREPOINT/
FileStreamDB', which contains data in:
   C:\temp\FileStreamDB\FileStreamDB.mdf
   C:\temp\FileStreamDB\FileStreamDB_log.ldf
   C:\temp\FileStreamDB\FileStreamData\*
```

Similarly, a restore session ends unexpectedly with the following error message:

```
[Major] From: OB2BAR_VSSBAR@computer.company.com
<mailto:OB2BAR_VSSBAR@computer.company.com>
"MSVSSW"  Time: 12/8/2010 3:23:16 PM
[145:298]   Writer 'SqlServerWriter(SQL Server 2008 R2:SQLWriter)'
 failed to prepare files for restore:
              Reported state:     VSS_WS_FAILED_AT_PRE_RESTORE
              Expected state:     VSS_WS_STABLE
              Failure code:       VSS_E_WRITERERROR_NONRETRYABLE
```

The issues appear if FILESTREAM access level is set to `Disabled`.

Action

Ensure that FILESTREAM access level is set to `Full access enabled` or `Transact-SQL access enabled`.

For details of how to configure RBS and FILESTREAM, see the Microsoft SQL Server 2008 documentation.

Problem

Microsoft Exchange Server restore or instant recovery fails

A Microsoft Exchange Server restore or instant recovery session fails with a message similar to the following:

```
[Major] From: OB2BAR_VSSBAR@tpc202.company.com "MSVSSW"
 Time: 19.02.2011 21:02:37
              Post Restore for backup '2011/02/19-1' failed.
```

This may happen if a recovery of a Microsoft Exchange Server database lasts longer than the timeout for post-restore operations (default is two hours). When the timeout is reached, Data Protector aborts the session.

Action

Increase the timeout using the OB2VSS_WAIT_TIMEOUT omnirc option and restart the session. For details on how to set the option, see the *HP Data Protector Help* index: "omnirc options".

Problem

The VSS integration can use only 5 concurrent threads for backup or restore

During a backup or restore session, the VSS integration always uses only 5 concurrent threads even if both the device and the application system are capable of handling more threads at once. This limitations appears regardless of device concurrency settings.

Action

Modify the limit by setting the omnirc option OB2VSS_MAX_CONCURRENT_WORKER_THREADS to a higher number. The maximum number of concurrent threads is 64.

ZDB related problems

Problem

Backup or instant recovery aborts due to VDS problems

A backup or instant recovery session aborts with the following error:

```
Failed to load VDS service.
```

This error can appear as a result of abnormal termination of a VDS service.

Actions

1. Stop the Virtual Disk Service (VDS):

   - Check that Virtual Disk Service is started, stop it using the command `net stop vds` or via the Control Panel.

   - If the above step does not help, stop the Virtual Disk Service by terminating the process `vds.exe` using the Task Manager. VDS will be started automatically as needed.

   - Alternatively, you could also log in to the system again and check if the system requests your confirmation to stop an abnormally terminated VDS. If it does, reconfigure the debugger to launch automatically to avoid similar problems in the future. This can be done by setting the `HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\AeDebug` value `Auto` to `1`.

2. Check your configuration.

Successful Data Protector backup or restore completion is influenced by a number of factors imposed by non-Data Protector components, including Microsoft Exchange Server and VSS. One of the factors is timing, which in turn is impacted by environmental conditions such as I/O loads, storage device activity and concurrent Data Protector actions.

To inhibit potential issues, it is recommended that you balance Data Protector activity across the available activity or maintenance window. If you encounter issues, *review and possibly reduce concurrent Data Protector activities* such as concurrent backups or restores.

Problem

Snapshots (volume shadow copies) creation fails with P6000 EVA Array provider

During a VSS backup, the following error is displayed:

```
[Critical] It was not possible to create Volume Shadow Copies.
```

If there are too many source volumes selected for backup (more than four) it may happen that the creation of snapshots takes too long and VSS fails the snapshots creation process.

Action

Do not select too many objects in a backup specification because they might reside on more than four source volumes. This means that for example a Microsoft Exchange Server storage group should not reside on more than four source volumes and only one storage group should be specified for backup in this case. The same applies to a Microsoft SQL Server database.

Problem

No HP SMI-S P6000 EVA Array provider login entries are configured within SMISDB

Action

Add the login information for HP SMI-S P6000 EVA Array provider by executing:

```
omnidbsmis -ompasswd -add ClientName [-ssl] [-port PortNumber] [-user
Username] [-passwd Password]
```

Problem

Configuration of an HP SMI-S P6000 EVA CIMOM failed

Action

Execute the following command to perform a health check of your environment, which may help identify any potential problems that occurred during the configuration:

```
omnidbsmis -ompasswd -check [-host ClientName]
```

Problem

Volume shadow copies cannot be imported due to low space in registry and consequently backup fails
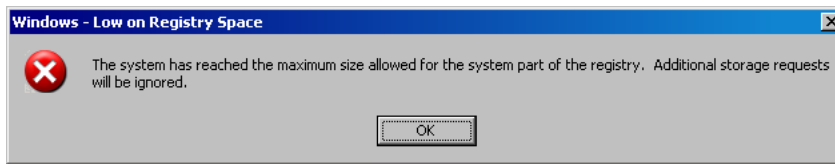
On a Windows system, one of the following errors is reported during a ZDB session:

- By Data Protector:

  ```
  [Critical] It was not possible to import Volume Shadow Copies from
  application host.
  ```

- By the operating system:

Figure 37 Error reported by the Windows Registry

These symptoms indicate that Windows Registry contains too many entries and ran out of space.

Such circumstances result in Data Protector ZDB session failures. Although Data Protector does not write anything to the Registry directly, the Registry contains entries for all volumes ever presented to the system. These entries are written by the storage device driver.

Action

Clean the Windows Registry up.

Problem

Data Protector reports that a volume was not removed

During a backup session, the following error is reported:

```
Volume 'StorageID', which is part of backup 'backupID', was not removed.
```

This error may appear if a replica on the disk array is removed, but its entry in the VSS database (VSSDB) remains intact.

This may happen when a backup session fails without creating a replica, but the reference that was created for it could not be removed due to network problems.

Action

To remove the false entry from the VSSDB:
1. In the error messages, find the ID of the session which could not remove the target volume.
2. Execute the following command, where *SessionID* is the session whose target volume(s) could not be removed:

   ```
   omnidbvss -remove session SessionID -reference
   ```

   Upon successful removal, the command should display the confirmation message `Removing references of session SessionID from VSSDB`.

Problem

Instant recovery of the Microsoft Exchange Server Writer fails

This problem may occur if the Microsoft Exchange Writer is not in the stable state. Check this by executing `VSSadmin list writers` from the command prompt.

Action

Bring the Exchange Server Writer to a stable state by restarting the Microsoft Exchange Information Store.

Problem

With a P6000 EVA Array, a copy-back instant recovery session completes with warnings

When you start a copy-back instant recovery session on P6000 EVA Array, the session completes with warnings similar to the following:

```
[Warning] From: SMISA@seven.e2008.company.com "SMISA"
Time: 5/13/2010 11:07:13 AM [236:8040] A target volume is located
in a different folder than the source volume.
     Source storage volume folder: Windows/VSSQA/Pump
     Target storage volume folder: Windows/VSSQA/Seven
```

This warning is displayed when you perform an instant recovery of a storage volume that is located in a different Command View EVA folder than the replica storage volume. As a result, when the instant recovery completes, the newly-restored storage volume is automatically moved to the same folder in which the replica storage volume is located.

Action

Using the Command View EVA, move the newly-restored storage volume back to its original location to avoid inconsistencies with the Data Protector ZDB database.

Problem

After restore, the system restart error is displayed

In some cases, if the HBAs are changed on the application system, the software driver from previous HBAs can cause this error after restore.

Action

Remove older drivers of uninstalled HBAs from the application system.

Problem

After an SQLServer writer instant recovery is restarted, the database cannot be brought online

If an instant recovery of the SQLServer writer is aborted or fails, you can restart the session. However, the SQLServer writer may report an error, stating that the files cannot be prepared during the instant recovery and the session completes with errors.

The issue can appear if the disk was dismounted before the session was aborted and the files are not visible to the writer. As a result, an error is reported and the database cannot be brought online after restore.

Action

To recover the database after such session:
1. Detach from the database.
2. Attach to the database.

Problem

Shadow copy creation fails with the VSS system provider

When using the VSS system provider for a hardware LUN disk, the shadow copy creation fails.

Action

Install the latest vendor drivers and their support packs, including the HBA adapter and MPIO driver.

Problem

During a backup session, the volumes are imported and then immediately deleted

When you perform a backup using a hardware provider, the volumes that are being backed up are imported and then immediately deleted with the following error:

```
[Normal] From: OB2BAR_VSSBAR@comp.company.com "MSVSSW" Time: 9/4/2010
1:37:41 PM Imported Volume Shadow Copy with the properties: ...

[Normal] From: OB2BAR_VSSBAR@comp.company.com "MSVSSW" Time: 9/4/2010
1:37:41 PM Deleting Volume Shadow Copies and releasing the volumes.

[Critical] From: OB2BAR_VSSBAR@comp.company.com "MSVSSW" Time: 9/4/2010
1:37:51 PM Backup failed.
```

Action

Ensure that the client security settings are correct. If your cell is secured, the following systems must be listed in the `allow_hosts` file:

- For a local or network backup, the application system
- For a transportable backup, both, the application and the backup system

For details, see the *HP Data Protector Help* index: "securing, client systems".

Problem

Instant recovery is not possible without P9000 XP Array VDS Hardware Providers

When the P9000 XP Array VSS hardware provider is in VSS compliant mode, the backup session ends abnormally with the following error message:

```
VDS Swap Instant Recovery is not possible without VDS Hardware providers.
Disk backup will be aborted.
```

Action

To be able to perform instant recovery, install the VDS hardware provider:

1. Install the VDS hardware providers on the application and backup system.
2. Resolve the application system:

   ```
   omnidbvss -resolve -apphost appSystem
   ```

3. Restart the session.

Alternatively, if the ZDB backup will not be used for instant recovery, you can set the `omnirc` option `OB2VSS_ALWAYS_ALLOW_DISK_BACKUP_WITHOUT_VDS` to 1.

ⓘ **IMPORTANT:** If you perform a backup without the P9000 XP Array VDS hardware provider, instant recovery is not possible.

Problem

Similarly, an instant recovery session ends abnormally with the message:

```
To perform VDS Swap Instant Recovery, the source and target LUNS need
to be resolved with VDS Hardware Providers.
```

Action

1. Install the VDS hardware providers on the application and backup system.
2. Resolve the application system:

   ```
   omnidbvss -resolve -apphost appSystem
   ```

3. Resolve the backup system:

   ```
   omnidbvss -resolve session SessionID
   ```

4. Restart the session.

Problem

Zero downtime backup sessions fail after updating the 3PAR StoreServ Storage firmware

After updating the firmware on a storage system of the HP 3PAR StoreServ Storage family and invoking a zero downtime backup session on this system, the session fails with an error similar to the following:

```
[Critical] From: SMISA@appsys.company.com "SMISA" Time: 06/15/2012
2:23:28 PM Replicator aborted with exception "No CIMOM found for ID
2FF70002AC000B6C".
```

Action

The cause of the problem is outdated information about the source volumes in the VSS database after the storage system's firmware update. To solve the problem, update the VSS database by executing the following command for the problematic application system:

```
omnidbvss —resolve —apphost ApplicationSystem
```

Glossary

access rights   *See* user rights.

ACSLS   *(StorageTek specific term)* The Automated Cartridge System Library Server (ACSLS) software that manages the Automated Cartridge System (ACS).

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

AES 256-bit encryption   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 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 and 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)* SAP R/3 backup programs can call the Data Protector backint interface program via an open interface, which enables them to communicate with Data Protector software. For backup and restore, SAP R/3 programs issue orders for the Data Protector backint interface.
backup API	The Oracle interface between the Oracle backup/restore utility and the backup/restore media management layer. The interface defines a set of routines to allow the reading and writing of data to the backup media, the creation, searching and removing the backup files.
backup chain	*See* restore chain.
backup device	A device configured for use with Data Protector, which can write data to and read data from storage media. This can be, for example, a standalone DDS/DAT drive or a library.
backup generation	One backup generation includes one full backup and all incremental backups until the next full backup.
backup ID	An identifier of an integration object that equals the session ID of the backup of this object. The backup ID is preserved when an object is copied, exported, or imported.
backup object	A backup unit that contains all items backed up from one disk volume (logical disk or mount point). The backed up items can be any number of files, directories, or the entire disk or mount point. Additionally, a backup object can be a database/application entity or a disk image.

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 (for example, SAP or Lotus).

• Type: Backup object type. For filesystem objects — filesystem type (for example, WinFS). 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. All clients configured in one backup specification are backed up together in one backup session using the same backup type. 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, full backup, and incremental 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.

backup system	*(ZDB specific term)* A system connected to a disk array together with one or multiple application systems. The backup system is typically connected to a disk array to create target volumes (a replica) and is used for mounting the target volumes (the replica). *See also* application system, target volume, and replica.
backup types	*See* incremental backup, differential backup, transaction backup, full backup, and delta backup.
backup view	Data Protector provides different views for backup specifications: By Type - according to the type of data available for backups/templates. Default view. By Group - according to the group to which backup specifications/templates belong. By Name - according to the name of backup specifications/templates. By Manager - if you are running MoM, you can also set the Backup view according to the Cell Manager to which backup specifications/templates belong.
BC	*(EMC Symmetrix specific term)* Business Continuance are processes that allow customers to access and manage instant copies of EMC Symmetrix standard devices. *See also* BCV.
BC Process	*(EMC Symmetrix specific term)* A protected storage environment solution that has defined specially configured EMC Symmetrix devices as mirrors or Business Continuance Volumes to protect data on EMC Symmetrix standard devices. *See also* BCV.
BCV	*(EMC Symmetrix specific term)* Business Continuance Volumes, or BCV devices, are dedicated SLDs that are pre-configured in the ICDA on which the business continuation operation runs. BCV devices are assigned separate SCSI addresses, differing from the addresses used by the SLDs they mirror. The BCV devices are used as splitable mirrors of the primary EMC Symmetrix SLDs that need to be protected. *See also* BC and BC Process.
Boolean operators	The Boolean operators for the full text search functionality of the 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 with files required for the initial step of the boot process. Microsoft terminology defines the boot volume/disk/partition as a volume/disk/partition containing the operating system files.
BRARCHIVE	*(SAP R/3 specific term)* An SAP R/3 backup tool that allows you to archive redo log files. BRARCHIVE also saves all the logs and profiles of the archiving process. *See also* 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* BRBACKUP and BRARCHIVE.
BSM	The Data Protector Backup Session Manager controls the backup session. This process always runs on the Cell Manager system.

C

CAP	*(StorageTek specific term)* Cartridge Access Port is a port built into the door panel of a library. The purpose is to enter or eject media.

Catalog 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 or SAN. 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	*(Windows specific term)* 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'

- 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 MC/ServiceGuard), application services, IP names and addresses, and so on).
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 application attributes, class attributes, and computer-level attributes. This guarantees 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 of 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)* An Oracle 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 Cell Request Server process (service), which runs on the Data Protector Cell Manager, and starts and controls the backup and restore sessions. The service is started as soon as Data Protector is installed on the Cell Manager. On Windows systems, the CRS runs under the account of the user specified at installation time. On UNIX systems, it runs under the account `root`.
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.

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 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 8, 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 8, 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 blobspace, dbspace, or logical log file.
