



Hewlett Packard
Enterprise

HPE Data Protector

Software Version: 9.09

Integration guide for Microsoft Volume Shadow Copy Service

Document Release Date: March 2017
Software Release Date: March 2017

Legal Notices

Warranty

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

The information contained herein is subject to change without notice.

Restricted Rights Legend

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

Copyright Notice

© Copyright 2017 Hewlett Packard Enterprise Development LP

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

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

UNIX® is a registered trademark of The Open Group.

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

Documentation Updates

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

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

To check for recent software updates: <https://softwaresupport.hpe.com/patches>.

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

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

Or click the **Register** link at the top of the HPE Software Support page.

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

Support

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

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

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

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

Most of the support areas require that you register as an HPE Passport user and sign in. Many also require a support contract.

To register for an HPE Passport ID, go to:

<https://hpp12.passport.hpe.com/hppcf/createuser.do>

To find more information about access levels, go to:

<https://softwaresupport.hpe.com/web/softwaresupport/access-levels>

HPE Software Solutions Now accesses the HPESW Solution and Integration Portal Web site. This site enables you to explore HPE Product Solutions to meet your business needs, includes a full list of Integrations between HPE Products, as well as a listing of ITIL Processes. The URL for this Web site is <https://softwaresupport.hpe.com>.

Contents

Chapter 1: Introduction	9
Chapter 2: Integration concepts	10
Volume Shadow Copy Service	10
Virtual Disk Service	12
VSSBAR agent	12
Backup types	13
VSS database	14
Querying the VSSDB	15
Deleting backup sessions	15
Enabling and disabling replicas	16
Changes in the environment	16
Chapter 3: Configuration	17
Prerequisites and limitations	17
Prerequisites	17
Transportable backup prerequisites	17
Limitations	18
ZDB and instant recovery limitations	18
HPE P6000 EVA Disk Array Family hardware provider limitations	18
HPE P9000 XP Disk Array Family hardware provider limitations	19
Configuring the application system for instant recovery-enabled backup sessions	19
Configuring HPE P4000 SAN Solutions	20
Prerequisites	20
Configuring the integration	21
Configuring the VSS hardware provider	22
Configuring HPE P6000 EVA Disk Array Family	22
Prerequisites	22
Configuring the integration	23
Considerations	23
Configuring the VSS hardware provider	24
Configuring HPE P9000 XP Disk Array Family	24
Prerequisites	24
Configuring the VSS hardware provider	25
Configuring the user authentication data	25
Configuring HPE 3PAR StoreServ Storage	25
Prerequisites	26

Configuring the Data Protector HPE 3PAR StoreServ Storage integration	26
Configuring the HPE 3PAR VSS hardware provider	26
Chapter 4: Backup	29
Backup types	29
Backup flow	29
Zero downtime backup with the Data Protector VSS integration	30
HPE P4000 SAN Solutions	30
HPE P6000 EVA Disk Array Family	30
HPE P9000 XP Disk Array Family	31
HPE 3PAR StoreServ Storage	31
Considerations	32
Microsoft Cluster environments and transportable backups	32
HPE P4000 SAN Solutions	32
HPE P6000 EVA Disk Array Family	33
Replica creation and reuse	33
HPE P9000 XP Disk Array Family	33
HPE 3PAR StoreServ Storage	34
Configuration check	34
Creating backup specifications	35
Backup options	39
ZDB options	40
Scheduling backup sessions	42
Starting backup sessions	44
Chapter 5: Restore	46
Standard restore	46
Restore modes	46
Component restore	46
File restore	46
Restoring using the GUI	47
Restore options	48
Restoring using the CLI	49
Instant recovery	50
Instant recovery methods	51
Switch of disks	52
Copy of replica data	52
Copy of replica data with the source volume retained	52
Copy of replica data with the source volume not retained	53
Restore snapshot data to the source volume	53
Restore snapshot data to the source volume with the source volume retained	53
Restore snapshot data to the source location with the source volume not retained	54

Resync of volumes	54
Limitations	54
HPE P4000 SAN Solutions considerations	55
HPE P6000 EVA Array considerations	55
HPE P9000 XP Array considerations	56
HPE 3PAR StoreServ Storage considerations	58
Instant recovery procedure	58
Limitations	58
Procedure	58
Chapter 6: Writer specifics	64
Microsoft Data Protection Manager writer specifics	67
Backup	67
Restore	69
Restore the DPM server first	69
Restore the DPM clients directly	70
Microsoft Exchange Server 2003 writer specifics	71
Backup	71
Restore	72
Rollforward recovery from the loss of one or more databases	72
Point-in-time restore after loss of a log file	73
Limitations	74
Instant recovery	74
Prerequisites	75
Post-instant recovery steps	75
Database recovery	75
User scenario for Microsoft Exchange Server 2003 backup and restore	75
Example - VSS transportable backup	76
Example restore scenario for Microsoft Exchange Server 2003	78
Microsoft Exchange Server 2007 writer specifics	78
Concepts	78
Continuous replication	79
Restore to original or another location	79
Backup	79
LCR and CCR environments	81
Restore	83
Standard restore	83
Rollforward recovery from the loss of one or more databases	83
Point-in-time restore after loss of a log file	84
Restoring individual mailboxes	85
Restoring a LCR or CCR copy to the original location	88
Instant recovery	90
Prerequisites	91
Limitations	91
Restoring an LCR or CCR copy to the original location	93

Post-instant recovery steps	93
Troubleshooting	93
Microsoft Exchange Server 2010 writer specifics	94
Introduction	94
Microsoft Exchange Server 2010 concepts	94
Integrating Data Protector Microsoft Volume Shadow Copy Service integration with Exchange Server 2010	94
Configuration	95
Prerequisites	95
Licensing	95
Configuring	95
Backup	95
Limitations	96
Creating a backup specification	96
Restore	96
Restore scenarios	96
Restoring a passive copy	96
Restoring an active copy	99
Point-in-time restore	100
Instant recovery	101
Limitations	101
Instant recovery of a P6000 EVA Array snapclone to a different client using the CLI ..	101
Troubleshooting	102
Microsoft Hyper-V writer specifics	102
Concepts	102
Prerequisites	103
Backup	103
Prerequisites	104
Limitations	104
Backup from a physical cluster node	105
Restore	105
Restore from a physical cluster node	105
Troubleshooting	106
Microsoft SharePoint Services writer specifics	106
Concepts	106
Backup	107
Restore	108
MSDE writer specifics	110
Restore	110
Chapter 7: Troubleshooting	111
Before you begin	111
Checks and verifications	111

Problems	111
ZDB related problems	116
Send Documentation Feedback	121

Chapter 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 <https://softwaresupport.hpe.com/>.
- *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.

Chapter 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.

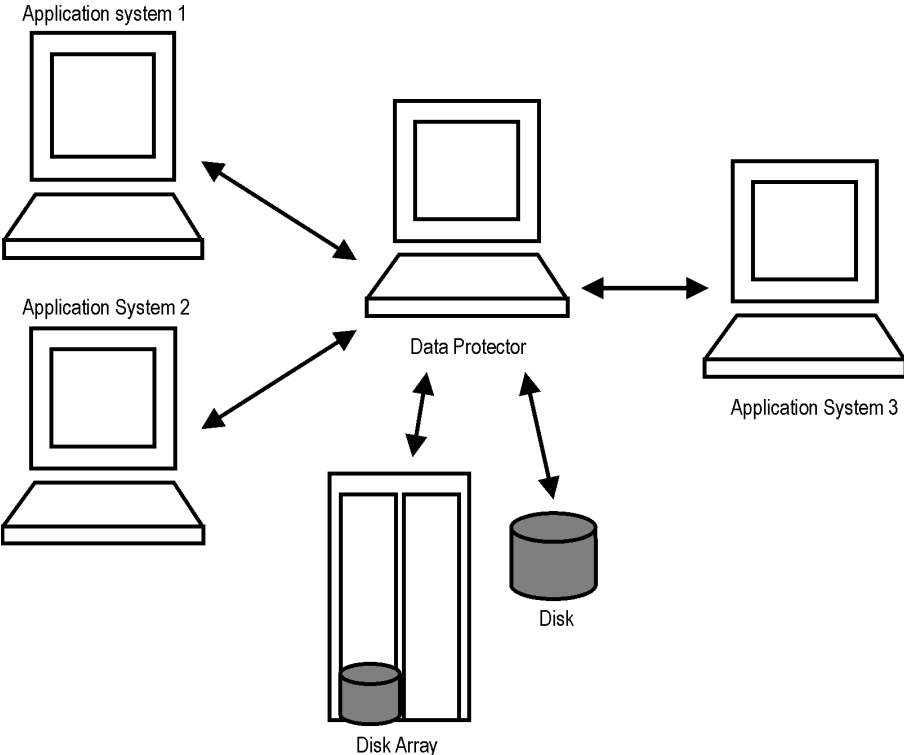
During backup, the service collaborates with a backup application (for example, Data Protector), applications to be backed up (usually database applications), shadow copy providers, and the operating system kernel to create a consistent, point-in-time shadow copy set. These shadow copy sets can be backed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

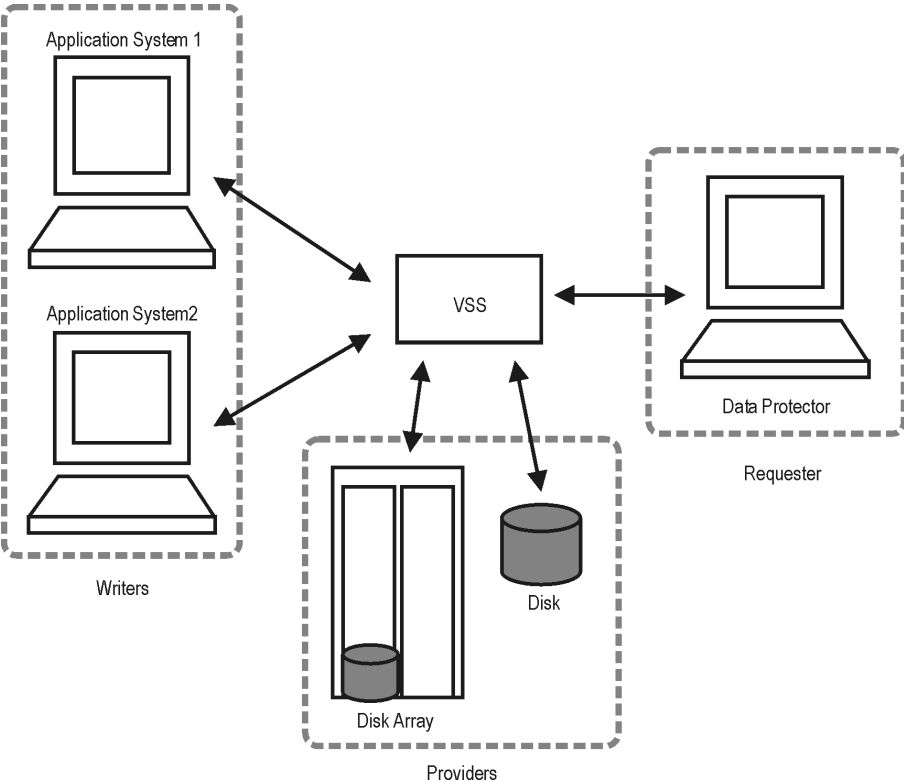
The Data Protector Volume Shadow Copy integration provides a unified communication interface that can coordinate backup and restore of an application regardless of their specific features. With this approach, a backup application does not need to handle each application to be backed up specifically.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

["Actors of the traditional backup model" on the next page](#) and ["Actors of the Data Protector VSS integration backup model" on the next page](#) show the differences between the traditional backup model and the model with the Data Protector Microsoft Volume Shadow Copy Service integration.

