



Hewlett Packard
Enterprise

HPE Data Protector

Software Version: 9.08

Product Announcements, Software Notes, and
References

Document Release Date: October 2016
Software Release Date: October 2016

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 2016 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: Announcements	9
What is supported?	9
Support for earlier agent versions	9
Updated information	10
Chapter 2: Product features and benefits	11
New Features in Data Protector 9.08	11
Support for object copy of file system data to Microsoft Azure	11
Support for CDB and PDB mode in Oracle 12c	11
Previous Releases	11
Chapter 3: Limitations and recommendations	26
Limitations	26
Scalability limitations	26
Backup infrastructure scalability	26
Internal Database scalability	26
Concurrency limitations	27
Prerequisites for increasing the limit on concurrent backup sessions	28
Enhanced incremental backup	29
Size of file depots used for file library	29
Installation limitations	29
Upgrade limitations	30
Migration limitations	30
Platform limitations	30
UNIX and Linux limitations	30
HP-UX limitations	30
Solaris limitations	31
Linux limitations	31
Mac OS X limitations	31
Windows limitations	32
Windows 32-bit limitations	32
Windows 64-bit limitations	33
Windows XP and Windows Server 2003 limitations	33
Windows Server 2012 limitations	33
Novell Open Enterprise Server (OES) limitations	33
HP OpenVMS limitations	33
Limitations on clusters	36
HPE Serviceguard limitations	36
Limitations on licensing	36

General licensing limitations	36
License upgrading limitations	37
Internet Protocol version 6 (IPv6) networking limitations	37
Limitations on license reporting in a traditional licensing model	37
Limitations on encryption	37
Limitations on data encryption	37
Limitations on encrypted control communication	38
Limitations on Data Protector MoM environments	38
Device and media limitations	38
NDMP limitations	39
NetApp filer	40
Celerra	40
Limitations on enhanced incremental backups	40
Limitations on virtual full backups	41
Limitations on object copy and consolidation	41
Limitations on object verification	41
General functionality limitations	41
Application integration limitations	42
Limitations on application integrations	42
General limitations	42
Oracle limitations	42
MySQL limitations	43
SAP R/3 limitations	43
Informix Server limitations	43
Microsoft SQL Server limitations	43
Microsoft Exchange Server limitations	43
Microsoft Volume Shadow Copy Service limitations	44
Common VSS limitations	44
Microsoft Exchange Server 2003	44
Microsoft Virtual Server 2005	44
Microsoft SQL Writer	44
Data Protector Virtual Environment Integration limitations	44
VMware limitations	44
Lotus limitations	45
Limitations on disk array integrations	45
HPE P4000 SAN Solutions limitations	45
HPE P6000 EVA Disk Array Family limitations	45
HPE P9000 XP Disk Array Family limitations	46
HPE 3PAR StoreServ Storage limitations	47
EMC Symmetrix disk array limitations	48
NetApp Storage limitations	48
EMC VNX limitations	48
EMC VMAX limitations	48
Disaster recovery limitations	49
User interface limitations	49
Reporting limitations	49

Other limitations	50
Recommendations	53
Using Data Protector with hierarchical storage management applications	53
Organizing Data Protector clients into cells	53
Support for NIS+	53
Large file support	53
Encrypted control communication recommendations	54
Pre-exec and Post-exec scripts	54
Enhanced incremental backup	54
Object consolidation	55
Microsoft Exchange Single Mailbox integration	55
Microsoft Volume Shadow Copy Service integration	55
Shadow copy storage area and disk space recommendations	55
Regular maintenance of the VSS part of the registry	55
Network Data Management Protocol Server integration	56
Windows Server 2008 clients	56
Windows Server 2012 clients	57
UNIX system clients	57
Chapter 4: Recognized issues and workarounds	58
Known Data Protector issues and workarounds	58
Installation and upgrade related issues	58
User interface related issues	61
Disk Agent related issues	63
Media Agent related issues	66
Integration related issues	68
Microsoft Exchange Server	68
Microsoft Exchange Single Mailbox	68
Microsoft SQL Server	69
Microsoft Volume Shadow Copy Service	69
SAP R/3	70
Oracle Server	70
VMware vSphere	71
Lotus Notes	72
Disk array integrations	73
Granular Recovery Extension issues	73
VMware vSphere	73
Microsoft SharePoint Server	74
Disaster recovery issues	74
Cluster related issues	75
Common issues	75
Issues in HPE Serviceguard	76
Issues in Microsoft Cluster Server	76
Reporting related issues	77
Other known issues	77

Known non-Data Protector issues and workarounds	83
Non-Data Protector issues related to installation or upgrade	83
Non-Data Protector issues related to user interface	84
Non-Data Protector issues related to Disk Agent	85
Non-Data Protector issues related to Media Agent	86
Non-Data Protector issues related to integrations	87
Microsoft Exchange Server	87
Microsoft SQL Server	87
Microsoft Volume Shadow Copy Service	87
Microsoft SharePoint Server	89
SAP MaxDB	89
SAP HANA Appliance	90
Oracle Server	90
VMware vSphere	91
Sybase Server	92
Disk array integrations	92
Non-Data Protector issues related to Granular Recovery Extensions	95
VMware vSphere	95
Microsoft Exchange Server	95
Non-Data Protector issues related to disaster recovery	95
Non-Data Protector issues related to reporting	96
Other known non-Data Protector issues	97
Chapter 5: Resolved issues and enhancements	99
List of enhancements implemented and resolved issues in Data Protector	99
Resolved issues	99
Chapter 6: Data Protector documentation	107
Documentation map	107
Abbreviations	107
Integrations	110
Localized documentation	112
Send Documentation Feedback	113

Chapter 1: Announcements

HPE Data Protector automates high performance backup and recovery, from disk or tape, over unlimited distances, to ensure 24x7 business continuity, and seamless integration with HPE storage hardware solutions. Data Protector delivers innovation and performance at a much lower cost than competitive solutions, while offering flexibility, scalability, and performance. Data Protector is an important member of the fast-growing HPE Software portfolio and offers the unique advantage of being able to source hardware, software, and award winning service offerings from a single, trusted source. Data Protector is both easy to deploy and use. It has a simple installation, automated routine tasks, and centralized licensing facility that reduces costs and data center complexity.

Now announcing its latest version: Data Protector 9.08.

What is supported?

Detailed information about supported platforms, devices, and integrations is available in the support matrices, which can be found on any Data Protector installation DVD-ROM in the \DOCS\support_matrices directory. The following support matrices are available in Portable Document Format (PDF):

- *HPE Data Protector 9.08 3PAR Support Matrix*
- *HPE Data Protector 9.08 Device Support Matrix*
- *HPE Data Protector 9.08 Disaster Recovery Support Matrix*
- *HPE Data Protector 9.08 Network Attached Storage (NAS) Support Matrix*
- *HPE Data Protector 9.08 Platform and Integration Support Matrix*
- *HPE Data Protector 9.08 Virtualization Support Matrix*
- *HPE Data Protector 9.08 VSS Integration Support Matrix*
- *HPE Data Protector 9.08 Zero Downtime Backup and Instant Recovery Support Matrix for HPE P6000 EVA Disk Array Family Using SMI-S Agent*
- *HPE Data Protector 9.08 Zero Downtime Backup and Instant Recovery Support Matrix for HPE P9000 XP Disk Array Family*
- *HPE Data Protector 9.08 Zero Downtime (Split-Mirror) Backup Support Matrix for EMC Arrays*
- *HPE Data Protector 9.08 Zero Downtime Backup Support Matrix for Non-HPE Storage Arrays*

For the latest version of support matrices on the Web, see <https://softwaresupport.hpe.com/>.

In the event of hardware or software failures on third-party products, please contact the respective vendor directly.

Commands of the Data Protector command-line interface (CLI) are documented in the *HPE Data Protector Command Line Interface Reference*.

Support for earlier agent versions

Wherever possible, Data Protector components on all clients in a Data Protector cell should be upgraded to version 9.08 during the regular upgrade process. This ensures that customers can benefit from the full feature

set of Data Protector 9.08 on all systems in a cell.

Nevertheless, Disk Agent and Media Agent components of an earlier Data Protector version (6.20, 7.00, or 8.00) are supported in a 9.08 cell with the following constraints:

- The earlier product version is still supported by HPE as an independent product. To check the announced end-of-support dates for HPE products, see the webpage <https://softwaresupport.hpe.com/>.
- Support is limited to the feature set of the earlier Data Protector version.
- If you are performing operations involving clients on different systems, all agents of the same type (for example Media Agents) must be of the same version.
- Earlier Media Agent component versions are not supported in combination with NDMP servers.
- If one Data Protector component on a client is upgraded to 9.08, all other components have to be upgraded to 9.08 as well.

If you encounter problems establishing a connection with agents of an earlier product version, consider upgrading to 9.08 as the first resolution step.

Note: Media Agents (8.0x and earlier clients) used as gateways are not compatible with the StoreOnce Software (SOS) server (8.1x or later).

Compatibility issues exist between the following:

- Clients before 7.03 and 7.03 server.
- 7.03 client and SOS server versions (8.0x and later).

Updated information

For the latest information about the product, see the Data Protector website <http://www.hpe.com/software/dataprotector>.

For the latest version of the Data Protector documentation set, including corrections and last-minute updates due to known issues, see <https://softwaresupport.hpe.com/>.

Chapter 2: Product features and benefits

The section lists all the new features and benefits.

New Features in Data Protector 9.08

The following key features are introduced in Data Protector 9.08:

- [Support for object copy of file system data to Microsoft Azure](#)
- [Support for CDB and PDB mode in Oracle 12c](#)

This section provides a brief description of the key Data Protector features that are introduced in the 9.08 release.

Support for object copy of file system data to Microsoft Azure

A new device is introduced to enable object copy to Microsoft Azure object store from Data Protector. This new device sends data to Cloud (Azure) containers and is created in a similar way as other devices in Data Protector. To create such a device in the GUI, the Backup to Disk device type with Cloud (Azure) interface is selected, and the parameters pertaining to account name and access keys are specified.

Support for CDB and PDB mode in Oracle 12c

A new feature is introduced in Oracle 12c for backup and restore of data called Pluggable Database (PDB). The PDB is located under container database (CDB) in Oracle 12c. The Backup Archive and Restore (BAR) GUI lists all the PDBs in the container database, selection of only one or more pluggable databases should be enabled.

Full and incremental backups are supported and from restore perspective recovery of CDB or any PDB are supported. Users can choose between point in time restore or until now restore.

Previous Releases

The following key features were introduced in the previous releases:

VMware/Virtual Environment Protection Agent (VEPA)

- [Support for SAN transport backups for VMware](#)
- [Support for VMware GRE with Windows Server 2012 Clustered Cell Managers](#)
- [Support for VDDK 6.0 Update 2](#)
- [Parallel VMDK backup](#)
- [VMware virtual machine restore from 3PAR Replica](#)

- Virtual Machine Power On from the Backup Image
- VMware Virtual Machine Live Migrate
- VEPA Application Aware Consistency
- Restore at VHD level in Hyper-V integration
- VEPA Platform and API Support
- VADP Reporting

Granular Recovery Extension

- Advanced GRE Web Plug-in for Data Protector
- Non-cached recovery using GRE for VMware vSphere
- VMware GRE from 3PAR snapshot for VMDK larger than 256GB
- VMware GRE GPT Support
- Cross Mounting
- Cached GRE, Power On and Live Migrate for VMware backups to StoreOnce Catalyst device

Zero Downtime Backups

- ZDB support for EMC VMAX
- ZDB support for EMC VNX
- Support for 3PAR ZDB/IR within a VMware VM
- Data Protector Virtual Environment ZDB integration
- 3PAR Remote Copy
- ZDB/IR support for 3PAR enhancements

Storage and Devices

- Support for backup or restore of Virtual Machine and Cinder Volume
- Support for Cloud (Helion) device
- Support for NetApp Storage
- Smart Cache device
- Support for StoreOnce Federated Catalyst
- Support for Data Domain Boost device replication
- Veritas Cluster 6.2 Support

SharePoint

- SharePoint Custom ID

MySQL

- Integration with MySQL
- Integration with PostgreSQL
- MySQL

Miscellaneous

- Support for Data Protector Telemetry
- Support for BTRFS
- Support for Storage Optimizer version 5.x
- Support for HPE Storage Systems with InForm OS 3.2.2
- Reconnect functionality for Copy Session Manager (CSM)
- StoreOnce Recovery Manager Central Integration
- Command Line Interface enhancements for bulk modification of backup specifications
- Support for Open SSL FIPS Object Module
- HPE Storage Optimizer is accessible from the Data Protector GUI
- Support for StoreOnce Application Source and Backup Server Deduplication using Catalyst
- Support for EMC Data Domain Boost Application Source and Backup Server Deduplication
- Enhanced Encrypted Control Communication
- New menu item under the Help menu
- HPE Storage Optimizer is accessible from the Data Protector GUI
- Scheduling and priority in Advanced Scheduler
- Omnicc command enhanced to include total data under protection
- Automated Replication Synchronization
- Site Specific Patch and Test Module installation enhancements
- Improved debug log gathering tools
- Platform coverage enhancements
- Enhanced device filtering
- Support for multiple interfaces to the same store
- Collection of high-level information for telemetrics
- Encrypted Control Communication enhancements
- NDMP integration enhancements
- Non-Changed Block Tracking (Non-CBT) backup

This section provides a brief description of the key Data Protector features that were introduced in the previous releases.

VMware virtual machine restore from 3PAR Replica

For zero downtime backups of virtual machines when **Disk** or **Disk+Tape** option is selected, the restore of virtual machines from the 3PAR replica is now supported.

Virtual Machine Power On from the Backup Image

Previously, virtual machines had to be powered on only after the complete data restore to the production data center. Virtual machines can now be powered on from the Data Protector backup image that resides on either the 3PAR replica or Smart Cache device. Use this feature if you want to verify the sanity of the backup.

Virtual Machine Live Migrate

This option will power-on the virtual machine from the backup image, and will simultaneously start the data migration to the destination data store. During this process, the virtual machine will continue to be accessible. Since the data movement is a back-end operation, it will have minimum impact on the usage and accessibility of the powered on virtual machine. Any modifications done to the virtual machine data will be consolidated, and the migrated virtual machine will have all the modified content on top of the restored image from the backup.

VEPA Application Aware Consistency

When the Quiescence option is enabled, backup of virtual machines are application consistent for MS SQL, MS SharePoint, MS Exchange and Oracle.

Restore at VHD level in Hyper-V integration

Data Protector introduces support for selecting individual disks while restoring the backed up Hyper-V virtual machines.

VEPA Platform and API Support

Data Protector VEPA component introduces:

- Support for Virtual Disk Development Kit (VDDK) 6.0
- Support for vSphere 6.0.
- Support for Virtual Volume (VVOL) datastore. vSphere 6.0 includes Virtual Volumes (VVols) for virtual machine storage at the VMDK level. From the backup and restore perspective, virtual machines on the VVOL datastore remains transparent, as they are handled by the VMware API.

VADP Reporting

This feature introduces changes in Data Protector reporting for VEPA objects, more specifically for VMware virtual machines (VCD and Hyper-V are not supported).

Previously, reports displayed ESXi Server and vCenter as clients and the object name contained the UUIDs of the Virtual Machines. With this enhancement, the reports now display additional information about virtual machines, such as DNS name of the Guest OS hosted by the VM, its IP address, path in datacenter, and operating system.

This enhancement is available for the following commands:

- omnirpt
- omnicellinfo
- omnidb
- omnimm

For more information on these Data Protector commands and command-line options, see the *HPE Data Protector Command Line Interface Reference*. Also, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

Advanced GRE Web Plug-in for Data Protector

The Data Protector GRE is now accessible through a new user interface; Advanced GRE Web Plug-in. This plug-in enables you to perform restore, presentation and recovery operations for both cached and

non-cached backups. Backups done to Smart Cache devices or backups from array snapshots residing on the HPE 3PAR storage are referred to as cached backups and the other backups are referred to as non-cached backups.

With the Advanced GRE Web Plug-in you can now import multiple vCenters into a single Cell Manager. You can also import a single vCenter into multiple Cell Managers.

For more information, see the HPE Data Protector Granular Recovery Extension User Guide.

For details on the supported operating systems, see the latest support matrices at: <https://softwaresupport.hpe.com/>.

Non-cached recovery

The non-cached recovery feature in GRE for VMware vSphere introduces presentation and recovery of files from the virtual machine image that is backed up to the 3PAR storage or Smart Cache device. With non-cached recovery, it is possible to control file recovery from a virtual machine backup, which is performed to a Smart Cache device or from a 3PAR ZDB backup, without restoring the VMDK files into a temporary location (mount proxy system).