DC directory	A directory that contains DC binary files, one for each configured Data Protector backup medium. DC directories constitute the Detail Catalog Binary Files part of the Data Protector Internal Database. *See also* Detail Catalog Binary Files (DBCF) 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	A physical unit which contains either just a drive or a more complex unit such as a library.
device chain	A device chain consists of several standalone devices configured for sequential use. When a medium in one device gets full, the backup automatically continues on a medium in the next device in the device chain.
device group	*(EMC Symmetrix specific term)* A logical unit representing several EMC Symmetrix devices. A device cannot belong to more than a single device group. All devices in a device group must be on the same EMC Symmetrix unit. You can use a device group to identify and work with a subset of the available EMC Symmetrix devices.
device streaming	A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little and resumes to write to the tape, and so on. In other words, if the data rate written

to the tape is less or equal the data rate which can be delivered to the device by the computer system, then the device is streaming. Streaming significantly improves the performance of the device and use of space.

DHCP server	A system running the Dynamic Host Configuration Protocol (DHCP) providing dynamic 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 backup	A high-speed backup where Data Protector backs up files as bitmap images. A disk image 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.

domain controller	A server in a network that is responsible for user security and verifying passwords within a group of other servers.
DR image	Data required for temporary disaster recovery operating system (DR OS) installation and configuration.
DR OS	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 (typically a tape drive). It can also read the data from the medium and send it to the computer system.
drive index	A number that identifies the mechanical position of a drive inside a library device. This number is used by the robotic control to access a drive.
drive-based encryption	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 Secure Socket Layer (SSL) that uses 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 several `StoreID`s 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 events are logged on the Cell Manager into the file `Data_Protector_program_data\log\server\Ob2EventLog.txt` (Windows systems),

or `/var/opt/omni/server/log/Ob2EventLog.txt` (UNIX systems). The Event Log is accessible only to users of the Data Protector `Admin` user group and to users who are granted the Data Protector `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	Also referred to as SCSI Exchanger. *See also* library.
exporting media	A process that removes all data about backup sessions, such as systems, objects, and 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	*(Windows specific term)* 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 (medium ID, description, and location) is saved in the IDB as well as on the respective media (media header). Data Protector media with protected data are not formatted until the protection expires or the media are unprotected/recycled.
free pool	An auxiliary source of media for use by media pools when they run out of media. The media pools must be configured to use free pools.
full backup	A backup in which all selected objects are backed up, whether or not they have been recently modified. *See also* backup types.
full database backup	A backup of all data in a database, not only the data that has been changed after the last (full or incremental) backup of the database. A full database backup does not depend on any other backup.
full mailbox backup	A full mailbox backup is a backup of the entire mailbox content.
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 on the Cell Manager in the directory `Data_Protector_program_data\Config\Server\holidays` (Windows systems), or `/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 P9000 XP Disk Array Family 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, 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, and 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 Business Copy (BC) P9000 XP, Main Control Unit, 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 Business Copy (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 HP P6000 / HP 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 component that executes all tasks needed by the Data Protector HP P9000 XP Disk Array Family integration. It uses RAID Manager Library for communication with a P9000 XP Array storage system. *See also* 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 session, 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 or 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 DB.
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 stores its data in 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 (DBCF).
Internet Information Services (IIS)	*(Windows specific term)* Microsoft Internet Information Services 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	*(EMC Symmetrix specific term)* A Logical Backup Object (LBO) is an object of data storage/retrieval in the EMC Symmetrix environment. It is stored/retrieved by EMC Symmetrix as one entity and can only be restored as a whole.
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 uses to log on to Microsoft SQL Server. A login ID is valid if Microsoft SQL Server has an entry for that user in the system table syslogin.
login information to the Oracle Target Database	*(Oracle and SAP R/3 specific term)* The format of the login information is *user_name/password@service*, where:
	• *user_name* is the name by which a user is known to Oracle Server and to other users. Every user name is associated with a password and both have to be entered to connect to an Oracle Target Database. This user must have 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 or HP CA+BC P9000 XP 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 Data Protector 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 (MoM) *See* MoM.

MAPI *(Microsoft Exchange Server 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 determine the state 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) 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.

O