Actors of the traditional backup model



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:

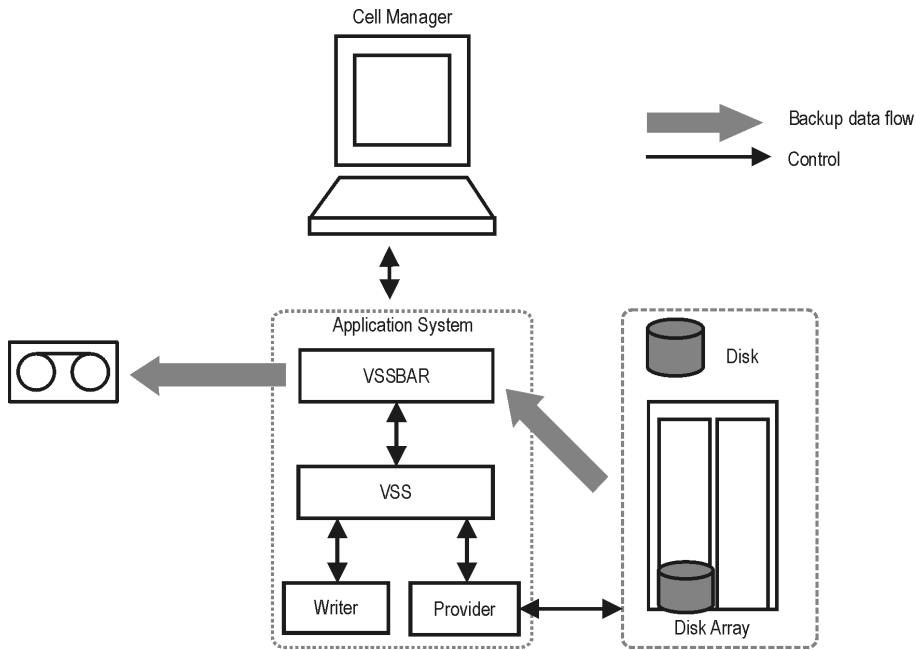
- With the *software provider*, you can back up the data to tape (**standard Data Protector backup to tape**). This can be done only on the same system where the shadow copy is created.
- With the *hardware provider*, you can create target volumes, which can be used for backing up to tape or for instant 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 *HPE Data Protector Concepts Guide*.

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

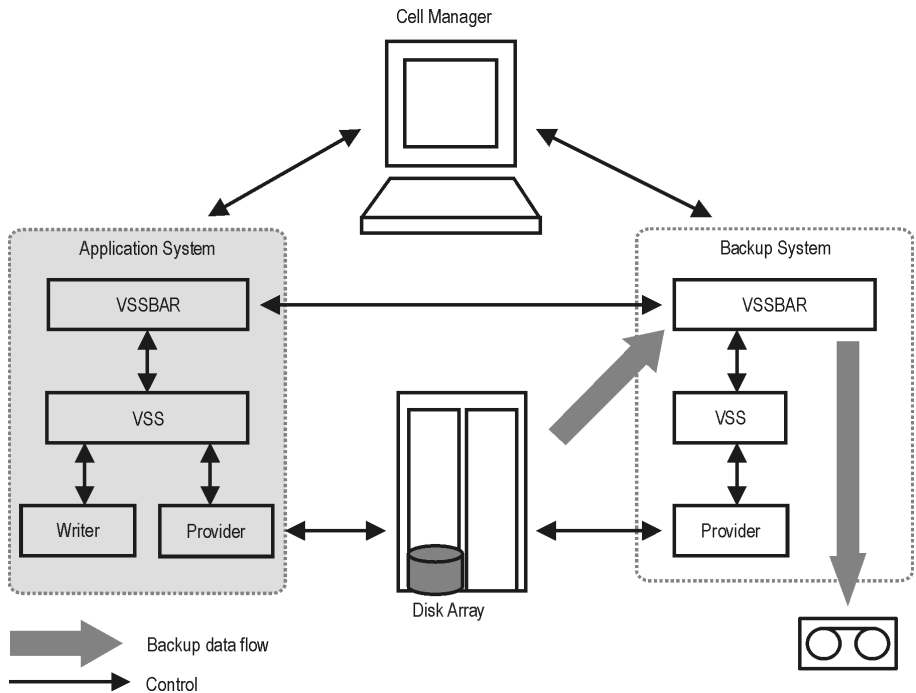
- *A dedicated system* (Data Protector ZDB terminology: a dual-host configuration).
In VSS terminology, a backup performed in such an environment is a **transportable backup**. It is available only in combination with a hardware provider.
- *The same system as the application system* (Data Protector ZDB terminology: a single-host configuration).
In VSS terminology, a backup performed in such an environment is a **local or network backup**.

["Local VSS backup" on the next page](#) shows the relationships between the components of a local VSS backup. ["Transportable VSS backup" on the next page](#) shows the relationships between the components of a transportable VSS backup.

Local VSS backup



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 and their

replicas. 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 unrepresenting) 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 *HPE Data Protector Command Line Interface Reference*.

The following tasks can be performed using the `omnidbvss` command:

- ["Querying the VSSDB" below](#)
- ["Deleting backup sessions" below](#)
- ["Enabling and disabling replicas" on the next page](#)

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 recorded in the persistent part of the VSSDB, 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 persistent part of 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 unmount (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) unmounted the replica from the backup system and mounted 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.

Chapter 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 *HPE 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.
<https://softwaresupport.hpe.com/>
 - *HPE Data Protector Installation 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 <https://softwaresupport.hpe.com/>.

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 HPE 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 HPE 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](#)" on page 50.

Limitations

For a list of general Data Protector limitations, see the *HPE 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.
- Data Protector MS Volume Shadow Copy integration skips reparse points. Exceptions are for the Microsoft Volume Deduplication deduplicated files, with reparse tag `IO_REPARSE_TAG_DEDUP`, which are backed up as normal file.

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 <https://softwaresupport.hpe.com/>.
- 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 systems, the created snapshots are unrepresented on the application system after the backup.
 - On Windows Server 2008 systems, the created snapshots are only put offline after the backup, while they remain presented on the application system.

HPE P6000 EVA Disk Array Family hardware provider limitations

- On Windows Server 2003 systems if a disk array from the HPE 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.

HPE P9000 XP Disk Array Family hardware provider limitations

- If a disk array from the HPE P9000 XP Disk Array Family is used and the P9000 XP Array VSS provider 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 HPE 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 than 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 HPE 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 *HPE 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 HPE P4000 SAN Solutions

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

Prerequisites

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

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

For information on supported product versions, see the latest support matrices at <https://softwaresupport.hpe.com/>.

- Obtain or install **Data Protector licenses and components**:
 - An instant recovery license.
 - HPE P4000 VSS Agent.
For installation and licensing information, see the *HPE Data Protector Installation Guide*.
- Make sure the same operating system (and its version) is installed on the application and backup systems.
- Connect an HPE 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 HPE 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 *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

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

Configuring the integration

Before you start configuration, make sure you met the prerequisites described in "[Prerequisites](#)" on the [previous page](#). Then set the login information for the SMI-S HPE 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 *HPE P4000 SAN Solutions* part of the *HPE Data Protector Zero Downtime Backup Administrator's Guide* and the `omnidp4000` reference page in the *HPE Data Protector Command Line Interface Reference*.

The information you provide is kept in the HPE 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 HPE P4000 VSS Agent communicates (default - 5988). If the SMI-S HPE P4000 SAN Solutions provider accepts the SSL-based connection, the default port number is 5989.

NOTE: When using StoreVirtual VSA software version 12.5, use SSL-based connection.

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

The HPE 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).

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 HPE 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 HPE 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 HPE P4000 SAN Solutions VSS provider documentation for details.

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

Configuring HPE P6000 EVA Disk Array Family

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

Prerequisites

- Obtain or install **HPE storage licenses and components**:
 - HPE 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 <https://softwaresupport.hpe.com/>.
 - HPE Business Copy (BC) P6000 EVA microcode and license.
 - HPE MPIO Full Featured DSM (Device Specific Module) for HPE 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.
 - HPE P6000 / HPE 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 *HPE Data Protector Installation 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 *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

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

Configuring the integration

Before you start configuration, make sure you met the prerequisites described in ["Prerequisites" on the previous page](#). 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 *HPE Data Protector Command Line Interface Reference*.

The information you provide is kept in the ZDB database for the HPE 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).

Considerations

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.

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 HPE 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 HPE 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 HPE P9000 XP Disk Array Family

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

Prerequisites

- Obtain or install **HPE storage 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 HPE sales representative for information on which version of RAID Manager Library to use.
 - HPE Business Copy (BC) P9000 XP microcode and a license for at least basic and HPE Business Copy operation.
 - P9000 XP Array VSS and VDS provider.
 - HPE MPIO Full Featured DSM (Device Specific Module) for HPE P9000 XP Disk Array Family

For information on supported product versions, see the latest support matrices at <https://softwaresupport.hpe.com/>.

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

- An instant recovery license.
- The HPE 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 *HPE Data Protector Installation 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 HPE 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 "[HPE P9000 XP Array considerations](#)" on page 56.

- 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 HPE 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 *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

Configuring HPE 3PAR StoreServ Storage

This section describes how to configure the Data Protector HPE 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 **HPE storage licenses and components**:

- HPE 3PAR VSS Provider Software.

For installation instructions, see the HPE 3PAR StoreServ Storage documentation. For information on supported product versions, see the latest support matrices at <https://softwaresupport.hpe.com/>.