For more information, see the *HPE Data Protector Granular Recovery Extension User Guide*.

For a list of supported file systems, operating systems on mount proxy systems, and partition types, see the latest support matrices at <https://softwaresupport.hpe.com/>.

VMware GRE from 3PAR snapshot for VMDK larger than 256GB

Granular Recovery Extension can be performed from the 3PAR snapshots for disks greater than 256 GB. The maximum disk size is 16 TB.

VMware GRE GPT Support

VMware Granular Recovery Extension (GRE) supports GUID Partition Table (GPT) disks (16TB max). VMware supports GPT disk size up to 64TB. Data Protector supports GRE for GPT disks up to 16TB.

Cross Mounting

In case of cached recovery operations, a Windows virtual machine can now be recovered using a Linux mount proxy. The ZDB recovery of Windows virtual machines on the 3PAR replica using a Linux mount proxy is also possible.

Cached GRE, Power On and Live Migrate for VMware backups to StoreOnce Catalyst device

Data Protector now supports the following restore operations from StoreOnce Catalyst backups:

- **Granular Recovery Extension:** The cached and non-cached recovery feature in GRE for VMware vSphere introduces presentation and recovery of files from the virtual machine image that is backed up (full, incremental and differential) to StoreOnce Catalyst device. For more information, see the *HPE Data Protector Granular Recovery Extension User Guide*.
- **Power On:** Virtual machines can now be powered on from the Data Protector backup image that resides on StoreOnce Catalyst device. Full, incremental and differential backups are supported. For more information, see the HPE Data Protector Integration Guide.
- **Live Migrate:** This option will power on the virtual machine from the backup image that resides on StoreOnce Catalyst device, and will simultaneously start the data migration to the destination datastore. During this process, the virtual machine will continue to be accessible. Since the data

movement is a back end operation, it will have minimum impact on the usage and accessibility of the powered on virtual machine. Any modifications done to the virtual machine data will be consolidated, and the migrated virtual machine will have all the modified content on top of the restored image from the backup. For more information, see the *HPE Data Protector Integration Guide*.

For more information on configuring a StoreOnce Catalyst device, see the *Configuring a backup to disk device - StoreOnce* page in *HPE Data Protector Help*.

For a list of supported file systems, operating systems on mount proxy systems, and partition types, see the latest support matrices at <https://softwaresupport.hpe.com/>.

Data Protector Virtual Environment ZDB integration

Data Protector offers Virtual Environment ZDB integration for VMware. This integration enables you to perform zero downtime backup of your virtual environment. The zero downtime backup (ZDB) functionality offers online backup capabilities with minimal degradation of the application system performance and the load on the application system is significantly reduced.

Virtual Environment ZDB integration for VMware supports environments where ESX and/or ESXi Server systems are set up with NetApp or 3PAR Storage System and managed through a vCenter Server (vCenter environments). Environments with standalone ESX(i) Server systems and mixed environments, in which some of the ESX(i) Server systems are managed through a vCenter Server system and some are standalone, are also supported.

For more information, see the *HPE Data Protector Zero Downtime Backup Integration Guide*.

Support for Cloud (Helion) device

Data Protector supports new Backup to Disk (B2D) device known as a Cloud (Helion) device configured with Cloud credentials. The Cloud (Helion) device is configurable on Windows and Linux media agent hosts. The Media Agent has been enhanced to act as a Cloud gateway to transmit data from on-premises backup devices to the Cloud (Helion) in object copy operations.

These object copy operations can be configured to run interactively, automated based on a schedule or as a post backup job action. Object copy to HPE's Public Cloud enables backup administrators to replicate backup jobs offsite electronically without the need for tape and can also be used as an offsite archive copy. All data transmitted to the cloud is encrypted and compressed before being sent and remains in this state when stored in the cloud.

For more information on configuring the Cloud devices, see the *HPE Data Protector Help*. For details on the supported operating systems, see the latest support matrices at: <https://softwaresupport.hpe.com/>.

Support for NetApp Storage

Data Protector introduces a NetApp Storage Provider to integrate with NetApp Storage and enable you to perform zero downtime backups. Data Protector NetApp Storage integration is supported for Windows and Linux systems. With NetApp Storage Provider, a ZDB to tape backup type is available, meaning you can restore the backed up data using conventional restore. Support is available for a filesystem, some ZDB application integrations (Oracle Server, SAP/R3, and Microsoft SQL Server), and Virtual Environment backup and restore.

For more information, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide* and the *HPE Data Protector Help*.

For details on the supported operating systems, storage system models, and versions of the storage systems' administrative software, see the latest support matrices at:

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

Smart Cache device

Data Protector introduces a new Backup to Disk (B2D) device known as Smart Cache. The Smart Cache device is configurable on Windows and Linux. To utilize the non-cached recovery feature in GRE for VMware vSphere, you must use a Smart Cache device for backups.

For more background information on Smart Cache devices, see the *HPE Data Protector Concepts Guide*. For more information on configuring the Smart Cache device, see the page "Configuring a Backup to Disk Device" in *HPE Data Protector Help*.

Support for StoreOnce Federated Catalyst

The StoreOnce Federated Catalyst feature enables distributed device presentation and deduplication. Data Protector supports StoreOnce software version 3.12.x and earlier. Data Protector can interface with:

- Expanded or contracted StoreOnce stores.
- Non-federated or federated stores over Fibre Channel.

For the StoreOnce Backup system, in the Deduplication System box, you can enter an IPv4 or IPv6 address, a fully qualified domain name (FQDN), or a Fibre Channel global identifier.

Support for Data Domain Boost device replication

Data Domain Boost devices provide functionality to replicate data from one device to another. Data Protector support this functionality through object copy sessions.

For more information on supported platforms, see the latest support matrices at

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

Support for EMC VNX and EMC VMAX Storage families

Data Protector introduces a EMC VNX Storage Provider and EMC VMAX Storage Provider to integrate with respective disk arrays and enable you to perform zero downtime backup. Both integrations are supported for Windows and Linux systems and enable ZDB to tape backup type. Support is available for a filesystem and ZDB application integrations (Oracle Server and Microsoft SQL Server) backup and restore. With EMC VNX Storage, Data Protector supports VNX Snapshot product. With EMC VMAX Storage, Data Protector supports TimeFinder/Clone and TimeFinder VP Snap products.

For more information, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide* and the *HPE Data Protector Help*.

For details on the supported operating systems, storage system models, and versions of the storage systems' administrative software, see the latest support matrices at:

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

Veritas Cluster 6.2 Support

Data Protector supports Symantec Veritas Cluster Server for Red Hat Enterprise Linux 6.6 operating system.

Symantec or Veritas Cluster Server (VCS) connects multiple, independent systems into a management framework for increased availability. Each system, or node runs its own operating system and cooperates at the software level to form a cluster.

Support for backup or restore of Virtual Machine and Cinder Volume

Data Protector introduces support for backing up the Nova Instance VM and Cinder Volume (Shadow VM) in the HPE Helion OpenStack 2.1 cloud system with VMware vCenter.

3PAR Remote Copy

You can now perform backup, restore, and instant recovery procedures from storage volumes that are part of the 3PAR Remote Copy groups.

See the *HPE Data Protector 3PAR Support Matrix* for information on the supported operating systems and applications.

ZDB/IR support for 3PAR enhancements

Data Protector introduces support for the following 3PAR features:

- ZDB/IR remote copy for virtual environment integration.
- Host set and Port present presentation types.
- Volume set storage volumes.
- Advanced GRE, Power On, and Live Migrate with remote copy snapshots.

SharePoint Custom ID

Granular recovery of SharePoint lists, which are customized with a Custom ID value greater than 10000 is now supported.

Integration with MySQL

Data Protector introduces integration with MySQL database server. Data Protector MySQL integration takes advantage of MySQL Enterprise Backup (MySQL tool) to provide enterprise-grade backup and recovery for MySQL through a System Backup to Tape (SBT) interface. Data Protector integrates with MySQL using the MySQL Integration agent, which provides the Data Protector interface for the MySQL tool.

For more information, see the *HPE Data Protector Integration Guide*.

MySQL

MySQL integration takes advantage of MySQL Enterprise Backup (MySQL tool) to provide enterprise-grade backup and recovery for MySQL through a System Backup to Tape (SBT) interface. Data Protector integrates with MySQL using the MySQL Integration agent, which provides the interface for the MySQL tool.

MySQL 5.5/5.6 is supported on Windows 2008/2012, on RHEL 5/6 and on SUSE 11 operating systems.

Integration with PostgreSQL

Data Protector introduces integration with PostgreSQL database server. PostgreSQL integration takes advantage of the PostgreSQL Write-Ahead Logging technology to provide enterprise-grade backup and recovery for the PostgreSQL data. Data Protector integrates with PostgreSQL using the PostgreSQL Integration agent, which provides the interface for PostgreSQL.

For more information, see the *HPE Data Protector Integration Guide* and *HPE Data Protector Command Line Interface Reference*.

Reconnect functionality for Copy Session Manager (CSM)

Data Protector attempts to reconnect broken connections between Media Agent (BMA, RMA, or MMA) and CSM, when running copy, consolidation, or replication sessions.

By default, the Data Protector reconnect broken connections functionality is enabled.

To disable the reconnect broken connections functionality, set the `omnirc` option `OB2_CSM_NORECON` to 1 on the cell server.

Data Protector tries to reconnect for 20 minutes by default. To modify this time-out period, set the `omnirc` option `OB2RECONNECT_RETRY` on the cell server and client(s). The cell server and client values must be in sync.

Note: The `OB2RECONNECT_RETRY` option on client refers to the duration covered by client in attempting to reconnect after an error. The `OB2RECONNECT_RETRY` option on server refers to the duration covered by server in waiting for the client to reconnect.

StoreOnce Recovery Manager Central Integration

StoreOnce Recovery Manager Central (RMC) software integrates HPE 3PAR StoreServ primary storage with HPE StoreOnce Backup systems. HPE StoreOnce RMC for VMware enables you to protect VMware Virtual Machine Disks (VMDKs) and data stores using application-consistent snapshots for rapid online recovery.

With Data Protector, you can backup virtual machines snapshots created by RMC to Data Protector supported secondary storage devices. You can then perform a restore operation to the required destination. Note that Granular Recovery Extension (GRE) operations are possible only for Snapshot+Tape backups.

For more details, see the *HPE Data Protector Integration Guide*.

Command Line Interface enhancements for bulk modification of backup specifications

The new CLI command `omniwl.pl` introduced as part of this feature allows you to modify a backup specifications. The input for modifying specifications must be present in a CSV file, with semicolon (;) as the delimiter. When you execute the `omniwl.pl` command, the CSV file is read, then the specified backup specifications are modified, and finally the new specifications files are placed in the specified directory.

For more details, see the *HPE Data Protector Command Line Interface Reference Guide*.

Support for Open SSL FIPS Object Module

Data Protector now uses Open SSL version 1.0.2d. Open SSL FIPS Object Module is added to the `libcrypto` library. For more information on the FIPS 140-2 Consolidated Validation Certificate No. 0018, see <http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140crt/FIPS140ConsolidatedCertList0018.pdf>

When AES-256 is selected for data security, the Disk Agent will switch to FIPS mode for data encryption. FIPS mode of operation is used by default on various platforms. For details on the supported platforms, see the Open SSL FIPS Object Module v2.0 guide at: <https://openssl.org/docs/fips/UserGuide-2.0.pdf>.

FIPS mode can be switched off using the `Omnirc` option. The `OB2_NO_FIPS_MODE` value should be set to true. For more information on the `Omnirc` Options, see the *HPE Data Protector Command Line Interface Reference Guide*.

HPE Storage Optimizer is accessible from the Data Protector GUI

The HPE Storage Optimizer user interface can be launched from the Data Protector GUI. The variable `StorageOptServer` must be enabled in the Data Protector Global file. You can launch Storage Optimizer by navigating to **Backup > Actions > HPE Storage Optimizer**.

For more information, see the HPE Data Protector Help.

Support for StoreOnce Application Source and Backup Server Deduplication using Catalyst

Data Protector includes support for StoreOnce Application Source and Backup Server Deduplication using Catalyst with Solaris (SPARC) and AIX platforms.

For more details on the supported operating systems, see the latest support matrices at: <https://softwaresupport.hpe.com/>.

For more information on Backup performance with StoreOnce devices, see the *Deduplication White Paper*.

Support for EMC Data Domain Boost Application Source and Backup Server Deduplication

Data Protector includes support for EMC Data Domain Boost Application Source and Backup Server Deduplication with Solaris (SPARC) and AIX platforms.

For more details on the supported operating systems, see the latest support matrices at: <https://softwaresupport.hpe.com/>.

Enhanced Encrypted Control Communication for Data Protector

Data Protector has the option to enable Encrypted Control Communication to raise security. With the latest release, additional features are being supported.

You can now, do the following:

- enable ECC on clients by default during remote installation or import.
- disable ECC for a client.
- use one CA per cell, reusing the CA currently used for JBoss.
- use the simplified GUI or the CLI to decide if you want to use the existing certificates or regenerate them.
- add your own external script (predefined name `gencert.pl`) to incorporate your own certificate generation process.
- replace the CA file for the whole cell.
- enable ECC in the MoM environment, the MoM server and then all the member Cell Managers.

For more information on encrypted control communication, see the HPE Data Protector Installation Guide, HPE Data Protector Command Line Interface Reference, HPE Data Protector Troubleshooting Guide, and HPE Data Protector Help.

For more information on VMware GRE and encrypted control communication, see the HPE Data Protector Granular Recovery Extension User Guide.

For details on the supported operating systems, see the latest support matrices at: <https://softwaresupport.hpe.com/>.

New menu item under the Help menu

The Help menu now includes an additional (new) menu item, "Latest Guides on the Web". You can use this menu item to download the latest Data Protector documentation from the HPE Software support site: <https://softwaresupport.hpe.com/manuals>.

Scheduling and priority in Advanced Scheduler

In the Advanced Scheduler, priority can be set for each schedule. In case multiple running sessions request access to a specific device at the same time, the priority determines the order in which the sessions will be queued.

In the Advanced Scheduler, you can specify that a scheduled session have the ability to pause other sessions if they are a lower priority than the selected session. For more information, see the *HPE Data Protector Help*.

- The ability to pause and resume from where the session left off is available for filesystem, VMware and Oracle Server integration sessions. For other integrations, after being paused, the backup session restarts.
- Backup sessions to Disk (B2D) devices are not subject to pausing due to priority.
- For backup sessions that contain a mix of backup device types, for example, file library and B2D, the pausing functionality will only apply to the non-B2D devices.

Omnicc command enhanced to include total data under protection

The `omnicc -check_licenses` command reports licensing related information from the cell. If the `-detail` option is specified, a detailed report is produced. Both the summary and detailed reports now include the total data under protection for the cell, in TB.

For more information on the command, see the *HPE Data Protector Command Line Interface Reference*.

Automated Replication Synchronization

Enables replication of meta data between Cell Managers for appliance to appliance replication via Catalyst or Boost.

Site Specific Patch and Test Module installation enhancements

The functionality for pushing Site Specific Patches (SSPs) and Test Modules (TMs) existed in Data Protector 8.0. However, in DP 9.04, this feature has been enabled to use the installation server for deployment of SSPs or TMs.

Improved debug log gathering tools

The following enhancements have been made for improved debug log collection:

- Specifying multiple modules for which you want the debug files.
- Compressing the debug logs on the fly.
- Collecting all DP logs using the `omnidlc` command by specifying multiple filter options.

Platform coverage enhancements

- EADR for SLES 11.3 with BTRFS
- Extended ACL for Linux
- EADR for dynamic disks
- Lustre file system
- IBM TS3500 (3584) library

Enhanced device filtering

Device filtering has been enhanced to select device targets for backup based on the active node of the cluster setup. Device filters can be used to assign backup targets to clients located in close proximity to reduce network traffic.

For more information on enhanced device filtering, see the *HPE Data Protector Concepts Guide*.

Support for multiple interfaces to the same store

Data Protector supports IP as well as a fiber channel connection to the same Catalyst / DDBoost store without the need to configure a separate store. The store is accessible simultaneously over both interfaces.

For example, sometimes a single Catalyst / DDBoost store can be accessed by local clients over fiber channel for faster backup while remote clients can access the same store over the WAN for slower backup.

This feature is not available in the Solaris environment or if FC is configured as the identifier for the deduplication target. This option applies to StoreOnce backup systems and DD Boost only.