OBDR capable device	A device that can emulate a CD-ROM drive loaded with a bootable disk and can thus be used as a backup or boot device for disaster recovery purposes.
obdrindex.dat	*See* IDB recovery file.
object	*See* backup object.
object consolidation	The process of merging a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object. The process is a part of the synthetic backup procedure. The result is a synthetic full backup of the specified backup object.
object consolidation session	A process that merges a restore chain of a backup object, consisting of a full backup and at least one incremental backup, into a new, consolidated version of this object.
object copy	A copy of a specific object version that is created during an object copy session or a backup session with object mirroring.
object copy session	A process that creates an additional copy of the backed up data on a different media set. During an object copy session, the selected backed up objects are copied from the source to the target media.
object copying	The process of copying selected object versions to a specific media set. You can select object versions from one or several backup sessions to be copied.
object ID	*(Windows specific term)* The object IDs (OIDs) enable access to NTFS 5 files no matter where in the system the files reside. 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. In an offline backup session, the database is generally put into a quiescent state that allows use by the backup system, but not the application, for the time period of the data replication process. For instance, for backup to tape, until streaming of data to the tape is finished. Normal database operation is resumed before potential post-backup operations are started. *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
	• ON-Bar catalog tables, which are used to back up dbobjects and track instances of dbobjects through multiple backups.
ONCONFIG	*(Informix Server specific term)* An environment variable that specifies the name of the active ONCONFIG configuration file. If the ONCONFIG environment variable is not present, Informix Server uses the configuration values from the `onconfig` file 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 of the data replication process. For instance, for backup to tape, until streaming of data to tape is finished. During this period, the database is fully operational, but there may be a small performance impact and log files may grow very quickly. Normal database operation is resumed before potential post-backup operations are started.

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.
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 Backup ownership affects the ability of users to see and restore data. Each backup session and all the data backed up within it is assigned an owner. The owner can be the user that starts an interactive backup, the account under which the CRS process is running, or the user specified as the owner in the backup specification options.

If a user starts an existing backup specification without modifying it, the backup session is not considered as interactive.

If a modified backup specification is started by a user, the user is the owner unless the following is true:

- The user has the Switch Session Ownership user right.
- The backup session owner is explicitly defined in the backup specification, where the username, group or domain name, and the system name are specified.

If a backup is scheduled on a UNIX Cell Manager, the session owner is root:sys unless the above conditions are true.

If a backup is scheduled on a Windows Cell Manager, the session owner is the user specified during the installation, unless the above conditions are true.

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/server/dr/p1s` (UNIX systems) with the filename `recovery.p1s`.

package *(MC/ServiceGuard and Veritas Cluster specific term)* A collection of resources (for example volume groups, application services, IP names, and addresses) that are needed to run a specific cluster-aware application.

pair status *(HP 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.
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.
prealloc list	A subset of media in a media pool that specifies the order in which media are used for backup.
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 also 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
- 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 snapclone), or a virtual copy (for example, a snapshot). From perspective of a basic operating system, the complete physical disk containing backup objects is replicated. However, if a volume manager is used on a UNIX system, the whole volume or disk group containing a backup object (logical volume) 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), mirror unit (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

SAPDBA	*(SAP R/3 specific term)* An SAP R/3 user interface that integrates the BRBACKUP, BRARCHIVE, and BRRESTORE tools.
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-exec 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` 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, and high-speed databases. If sparse file processing is not enabled during restore, it might be impossible to restore this file.

split mirror
(EMC Symmetrix Disk Array 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	*(disaster recovery specific term)* A text file in the Unicode (UTF-16) format, generated during CONFIGURATION backup of a Windows or Linux 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)* An object that may be presented to an operating system or some other entity (for example, a virtualization mechanism) upon which volume management systems, filesystems, or other objects may exist. The volume management systems, file systems are built on this storage. Typically, these can be created or exist within a storage system such as a disk array.
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. 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	*(Oracle specific term)* An Oracle interface that handles the actions required to load, label, and unload correct backup devices when Oracle issues a backup or restore request.
system databases	*(Sybase specific term)* The four system databases on a newly installed Sybase SQL Server are the:

- master database (master)
- temporary database (tempdb)
- system procedure database (sybsystemprocs)