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

- Appropriate zero downtime backup extension and instant recovery extension licenses-to-use (LTU).
- HPE 3PAR VSS Agent installed on both the application system and the backup system

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

- Make sure the same operating system (and its version) is installed on the application and backup systems.
- Connect a storage system of the HPE 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 HPE 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 *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

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

Configuring the Data Protector HPE 3PAR StoreServ Storage integration

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

Configuring the HPE 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 HPE 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 HPE 3PAR StoreServ Storage documentation.

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

Chapter 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" on page 13](#) 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" on page 64](#).

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" on page 14](#) 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 *HPE Data Protector Concepts Guide*.

HPE P4000 SAN Solutions

There is one replica type available when backing up using the HPE 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.

HPE 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). Vsnap

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**

HPE 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.

HPE 3PAR StoreServ Storage

There is one replica type available when backing up using the HPE 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).

HPE P4000 SAN Solutions

When a new snapshot is created with HPE 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.

HPE 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.

HPE 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 "[HPE P9000 XP Array considerations](#)" on page 56.
- 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.

HPE 3PAR StoreServ Storage

- After a new snapshot is created with HPE 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.
- HPE 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 HPE 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.

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](#)" on page 64.


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

1. In the **HPE 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.

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:
Replica types

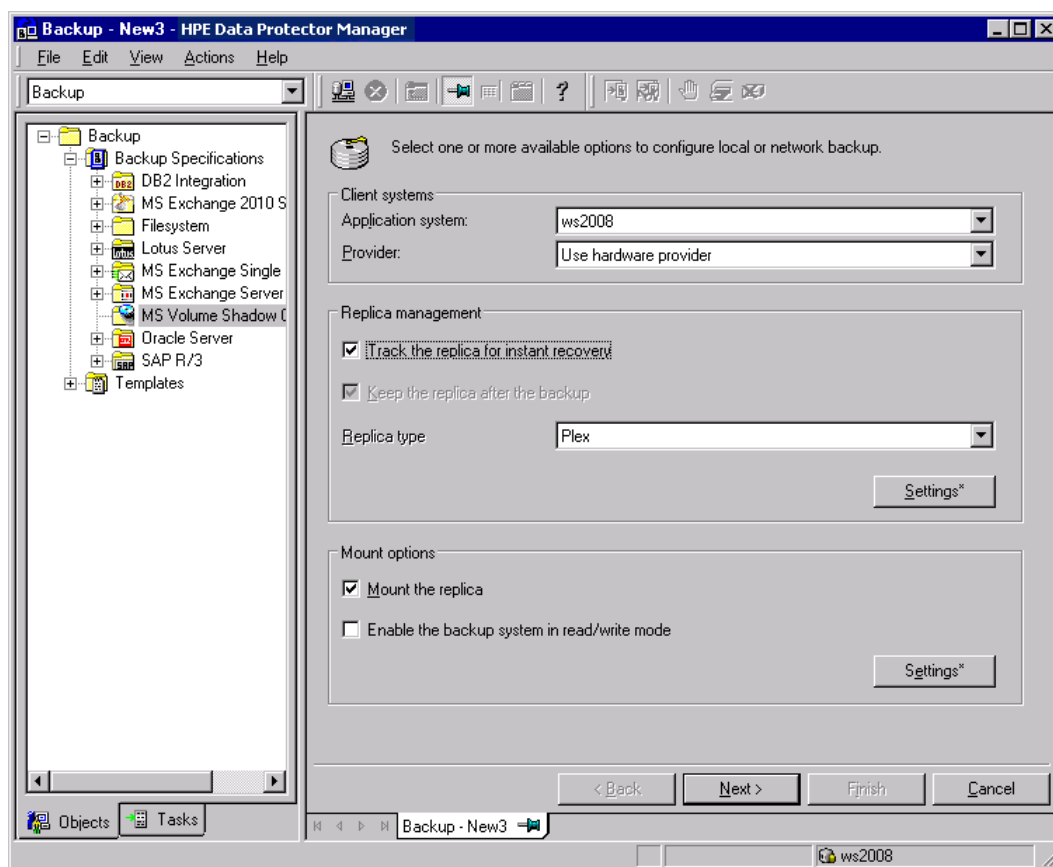
Replica Type	Description
Plex (Clone/Mirror)	A shadow copy, independent from its source volume is created: <ul style="list-style-type: none"> ○ This replica type is not supported by the HPE P4000 SAN Solutions and HPE 3PAR StoreServ Storage hardware providers. ○ With the HPE P6000 EVA Disk Array Family hardware provider, a snapclone is created. ○ With the HPE 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. <ul style="list-style-type: none"> ○ With the HPE P4000 SAN Solutions or HPE 3PAR StoreServ Storage hardware provider, a snapshot is created. ○ With the HPE 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 " on page 40. ○ With the HPE 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"](#) on page 30.

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

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 ["Replica management options "](#) on page 40 and ["P6000 EVA specific options "](#) on page 40 or press **F1**.

For descriptions of mount options, see ["Mount options" on page 41](#) 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.

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.

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 *HPE 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 *HPE 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

VSS-specific backup options

Option	Description
Pre-exec	Specify a command that will be started by <code>vssbar.exe</code> on the application

Option	Description
	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 the default Data Protector commands directory.
Post-exec	Specify a command that will be started by <code>vssbar.exe</code> 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 the default Data Protector commands directory.

ZDB options

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" on page 34 .
Number of replicas rotated	<p>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.</p> <p>Default: 3</p> <p>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.</p>

P6000 EVA specific options

Option	Description
Replica sub type	<p>Available if Keep the replica after the backup is selected.</p> <p>Sets the replica subtype. Depending on the Replica type you selected, the following subtypes can be selected:</p> <ul style="list-style-type: none"> • Plex (Clone/Mirror): <ul style="list-style-type: none"> • SnapClone • Differential (Snapshot): <ul style="list-style-type: none"> • Snapshot • MirrorClone Snapshot

Option	Description
	Disabled if Replica type is set to Default .
Wait for data copy to complete	<p>Available if Replica subtype is set to SnapClone.</p> <p>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.</p>
Attempt to minimize wait	<p>This option resumes the backup process after the copy process has finished if it is finished before the specified time.</p> <p>Selected by default if Wait for data copy to complete is selected.</p>

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.

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 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 <code>c:\mnt</code> .
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.

Option	Description
	<p>Example:</p> <p>Root directory: c:\mnt</p> <p>Application system: app.comp.com</p> <p>Backup session ID: 2008-02-22-4</p> <p>Mount path on the application system: E:\disk1</p> <p>If Hostname is selected: c:\mnt\app.comp.com\E\disk1</p> <p>If Hostname + session ID is selected: c:\mnt\app.comp.com\2008-02-22-4\E\disk1</p> <p>If Session ID is selected: c:\mnt\2008-02-22-4\E\disk1</p> <p>If Session ID + hostname is selected: c:\mnt\2008-02-22-4\app.comp.com\E\disk1</p> <p>Default: Hostname + session ID.</p>
<p>Automatically dismount the filesystems at destination mount points</p>	<p>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.</p> <p>Default: not selected.</p>


Scheduling backup sessions

For more detailed information on scheduling, see the *HPE Data Protector Help* index: “scheduled backups”.

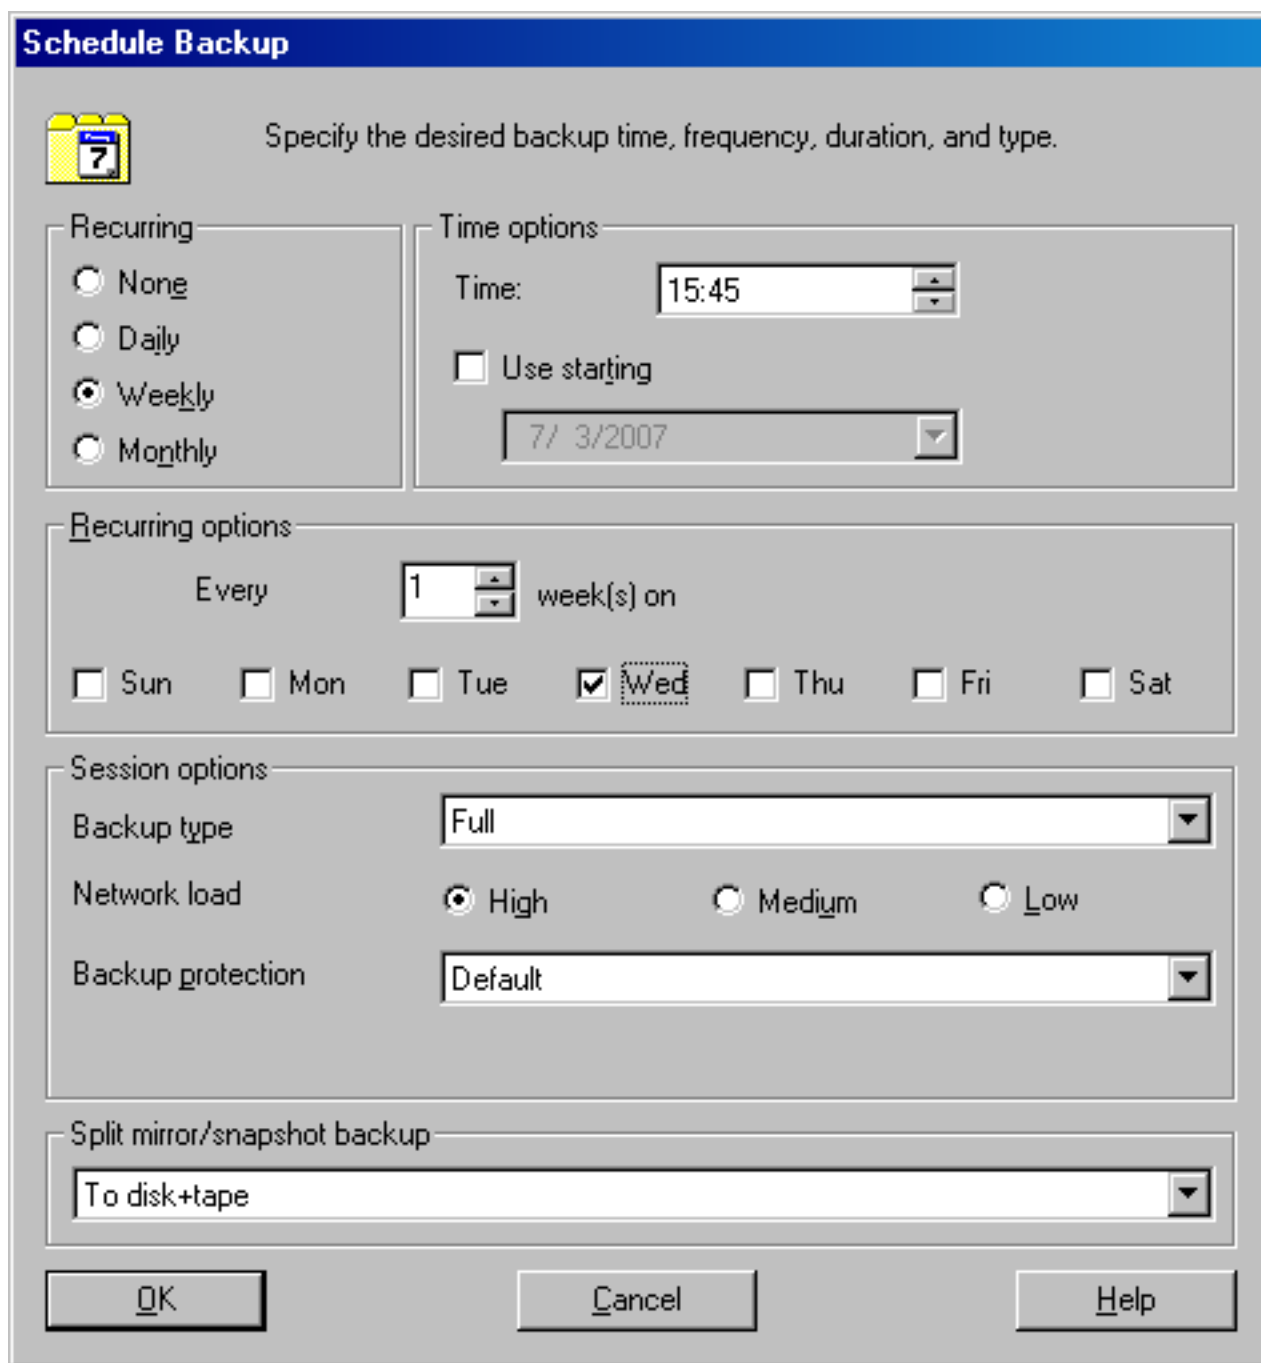
To schedule a Microsoft Volume Shadow Copy Writers backup specification, perform the following steps in the Data Protector GUI:

1. In the **HPE 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.

Scheduling a backup



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 **HPE 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 *HPE Data Protector Help*.

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

Chapter 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](#)" below.
- Using the instant recovery functionality. See "[Instant recovery](#)" on page 50.

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.

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.

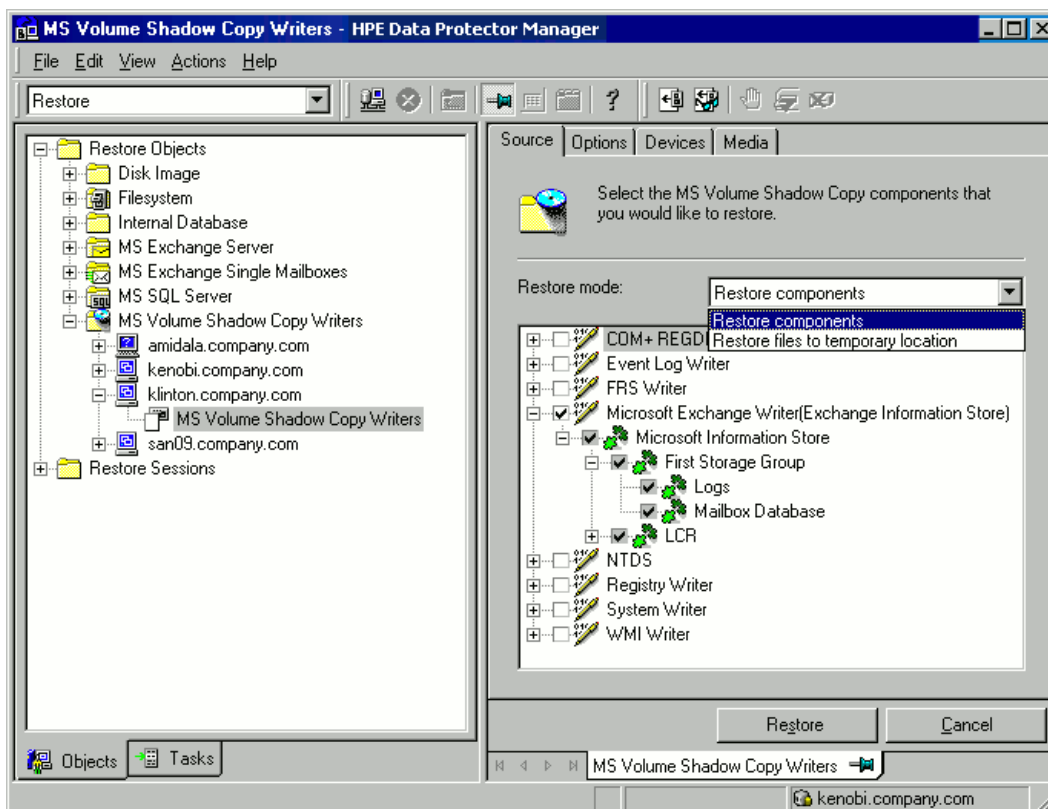
Restoring using the GUI

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](#)" on page 64.

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).

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" on page 64](#).

5. In the **Options** property page, select the MS Volume Shadow Copy specific restore options. See ["Restore options" below](#).
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 *HPE 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 Options

Restore Options	Description
Restore to another client	<p>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</p> <p>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.</p>
Restore into the following directory	<p>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).</p> <p>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.</p>

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 *HPE 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 *HPE 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.

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 *HPE Data Protector Concepts Guide* and *HPE 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 <https://softwaresupport.hpe.com/>.

- **Using a Data Protector disk array integration agent** (HPE P4000 VSS Agent, HPE P6000 / HPE 3PAR SMI-S Agent, HPE P9000 XP Agent, or HPE 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 HPE P4000 SAN Solutions](#)" on page 55, "[Supported instant recovery methods with HPE P6000 EVA Disk Array Family](#)" on page 55, "[Supported instant recovery methods for HPE P9000 XP Disk Array Family](#)" on page 56, or "[Supported instant recovery methods for HPE 3PAR StoreServ Storage](#)" on page 58.

Switch of disks

With this method, the source volumes are unrepresented 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 HPE P6000 / HPE 3PAR SMI-S Agent or HPE 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, HPE 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.

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 *HPE 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, HPE P4000 SAN Solutions, HPE P6000 EVA Disk Array Family, HPE P9000 XP Disk Array Family, and HPE 3PAR StoreServ Storage restore options are not available.

For additional writer specific limitations, see "[Writer specifics](#)" on page 64.

HPE P4000 SAN Solutions considerations

- The following table lists instant recovery methods available with the HPE P4000 SAN Solutions
Supported instant recovery methods for HPE P4000 SAN Solutions

Instant recovery method		HPE P4000 VSS Agent
Restore snapshot data to the source volume	with the source volume retained	No
	with source volume not retained	Yes

- HPE 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.

HPE P6000 EVA Array considerations

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

Instant recovery method		VDS hardware provider	HPE P6000 / HPE 3PAR SMI-S Agent
Switch of disks		Yes ¹	Yes ²
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

- 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).

HPE P9000 XP Array considerations

- The following table lists instant recovery methods available with the HPE P9000 XP Array:
 Supported instant recovery methods for HPE P9000 XP Disk Array Family

¹Available if the backup was performed using snapclone as subtype.
²Available if the backup was performed using snapclone as subtype.

Instant recovery method		VDS hardware provider	HPE P9000 XP Agent
Switch of disks		Yes ¹	No
Resync of volumes		No	Yes ²
Restore snapshot data to the source volume	with the source volume retained	No	No
	with source volume not retained	No	Yes ³

- 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,

¹ hardware provider in the VSS compliant mode

² hardware provider in the resync mode

³ hardware provider in the resync mode

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 volumes on the application system: `omnidbvss -resolve -apphost ApplicationClient`
- To resolve the target volumes created in the backup session: `omnidbvss -resolve -session SessionID`

The following restore considerations apply to restore using the HPE 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.

HPE 3PAR StoreServ Storage considerations

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

Instant recovery method		HPE 3PAR VSS Agent
Restore snapshot data to the source volume	with the source volume retained	No
	with source volume not retained	Yes

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

Instant recovery procedure

Limitations

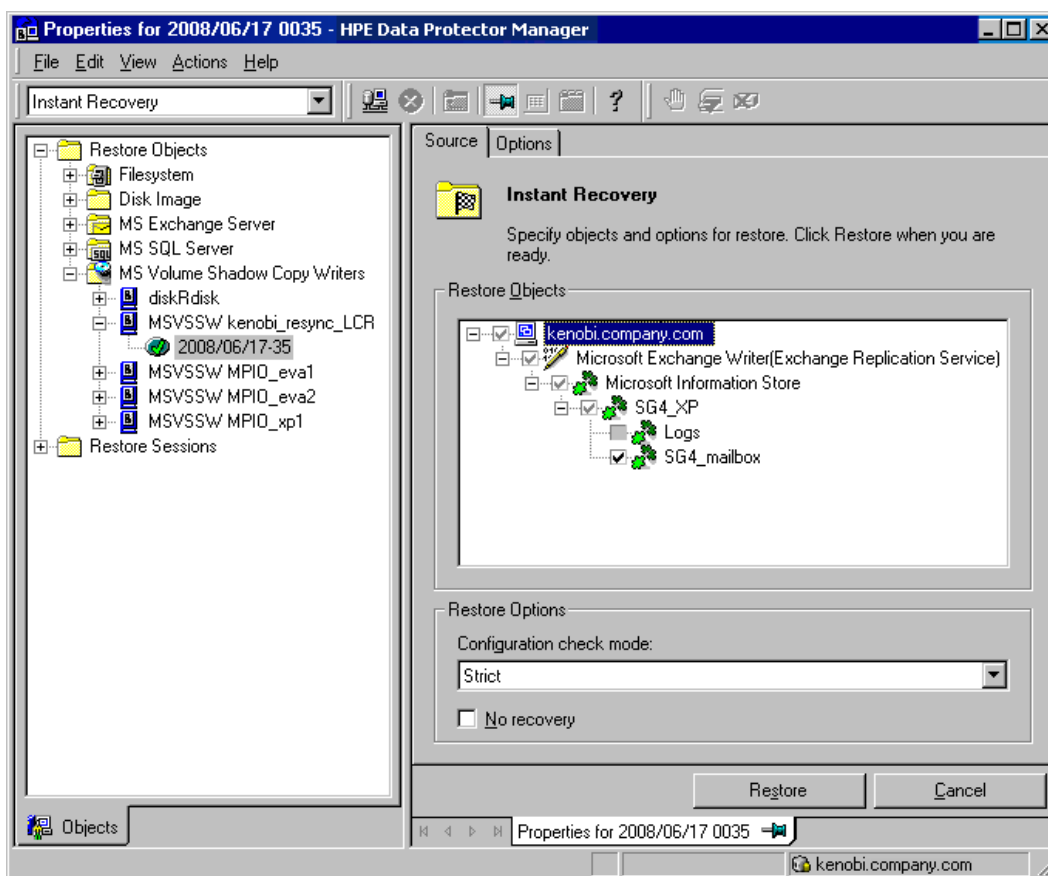
- Instant recovery of Cluster Shared Volumes is not supported.
- Snapshots can be restored only to the original LUN (the LUN from where the replica was created) and to the same client.