For more details on the working of this feature, see the *HPE Data Protector Administrator's Guide*, *HPE Data Protector Help* and *HPE Data Protector Command Line Interface Reference*

Collection of high-level information for telemetrics

Data Protector gathers and collects the following high-level information for telemetrics:

- Host OS version
- Data Protector components and its versions
- Devices or Media Servers - Are associated to a client in the Cell Manager. It includes the host name details where the device is attached, name of the device, library name, pool name where the media is placed, and device type.
- Schedules - The schedule telemetry exposes information grouped by backup and session types. It represents the number of full and incremental backup processes scheduled every year by backup and session types.
- Capacity Based Licensing (CBL) - CBL is leveraged to gather information on capacity. For more information, see the *HPE Data Protector Installation Guide*.
- License categories - Lists the number of licenses available in Data Protector.

After the telemetric data is collected, the data is uploaded to Support using the debug logs. For further information, see the *HPE Data Protector Troubleshooting Guide*, or *HPE Data Protector Command Line Interface Reference*.

Note: The Cell Manager performance will not be impacted significantly during the collection of telemetry data.

Encrypted Control Communication enhancements

Data Protector introduces support for the following:

- Display of expiry date for Encrypted Control Communication certificates - The expiration date for Encrypted Control Communication certificates is now visible in the Data Protector GUI.
- Review the expired Encrypted Control Communication certificates - A notification has been added to warn about the expired Encrypted Control Communication certificates.
- Configure the TLS for Encrypted Control Communication - The default TLS version used by

Encrypted Control Communication has been modified from 1.0 to 1.1. The binaries that are not released in patch will still use only TLS 1.0. Also, the client can be configured to support specific TLS versions using CLI.

NDMP integration enhancements

Data Protector introduces support for the following:

- NDMP 3-way backup (Refers to the backup flow where the data from one filer is streamed to another. The second filer stores data to its local backup device).
- NDMP NetApp Cluster Aware Backup (CAB).
- NDMP 64-bit Media Agent support for Linux 64-bit environments.

Non-Changed Block Tracking (Non-CBT) backup

Non Changed Block Tracking (Non-CBT) backup is a feature which does not depend on block level changes for backup.

Allow fallback to non-CBT backups option is enabled when change block tracking backup fails.

The non-CBT backup can be used in the following scenarios:

- When the hardware version of virtual machine is lesser than 7.
- When backing up virtual machines without the older version of operating system installed (example, Windows 2003).
- When snapshots are available on virtual machine and CBT is not enabled.

Only full backups are available for non-CBT backups.

For more information on Non-CBT backup, see the HPE Data Protector Integration Guide and HPE Data Protector Command Line Interface Reference.

Support for Data Protector Telemetry

Data Protector gathers and collects the following high-level information for telemetrics:

- Host OS version
- Data Protector components and its versions
- Devices or Media Servers - Are associated to a client in the Cell Manager. It includes the host name details where the device is attached, name of the device, library name, pool name where the media is placed, and device type.
- Schedules - The schedule telemetry exposes information grouped by backup and session types. It represents the number of full and incremental backup processes scheduled every year by backup and session types.
- Capacity Based Licensing (CBL) - CBL is leveraged to gather information on capacity. For more information, see the *HPE Data Protector Installation Guide*.
- License categories - Lists the number of licenses available in Data Protector.
- Patches - Describes the patch version installed.

Note: The customer related internal information is gathered, but the Host information is masked or replaced with a character numeric format.

Once the telemetric data is collected, the data is uploaded to Support using the debug logs. For further information, see the *HPE Data Protector Troubleshooting Guide*, *HPE Data Protector Online Help*, or *HPE Data Protector Command Line Interface Reference*.

Note: The Cell Manager performance will not be impacted significantly during the collection of telemetry data.

Support for BTRFS

Data Protector introduces backup of BTRFS volumes and subvolumes without the need of mounting them.

Support for SAN transport backups for VMware

Data Protector introduces support for VMware VDDK 6.0, Update 2. It supports SAN transport backups.

Support for VMware GRE with Windows Server 2012 Clustered Cell Managers

Data Protector introduces support for VMware GRE with Windows Server 2012 Clustered Cell Managers.

Support for Storage Optimizer version 5.x

Data Protector is now integrated with Storage Optimizer version 5.x. To backup the archived data, you should set up regular backups of targets on the archiving storage systems. Use Data Protector snapshot capabilities (enabled by default for Windows filesystem backup) when creating backups of Storage Optimizer repositories and archive targets.

Set the Data Protector `OB2HSMBACKUPALL` variable to enable backup of offline files which are generated by Storage Optimizer. For more information, see *HPE Data Protector Help*

Ensure that you enable the Enhanced Incremental option when creating a datalist to backup the offline files.

Support for HPE 3PAR Storage Systems with InForm OS version 3.2.2

Data Protector introduces support for HPE 3PAR Storage Systems with InForm OS version 3.2.2.

Support for HPE 3PAR ZDB instant recovery within a VMware VM

Data Protector offers HPE 3PAR ZDB instant recovery for agents within a VMware virtual machine. For more information, see the *HPE Data Protector Zero Downtime Backup Integration Guide*.

ZDB support for EMC VMAX

Data Protector introduces an EMC VMAX Storage Provider to integrate with the EMC VMAX Storage families and enables you to perform zero downtime backups. With an EMC VMAX Storage Provider, a ZDB to tape backup type is available, meaning the backed up data can be restored using conventional restore. Support is available for a filesystem and ZDB application integrations (Oracle Server, SAP R/3, and Microsoft SQL Server) backup and restore.

EMC VMAX Storage Provider is supported on Windows 2008/20012, on RHEL 5/6, on SLES 10/11, and on OEL 5 (filesystem backup also on OEL 6) operating systems.

ZDB support for EMC VNX

Data Protector introduces an EMC VNX Storage Provider to integrate with the EMC VNX Storage families and enables you to perform zero downtime backups. With an EMC VNX Storage Provider, a ZDB to tape backup type is available, meaning the backed up data can be restored using conventional restore. Support is available for a filesystem and ZDB application integrations (Oracle Server, SAP R/3, and Microsoft SQL Server) backup and restore.

EMC VNX Storage Provider is supported on Windows 2008/20012, on RHEL 5/6, on SLES 10/11, and on OEL 5 (filesystem backup also on OEL 6) operating systems.

Chapter 3: Limitations and recommendations

Limitations

Scalability limitations

Backup infrastructure scalability

Backup infrastructure metric	Limit
Clients in a Data Protector cell	5000
Cell Managers (cells) in a Data Protector Manager-of-Managers (MoM) cell	50
Total number of clients in a MoM environment	50 000 ¹

Internal Database scalability

Specific limits, which are marked as such in the following tables, can be reconfigured by adjusting the Data Protector global options. For details, see the *HPE Data Protector Troubleshooting Guide* and the *HPE Data Protector Help*.

Basic Internal Database capacity	Limit
Data Protector sessions stored in the Internal Database (IDB)	100 million (100 000 000)
Filenames with metadata referenced in the IDB	1 trillion (10 ¹²)
Backup objects referenced in the IDB	1 million (1 000 000)
Backup object versions referenced in the IDB	50 million (50 000 000)

Detail Catalog Binary Files capacity	Maximum configurable limit (Predefined default limit)
Detail Catalog (DC) directories	100 (50)

¹ In MoM environments, the total number of clients does not scale linearly.

Detail Catalog Binary Files capacity	Maximum configurable limit (Predefined default limit)
Size per DC directory	2047 TB ^{1 2} (200 GB)
Files per DC directory	500 000 (100 000)
DC binary file size	n/a ³
DC directory low space (minimum difference to the effective limit on the directory size)	n/a ⁴ (2 GB)

Media Management Database capacity	Limit
Backup media in the Media Management Database (MMDB)	50 million (50 000 000)
Total number of backup media in all MMDBs (or in the CMMDB ⁵) in a MoM environment	2.5 billion (2 500 000 000)
Backup media in a media pool	200 000

Concurrency limitations

Backup session concurrency metric	Maximum configurable limit (Predefined default limit)
Concurrent ⁶ backup sessions	1000 (100)
Total number of concurrent backup sessions in a MoM environment	50 000 ⁷
Backup sessions in a day	99 999 ⁸

¹The underlying file system limitations or settings may prevail over this limit.

²In the Data Protector GUI, you are only allowed to set amounts up to 10 240 GB (10 TB). You can use the `Data Protector omnidbutil` command for setting bigger amounts.

³The effective limit depends on the underlying file system limitations or settings.

⁴The effective limit for a particular DC directory depends on the configured Data Protector limit for its maximum size.

⁵Centralized Media Management Database.

⁶In this context, “concurrent” means “running concurrently in a Data Protector cell”.

⁷This limit cannot be changed.

⁸This limit cannot be changed.

Backup device concurrency metric	Limit
Disk Agent concurrency (device concurrency)	32 ¹
Backup devices (drives) used in a backup, object copy, object consolidation, or restore session	128 ²
Concurrent ³ physical drives (DLT7000 and lower performing models)	1000
Concurrent physical drives (DLT8000, SDLT, LTO)	500
Concurrent virtual drives (LTO—where drive concurrency is set to 1)	1000

Intrasession concurrency metric	Limit
Backup objects processed concurrently in a session	4096 ⁴
Backup media imported at the same time	100

Prerequisites for increasing the limit on concurrent backup sessions

If you increase the value of the `MaxBSessions` global option to a certain amount (for example, near 1000), you may have to modify specific system parameters on the Cell Manager that limit the number of concurrent sessions in general. After such modifications, restart the local Data Protector services with the Data Protector `omnisv` command. The modifications are as follows, and depend on the Cell Manager's operating system.

Windows systems:

The default non-interactive desktop heap size is sufficient for approximately 100 parallel sessions. Therefore, you need to increase the size of the non-interactive desktop heap.

Update the `Windows` value in the Windows Registry key `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SessionManager\SubSystems` by changing the third numerical value in the parameter `SharedSection=1024,20480,768` from 768 to 10240.

For example:

```
Windows="%SystemRoot%\system32\csrss.exe ObjectDirectory=\Windows
SharedSection=1024,20480,10240 Windows=On SubSystemType=Windows ServerDll=basesrv,1
ServerDll=winsrv:UserServerDllInitialization,3
```

¹ This is the maximum configurable limit. The effective limit for a particular device depends on the configured concurrency for that device in the device configuration or session specification.

² This is the maximum configurable limit.

³ In this context, "concurrent" means "used concurrently in a Data Protector cell".

⁴ This is the maximum configurable limit. It depends on the effective limits for the Disk Agent concurrency and the number of backup devices that can be used in a session.

```
ServerDll=winsrv:ConServerDllInitialization,2 ProfileControl=Off  
MaxRequestThreads=16"
```

HP-UX systems:

You need to adjust the operating system kernel parameters `nproc` and `nkthread`, and possibly other dependent parameters as well. You can perform the adjustments using the `kmtune` command or the System Administration Manager (SAM) application. For instructions, see the HP-UX operating system documentation.

Linux systems:

You need to add the following line into the preconfigured service `omni` parameter group in the file `/etc/xinetd.d/omni`:

```
cps = 1100 10
```

Enhanced incremental backup

- Each new enhanced incremental database can support a maximum of 40 billion files per mount point and a maximum of 40 million files per directory.
- The maximum memory consumption is determined by the largest number of files within one single directory. The approximate maximum memory consumption is 130 MB per 1 million files within one directory.
- Data Protector supports enhanced incremental backup of the following number of files per directory:

Windows systems (32-bit):

10 million files

HP-UX systems:

5 million files

Linux systems (32-bit):

5 million files

Size of file depots used for file library

HPE recommends that you use the default file depot size (5 GB). Increasing the size may cause certain performance degradation. The maximum supported file depot size is 2 TB.

Installation limitations

Data Protector cannot be installed if the installation path:

- Contains non-ASCII characters
- Contains the characters "@" or "#" or "&"
- Contains a directory that ends with the character "!"
- Is longer than 80 characters

Spaces are not supported in the Omniback data directory during an upgrade procedure.

If you are upgrading from a path that has any of the above characteristics, you must migrate the installation to a different directory. See the troubleshooting chapter of the *HPE Data Protector Installation Guide*.

Upgrade limitations

- A backup of the Internal Database, created with previous versions of Data Protector, cannot be restored. After upgrading the Cell Manager, back up the Internal Database before you continue using Data Protector.
- If the stores created using an older Data Protector version are not visible after upgrading, then you should restart the client (where stores reside).

Migration limitations

- Cell Manager can only be migrated to the same Data Protector version.
To use a new Data Protector version on the system you want to migrate to, upgrade the existing Cell Manager installation to the new version before you start migration.
- Cross-platform migration, for example from a Windows system to an HP-UX system, is not supported.
- Migrating the Data Protector Cell Manager from a clustered environment to a single server environment is not supported.

Platform limitations

UNIX and Linux limitations

- LOFS filesystems are fully supported. However, Data Protector does not recognize directories that are lofs-mounted if they are mounted within the same filesystem. This will result in additional data being backed up.
- Cross-filesystem restore of ACLs (file permission attributes) is not supported. For example, ACLs backed up from the VxFS filesystem cannot be restored to a UFS filesystem and the other way round. File objects however, can be restored to a different filesystem without ACLs.
- Cross-platform restore of ACLs is not supported. This limitation is due to different internal ACL data structures on different operating systems.
- Modification of ACL entries does not affect the modification time of the file object, so the file object (and the modified ACL) is not backed up during an incremental backup.
- The GUI can display a maximum 64 000 items (files in one directory, slots in a library, and so on) in a tree view.
- File names containing quotation marks are not supported.

HP-UX limitations

- Restore of a single file from a disk image is not supported.
- On HP-UX 11.31 that uses new persistent multi-pathing and path-independent Device Special Files

(DSFs), backup specifications referring to the old DSF may not work if the old DSF is disabled on the system. In this case, reconfigure the devices and update backup specifications to use the new-style DSF.

Solaris limitations

- If a `csh` script is used for `pre-` or `post-exec`, the `-b` option must be specified in the interpreter specification line: `#!/bin/csh -b`
- On Solaris, `/tmp` is a virtual filesystem in the swap area. If the `/tmp` directory is included in a backup specification, it is backed up as an empty directory. If restoring such backup, a swap area must be configured on the client prior the restore, otherwise the `/tmp` directory cannot be re-created.
- Backup and restore of access control lists (ACLs) on Veritas Cluster File System (CFS) is not supported.
- On Solaris, detection of media types other than Data Protector media is not reliable, due to the use of a number of different block sizes. Do not rely on Data Protector to recognize foreign media.
Workaround: To prevent Data Protector from automatically initializing a medium it does not recognize correctly, set `InitOnLoosePolicy` global option to `0`. All media then have to be initialized manually.
- Cleaning tape recognition in DDS libraries does not function.

Linux limitations

- After the transition from the `ext2` to the `ext3` filesystem on Linux systems, the journal will be visible as the `.journal` file in the `root` directory of the filesystem. If the filesystem is not mounted, the journal will be hidden and will not appear in the filesystem.
Due to the Linux operating system limitations, do not delete this `.journal` file, do not back it up, and do not restore it from backup.
- If you use access control lists (ACLs) and perform backup and restore between 32-bit and 64-bit Linux systems (for example, you perform a backup on a 32-bit Linux system and restore this backup to a 64-bit Linux system), the ACL entries are not restored.
- Cross-platform restore of ACLs between 32-bit and 64-bit Linux operating systems is not supported.
- On Linux systems, before you restore a symbolic link whose owner is not the `root` user, ensure that all directories in the path where the link will be restored have `execute` permission set for the link owner. Otherwise, the restore session will fail.
- Disaster recovery (Enhanced Automated Disaster Recovery or One Button Disaster Recovery) is not supported if SELinux is enabled.

Mac OS X limitations

- The Internet Protocol version 6 (IPv6) is not supported on Mac OS X operating system.
- Cross file system restore of ACL (Access Control List), extended ACLs, and file attributes is not supported (for example, ACL's backed up from the HFS+ file system cannot be restored to the UFS file system and the other way round).
- Encrypted control communication is not supported on Mac OS X operating system.