- model database (model).

System Recovery Data file	*See* SRD file.
System State	*(Windows specific term)* The System State data comprises the Registry, COM+ Class Registration database, system startup files, and the Certificate Services database (if the server is a certificate server). If the server is a domain controller, Active Directory services and the SYSVOL directory are also contained in the System State data. If the server is running the Cluster service, the System State data also includes resource registry checkpoints and the quorum resource recovery log, which contains the most recent cluster database information.
system volume/disk/ partition	A volume/disk/partition containing operating system files. Microsoft terminology defines the system volume/disk/partition as the volume/disk/partition containing files required for the initial step of the boot process.
SysVol	*(Windows specific term)* A shared directory that stores the server copy of the domain's public files, which are replicated among all domain controllers in the domain.

T

tablespace	A part of a database structure. Each database is logically divided into one or more tablespaces. Each tablespace has data files or raw volumes exclusively associated with it.
tapeless backup (ZDB specific term)	*See* ZDB to disk.
target (R2) device	*(EMC Symmetrix specific term)* An EMC Symmetrix device that participates in SRDF operations with a source (R1) device. It resides in the remote EMC Symmetrix unit. It is paired with a source (R1) device in the local EMC Symmetrix unit and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF2 group type. *See also* source (R1) device.
target database	*(Oracle specific term)* In RMAN, the target database is the database that you are backing up or restoring.
target system	*(disaster recovery specific term)* A system after a computer disaster has occurred. The target system is typically in a non-bootable state and the goal of disaster recovery is to restore this system to the original system configuration. The difference between a faulty system and a target system is that a target system has all faulty hardware replaced.

target volume	*(ZDB specific term)* A storage volume to which data is replicated.
Terminal Services	*(Windows specific term)* Windows Terminal Services provide a multi-session environment that allows clients to access a virtual Windows desktop session and Windows-based programs running on the server.
thread	*(Microsoft SQL Server specific term)* An executable entity that belongs to only one process. It comprises a program counter, a user-mode stack, a kernel-mode stack, and a set of register values. Several threads can run at a time within one process.
TimeFinder	*(EMC Symmetrix specific term)* A business continuation process that creates an instant copy of single or multiple Symmetrix Logical Devices (SLDs). The instant copy is created on specially preconfigured SLDs called BCVs and is accessible via a separate device address to the system(s).
TLU	Tape Library Unit.
TNSNAMES.ORA	*(Oracle and SAP R/3 specific term)* A network configuration file that contains connect descriptors mapped to service names. The file may be maintained centrally or locally, for use by all or individual clients.
transaction	A mechanism for ensuring that a set of actions is treated as a single unit of work. Databases use transactions to keep track of database changes.
transaction backup	Transaction backups generally use fewer resources than database backups, so they can be created more frequently than database backups. By applying transaction backups, you can recover the database to a specific point in time prior to when a problem occurred.
transaction backup	*(Sybase and SQL specific term)* A backup of the transaction log providing a record of changes made since the last full or transaction backup.
transaction log backup	Transaction log backups generally use fewer resources than database backups so they can be created more frequently than database backups. By applying transaction log backups, you can recover the database to a specific point in time.
transaction log files	Files that record transactions of the database modifications, and provide fault tolerance in case of a database disaster.
transaction 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 where a backup can be performed. *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, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012 that limits application software to standard user privileges until an administrator authorizes an increase in privilege level.
user disk quotas	NTFS quota-management support enables an enhanced tracking mechanism and control over disk space usage on shared storage volumes. 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	*(VLS specific term)* 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)	*(VLS specific term)* An emulated tape library that provides the functionality of traditional tape-based storage. *See also* Virtual Library System (VLS).
VMware management client	*(VMware (Legacy) integration specific term)* The client that Data Protector uses to communicate with VMware Virtual Infrastructure. This can be a VirtualCenter Server system (VirtualCenter environment) or an ESX Server system (standalone ESX Server environment).
volser	*(ADIC and STK specific term)* A VOLume SERial number is a label on the medium to identify the physical tape used in very large libraries. A volser is a naming convention specific to ADIC/GRAU and StorageTek devices.
volume group	A unit of data storage in an LVM system. A volume group can consist of one or more physical volumes. There can be more than one volume group on the system.
volume mountpoint	*(Windows specific term)* An empty directory on a volume that can be used to mount another volume. The volume mount point acts as a gateway to the target volume. Provided that the volume is mounted, users and applications can refer to the data on the mounted volume by the full (merged) filesystem path as if both volumes are one integral part.

Volume Shadow Copy Service	*See* Microsoft Volume Shadow Copy Service (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 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).
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 in zero downtime backup, instant recovery, and split mirror restore sessions. *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