Procedure

1. In the **HPE 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.

Selecting writers components for instant recovery



- In the Source property page, specify writers and/or components for recovery.
Select the configuration check mode. For more information, see ["Configuration check" on page 34](#).
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" on page 64](#).
- 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 agentSee ["Instant recovery methods" on page 51](#) 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 HPE P4000**

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

For details on these methods see ["Instant recovery methods" on page 51](#).

- **Restore using HPE 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" on page 51](#).

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 HPE 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" on page 51](#).

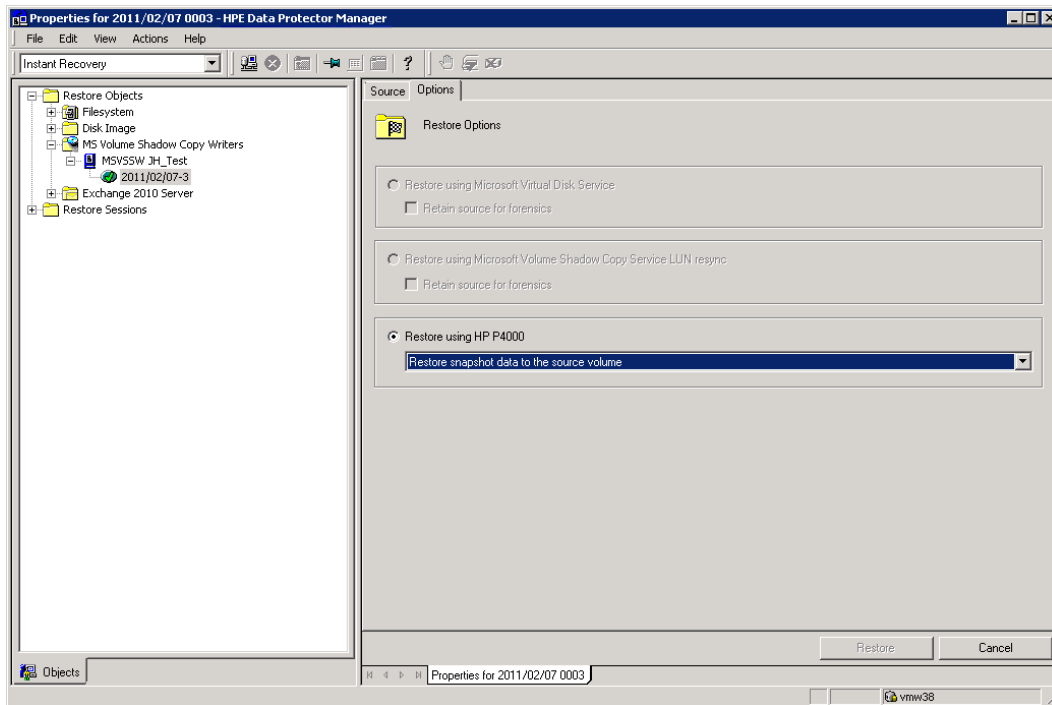
- **Restore using HPE P10000 3PAR**

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

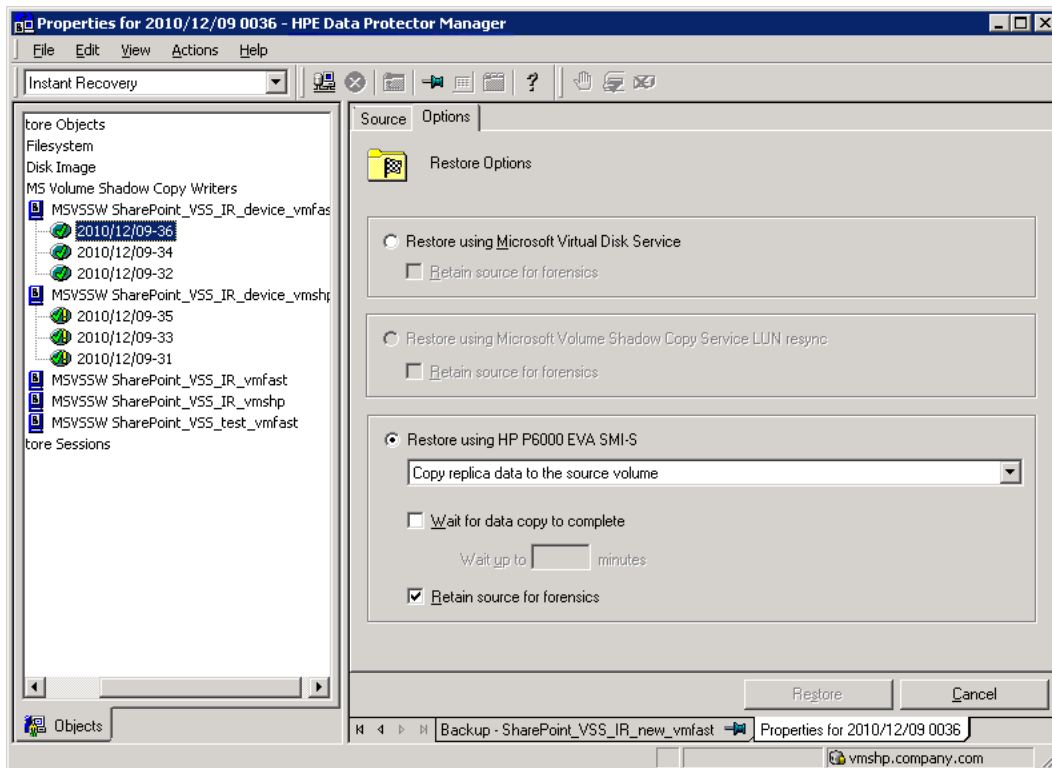
For details on these methods see ["Instant recovery methods" on page 51](#).

For details, press **F1**.

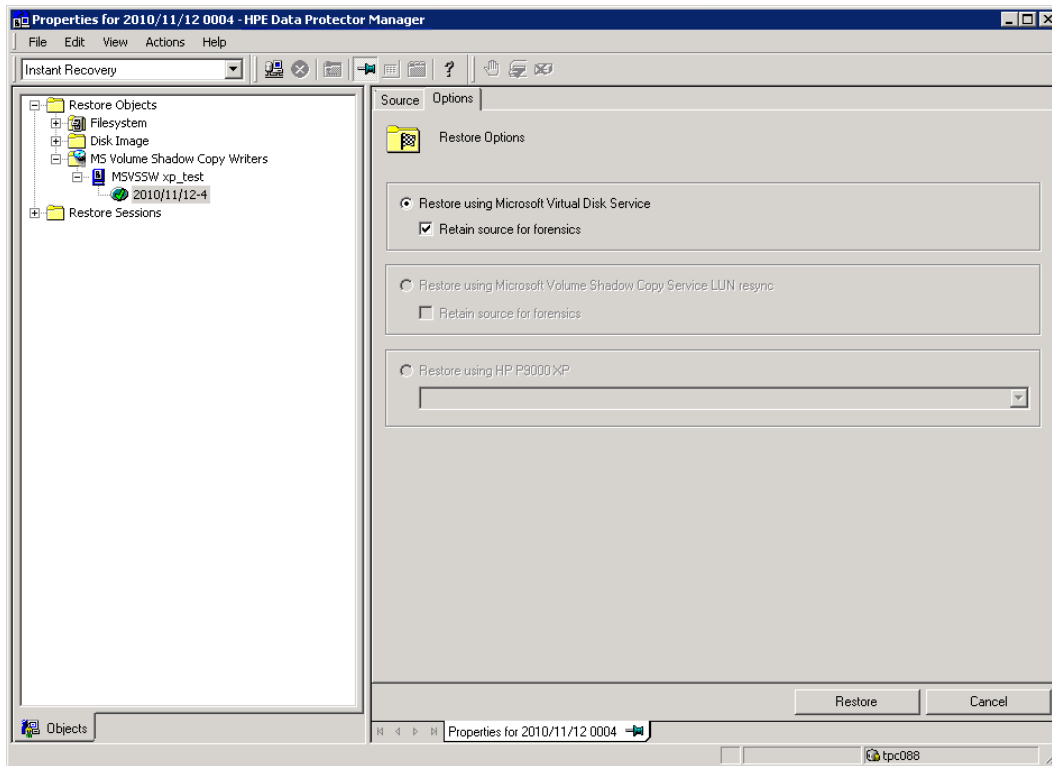
Selecting instant recovery options (P4000 SAN Solutions integration)



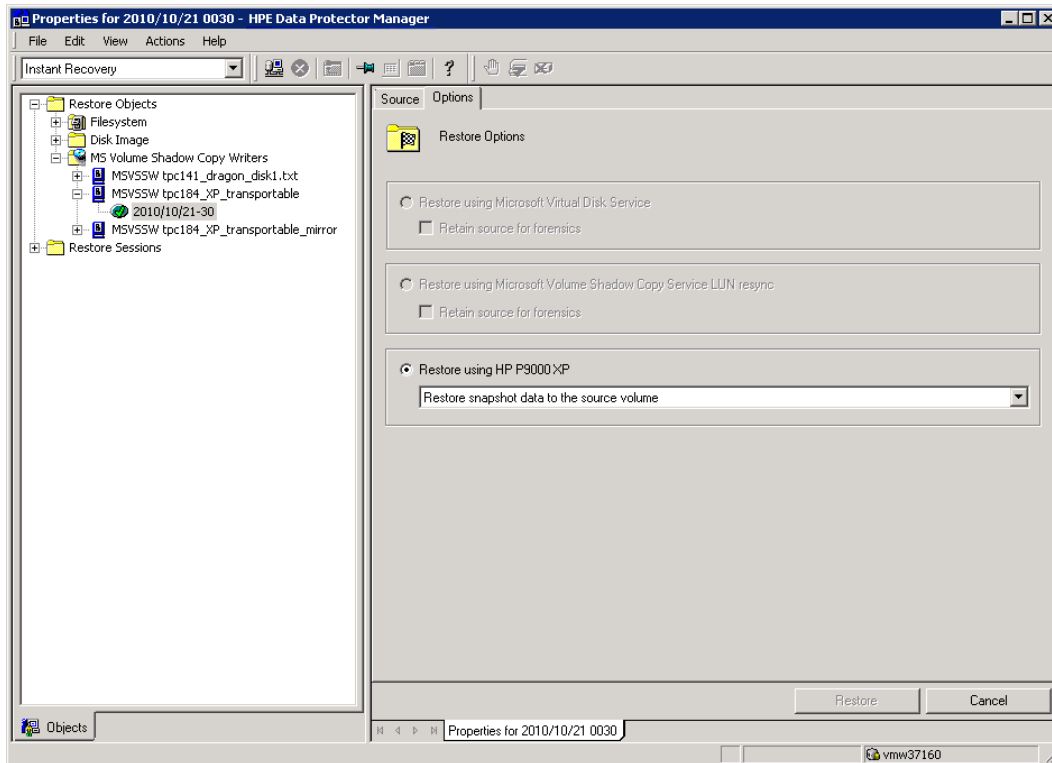
Selecting instant recovery options (P6000 EVA Array integration)




Selecting instant recovery options (P9000 XP Array integration, split mirror)



Selecting instant recovery options (P9000 XP Array integration, snapshot)



Selecting instant recovery options (3PAR StoreServ integration)



5. Click **Restore**.
6. Perform any additional writer specific steps. See ["Writer specifics"](#) on page 64.

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 *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

Chapter 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 <https://softwaresupport.hpe.com/>.