Windows limitations

- Windows directory share information can only be restored to a Windows system with the Data Protector Disk Agent. If this requirement is not met, the directory will still be restored, but the Disk Agent will ignore the directory share information.
- Only one CONFIGURATION backup can run on a Windows client at a time.
- Data Protector requires the same name for both, the computer name and the resolving hostname.
- Remote installation using secure shell (SSH) is not supported on Windows platforms.
- Local secure shell installation supports key-based authentication. It does not support other authentication modes.
- Backing up network shared volumes using the VSS functionality is not supported.
- The GUI on Windows systems can display a maximum 64000 items (files in one directory, slots in a library, and so on) in a tree view.
- The name of the file cluster resource used during the installation of the Data Protector Cluster Integration on Windows must not be omniback. For details, see the *HPE Data Protector Installation Guide*.
- When browsing with the backup specification editor a Windows client, the Windows user interface lists both online and offline Informix Server dbspaces. To check for databases, use the `onstat -d` command. Available databases are marked with the PO flag.
- On Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012 systems, the user performing a network share backup must be a member of the operating system Backup Operators user group and must be added to the Inet configuration on the system where Disk Agent is running (using `omniinetpasswd -add`). In a cluster environment, users must be configured on both nodes.
- On Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012, the broadcast message send method is not supported.
- Backed up directory share information of directories which were located on a 32-bit Windows system cannot be restored to a 64-bit Windows system and the other way round. In such restore scenarios, the selected directories and their contents will be restored as expected, but without their share information.
- VSS disk image backup of the logical volumes can be used for disaster recovery only on Windows Vista, Windows 7, and Windows Sever 2008 systems.
- You can boot the target system over the network only on Windows Vista, Windows 7, and Windows 2008 Server systems.
- The HPE Data Protector Disaster Recovery GUI is available only on Windows Vista, Windows 7, and Windows 2008 Server systems. On other Windows systems, a console interface is available.
- IPv6 addresses cannot be used in share names when backing up network share volumes.
- Data Protector Inet service cannot be started if the Windows system is started in Safe Mode with Networking.

Windows 32-bit limitations

- On Windows systems, the native robotics driver (Removable Storage Manager) is automatically loaded to enable tape libraries. To use the library robotics with Data Protector on 32-bit Windows

systems, disable the Windows medium changer (robotics) driver before you configure the system with the Data Protector Media Agent.

Windows 64-bit limitations

- The original Microsoft Windows installation CD-ROM is supported for Automated System Recovery (ASR). The Windows XP 64-bit Edition Recovery DVD that comes with Itanium systems *cannot* be used for ASR.
- It is not possible to integrate the Data Protector GUI with the Microsoft Management Console (MMC) using the Data Protector OB2_Snap snap-in.
- Data Protector does not support Java web reporting on Windows systems based on the Itanium 2 processor architecture, since Java runtime environment is not supported on this platform.
- On AMD64/Intel EM64T systems, sending notifications and reports by e-mail using MAPI is supported only with Microsoft Outlook Express, and not Microsoft Outlook.

Windows XP and Windows Server 2003 limitations

- In order to perform Data Protector remote installation if any of the clients are running Windows XP or Windows Server 2003, the Installation Server and the clients must have IPv4 protocol enabled.

Although both systems natively support IPv6, there is a limitation:

- The Windows Remote Procedure Call (RPC) provider does not provide IPv6 support on these systems. As a consequence, accessing remote network shares on systems using IPv6-only configuration may not be possible.

Network shares are used by Data Protector remote installation in order to install the initial services when performing a clean client installation, as well as accessing the installation depot from the client.

Windows Server 2012 limitations

- Filesystem backup of volumes formatted with the Resilient File System (ReFS) is not supported. Use disk image backup instead.
- Backup of network shared disks is supported without using the VSS for SMB File Shares functionality.

Novell Open Enterprise Server (OES) limitations

- Back up or restore of any GroupWise system files is not supported.
- The Internet Protocol version 6 (IPv6) is not supported for OES cluster configurations.

HP OpenVMS limitations

- The OpenVMS client must be installed locally on the OpenVMS system. There is no support for remote installation from an Installation Server.
- The product can only be installed on the system disk in `SYS$COMMON:[OMNI]`.
- Any file specifications that are passed to the CLI must conform to a UNIX-style syntax:

/disk/directory1/directory2/filename.ext.n

- The string should begin with a slash, followed by the disk, directories, and file name, separated by slashes.
- Do not place a colon after the disk name.
- A period should be used before the version number instead of a semi-colon.
- File specifications for OpenVMS files are case insensitive, except for the files that reside on ODS-5 disks.

For example:

An OpenVMS file specification of:

```
$1$DGA100:[USERS.DOE]LOGIN.COM;1
```

must be specified in the form:

```
/$1$DGA100/USERS/DOE/LOGIN.COM.1
```

- Patch level display is not available on OpenVMS.
- There is no implicit version number. You always have to specify a version number. Only file versions selected for the backup will be backed up. If you wish to include all versions of the file, select them all in the GUI window, or, using the CLI, include the file specifications under the `Only (-only)` option, including wildcards for the version number, as follows

```
/DKA1/dir1/filename.txt.*
```

- If the `Do not preserve access time attributes` option is enabled during a backup, the last accessed date will be updated with the current date and time on ODS-5 disks. On ODS-2 disks, this option has no effect, and all the dates remain unchanged.
- Disk backup is not available on OpenVMS. There is no equivalent to “BACKUP/IMAGE” or “BACKUP/PHYSICAL”.
- When the data backed up from an OpenVMS Alpha system is restored or migrated to an OpenVMS Integrity system using Data Protector, some of the default file attributes (such as creation time, last revised time, version limit and some of the file record attributes) may get lost. This also applies to the data restore or migration from Itanium to Alpha.

Workaround: Manually reset the attributes using the DCL command line.

- The `Backup POSIX hard links as files (-hlink)` option is not available on OpenVMS. Files with multiple directory entries are only backed up once using the primary path name. The secondary path entries are saved as soft links. During a restore, these extra path entries will also be restored.

For example, system specific roots on an OpenVMS system disk will have the `SYSCOMMON.DIR;1` path stored as a soft link. The data for this path will be saved under `[VMS$COMMON...]`.

- Files being backed up or restored are always locked regardless of whether the `Lock files during backup (-lock)` option is enabled or disabled. With the `-lock` option enabled any file opened for write is not backed up. With the `-lock` option disabled any open file is backed up as well. No message is issued when an open file is saved.
- The default device and directory for pre- and post-exec command procedures is `/omni$root/bin`. To place the command procedure anywhere else the file specification must contain the device and

directory path in UNIX-style format. `/SYS$MANAGER/DP_SAVE1.COM` is an example of a valid specification.

- If you restore to a location other than the original location, only the disk device and starting directory are changed. The original directory path is added to the destination path to form the new restore location.
- To successfully back up write-protected and shadow disks, enable the `Do not preserve access time attributes` option in the backup specification.
- If the `Do not preserve access time attributes` option is disabled during a backup and if the `Restore Time Attributes` option is disabled during a restore, the last accessed date will be updated with the current date and time on ODS-5 disks. On ODS-2 disks, the original dates will be set on the files.
- The `Move Busy Files (-move)` and `Restore Sparse Files (-sparse)` options are not available on OpenVMS.
- Files backed up from an ODS-5 disk on an OpenVMS system that have extended filesystem names (for example upper and lower case letters, Unicode characters, etc.) may not be restored to an ODS-2 disk.
- If the `Restore Protection Attributes (-no_protection)` option is disabled, the files are created with the default owner, protection and ACL.
- There is no support for a BACKUP/IMAGE equivalence. To make a restored copy of an OpenVMS system disk bootable, the OpenVMS WRITEBOOT utility has to be used to write a boot block onto the restored disk.
- The `omnicheck -patches -host` command is not supported on OpenVMS.
- The `omnirpt -email` command is not supported on OpenVMS. You can use the `-log` option to create a local dump of a report file and use the native OpenVMS mail utility to send an e-mail with this file as an attachment.
- 16-bit Unicode filenames on an ODS-5 disk volume will be displayed in VTF7 (OpenVMS specific) notation on the Cell Manager in the form of `^Uxxyy` for a Unicode character where `xx` and `yy` are the Unicode hex codes for this character. Other valid characters for files on ODS-5 volumes can be specified using the OpenVMS guidelines for extended file specification syntax.
- If an OpenVMS file is restored to a non-OpenVMS platform, file attributes specific to OpenVMS may not be retained (for example record format, backup date, ACL).
- Files that have been saved on non-OpenVMS platforms and are to be restored to an OpenVMS system may lose some file attributes. No ACL will be restored in this case.
- No qualification is done for tape drives which are not supported by OpenVMS. For a complete list of tape drives, see the OpenVMS Software Product Description (SPD).
- HSJ connected tape libraries cannot be autoconfigured. Use manual configuration methods to add these devices to Data Protector.
- Maximum block size for Media Agent on OpenVMS is 63.5 kB. If a device/drive is configured with a bigger block size, it will be changed to 63.5 kB.
- Data Protector file library is not supported on OpenVMS ODS-2 disks.
- All tape media initialized by the Media Agent starts with an ANSI VOL1 label having a non-blank Volume Accessibility character. To mount such a tape volume on OpenVMS, use the `/OVERRIDE=ACCESSIBILITY` qualifier. However, the tape volume does not comply with ANSI tape labeling and can therefore not be used with OpenVMS utilities like DCL-COPY.

- Restore file to original location with the `-no_overwrite` option will not restore any files.
- Incremental backup will work at the directory level only, because OpenVMS creates a new file with a new version number upon modification of an existing file. Data Protector on OpenVMS allows to create incremental backup at file level only if the filename is exactly the same as the previous, including the version number.
- On the OpenVMS client with the Oracle integration installed, you have to configure a Data Protector admin user with the username `<Any>` and the group name `<Any>`. This limitation is due to the lack of the user group name concept on OpenVMS.
- If you run the Media Agent and the Data Protector Oracle integration agent on the same OpenVMS client, modify the group ID of the `omniadmin` user as `DBA` using the `MCR AUTHORIZE` utility.
- When a debug and log file collector is used on OpenVMS, the following applies:
 - The OpenVMS ODS-2 disk structure file name can contain a maximum of 39 characters.
 - As OpenVMS systems do not have the `get_info` utility, the `get_info.out` file is blank and is not collected.
 - The `omnidlc` command executed with the `-session` parameter does not collect the debug files produced during specified session, because session names are not part of the OpenVMS debug filename. All available logs are collected instead.
- The Oracle environmental variables and `omnirc` options `OB2_RMAN_COMMAND_TIMEOUT` and `OB2_SQLP_SCRIPT_TIMEOUT`, which help improving Oracle Server backup session handling, are not supported on OpenVMS systems.
- The Internet Protocol version 6 (IPv6) is not supported on HP OpenVMS.
- Encrypted control communication is not supported on HP OpenVMS.
- Enhanced incremental backup is not supported.

Limitations on clusters

HPE Serviceguard limitations

- When adding components on HPE Serviceguard, add the components on the active node. Then start the package on the other node, and add the components on this node too.

Limitations on licensing

General licensing limitations

- Both licensing models, the capacity based and traditional cannot coexist in the same Data Protector cell.
- In a MoM environment, you cannot mix both licensing models — all Cell Managers must use the same licensing method.

License upgrading limitations

- Limitations on advanced backup to disk licensing:
 - The library capacity (VTLCAPACITY) of a virtual tape library, which was created with a previous version of Data Protector, is by default set to 1 TB after the upgrade to the latest version. You must enter the estimated library capacity value manually through the graphical user interface (GUI) or via the command-line interface (CLI). See the advanced backup to disk example in the *Data Protector licensing* chapter of the *HPE Data Protector Installation Guide*, and the `omniupload` man page or the *HPE Data Protector Command Line Interface Reference*.

Internet Protocol version 6 (IPv6) networking limitations

- Licenses acquired with versions earlier than Data Protector 9.00 only support IPv4 and dual IP stack environments and require the Cell Manager to have an IPv4 address. In order to use Data Protector in an IPv6 only environment, the licenses acquired with versions earlier than Data Protector 9.08 should be converted to the Data Protector licensing technology. To convert the keys, follow the procedures provided by the password delivery center.

This limitation does not apply to licenses acquired with or after Data Protector 9.00 or later.

Limitations on license reporting in a traditional licensing model

- In a cell if the Data Protector Cell Manager and a client is not upgraded, Media Agent on a client cannot send the information about the used disk capacity to the Cell Manager. Consequently, the license checker does not receive the needed information about disk space that is used and cannot report the actual license capacity in use. Therefore, the license checker reports an additional Advanced backup to disk for 1 TB license-to-use is required for such file library.
- Due to the migration of multi-drive server licenses to single-drive licenses, the license enforcement is stronger than the license checking. If a multi-drive server license is installed on a system that is not a drive server, the multi-drive license is not used and the backup may not be possible, although the license checker reports enough appropriate licenses installed.
- Due to the platform independent licenses for slot libraries, the license enforcement is stronger than the license checking. During the backup, Data Protector is checking the licenses for different platforms and the backup may not be possible because of the missing licenses for a specific platform, although the license checker reports enough appropriate licenses installed.
- Since the legacy ZDB and the IR licenses respectively are grouped into one generic license, the license enforcement is stronger than the license checking. During the ZDB backup, Data Protector is checking licenses for different storage arrays and the backup may not be possible due to the missing licenses for a specific storage array, although the license checker reports enough of the zero downtime backup extension and instant recovery extension licenses-to-use (LTU) installed.

Limitations on encryption

Limitations on data encryption

- Consolidation of objects backed up with software encryption is not supported.

Limitations on encrypted control communication

- Communication between the client, which is using plain control communication and the client with enabled encrypted control communication is not supported. This means, that Data Protector operations will not be executed (for example, remote installation from an Installation Server, which is using plain control communication to the client with enabled encrypted control communication will not succeed).

However, the Cell Manager can communicate with both types of clients in the Data Protector cell.

- End user authentication is not supported.
- To satisfy U.S. Export Regulations, encrypted control communication uses only export ciphers. Key lengths are limited to 64 bits for symmetric and 512 bits for asymmetric encryption. These regulations are enforced on a code level.

Limitations on Data Protector MoM environments

- Debug log collection is not supported in MoM environments.

Device and media limitations

- Device filtering during a backup session is supported for Data Protector File System, Data Protector Oracle Server integration, Data Protector Microsoft Exchange Server 2010 integration, and Data Protector Microsoft SQL Server integration.

Device filtering can be enabled by setting the `global` option `EnableDeviceFilters` to 1.

For details on setting the `global` options, see the *HPE Data Protector Help*.

If you are an existing customer using the `omnirc` variable `OB2DEVICEFILTER` for the device filter feature, you need to migrate the filter tags using the `omnicc -migrate_devfilter [HostName] [-delete_old_devfilter]` command.

For information on migrating the existing `omnirc` variable based `OB2DEVICEFILTER` tags, see the *HPE Data Protector Command Line Interface Reference*.

- Source-side deduplication backup to StoreOnce Backup system devices configured with Fiber Channel (FC) can be performed on only those systems that are connected to FC. Therefore, before performing this backup, you must ensure that the systems meet the following requirements:
 - Data Protector Disk Agent is installed.
 - Data Protector Media Agent is installed.
 - Fiber Channel connection is configured.

During backup, you can use the **Systems ready for source-side deduplication** option to filter out systems that do not support the source-side deduplication. However, this option does not filter out systems that do not have the FC connection, which is one of requirements for performing the source-side deduplication backup to StoreOnce Backup system devices configured with FC.

To verify whether systems have the FC connection, click **Check** to validate the Gateways, while adding the StoreOnce Backup system device.

- Replication between Backup to Disk (B2D) is not supported if the source and target store does not have the appropriate credentials. The replication session fails with the following error message:
Permission Denied
If the original client with the source store does not have access rights on the target store, the client information is written in the source store which cannot be changed when imported to other clients. This problem cannot be resolved even if the new client has access to the target store. The original information on the source host is written on the store which is a crucial one.
- Reconstruction of broken media soon after a failed backup fails with the following error message:
The media cannot be loaded or open . This problem occurs on StoreOnce or EMC Data Domain Boost Backup system devices. Therefore, you must wait for some time before triggering the reconstruction.
The wait time for StoreOnce software deduplication is 2 hours, and the wait time for EMC Data Domain Boost device is 3 hours.
- StoreOnce catalyst does not support reconnect to Data Protector Media Agent when the connection is lost.