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](#) " below provides a description of VSS writers.

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,	Files are restored after a system restart.

Writer name	Description	Restore method
	such as starting and stopping services or the logging on and logging off of a user.	
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 system that can respond to queries from client machines formatted in the SQL language.	See "MSDE writer specifics" on page 110.
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" on page 67.
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" on page 71, "Microsoft Exchange Server 2007 writer specifics" on page 78, or "Microsoft Exchange Server 2010 writer specifics" on page 94.

Writer name	Description	Restore method
Microsoft Virtual Server 2005 Writer	<p>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.</p> <p>Cluster configurations are not supported, only individual nodes can be backed up.</p>	Standard VSS restore and instant recovery.
Microsoft Hyper-V Writer	<p>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.</p> <p>Cluster-aware backups are not supported.</p>	Standard VSS restore and instant recovery. For writer specifics, see "Microsoft Hyper-V writer specifics" on page 102.
SharePoint Services Writer	<p>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.</p> <p>Only a single server configuration (farm) is supported.</p> <p>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.</p>	Standard VSS restore. For writer specifics, see "Microsoft SharePoint Services writer specifics" on page 106.
NTDS Writer	<p>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.</p>	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	Custom restore method

Writer name	Description	Restore method
	back up and restore Windows Registry. Windows Registry is a database 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

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.

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" below](#)
- 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" on the next page](#).


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" on 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*.
Restoring the Microsoft Data Protection Manager client



6. Use the DPM to restore individual clients.

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 "[Standard restore](#)" on page 46 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.

This backup type is not available with the VSS hardware providers.

- *Differential* - similar to incremental backup, but the transaction logs are not truncated.

This backup type is not available with the VSS hardware providers.

- *Copy* - similar to full backup, with the difference that the transaction logs are not truncated. This backup type is not intended for use in recovering failed systems.

Limitations

- A combination of VSS snapshot backups and non-VSS backups (for example, stream incremental 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 HPE 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.

Selecting Microsoft Exchange Server stores for rollforward recovery



3. Proceed as with general VSS writer restore. See ["Standard restore" on 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.

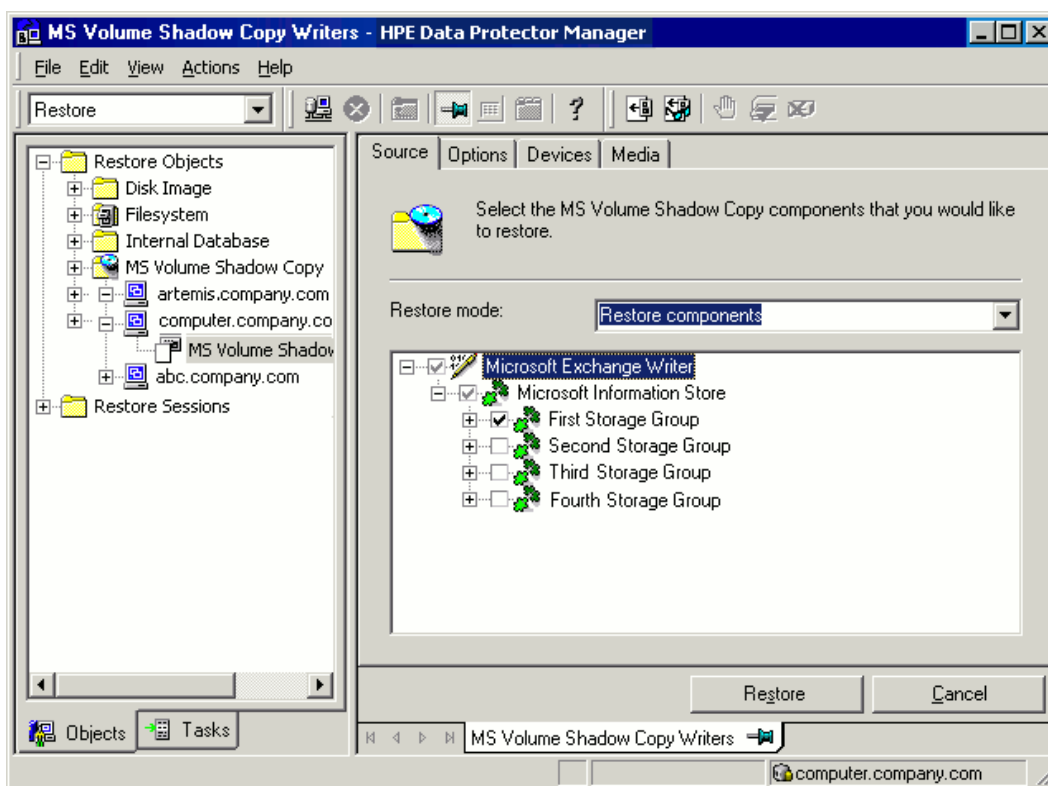
Point-in-time restore after loss of a log file

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.

Selecting Microsoft Exchange Server stores for point-in-time restore



3. Proceed as with general VSS writer restore. See ["Standard restore" on 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" on page 78](#).

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 HPE 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 HPE 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](#)" on [page 116](#) for details of how to avoid issues with too long copy creation.

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](#)" on page 76 and the backup policy given in the example was implemented.

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 - 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.

This backup type is not available with the VSS hardware providers.

- *Differential* - similar to incremental backup, but the transaction logs are not truncated.
This backup type is not available with the VSS hardware providers.
- *Copy* - similar to full backup, with the difference that the transaction logs are not truncated. This backup type is not intended for use in recovering failed systems.

Limitations

- A combination of VSS snapshot backups and non-VSS backups (for example, stream incremental 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 HPE 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 "[Selecting a replicated Microsoft Exchange Server 2007 storage group](#)" below.

Selecting a replicated Microsoft Exchange Server 2007 storage group



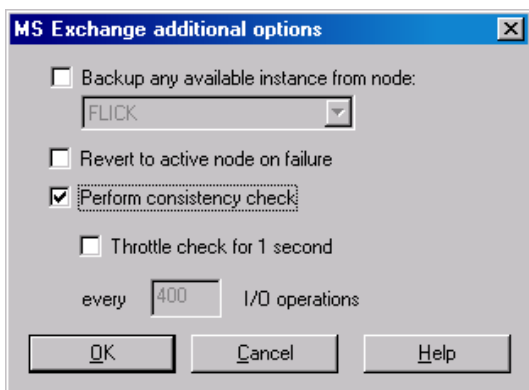
Cluster support in a CCR environment

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 "[Additional options for Microsoft Exchange Server 2007 in a CCR environment](#)" below.

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.

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 "[HPE P9000 XP Disk Array Family](#)" on page 33 and "[Limitations](#)" on page 91. 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" below](#).
- See ["Instant recovery" on page 90](#).

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" on page 85](#).

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.

Selecting Microsoft Exchange Server stores for rollforward recovery



3. Proceed as with general VSS writer restore. See ["Standard restore" on 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.

Selecting Microsoft Exchange Server stores for point-in-time restore



3. Proceed as with general VSS writer restore. See ["Standard restore" on 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"
```

- 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 RSG\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 Group/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
```

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

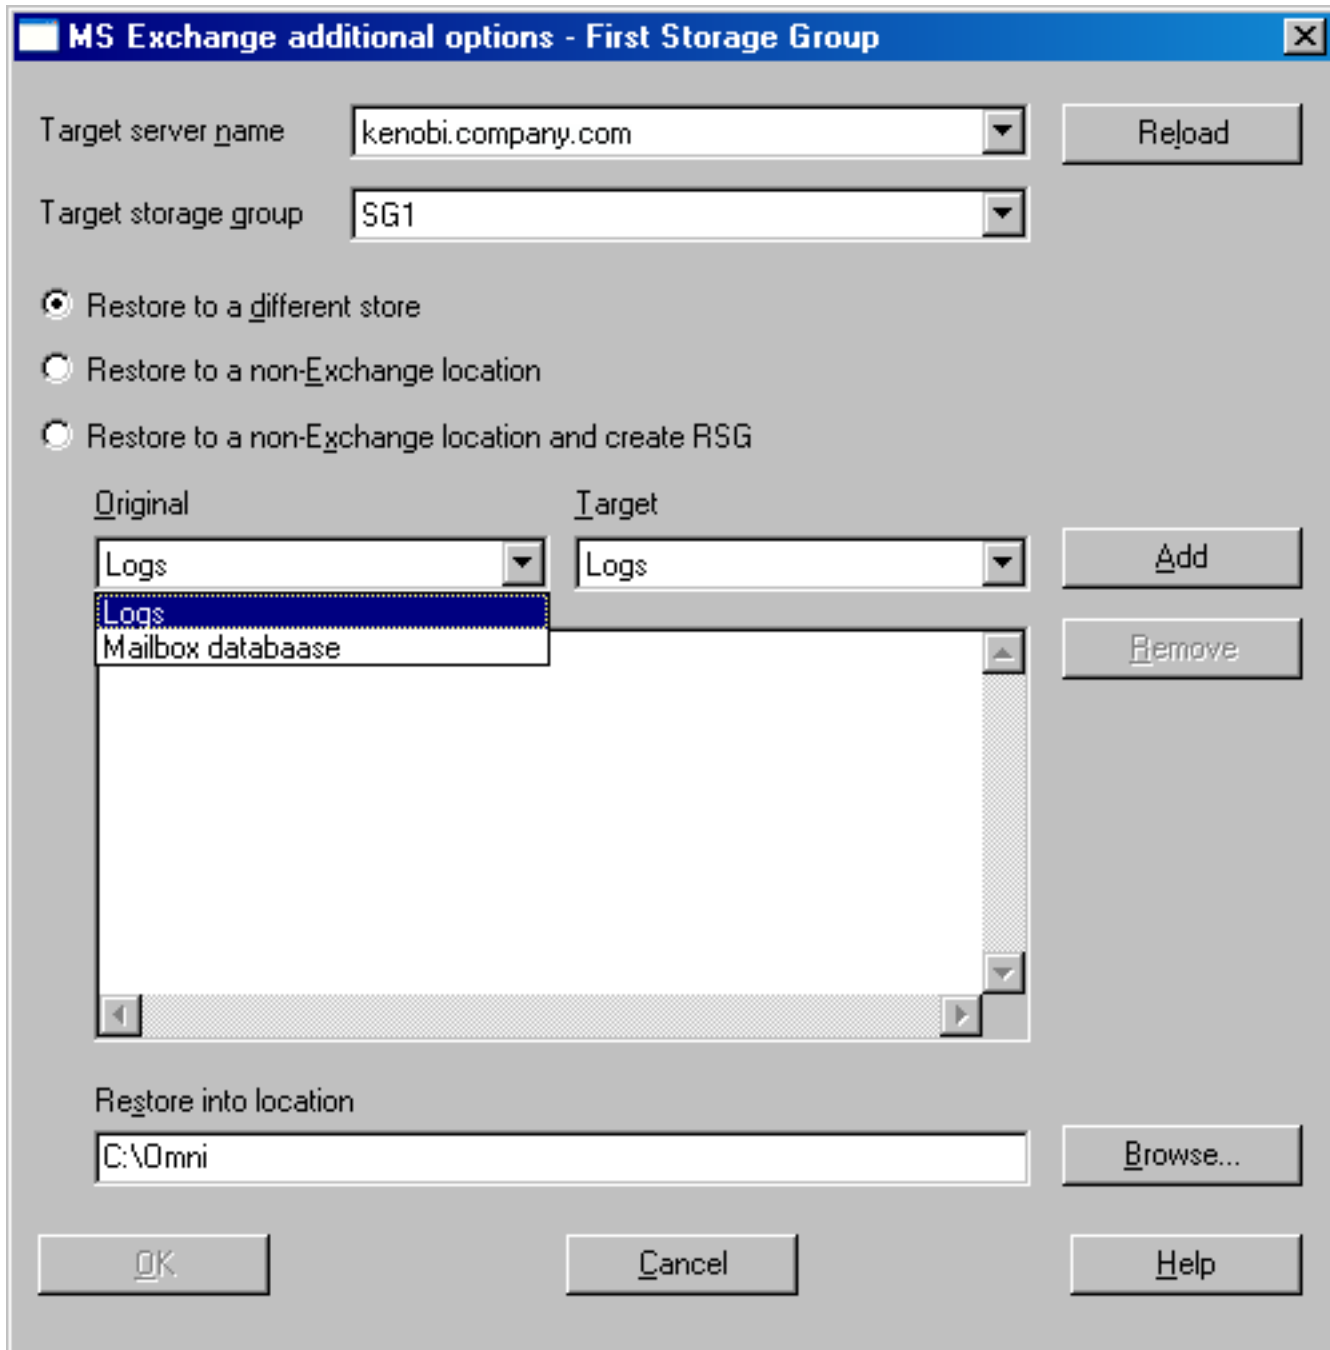
server name drop-down list, and then choose the desired stores using the Original drop-down list.

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.

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.

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](#)" on page 78.

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](#)" on page 85). 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**.

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 unrepresents 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 `omniarc` 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:

- a. Fail over Exchange Server. This will cause your production database to reside on the former backup system, where the database copy resided before failover.
- b. Perform instant recovery.
- c. 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" on the previous page](#).

Restoring an LCR or CCR copy to the original location

See "[Restoring a LCR or CCR copy to the original location](#)" on page 88.

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:

- i. On the application system, mount the restored store.
- ii. On the passive node, where the copy of the database exists, delete transaction logs.
- iii. 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.
- iv. 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](#)" on page 111.

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 *HPE 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 *HPE Data Protector Integration Guide* and the *HPE 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](#)" on [page 96](#) 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 VSS software provider)
- Differential (only with the VSS 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 ["Creating backup specifications" on page 35](#).

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**.

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 ["Restore" on page 46](#). 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**.

Enabling restore for a database



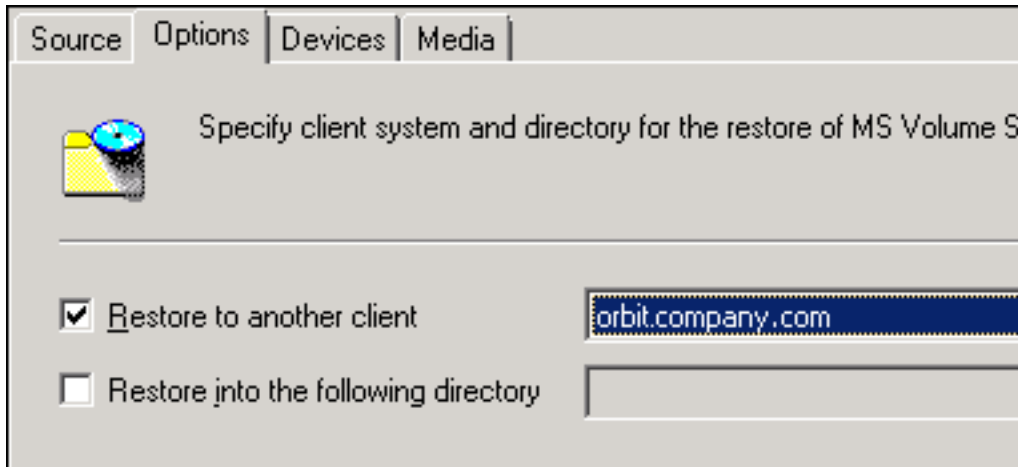
3. In the Data Protector GUI, select the **Restore** context. In the Scoping Pane click **MS Volume Shadow Copy** and select the client from which the database was backed up. You can select any client, not only the one to which you restore.
4. Depending on whether you backed up an active or passive copy,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

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.

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](#)" below.

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](#)" on the previous page.

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.

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](#)" on page 97
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](#)" on page 97.

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 "[Enabling restore for a database](#)" on page 97.
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](#)" on page 97.

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](#)" on page 97
3. Suspend the database replication.
4. Restore the active copy. See "[Restoring an active copy](#)" on the previous page.

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.

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](#)" on page 96.

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" on page 29](#). For general prerequisites and conditions (for example dismounting the Microsoft Exchange Server databases, suspending the replication), see ["Restore scenarios" on page 96](#). 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 HPE 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 *HPE Data Protector Command Line Interface Reference* for details on the `omndbvss` 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 *HPE 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" on page 111](#).

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](#)" below and "[Restore](#)" on page 105.

Prerequisites

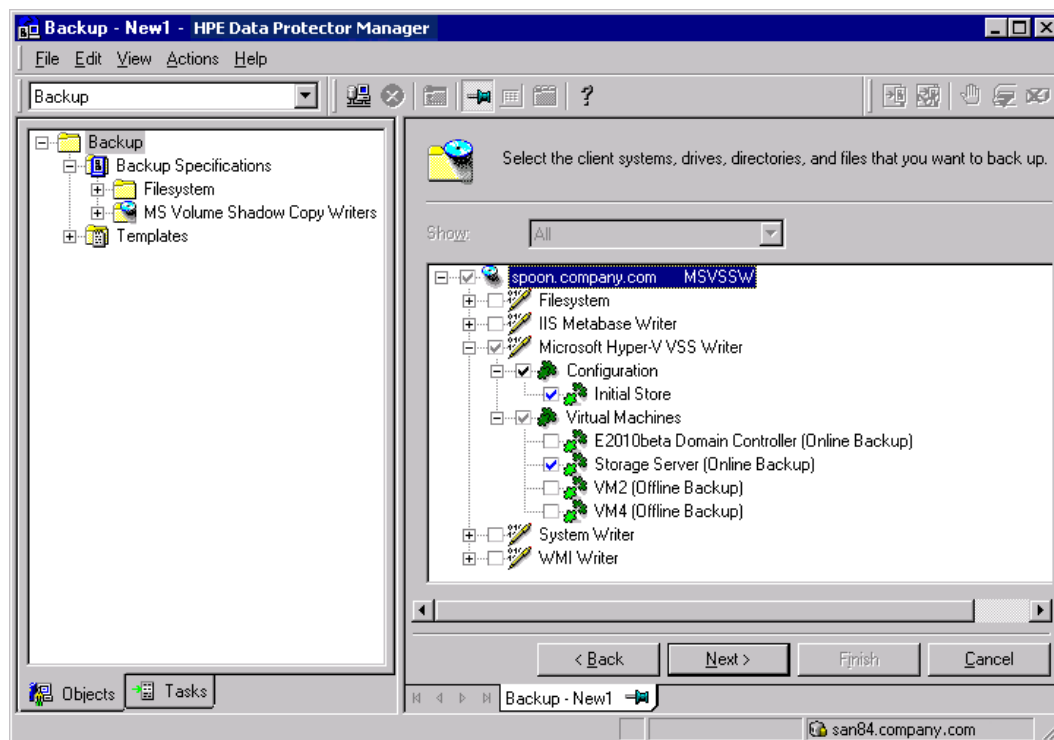
- Windows Server 2008 Service Pack 2 must be installed. For more information, see <http://support.microsoft.com/kb/948465>.