NDMP limitations

- Only filesystem backup and restore are available.
- Only Full and Incr1 backup types are supported.
- Maximum device concurrency is 1.
- Device selection as well as filesystem browsing is not possible.
- NDMP devices must use dedicated media pools.
- Localization for the NetApp-specific messages is not possible.
- It is not possible to deselect a subtree of the selected tree to be restored.
- It is not possible to perform a restore of the selected fileset as a tree with a different path name.
- Copying of NDMP backup objects and object mirroring in NDMP backup sessions are not supported.
- NDMP object copy to more than one target device is not supported.
- Medium header sanity check is not supported on NDMP clients.
- The data that was backed up from an NDMP Server of a particular type (for example, NDMP-NetApp) cannot be restored to an NDMP Server of another type (for example, NDMP-Celerra).
- When restoring to another NDMP Server, the device to restore from must be connected directly to the target NDMP Server, must be of the same type, and selected or specified as the restore device in the Data Protector GUI or CLI.
- Restore preview is not supported.
- Restoring data using the Data Protector Restore by Query functionality is not supported.
- Data Protector does not support IPv6 for NDMP backup sessions, therefore the NDMP servers should have IPv4 protocol enabled.
- On 64-bit Linux systems, the Data Protector NDMP Media Agent does not support ADIC/GRAU DAS library devices.
- Three-way backup or restore is only supported in filers with the same major version of firmware (For

eg. ONTAP 8.x).

- Three-way backup does not offer remote copy facility wherein the data from NDMP filer is sent to the Media Agent client independently, which then backups the data to the target.
- Cluster Aware Backup (CAB) or restore is supported only on the same cluster filer. Data Protector supports only backup and restore of volumes and files; which are on the same cluster filer as the devices used for backup or restore operation.
- In a non Cluster Aware Backup (CAB) environment, you cannot combine the three-way objects with the local objects during backup specification.

NetApp filer

- On NetApp filers running Data ONTAP version prior to 6.4, direct access restore (DAR) is not supported for directories; a standard restore will be performed instead. This has performance implications only.
- With the SMTape backup type, a backup image of a volume in a particular aggregate type cannot be used for restore to a volume in a different aggregate type.
- With the SMTape backup type, a backup image of a volume in a regular aggregate cannot be used for restore to a volume in a larger aggregate, and the other way round.
- The SMTape backup type offers only full backup (level-0 backup).
- The SMTape backup type enables you to only back up entire file systems. For example, you can back up /ufs1, but not /ufs1/dir1.

Celerra

- Media copying is not supported for NDMP-Celerra backup sessions.
- If you select both a directory and individual files from another directory and start the restore, only the selected files are restored. To restore both, use standard restore (set the NDMP environment variable DIRECT to N).
- Directory direct access restore (DDAR) cannot be used with backup images created with the NDMP volume backup (NVB) option selected.
- The NVB backup type enables you to only back up entire file systems. For example, you can back up /ufs1, but not /ufs1/dir1.
- The NVB backup type and file or directory filtering cannot be used together. If both are used, NVB takes precedence and the filters have no effect.

Limitations on enhanced incremental backups

- Limitations on the enhanced incremental database:
 - Hard link detection is not supported with enhanced incremental backups. Hard links for selected object will be backed up as files.
 - To maintain the optimal size of a new enhanced incremental database, Data Protector by default performs a regular check every 30 days. The objects that were deleted from the source volume or were not backed up for a period of 30 days are removed from the database. Thus, the objects that

were not backed up for 30 days will be backed up in the Full mode. This is applicable only to HP-UX, Windows, and Linux systems.

- Limitations on incremental backups using Windows NTFS Change Log Provider:
 - Hard link detection is not supported with Windows NTFS Change Log Provider incremental backups. Hard links for selected object will be backed up as files.
 - Backup of **FAT16** and **FAT32** filesystems is not supported.
 - Data Protector does not have private access to the Windows Change Journal which means that other applications might turn it off while Data Protector is using it.

Limitations on virtual full backups

- Virtual full backup is only supported for filesystem data. Data Protector integrations do not support the functionality.

As the distributed file media format is optimized for virtual full backups, do not use it when performing, for example, an Oracle or Microsoft SQL Server database integration backup. It will only reduce backup performance, without providing any beneficial results.

Limitations on object copy and consolidation

Due to dynamic expansion of gateways and gateway or store or device connection limitations, object copy and consolidation must ensure the following:

- When B2D devices are used as sources, at least one connection should be available for object copy and at least N for object consolidation, where N is the number of source media used for consolidation.
- When B2D devices are used as targets, at least M connections should be available – M being the MIN device setting in the copy or consolidation specification, if only B2D devices are used as targets. If other types of devices are used in parallel, CSM will balance them so that MIN setting is reached, or abort the session otherwise.

Limitations on object verification

General functionality limitations

- Object verification is applicable to backups stored in Data Protector tape format that can be restored using standard Data Protector network restore. It is not applicable to ZDB-to-disk, or the disk part of ZDB-to-disk+tape, which use instant recovery for restore.
- While the source media are being read for object verification, they are unavailable for restore.
- The use of Web Reporting with object verification is not supported.

Application integration limitations

- Object verification only verifies application integration objects from the Data Protector point of view: It can verify object data and delivery of that data to the required destination host. The object verification process does not communicate in any way with integrated applications and so cannot verify restore capability by the applications concerned.

Limitations on application integrations

For integration-specific limitations, see the *HPE Data Protector Integration Guide* and *HPE Data Protector Zero Downtime Backup Integration Guide*.

General limitations

- With database integrations that support restore by starting the integration agent via the CLI, starting such a restore is not supported if you access the client through Remote Desktop Connection and the Media Agent to be used is on the same client.

Oracle limitations

- When using RMAN scripts in Oracle backup specifications, double quotes (") must not be used, single quotes (') must be used instead.
- Data Protector does not check whether database objects to be restored were backed up and exist in the Data Protector Internal Database. The restore procedure simply starts.
- When restoring tablespaces to point in time the RMAN interface has to be used.
- Only the Oracle Restore GUI and Oracle RMAN can be used to recover the Oracle recovery catalog database.
- When restoring a database using the Data Protector GUI to a client system other than the one where the database originally resided, the instance name chosen on the new client system must be the same as that of the original instance name.
- On Windows platforms, a proxy copy backup of an Oracle database is not possible if the database is on raw disks even if the backup seems to complete without any problems reported.
- If an object is deleted from the RMAN Recovery Catalog database, these changes will not be propagated automatically to the IDB and the other way round.
- The Oracle backup set ZDB method is not supported if the database is installed on raw disks.
- Configuration of multiple Oracle databases using user created XLS (Microsoft Office Excel) and CSV (comma separated values) files is not supported on HP OpenVMS clients. Also, this feature cannot be used to configure standby databases and Oracle databases in ZDB environment. The Microsoft Office Excel 2007 Open XML Format is also not supported.
- Backing up Oracle control files with Oracle backup set ZDB method on IPv6-only clients is not supported.
- You cannot use the Data Protector GUI to configure an Oracle database whose files are managed by Automatic Storage Management (ASM) and for which any of the following ASM properties differs

from their default values: home directory of the ASM instance, the authentication mode used by the Data Protector Oracle integration agent to connect to the ASM instance.

MySQL limitations

- MySQL configurations with circular binary log implementation cannot be backed up with Data Protector.
- The incremental backup type is only available for InnoDB database tables. When tables using other MySQL storage engines, are included in an incremental backup session, full backup will be performed for such tables.
- MySQL data can only be restored to a system that hosts the same MySQL version as the system from where the data was backed up.
- Import (restore with the **Import tables to target instance** option) of the database or database tables that use the InnoDB storage engine can be only performed if the MySQL file-per-table setting was enabled for the backup.

SAP R/3 limitations

- If ZDB to tape is used to back up a tablespace in a ZDB environment on Windows systems, and the `ZDB_ORA_INCLUDE_CF_OLFOmnirc` option is not set to 1, the backup will fail if the control file is not on the mirrored disk or in the snapshot that will be backed up.
- A split-mirror restore of the SAP R/3 data using the Data Protector GUI on the backup system is executed as a regular filesystem restore, during which ZDB agents (SYMA, SSEA) mount disks on `/var/opt/omni/tmp` (the default mount point). Since this is a restore of application data, VRDA restores files to the original mount points. Therefore, the data is not restored to EMC Symmetrix or P9000 XP Array disks, but to the root partition instead.

Informix Server limitations

- On Windows systems, cold restore of non-critical dbspaces is not possible.

Microsoft SQL Server limitations

- Backup preview is not supported.
- Backup compression is supported only by SQL Server 2008 Enterprise and later.
- A data file belonging to a database which was backed up with the option **Concurrent streams** set to more than 1 cannot be restored. You can only restore the entire database.

Microsoft Exchange Server limitations

- Backup preview is not supported.

Microsoft Volume Shadow Copy Service limitations

Common VSS limitations

- Preview is not supported for any type of VSS sessions: backup, restore, zero downtime backup, and instant recovery.
- The **Restore files to temporary location** mode is not available for the DPM *Database* writer components. Because the files were backed up by another writer (MSDE Writer in this case), they are not shown in the restore page. Only the **Restore components** mode is available in such cases.

Microsoft Exchange Server 2003

- Due to a Microsoft Exchange Server 2003 writer issue, non-Latin characters (for example, Japanese characters) for Exchange store or storage group names are not supported.

Microsoft Virtual Server 2005

- Cluster backup of Microsoft Virtual Server 2005 is not supported. You can back up only individual nodes.

Microsoft SQL Writer

- Microsoft SQL Writer does not support Microsoft SQL database restore to another system. If you try to perform a restore, only files are restored.

Data Protector Virtual Environment Integration limitations

VMware limitations

- After you upgrade to Data Protector [[[Undefined variable DP.DP_Version_0905]]] or later, you cannot restart the failed VMware backup sessions from earlier Data Protector versions.
- After you upgrade from 7.03 or earlier versions to , you cannot run incremental or differential backups without running a full backup.
- Power On and Live Migrate operations from StoreOnce catalyst are supported only with Linux backup hosts.
- The SAN transport mode falls back to NBDSSL on the vSphere 5.5 environment, after an upgrade to VDDK 6.0 Update 1.
- For Changed Block Tracking requirements, see the following URL:
http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1020128
- 3PAR Zero Downtime Backup (ZDB) and Instant Recovery of virtual machines on vSphere VVol (Virtual Volume) datastore are not supported. So, it is recommended to ensure that VMs are not hosted on VVol datastores before starting a ZDB backup. In case, VMs are hosted on VVols in the Datacenter, which is selected for backup, the following Warning message is displayed:

```
[Warning] From: VEPALIB_VMWARE@hostname "/Datacenter" Time: <Date> <Time>  
ZDB backups are not supported with vSphere Virtual Volumes. Skipping Virtual  
Machine 'vm_name'...
```

Lotus limitations

- On Solaris and AIX systems, offline restore is not available for Lotus Notes/Domino Server 7.0 and later versions.
- On Solaris systems, restore with recovery is not available for Lotus Notes/Domino Server 7.0 and later versions.

Limitations on disk array integrations

For additional integration-specific limitations not included in this section, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

HPE P4000 SAN Solutions limitations

- In a Microsoft server cluster environment, all volumes which are selected for zero downtime backup session must belong to the same cluster.
- Backup preview is not supported.
- Object copying and object mirroring are not supported for ZDB to disk.
- Although you can create replica sets, replica set rotation is not supported.
- A replica cannot be used for instant recovery under any of the following conditions:
 - A target volume of the replica has been automatically removed during an instant recovery session based on another ZDB backup specification.
 - Other entities exist on the disk array which depend on the source volume that was used to create a target volume of the replica:
 - A newer target volume exists, and a smartclone is attached to it.
 - A newer snapshot exists, and the snapshot was not created by Data Protector.
- The Data Protector `omnidbp4000` command that you should use for configuring access to the CIMOM provider of the HPE P4000 SAN Solutions is available only on Windows systems.

HPE P6000 EVA Disk Array Family limitations

- The single-host (BC1) configuration based on Linux platform is not supported. In such a configuration, a single Linux system acts as the application system and the backup system.
For a list of supported configurations, see the latest support matrices at <https://softwaresupport.hpe.com/>.
- Dynamic disks are not supported.
- Only one type of target volume per source volume can exist on a disk array at the same time. For example, a snapclone of a source volume cannot be created if a vsnap or a standard snapshot of the

same source volume already exists.

- A replica cannot be reused if any snapclone from this replica has a snapshot attached or if a target volume from this replica is presented to a system other than backup system.
- Data Protector does not allow ZDB to use an instant recovery object as a source volume.
- When cloning of a source volume is in progress, another snapclone of that source volume cannot be created.
- Backup preview is not supported.
- Object copying and object mirroring are not supported for ZDB to disk.
- When using the “switch of disks” instant recovery method with HPE P6000 EVA Disk Array Family, care must be taken when instant recovery is performed on objects located on lower performance disks, as this may result in undesired performance penalties. In such cases, a ZDB to the high performance disks and subsequent instant recovery will reverse the situation.
- During instant recovery, CRC check is not performed.
- Instant recovery from a ZDB-to-disk+tape session cannot be performed using the Data Protector GUI after exporting or overwriting the media used in the backup session. The backup media must not be exported or overwritten even after an object copy session. If the backup media have been exported or overwritten, perform instant recovery using the Data Protector CLI. For information, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.
- Routine maintenance tasks, including (but not limited to) hot-swapping HBAs/SCSI controllers, disk array controllers, FC switches, and online firmware upgrade during backup are not supported. Backup is a high-IO activity and should not be done at the same time as routine maintenance.
- The number of standard snapshots or vsnaps that can be created for a specific source volume is limited by the HPE P6000 EVA Disk Array Family storage system. The actual limitation is determined by the storage system's firmware revision. For details, see the HPE P6000 EVA Disk Array Family documentation. Consider the limitation when specifying a value for the option **Number of replicas rotated** of a zero downtime backup specification. Note that the limitation does not apply to snapclones.
- In zero downtime backup sessions using multisnapping, only two snapshot types are supported by default: standard snapshot and snapclone. For information on whether your HPE P6000 EVA Disk Array Family environment supports multisnapping using vsnaps, see your HPE Command View (CV) EVA documentation. For instructions on how to enable support for the vsnap snapshot type in multisnapping ZDB sessions in Data Protector, contact HPE technical support.
- The Data Protector `omnicreated1` command cannot be used for creating Microsoft Exchange Server ZDB backup specifications for ZDB sessions involving P6000 EVA Array or P9000 XP Array.

HPE P9000 XP Disk Array Family limitations

- Asynchronous HPE Continuous Access P9000 XP configuration is not supported.
- The single-host (BC1) configuration based on Linux platform is not supported. In such a configuration, a single Linux system acts as the application system and the backup system.

For a list of supported configurations, see the latest support matrices at <https://softwaresupport.hpe.com/>.

- With the single-host (BC1) configuration, only filesystem and disk image backup are supported.
- Split-mirror restore (restore of data from the backup medium to a secondary volume and restore of

data from the secondary volume to a primary volume afterwards) is supported for the filesystems and disk images in the HPE Business Copy P9000 XP configuration. Database (application) split-mirror restore is not supported.

- Instant recovery is only available in HPE Business Copy P9000 XP configurations.
- In case Microsoft Exchange Server is installed on the backup system, its Information Store (MDB) and Directory Store have to be installed on the HPE P9000 XP Disk Array Family LDEVs that are different than the mirrored LDEVs used for the integration. The drive letters assigned to these LDEVs have to be different from those assigned to the LDEVs that are used for the integration.
- Backup preview is not supported.
- Object copying and object mirroring are not supported for ZDB to disk.
- Instant recovery from a ZDB-to-disk+tape session cannot be performed using the Data Protector GUI after exporting or overwriting the media used in the backup session. The backup media must not be exported or overwritten even after an object copy session. If the backup media have been exported or overwritten, perform instant recovery using the Data Protector CLI. For information, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.
- When restoring filesystems in an instant recovery session, no object other than those selected for instant recovery should share the disks that are used by objects selected for the session.
- Routine maintenance tasks, including (but not limited to) hot-swapping any field replaceable components like, disk array controllers, FC switches, and online firmware upgrade during backup are not supported. Backup is a high-IO activity and should not be done at the same time as routine maintenance.
- The maximum number of secondary volumes (mirrors, volumes to be used for snapshot storage) that can be created for a specific primary volume is limited by the HPE P9000 XP Disk Array Family model used and its installed firmware revision. Note that the limitation for mirrors and the limitation for volumes to be used for snapshot storage differ. For details, see the HPE P9000 XP Disk Array Family documentation.

HPE 3PAR StoreServ Storage limitations

- Snapshot of snapshot is not supported.
- On Windows, only MBR formatted volumes are supported.
- On Linux, only two-host configurations are supported. The application system and the backup system must not be the same client.
- During instant recovery, the target volumes from which the replica data copy is to happen, should not be presented to any client.
- Data Protector does not support the HPE 3PAR StoreServ Storage iSCSI host interface.
- In a Microsoft server cluster environment, all volumes which are selected for zero downtime backup session must belong to the same cluster.
- Backup preview is not supported.
- Object copying and object mirroring are not supported for ZDB to disk.
- A replica cannot be used for instant recovery if it was created with a Data Protector version earlier than 7.00 or with Data Protector 7.00 without the patch bundle set installed.
- Oracle ASM is supported with firmware 3.1.2 MU2 or later.

- During Instant Recovery, all volumes that are part of the volume set are accessible until the restore is completed.

EMC Symmetrix disk array limitations

- Only ZDB to tape is supported which means that instant recovery is not.
- Backup preview is not supported.
- Routine maintenance tasks, including (but not limited to) hot-swapping any field replaceable components like, disk array controllers, FC switches, and online firmware upgrade during backup are not supported. Backup is a high-IO activity and should not be done at the same time as routine maintenance.

NetApp Storage limitations

- In cluster environments, the backup system must not be in the same cluster with the application system. Additionally, the backup system cannot be the cluster virtual server, it can only be a cluster node.
- Instant recovery is not supported.
- Oracle ASM is not supported.
- Oracle RAC is not supported.

EMC VNX limitations

- In cluster environments, the backup system must not be in the same cluster with the application system. Additionally, the backup system cannot be the cluster virtual server, it can only be a cluster node.
- Instant recovery is not supported.
- SnapView Snapshot is not supported.
- Oracle ASM is not supported.
- Oracle RAC is not supported.

EMC VMAX limitations

- In cluster environments, the backup system must not be in the same cluster with the application system. Additionally, the backup system cannot be the cluster virtual server, it can only be a cluster node.
- Instant recovery is not supported.
- TimeFinder/Snap snapshot type is not supported.
- TimeFinder/Consistency Groups are not supported.
- Oracle ASM is not supported.
- Oracle RAC is not supported.

Disaster recovery limitations

- One Button Disaster Recovery (OBDR) method is not available for Data Protector Cell Managers.
- The Cell Manager recovery wizard is no longer available in the Data Protector GUI. As the Internal Database can only be backed up separately from the files systems, when preparing for disaster recovery of a Cell Manager, you need to ensure the backup images are created in the correct order: the Cell Manager filesystem backup image first, the Internal Database backup image second.
- When recreating volumes during Phase 1 of an automated disaster recovery process (Enhanced Automated Disaster Recovery, One Button Disaster Recovery), the original volume-compression flag is not restored (always saved to non-compressed).

Workaround: Restore the volume compression flag manually after restore.

- Disaster recovery functionality is supported only if the platform on which the Data Protector GUI is used and the platform of the system which will be recovered are the same. This means, for example, that you cannot use the GUI launched on a Windows 7 system to perform a Windows Server 2008 backup for Enhanced Automated Disaster Recovery (EADR).
- One Button Disaster Recovery (OBDR) functionality is available only locally on the system to which the OBDR device is connected.
- The path of DR image file is limited to 250 characters, if it is saved on the Cell Manager during backup.
- The booting of Disaster Recovery image for Linux OS in physical machines with UEFI ROM can take a long time. Usually, it takes a minute for each physical core. But this does not affect the Disaster Recovery (DR) process, as once the image is booted, the DR continues to work properly.

User interface limitations

- The Data Protector GUI can display a limited number of backup specifications. The number of backup specifications depends on the size of their parameters (name, group, ownership information and information if the backup specification is dynamic or not). This size should not exceed 80 kB.
- The Data Protector command-line interface (CLI) does not support logging of user-triggered events to the Data Protector event log.
- On Linux systems, messages and notifications of the Data Protector CLI are only available in the English language.

Reporting limitations

- Information about physical devices, which is shown in the Device Flow report if the `RptDisplayPhysicalPath` global option is set to 1, is acquired from the current device configurations and may therefore be different from information at the time when the devices were actually used.
- In the Manager-of-Managers enterprise (multicell) Device Flow Web Report, devices are not sorted separately for each Cell Manager in the MoM.
- The following reports provide information only on destination media: Configured Devices not Used by Data Protector, Extended Report on Used Media, Report on Used Media, Session Media Report,

and Session Devices Report.

- The virtual machine must not contain the following special characters: & ^ \$! ~ , . ' ; () { } [], otherwise the `omnidb -session <sessionID> -detail` does not display the correct Object Name, VM Path, and VM Name, and the `omnidb -veagent <ObjectName>` report displays a syntax error or "No objects were found" error message.

Note: This is applicable for VADP reporting feature.

Other limitations

- For StoreOnce federated stores, all writing operations are performed in the low bandwidth mode (Application Source or Backup Server deduplication). Even if a gateway is configured as target-side deduplication (high bandwidth mode), it automatically switches to the low bandwidth mode.
- Only local shared storage (connected to cluster nodes via SCSI) is supported in cluster environments for ASR. Shared storage on disk arrays connected to cluster nodes via Fibre Channel (for example: P6000 EVA or P9000 XP disk arrays) is not supported unless appropriate device drivers are provided during the initial phase of ASR recovery (by pressing F6). This enables Windows Server 2003 Setup to correctly detect shared storage located on disk arrays.

It is necessary to execute a test plan. The operation is at your own risk.

- Data Protector does not support hostnames with non-ASCII characters.
- Do not export media which contain integration object copies made from platforms that support Unicode (for example, Windows) to non-Unicode platforms (for example, HP-UX) or the other way round.
- The STK - Horizon Library manager is not supported.
- The filestream database trans backups are not supported.
- You cannot select different condition factors for pools sharing the same free pool. All media pools using a free pool inherit the condition from the free pool.
- Device files for the spt driver cannot be created automatically by Data Protector. The device file needs to be created manually using the `mknod` command.
- Media pools with magazine support cannot use free pools.
- Data and catalog protection can only be set until the year 2037.

Workaround: Set protection period to 2037 or less and extend it with one of the future Data Protector releases that will support time settings past the year 2037.

- The network connections from the Cell Manager to Disk Agent clients must respond within 10 seconds or the session will be marked as failed.
- The name of a backup specification should not exceed 64 characters.
- The maximum length of text strings to identify or describe the properties of media and devices (for example, the media label applied to a medium when being initialized) is 80 characters.
- Session level restore is not available for the online database integrations.
- Automatic device selection during restore and/or object copy is limited to libraries. Only a device in a library can be automatically replaced with another device from the same library and of the same media type (for example, LTO).
- Automatic device selection during restore cannot be disabled for the Data Protector integrations that

cannot be restored using the Data Protector GUI or CLI (for example, Sybase integration).

- The minus symbol (–) must not be used as the first character in any Data Protector labels or descriptions.
- The word DEFAULT is a reserved keyword and must not be used in device names, backup specification names, and pool names.
- All media with barcode labels starting with the CLN prefix are treated as cleaning tapes. Labels with this prefix should only be used on cleaning tapes.
- Software data compression for online database backups, such as Oracle, Sybase, SAP R/3, Informix Server, and Microsoft SQL Server, is not supported.
- The eject/enter functionality for ATL 2640 and ATL 6/176 devices is not supported using the fast access port.
- Media of different format types are not compatible:
 - Data Protector (written by devices under direct Data Protector MA control)
 - NDMP NetApp (written by devices connected to NetApp filers)
 - NDMP Celerra

Media from these different format categories cannot reside in the same pool. Media from one format category cannot be recognized when subjected to one of the other environments using a different format category. In such a case, the media will be viewed as foreign and depending on the policy, unexpected overwrites might occur.

- From one backup object, only 1024 files and/or directories can be selected, otherwise select the entire object. For details about backup objects, see the *HPE Data Protector Help*.
- Some filesystems allow creation of deep directory structures (deeper than 100). Data Protector can only back up down to a depth of 100.
- When changing the `omnirc` file, it is required to restart the Data Protector services/daemons on the system. This is mandatory for the `crs` daemon on UNIX and recommended for Data Protector Inet and CRS services on Windows. On Windows systems, restarting is not required when adding or changing entries, it is required only when removing entries.
- If you use quotes (") to specify a pathname, do not use the combination of a backslash and quotes (\"). If you need to use trailing backslash at the end of the pathname, use double backslash (\\).
- Tape quality statistics functionality is not supported if the Media Agent runs on a Linux or AIX system.
- Automatic drive cleaning for library definitions with a shared cleaning tape is not supported. Each library definition needs to have its own cleaning tape configured.
- The maximum pathname length supported by Data Protector is 1023 characters.
- Devices of file library type are not supported for filesystems that have compression turned on.
- The length of the pathnames of the directories that can be used for configuring devices of the file library type cannot exceed 46 characters.
- The length of the pathnames for jukebox slots and standalone file devices cannot exceed 77 characters.
- Data Protector does not support copying a media copy. However, such a copy can be made if the original medium is exported and thus the copy becomes the original. If you export the second level

copy, you cannot import it again if the original medium is imported.

- The configuration of SNMP traps using the Data Protector Manager depends on the platform of the Cell Manager:
 - On HP-UX systems, the recipient system for the trap that is configured in the GUI receives the traps.
 - On Windows systems, the content of the recipient field in the GUI is ignored. The recipient must be configured on the Cell Manager in the Control Panel under **Network > Services > SNMP Services**.
- If Boot Configuration Data (BCD) is located on removable storage like floppy disk, USB flash drive, CD-ROM, or DVD-ROM, Data Protector cannot back up BCD registry entries.
- The Windows NTFS Change Log Provider cannot be used with Hierarchical Storage Management (HSM) solutions.
- The maximum size the Windows Change Journal is 4 GB. This space allows logging about 10,000,000 changes. After the maximum size is reached, a part of data is overwritten. In this time frame an incremental backup should be run.
- Automated System Recovery (ASR) cannot be used in IPv6-only environments. ASR can be used only in environments with a functioning DHCPv4 server.
- The Data Protector replication must not be configured, when:
 - a. Copying from two or more source Stores on an Appliance, to one target Store on an Appliance.
 - b. Copying from two or more source Stores on an Appliance, to two or more target Stores on an Appliance.
 - c. Copying from one source Store on an Appliance, to two or more target Stores on an Appliance.
- For Data Protector integration objects, the following actions cannot be restricted by using the user_restrictions file:
 - Start backup
 - Start backup specification
 - Start restore
- If the VADP reporting is enabled when a Virtual Machine (VM) is configured with IPv6 and DNS hostname (FQDN) is not known, reports display the VM name instead of the IP address.
- At the end of a Data Protector integration backup preview session, the backup statistics report that gets displayed contains irrelevant information. The following statistics are always zero: Completed Media Agents, Failed Media Agents, Aborted Media Agents, Media Agents Total, Mbytes Total, and Used Media Total.
- On UNIX systems, the original creation timestamp of a symbolic link is not preserved during a restore. The timestamp is set to the current system time. Due to a limitation of the system call `utime()`, the creation timestamp of a symbolic link cannot be changed after the link creation.

Recommendations

Using Data Protector with hierarchical storage management applications

If you are using an application such as file migration or hierarchical storage management that creates links, stubs, or shortcuts, it is recommended that Data Protector backup policies and file migrations do not execute against the same data set at the same time. This will avoid the situation where the backup may capture the stub and the migrated file in an inconsistent state.

Organizing Data Protector clients into cells

In small environments, the most simple approach is to manage all Data Protector clients within one Data Protector cell.

To efficiently hierarchically structure and manage large-scale environments, you can use the Data Protector Manager-of-Managers (MoM). An environment structured in such a way allows you to manage numerous clients from a central location; for MoM-related scalability limitations, see ["Backup infrastructure scalability" on page 26](#). Furthermore, multiple MoM cells can be centrally managed using the HPE System Management. Such a setup allows you to manage an unlimited number of Data Protector clients from one central location while distributing administrative and managerial rights to different Data Protector users and user groups.

Support for NIS+

NIS+ cannot be used as the primary name resolution for hosts when using Data Protector. However, you can run Data Protector on the hosts where NIS+ is configured if one of the following alternatives for name resolution with Data Protector is chosen:

- Using DNS. In this case, change the line starting with hosts in the `/etc/nsswitch.conf` file as follows:

```
hosts: dns [NOTFOUND=continue] nisplus
```
- Using hosts file. In this case, change the line starting with hosts in the `/etc/nsswitch.conf` file as follows:

```
hosts: files [NOTFOUND=continue] nisplus
```

In both cases, the Cell Manager must have fully qualified domain name registered in DNS or hosts file.

Large file support

- HPE recommends that the file system where DC directories reside supports files larger than 2 GB, especially if drives with large capacity, for example LTO 6, are used, and more than 10 million files are backed up on tape. In addition, on Windows systems HPE strongly recommends to use NTFS filesystem.

Encrypted control communication recommendations

- Use the generated keys and certificates instead of the default HPE Data Protector certificate `hdpcent.pem`.
- Consider the manual distribution of keys and certificates to clients.

For more information, see *Enabling encrypted control communication with manual distribution of certificates and keys* in *HPE Data Protector Installation Guide*.

Pre-exec and Post-exec scripts

The Pre- and post-exec commands for backup specifications must be located as follows:

- **Windows systems:** The scripts must be located in the `Data_Protector_home\bin` directory.
- **UNIX systems:** The exec commands for backup specifications must be located as follows:
 - **HP-UX, Solaris, and Linux systems:** `/opt/omni/lbin`
 - **Other UNIX systems:** `/usr/omni/bin`

Note: For the commands located in the `/opt/omni/lbin` or in the `/usr/omni/bin` directory, specify only the filename, otherwise, specify the full pathname. For more information, see *Pre- and Post-Exec Commands for a Backup Specification* in *HPE Data Protector Online Help*.

With the enhancements in Data Protector [DP.DP_Version_0905], if the Pre- and Post-exec scripts are not located in the specified directories, you will encounter an error message and the backup specification will fail.

Hence, it is highly recommended that you run `omnicellinfo -prepostinfocheck` to identify the locations that are not valid before performing backup. If the scripts are located in folders different than the ones specified, move the scripts to the specified locations and then continue.

Enhanced incremental backup

- To enable Data Protector Disk Agent to access more memory if needed for enhanced incremental backup on HP-UX systems, set the tunable kernel parameter `maxdsiz` as follows:

HP-UX 11.11 systems:

```
kmtune set maxdsiz=2147483648  
kmtune set maxdsiz_64bit=2147483648
```

HP-UX 11.23/11.31 systems:

```
kctune set maxdsiz=2147483648  
kctune set maxdsiz_64bit=2147483648
```

Object consolidation

- When consolidating a large number of objects from synthetic backup with very long restore chains, an error might occur. To prevent this, run object consolidation regularly, for example, when you would normally run a full backup, to keep the restore chain manageable.
- Before starting an object consolidation session, ensure that the order of the objects is kept the same. Changing the order of the backed up objects may result in object consolidation failure.

Microsoft Exchange Single Mailbox integration

- Microsoft Exchange Server single mailbox backup is not as space- and CPU-efficient as backup of the whole Microsoft Exchange Server. HPE recommends to use the Microsoft Exchange Single Mailbox integration only for backup of a small number of mailboxes. If you are backing up large numbers of mailboxes, use the Microsoft Exchange Server integration instead.

Microsoft Volume Shadow Copy Service integration

Shadow copy storage area and disk space recommendations

- When backing up volumes using VSS (either using the Disk Agent or the VSS integration), ensure that there is enough free space available for the shadow copy storage area.

By default, the initial size for the shadow copy storage area is set to 300 MB on Windows Server 2003 systems (100 MB if the hotfix KB826936 is not installed) and Windows Server 2008 systems, and 320 MB on Windows Server 2008 R2 systems. This means that for example on Windows Server 2008 R2 systems with the default settings there must be at least 320 MB of free space available on the volume that you are backing up.

If you encounter timeout errors during the shadow copy creation, you may also want to increase the initial size for the shadow copy storage area. For details, see the Microsoft Knowledge Base article at <http://support.microsoft.com/kb/826936>.

Regular maintenance of the VSS part of the registry

- Microsoft Windows operating systems maintain a record of mount operations in the registry. This process results in registry growth over time and eventually leads to volume shadow copy import problems. For details, see the *HPE Data Protector Zero Downtime Backup Integration Guide*, chapter *Integrating the Data Protector ZDB integrations and Microsoft Volume Shadow Copy Service*, section *Troubleshooting*.

To prevent the registry from growing excessively, HPE recommends that you periodically perform registry management tasks with Microsoft Registry Management Tool.

Network Data Management Protocol Server integration

- The maximum number of files and directories per NDMP backup specification should not exceed 20 million. The recommended number of files and directories per NDMP backup specification is 10 million.