Backup

With the Hyper-V VSS writer, it is possible to back up:

- The Hyper-V configuration
- Virtual machines

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, execute `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 dismantled 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 HPE 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](#)" on 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 executing 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

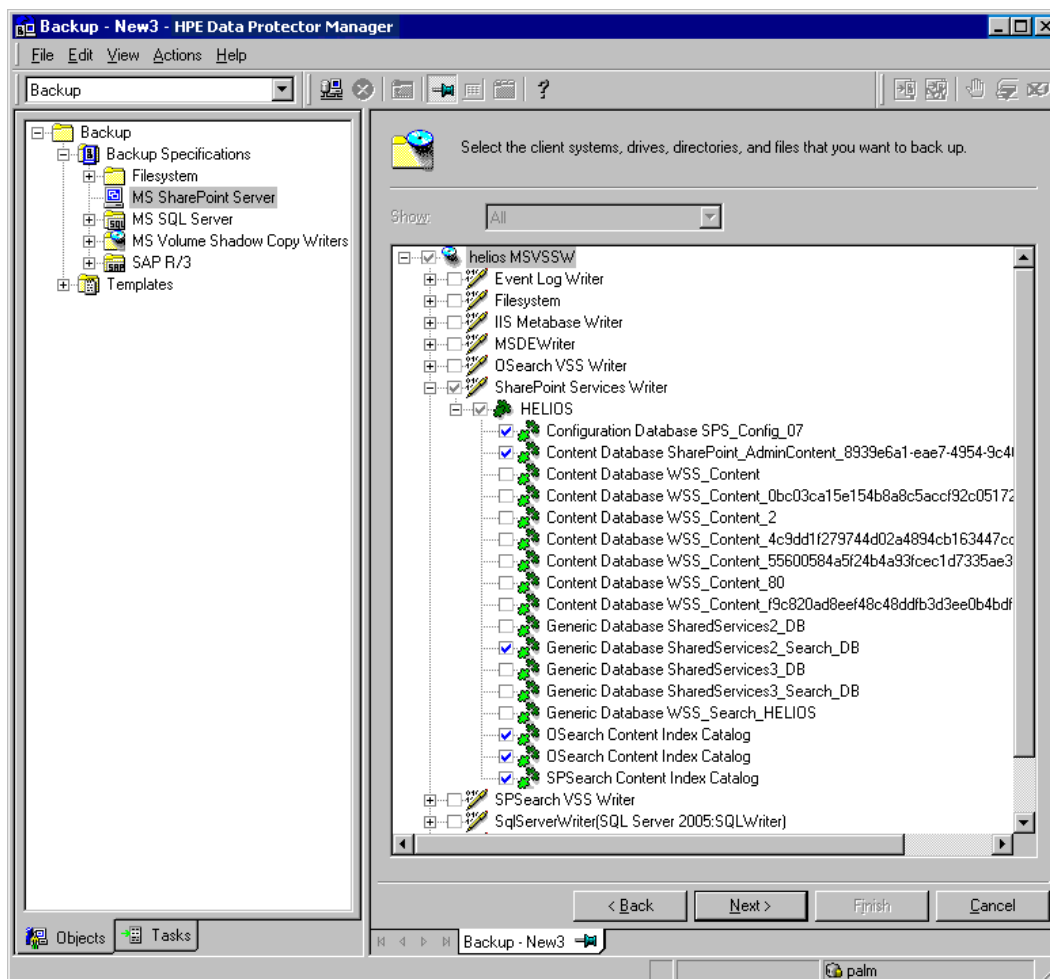
- If you back up individual writer components, to ensure that data is synchronized, the following combinations of items must be backed up in the same session:
 - The configuration database and central administration content database
 - Search databases and corresponding index files

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](#)" below for an example of how to select writer components.

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](#)" on the next page.

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.

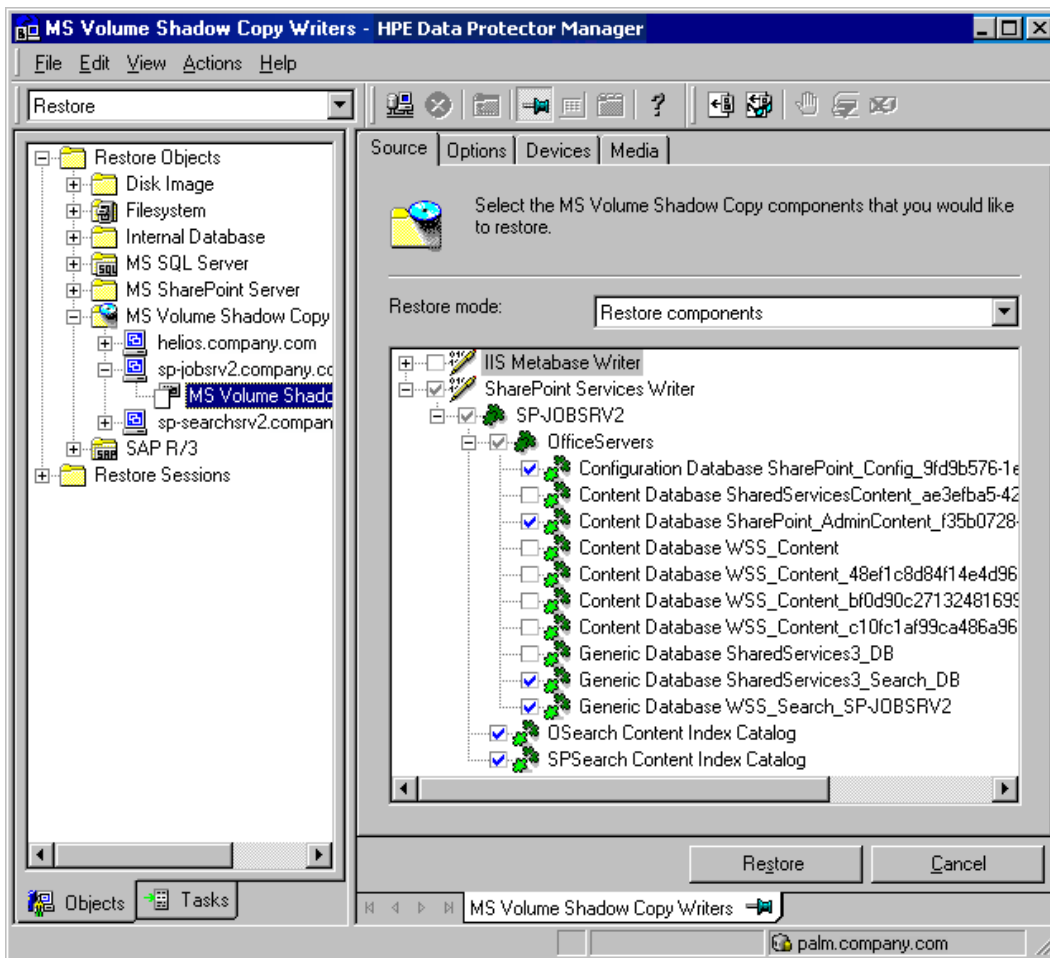
Prerequisites

- 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.

Selecting Microsoft SharePoint Services writer and the corresponding search writers for restore



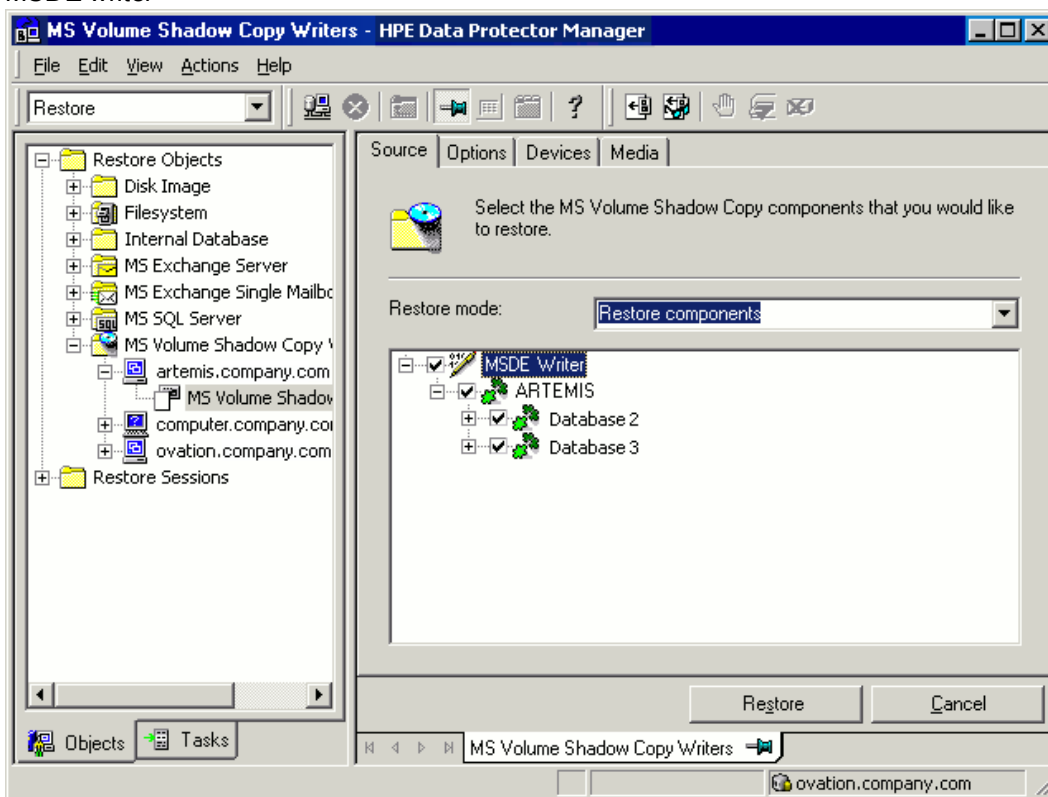
MSDE writer specifics

Restore

MSDE writer is used to back up and restore Microsoft SQL database.

Before restoring the SQL system databases (master, model, msdb and pub), you have to stop the SQL service.

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.

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.

Chapter 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 *HPE Data Protector Troubleshooting Guide*.

Before you begin

- Ensure that the latest official Data Protector patches are installed. See the *HPE Data Protector Help* index: “patches” for information of how to verify this.
- See the *HPE Data Protector Product Announcements, Software Notes, and References* for general Data Protector limitations, as well as recognized issues and workarounds.
- See <https://softwaresupport.hpe.com/> 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](#)" on page 71.

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:
 - a. On the passive node, suspend the replication by using `Suspend-StorageGroupCopy` cmdlet.
 - b. Delete all log files from the `Logs` directory of the database copy.
 - c. Seed the database copy or re-synchronize the original database and its copy by using the `Update-StorageGroupCopy` cmdlet.
 - d. Resume the database copy by using the `Resume-StorageGroupCopy` cmdlet.
 - e. 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-bfdb4c37890}
  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 *HPE 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 limitation 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

No HPE SMI-S P6000 EVA Array provider login entries are configured within SMISDB

Action

Add the login information for HPE SMI-S P6000 EVA Array provider by executing:

```
omnidbsmis -ompasswd -add ClientName [-ssl] [-port PortNumber] [-user Username] [-passwd Password]
```

Problem

Configuration of an HPE 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:

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.

Windows Server 2003, Windows Server 2008: In the Command Prompt window, execute the command `mountvol /R`.

It is recommended to periodically perform this cleanup task to prevent the Registry from running out of space and Data Protector ZDB sessions from failing. However, the task should be run under supervision and should not be automated.

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:

```
omnidbvs -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 *HPE 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 (the source volumes):
`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.

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 (the source volumes):
`omnidbvss -resolve -apphost AppSystem`
3. Resolve the backup system (the target volumes created in the backup session):
`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 HPE 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
```


Send Documentation Feedback

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

Feedback on Integration guide for Microsoft Volume Shadow Copy Service (Data Protector 9.09)

Just add your feedback to the email and click send.

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

We appreciate your feedback!