Windows Server 2008 clients

- **Server roles and services on Windows Server 2008**

Similarly to previous Windows Server operating system releases, Microsoft extended the concept of server roles and services in Windows Server 2008. To enable backup of data belonging to the server roles and services introduced with Windows Server 2008, Data Protector provides extended filesystem backup functionality for this platform. Among others, the following roles can be backed up using filesystem backup:

- Active Directory Certificate Services (AD CS)
- Active Directory Domain Services (AD DS)
- Application Server (requires IIS 6 compatibility)
- Dynamic Host Configuration Protocol (DHCP) Server
- DNS Server
- Network Policy and Access Services
- Terminal Services
- Web Services (IIS) (requires IIS 6 compatibility)

When configuring a backup specification for data belonging to a particular server role or service, you should select either the entire volume on which the data resides or the entire client system that hosts the server role or service. Moreover, you should select the **Use Shadow Copy** option on the **WinFS options** property page of the **Filesystem Options** window. When selected, this option provides a consolidated, consistent state of the backed up data.

Caution: Additionally, if configuring a backup specification for disaster recovery purposes, clear the option **Allow Fallback**. Failing to do so may result in backup data unusable for disaster recovery.

- **System State backup and the CONFIGURATION object**

To perform a System State backup on Windows Server 2008, you should follow the above instructions for filesystem backup of the relevant volumes or the entire client system, instead of backing up the CONFIGURATION object.

- **Active Directory Domain Services restore**

On Windows Server 2008, only offline restore of Active Directory Domain Services is supported, which must be performed in Directory Services Restore Mode. Since the Active Directory Domain

Services restore is a complete overwrite of the existing database, it does not preserve any new users which are created after the backup operation.

Windows Server 2012 clients

- When performing a filesystem backup or restore of a deduplicated volume, data is backed up and restored rehydrated (nonoptimized backup or restore). Therefore, consider the following:
 - Make sure that the target media and the restore target have enough storage space available:
 - Plan the storage space needs based on the original (logical) size of data, *not* the space occupied by deduplicated data on the volume. For example, if 100 GB of data occupies 40 GB of space on the volume, plan for 100 GB.
If you do not have enough free capacity, and you do not need to restore individual files, you can perform a disk image backup instead (optimized backup).
 - Data is always restored to the original (non-deduplicated) size, in the above example, it occupies 100 GB space after restore.
Schedule a periodic deduplication process at the restore target system to deduplicate the restored data again.
 - In case of a backup to a Backup to Disk device, HPE recommends to use the Data Protector deduplication functionality (either StoreOnce software deduplication or the method provided by the device) to reduce the amount of space consumed by the backed up data. To reduce the network load, use source-side deduplication.
 - If you back up data with a high deduplication ratio, an optimized disk image backup is faster because less data needs to be transferred over the network.

UNIX system clients

- When performing a disk image backup, HPE recommends to dismount disk partitions before the backup and mount them back later.

Chapter 4: Recognized issues and workarounds

This section lists known Data Protector and non-Data Protector issues and workarounds.

Known Data Protector issues and workarounds

Installation and upgrade related issues

- After upgrading to Data Protector 9.00 or later, if the backup specification is created using Data Protector 7.0, 7.01, or earlier versions, the VEAgent backup fails with the following error:

```
[Critical] From: VEPALIB_VMWARE@hostname; "<Datacenter>" Time: <Date Time>;  
No Objects found for backup
```

Workaround: After upgrading, re-create the backup specification with the same VM selection and options as earlier, and run the backup again.

- If you install the Installation Server directly from another installation server (where GR, and MMR patches have been applied), these patches will not be shown in the list of Installation Server patches for the newly installed Installation Server. However, patches are only not shown but the new Installation Server depot contains the same patches as the one from where it was installed.

Workaround: Hence if you need to see the installed patches, then copy the patch_* files from the original Installation Server system <DP_DATA_DIR>\Config\Server\install\patch_* to the newly installed Installation Server system: <DP_DATA_DIR>\Config\Server\install

- Encryption keys are not migrated correctly when migrating the Cell Manager from 32-bit to 64-bit Windows systems. As a result, restore of encrypted backups fails after the migration.

To ensure that encryption keys are correctly migrated, perform the following:

- a. Export all keys from the Key Management Server (KMS) on the 32-bit system using the `omnikeytool` command.
 - b. After you perform the migration, *delete* all data (DAT) files from all key store folders from the directory `Data_Protector_program_data\server\db80\keystore` on the 64-bit system, except from the `catalog` folder. Do not delete the index files.
 - c. Import all previously exported keys to the KMS on the 64-bit system. After the import, encrypted backup can again be restored.
- **HP-UX and Linux systems:** For Data Protector cluster-aware clients, the Data Protector Cell Manager will only update configuration information for their cluster virtual system during the upgrade process, but not for the corresponding cluster nodes (physical systems).

Workaround: The issue has no effect on the actual state of such clients, only their configuration data is not upgraded. After the upgrade, to update the configuration data and complete the upgrade process, execute the command `omnicc -update_host ClientName` for each cluster-aware client, where `ClientName` is the name of a particular cluster node.

- **HP-UX and Linux systems:** The Data Protector GUI enables you to remotely install the components to a

virtual host, even though the components must not be added to the virtual host.

Workaround: None. Do not remotely install the components to the virtual host, but install the clients locally as described in the *HPE Data Protector Installation Guide*.

- Import of the cluster virtual host with Data Protector installed will not finish successfully (cluster will be imported but offline virtual servers will not be imported) during the installation of cluster-aware Cell Manager if there is another cluster virtual server configured on Microsoft Cluster Server in any cluster group and is offline. If this virtual server is online during the Data Protector installation, the import of the Data Protector cluster virtual server will be successful.

Workaround: Put all virtual servers in your cluster online and import the Data Protector cluster virtual server manually after the installation.

- If you upgrade a Data Protector client on an HP-UX 11.23 or HP-UX 11.31 system, the binaries of the Data Protector components that are not supported on HP-UX 11.23 or HP-UX 11.31 (for example, EMC Symmetrix Agent, DB2 Integration) are not removed. If you later uninstall Data Protector, the binaries are left on the system.

Workaround: Uninstall the earlier version of Data Protector. Install Data Protector 9.00, and then install the Data Protector[[[Undefined variable DP.DP_Version_0905]]] patch.

- On Windows systems, desktop shortcuts for starting Data Protector that were created by the user, for example by dragging the menu item to the desktop, do not function after an upgrade.

Workaround: Recreate the desktop shortcuts after upgrading.

- In a cluster-aware Cell Manager configurations running on an HPE Serviceguard cluster or Veritas Cluster Server, the installation check fails on the non-active node although Data Protector is correctly installed, because only the active node can access the Cell Manager configuration.

If the cluster fails over, the check on the now active node succeeds.

- If a remote UNIX or Linux client installation fails, and you restart the installation using the **Restart failed clients** option, the installation is either skipped or fails again, although the issue that caused failure of the first installation session is resolved.

Workaround: Locally uninstall the client and repeat the remote installation. For uninstallation details, see the *HPE Data Protector Installation Guide*.

- On Windows systems, the Data Protector installation might fail with the following error:

```
Error 1601. The Windows Installer Service could not be accessed. This can occur if the Windows Installer is not correctly installed. Contact your support personnel for assistance.
```

The root cause of the problem is the Windows Installer service that could not be started at the beginning of the installation.

If the service cannot be started, the installation fails.

Workaround: In the **Control Panel > Administrative Tools > Services**, change the startup type for the Windows Installer service from Manual to Automatic, start the service, and restart the Data Protector installation.

- After upgrading the Data Protector Virtual Environment integration component without the installed patch bundle set 6.21 to the latest release version, passwords of all virtual environment hosts will no longer work. To solve this, execute the following command:

```
vepa_util.exe --upgrade -cell_info
```

This is needed due to a change in password encoding in the `cell_info` file. It will re-encode the passwords of all virtual environment hosts, first creating a `cell_info.bak` file.

- On HP-UX systems, the following message may be reported during installation of the Data Protector Cell Manager, but the installation succeeds nevertheless:

```
* "Hostname:/cdrom/hpux/DP_DEPOT": Cannot open the logfile on
this target or source. Possibly the media is read-only or
there is a permission problem. Check the daemon logfile and
"/var/tmp/swagent.log" on this host for more information
```

Workaround: None. You can safely ignore the message.

- During the installation, Data Protector setup will update only Private and Domain Windows Firewall profiles. The Public profile is not updated. As a result, communication between Data Protector components in the cell may not work properly, resulting in various issues.

Workaround: If you are using the Public profile, update the firewall rules manually. For more information about Data Protector and firewalls, see the *Data Protector Help index: "firewall support"*.

- On HP-UX systems, the HPE Software Assistant (SWA) may report errors after installing General Release Patches:

```
ERROR: Patch PHSS_XXXXX is not recognized.
```

This occurs in cases where General Release Patches have been correctly installed on top of a Data Protector patch bundle, and the patches contained inside the bundle are not present in the catalog referenced by SWA.

Workaround: None. You can safely ignore the messages.

- On HP-UX systems, after you install a patch on the Cell Manager system, the GUI connection to the proxy service fails with an exception:

```
16:38:04,534 SEVERE [org.jboss.resteasy.core.SynchronousDispatcher] (http--
0.0.0.0-7116-4)
```

```
Failed executing GET /backupspec: org.jboss.resteasy.spi.WriterException:
java.lang.IllegalStateException: Invalid JSON namespace:
http://www.hp.com/2011/software/im/dp/data_model at
org.jboss.resteasy.core.ServerResponse.writeTo(ServerResponse.java:262)
[resteasy-jaxrs-2.3.2.Final.jar:]
```

Workaround: After installing a patch, restart the Cell Manager.

- On Windows systems, the following message may be reported during installation of Data Protector, but the installation succeeds nevertheless:

```
{A37E26EF-E4F1-432B-ABA4-02268BC99B80}: related product unexpectedly found on the
system. or/and "{30692C3E-7A60-4BD4-B021-213055B1810F}: related product
unexpectedly found on the system."
```

The root cause of the problem are previously uninstalled Data Protector components VMware Granular Extension Web Plug-In and VMware Granular Recovery Extension Agent, which the uninstallation process did not remove.

Workaround: You can safely ignore the message. To properly uninstall the Data Protector Granular Recovery Extension for VMware vSphere, manually remove entire Data Protector (all installation

components) from your system. For details on removing Data Protector, see the *HPE Data Protector Help* index: “uninstalling, Data Protector software”.

- While upgrading to the latest Data Protector General Release patch bundle, the installation displays the following warning:

```
Timeout reached before Data Protector Application Server stopped.
```

Workaround: You can safely ignore the message

- Prior to an upgrade to Data Protector 9.04 and above, if you have used FUSE in Data Protector 9.02 and 9.03, and 3PAR VMware GRE, the older FUSE mount points and 3PAR replica presentations to the mount proxy must be cleaned up.

To clean up the older FUSE mount points and 3PAR replica presentations to the mount proxy, proceed as follows:

- Delete the request files from the Cell Server. The request files are available in the following location:

Windows Cell Manager:

```
C:\ProgramData\OmniBack\Config\Server\Integ\Config\Vmware\\
```

Linux Cell Manager:

```
/etc/opt/omni/server/integ/config/Vmware/<vCenterHost>/
```

- Unmount the vmfs fuse mount point (Specific to Linux only)
#umount <MountPoint>

The <MountPoint> format is as follows:

```
/var/opt/omni/tmp/VMWareGRE/<vCenterName>/<reqID>
```

- Unpresent replica from the 3PAR Array console, if requests are created using the ZDB backups.
- Remove the nfs share from the Media Agent, if requests are created using the Smart Cache backups.

User interface related issues

- When you are setting up a user account for the Data Protector Inet Service user impersonation using the Data Protector GUI, the configuration may fail with an error message similar to the following:

```
Failed to modify config information for user myuser@hostname.
```

Workaround:

- a. Connect to the client where the issue appears.
- b. Delete the user impersonation configuration for the specified client using the `omniinetpasswd` command:

```
omniinetpasswd -delete myuser@hostname
```

- c. Reconfigure the user impersonation for the specified client using the `omniinetpasswd` command:

```
omniinetpasswd -add myuser@hostname
```

For details of the `omniinetpasswd` command, see the *HPE Data Protector Command Line Interface Reference*.

- When using the Data Protector CLI on a Windows system to manage backups of data residing on clients running on other platforms, the filenames will only be displayed correctly for code page 1252. Characters from other code pages will appear corrupted. Even though a filename appears corrupted in the CLI, it will be backed up or restored properly. Data Protector CLI expects such "corrupted" filenames as input parameters. You can use copy and paste functionality to input filenames as they appear in code page 1252.

For internationalization limitations tables, see the *HPE Data Protector Help* index: "internationalization".

- On Windows Server 2003 systems, after saving backup specifications whose names include non-Latin characters (for example, Russian or Greek), names of the backup specifications may appear corrupted in the Data Protector GUI.

Workaround: Install Windows Server 2003 Service Pack 2 on the system where Data Protector GUI is installed.

- Backup or object operations sessions scheduled in the basic scheduler that experience problems and miss their scheduled execution times will not be displayed in the Missed Job Executions list.

Workaround: To schedule backup or object operation sessions so that they display in the Missed Job Executions list, use the advanced scheduler.

- The advanced scheduler feature is not available for the backup templates workflow.

Workaround: To schedule backups using backup templates, use the basic scheduler.

- The advanced scheduler feature is not available for replication specifications. Selecting the Advanced Scheduler menu item for a replication specification fails with the error "Unable to find operation".

The advanced scheduler will also no longer display existing object copy specifications originally scheduled in advanced scheduler, and subsequently edited to include replication.

Workaround:

- To schedule a new replication specification session, use the basic scheduler.
- To add replication to existing object copy specifications originally scheduled in the advanced scheduler, delete them in advanced scheduler and recreate them in the basic scheduler.
- The advanced scheduler feature is not available for replication specifications. Selecting the Advanced Scheduler menu item for a replication specification fails with the error "Unable to find operation".

The advanced scheduler will also no longer display existing object copy specifications originally scheduled in advanced scheduler, and subsequently edited to include replication.

Workaround:

- To schedule a new replication specification session, use the basic scheduler.
- To add replication to existing object copy specifications originally scheduled in the advanced scheduler, delete them in advanced scheduler and recreate them in the basic scheduler.
- The session messages for a paused backup session display an error from the Disk Agent when the backup session is scheduled from the Advanced Scheduler and has the ability to pause lower priority sessions.

[Normal] From: BSM@hostname "Bkp_Low" Time: <Date> <Time>

Session pause request received by user ADMINISTRATOR.123@hostname. Pausing the

session.

```
[Normal] From: VBDA@hostname "C:" Time: <Date> <Time>  
Received ABORT request from SM => aborting.
```

Workaround: None. This is the expected behavior when a backup session is being paused on a busy backup device. The warning can be safely ignored.

- The session messages for a paused backup session may contain some failure messages for a session that succeeds.

In certain scenarios where the same backup session is scheduled from the Basic Scheduler and from the Advanced Scheduler for the same device, a lower priority session is waiting for a device to start before it receives a session pause request due to a higher priority session. The session logs contain all of the success messages, warning and pause messages.

```
[Warning] From: BSM@hostname "Bkp_Low_FL1" Time: <Date> <Time>  
[61:2013] Some of the backup devices are occupied. Session is waiting for all the  
devices to get free.  
[Normal] From: BSM@hostname "Bkp_Low_FL1" Time: <Date> <Time>  
Session pause request received by user@hostname. Pausing the session.
```

Workaround: None. This is the expected behavior when a backup session is being paused on a busy backup device. The warning can be safely ignored.

Disk Agent related issues

- When attempting a parallel restore which uses more Disk Agents than the current Media Agent concurrency setting, some Disk Agents may fail with the following error:
Cannot handshake with Media Agent (Details unknown.) => aborting.
Workaround: Restart the restore objects of the failed Disk Agents.
- During restore, the volume restore Disk Agent (VRDA) displays the mount points of the application system in the monitor. For example, instead of the restore target mount point `/var/opt/omni/tmp/hostname/BC/fs/LVM/VXFS` it actually displays the corresponding application source mount point `/BC/fs/LVM/VXFS`.
- When restoring files to a different system via a UNC share, the restore fails with the following message in the session log:
Can not open: ([112] There is not enough space on the disk.) => not restored.
[Warning] From: VRDA@hostname "host2.test.com [/H]" Time: <Date> <Time> Nothing restored
Workaround: The Data Protector Inet logon user account must have the access to log on to the remote system, which is specified in the UNC path. You should also be the owner or have permission to write to the files you want to restore via UNC share.
- When trying to back up directory structure with more than 100 directories (on HP-UX systems, this number is equal to the maximum number of allowed open file descriptors), the following message is displayed twice instead of once:
[Major] From: VBDA@hostname "C:" Time: <Date> <Time>
[81:74] File system too deep: (100) levels.
- When backing up a mount point that resides on a Windows system, if a subdirectory is deselected

and therefore excluded from backup, the whole mount point might be backed up nevertheless.

- When trying to expand the empty Windows mount point in tree view, the following error is reported:
Cannot read directory contents.
- On Windows, the encrypt attribute of encrypted folders will be restored. However, only a user who logs on using the account under which the Inet service runs on the client or an Administrator will be able to remove the attribute.
- On Windows systems, when backing up Macintosh files, certain characters in file names may cause problems. If file names contain characters considered invalid on a Windows filesystem (typically '*' and '?'), or contain characters mapped to such invalid characters (for example, the Macintosh bullet character), it is possible that individual files are not backed up or that the Disk Agent terminates abnormally.

Workaround: Rename the problematic files.

- Data backed up from a shared network folder using Data Protector Disk Agent installed on a Windows Vista, Windows 7, Windows 8, Windows Server 2008, or Windows Server 2012 system cannot be restored to its original location, even though the user account which was used during the backup session is granted write permissions for the folder.

The problem occurs because Data Protector does not have impersonation capability for filesystem restore sessions.

Workaround: Using the `runas.exe` command, start the Data Protector GUI as the user whose account was used during the backup session, and only then start the restore session.

- On HP-UX systems, when performing a disk image backup, a warning message is displayed although the backup session succeeds:

Object is a mounted filesystem.

Workaround: None. Check if the disk or volume is mounted. If it is not mounted, you can ignore the warning message.

- If you schedule several replication sessions to run in parallel and the replication source are also replication sessions, the sessions may fail with an error similar to the following:

```
[Major] From: CSM@hostname "QCTP2A53730" Time: <Date> <Time>  
[65:99] Import failed with possible cause:  
this media already has valid copy in DB.
```

The issue is caused by objects with identical labels that appear in multiple backup specifications, for example, if you create multiple backup specifications for different directories of the same filesystem on the same client.

Workaround: Using the Data Protector GUI, provide different descriptions for the conflicting objects in the backup specifications that were the initial source for the replication specifications or make sure that replication sessions containing these objects are not started in parallel.

- It is possible to start a replication when the source and target of the replication is the same B2D device because the Data Protector GUI is not able to distinguish between active source devices.

If the user attempts to create an object copy specification using the **Capable of Replication** option, it is possible to select the same source and target B2D devices in the Data Protector GUI.

Workaround: Ensure that the replication source and targets are different B2D devices.

- When resuming a backup on a filesystem where too many changes took place from the time the session was aborted, Data Protector might not be able to use the Windows native change journal to

analyze what files have already been backed up. In this case, a filesystem scan will disregard all resume related information upon resume. The backup session will be displayed as resumed, however, all the files will be backed up again, which might result in a larger size of the resumed backup than expected.

- When trying to backup a system reserved partition and multiple full volume objects, the backup fails with either of the following error message:
 - Cannot read <number> bytes at offset <number>(:1): ([21] The device is not ready.).
 - Cannot open: ([2] The system cannot find the file specified.) => not backed up.

Note: The problem occurs only if the VSS option is enabled and if the system reserved partition does not have enough space to hold multiple snapshots.

Workaround:

Set the omnirc variable `OB2_DISABLE_REGLIST_FOR_FULL_VOLUME` to 1 and restart the backup. If the error persists, see the following Microsoft webpage for information on how to resolve this problem:

<http://support.microsoft.com/kb/2930294>

- When restoring a folder (C:\Program Files\WindowApps) on a Windows 8.1 client, the restore fails with an error similar to the following when this folder is overwritten:

Cannot write: ([13]) The data is invalid.) => not restored.

Workaround:

- a. Restore the folder to an alternate location by performing the following:
 - In the Context List, click **Restore**.
 - In the Scoping Pane, expand Filesystem and select the folder from which you backed up.
 - In the Destination page, select **Restore to new location** and specify an alternate location.
 - Click **Restore**.
 - b. Change the access rights on the folder to provide full access permissions to the Administrator.
 - c. Delete the folder and move the restored folder to this location.
- A ZDB filesystem backup with the **Enhanced incremental backup** option enabled will result in full backup, if the ZDB is configured to add the **session ID** directory to the mount path.

Workaround:

Use **Hostname** option for directories to be added to the mount path under **Backup system** option section. Ensure that mount path is free before the next session by using **Automatically dismount filesystems at destination mountpoints** option or be sure that **Leave the backup system enabled** is not selected.

- Disk Agent does not backup all the sub volumes directly, if accessed using its parent sub volume. Therefore, the sub volumes must be mounted and backed up separately. BTRFS (a new file system for Linux) as a feature enables creating sub volumes from one folder tree. So, you can have many sub volumes on one file system. Once such a sub volume is created, the Disk Agent will not able to backup files from that sub volume.

Workaround:

You can mount that volume as a new mount point and configure that in the backup specification to backup the mount point.

Media Agent related issues

- If during a backup session a shared StorageTek ACS tape library is used as a backup device, and intercommunication between a Disk Agent and a Media Agent is interrupted, the Utility Media Agent (UMA) may stop responding. Consequently, subsequent sessions that use the involved tape drive may fail.

Workaround: Use the `omniirc` option `OB2ACSUMATIMEOUT` to specify how long Data Protector should wait for the connection between the Disk Agent and the Media Agent to be restored before it terminates the UMA.

- If during a backup, copy, or restore session, SCSI read or SCSI write errors are intermittently reported, there may be intercommunication problems between a Media Agent and a SCSI device connected to SAN.

Workaround: The problem may be solved by configuring the following `omniirc` options on the affected Media Agent system: `OB2MAREADRETRY`, `OB2MAXREADRETRIES`, `OB2MAREADRETRYDELAY`, `OB2MAWRITERETRY`, `OB2MAXWRITERETRIES`, and `OB2MAWRITERETRYDELAY`.

- In a cell where the Cell Manager is not installed on the cluster, the devices are connected to cluster nodes, and a failover during backup activity occurs, the Media Agent may not be able to properly abort the session, which results in the medium no longer being appendable.
- Cleaning tape drive functionality functions correctly only when there is a cleaning tape present either in the library slot or in the repository slot. If the cleaning tape is not present, the mount request for the cleaning tape will not function properly.
- When importing a range of tapes, Data Protector normally skips all invalid tapes (such as tar tapes, blank tapes, and so on) and continues with the next slot. When importing a range of tapes on NetApp Filer (Celerra) and when a NetApp tape is detected, Data Protector reports a major error and ends abnormally.
- If ACSLS library mount request occurs during backup or restore session (in case that library ran out of usable media), do not format or scan additional tapes with the tape device currently being used by the session. Use a different tape device in the library to perform this operation and confirm the mount request.
- During a backup session, if you restart the system that hosts a Data Protector Media Agent, the medium to which data is backed up with this Media Agent becomes corrupted, although Data Protector does not report any errors. Consequently, you may not be able to restore any backup data from this medium. Subsequent backup sessions to the corrupted medium will fail, too.
- Backup sessions for backing up to a file library device ignore the media pre-allocation list.
- If the media of a file library device are unprotected, they are deleted at the beginning of the next backup session that is using this device. However, the session which was using the first medium of the file library device is still stored in the database. If you attempt to restore data by specifying this session, the restore fails and the following message is issued:
Object not found.
- If a disk becomes full during a backup session using a jukebox (with media of type file) as destination device, all slots configured on this disk which contain unprotected media will be marked

as empty.

Workaround:

- a. Rescan the slots which are marked as empty.

After the rescan, the media will be visible again in the slot.

- b. Free up space on the disk to avoid this problem from recurring.

After performing both steps, you can continue to work with the jukebox device.

- An object copy session containing numerous objects (more than 200) or complex object media relations (see below) may become unresponsive.

Workarounds:

- Change the device mapping so that only one device is used to read the copy source media per media type (DLT or LTO) and restart the session.
- Split the original object copy session into multiple sessions and restrict each session to copy objects from one backup session only.
- Split the original object copy session into multiple sessions and restrict the session to copy as few media as possible in a single session.

Unresponsiveness is commonly caused by copying objects from source media which were created by different backup sessions using different (logical) devices.

- An NDMP object copy session with more than one source device used for the object copy may become unresponsive.

Workaround:

Limit the number of objects to be copied. Additionally, limit the number of source devices to 1.

- When an external encryption controller is controlling encryption on a tape device, a failure to read the tape medium header of a previously encrypted medium can occur. This happens if the connection to the external encryption controller is not available or a decryption key is deleted from the external encryption controller.

Workaround:

Set the `OB2_ENCRYPT_FORCE_FORMAT` environment variable to force a format operation on the tape.

The following options are available:

- If the variable value is set to 0 then a format operation will be aborted.
- If the variable value is set to 1 then Data Protector Media Agent will force a format operation.

The default value is 0 (not set).

- When importing legacy NDMP media populated in a Data Protector version earlier than 7.x and 8.x, there is a problem with imported UNIX file system backup objects: ownership flags (owner, group) of such objects are set to `0 0`. The root cause of the problem is incorrect handling of the ownership of such objects in earlier Data Protector versions, where the ownership fields stored alongside the object data were misused for purposes other than object ownership.

Workaround: None.

- When a device required for restore is unavailable or disabled, the restore session fails.

Workaround:

Restore your data with an alternative device. For instructions, see the *HPE Data Protector Help* index: "selecting, devices for restore".

- When more than 4 TB of backup data is stored to a single medium, its capacity and usage are incorrectly recorded in the IDB. Consequently, invalid figures are presented.

Workaround: None. Nevertheless, all the Data Protector operations such as backup and restore, function properly.

- In previous releases, Data Protector incorrectly calculated the space used by file libraries. As a result, the TB based licensing database tracked file library usage as being many times lower than actual.

Workaround: None. In this release, customers may see a sudden increase in reported file library usage when file library usage is checked with `omnicc`.

- The gateway filter cannot distinguish between media agent versions of Data Protector 9.08 and subsequent General Release patches. When attempting to create a Cloud device with a gateway on a media agent, the GUI reports an error of unsupported. Install the latest `[[[Undefined variable DP.DP_Version_0905]]]` patch.

Workaround: Update all clients with media agents to the latest General Release patch bundle.

- General Data Protector backups use upto 64 connections but backups over Fiber Channel (FC) use only 16 connections, which is four times less.

Workaround: Increase the number of initiators on the Media Agent to 64.

Integration related issues

Microsoft Exchange Server

- In the Data Protector GUI, the tape device you want to use for a Microsoft Exchange Server restore cannot be changed from the device originally used by backup.

Workaround: To change the device for restore, in the Data Protector GUI, click the **Change** button. You cannot change the device by just deselecting the default device and selecting the desired device.

Microsoft Exchange Single Mailbox

- When configuring the Microsoft Exchange Single Mailbox integration, the following issues may occur:
 - The CLI configuration session finishes without errors, but the configuration actually fails. When creating a backup specification, the configuration dialog displays. If the backup is started from the CLI or from the GUI where the configuration was not performed in GUI, the session finishes immediately without backing up any data.
 - If the integration was configured using the GUI, and you run the configuration check from the CLI, the check will fail with `*RETVAl *8561`.

Workaround:

- Use the GUI to configure the integration and to check the configuration.
- Set or export the environment variable `OB2BARHOSTNAME` on the client system with the command `set OB2BARHOSTNAME=client_name` (on Windows systems) or `export OB2BARHOSTNAME=client_name` (on UNIX systems) and repeat the configuration from the CLI.

Microsoft SQL Server

- In the Data Protector GUI, the tape device you want to use for a Microsoft SQL Server restore cannot be changed from the device originally used by backup.
Workaround: To change the device for restore, in the Data Protector GUI, click the **Change** button. You cannot change the device by just unselecting the default device and selecting the desired device.
- When performing an object copy session where B2D devices are the source devices, Data Protector may not enforce the requirement that different source media are copied to different target media.
Workaround: To make sure that several data streams are not multiplexed to the same medium, configure a separate object copy specification for each database and set the following:
 - the number of devices must be the same as the number of streams of an object in the selected session
 - the minimum value for Load balancing must equal the number of the devices
 - target device concurrency to 1.

To calculate the number of devices you need, multiply the maximum number of streams for any database in a session with the number of parallel object copy sessions from these specifications. For example, if there are a maximum of five data streams in a session and two parallel object copy sessions based on this session, you need 10 devices.

Microsoft Volume Shadow Copy Service

- When backing up many writer components at once, the backup fails if a transportable backup is performed in parallel.
The issue appears if you select many writer components (for example SQL databases for `SqlServerWriter` or Hyper-V virtual machines for the Hyper-V writer) directly instead of their container (for example, the whole writer) and, at the same time, a transportable backup is performed.
Workarounds:
 - Perform a local backup instead of a transportable backup. You can still use the VSS hardware provider if necessary. The writer components (databases, virtual machines, and so on) can be selected directly one by one.
 - Select the whole container (for example an SQL instance for the `SqlServerWriter` or the entire Hyper-V writer) instead of individual components (for example SQL databases or virtual machines). In this case, you can use VSS transportable backup.
- Data Protector might incorrectly process restore chains of two backup objects (volumes, folders, or files) during the restore sessions. For example, if both full and incremental backup images exist for

objects A and B, and you select a full backup image of A and an incremental backup image of B for restore, Data Protector overrides your selection for A by also including its incremental backup image in the restore chain.

Workaround: To make Data Protector correctly process different restore chains for multiple objects backed up in the same sessions, restore the objects separately in several restore sessions.

SAP R/3

- Backup of SAP R/3 data fails when the `-u` option is specified in the command line for the `brbackup` or `branchive` command.

Workaround: If you specified `-u` in the command line of `brbackup` or `branchive`, it should be followed by `username/password`.

Oracle Server

- On Windows system, Oracle backup sessions wait for 20 seconds before they end. This waiting time occurs because Oracle does not notify that the API session is complete. If you run a backup from RMAN and use the Data Protector library (`orasbt.dll`) to perform that task, you must wait at least 20 seconds between two backup sessions using the same backup specification. In the opposite case, all the backup objects will be backed up within the same backup session.
- The `ZDB_ORA_INCLUDE_CF_OLF`, `ZDB_ORA_INCLUDE_SPF`, and `ZDB_ORA_NO_CHECKCONF_IR` omnirc options are set and database recovery after instant recovery fails with the following error:

```
ORA-00338: log Name of thread Num is more recent than control file
```

The above message indicates that the control file was overwritten during instant recovery. This happens if the Oracle control file location was specified for the `control_file_location` parameter which should define the location of the control file copy.

Workaround: Perform recovery using a backup of the control file.

Ensure that `control_file_location` does not point to the location where the Oracle control file is located.

- If you restore backup data created using the proxy-copy method and perform a database recovery, RMAN may try to use the channel allocated for restoring proxy-copy backups to recover the database. As a result, the recovery fails.

Workaround: Start a database recovery-only session from the Restore context or using RMAN scripts.

- In an Oracle Real Application Clusters (RAC) environment with Oracle version 11.2.0.2 or later, a Data Protector managed control file backup ends unexpectedly with the following message:

```
The database reported error while performing requested operation.
```

```
ALTER DATABASE BACKUP CONTROLFILE TO '/var/opt/omni/tmp/ctrl_dbpp.dbf' REUSE  
sqlcode 245 error occurred at line 1.
```

```
ORA-00245: control file backup operation failed
```

By default, Data Protector backs up the Data Protector managed control file from the default Data Protector temporary files directory .

Workaround: The directory to back up the control file from must be located on a shared disk and be accessible from all RAC nodes. Define an appropriate directory by setting the `OB2_DPMCTL_SHRLOC` environment variable for the control file copy, and restart the session.

- If you are configuring a backup set ZDB backup specification that involves a disk array of the HPE 3PAR StoreServ Storage family, Data Protector may incorrectly detect that the chosen database files are managed by an Oracle ASM instance, in case it is running on the same Oracle Server system. Since the atomic snapshot feature is a prerequisite for backing up Oracle ASM-managed data with Data Protector, but the HPE 3PAR StoreServ Storage disk arrays do not support it, sessions based on such a backup specification fail.

Workaround: Before you start configuring a Data Protector ZDB backup specification for Oracle Server backup, shut down the running Oracle ASM instance and start it up again after you have saved the backup specification. Data Protector will successfully back up correct data that you have selected for backup.

VMware vSphere

- When restoring a virtual machine with a virtual Raw Device Mapping (vRDM) disk to a directory using the option **Restore to Directory**, the restore succeeds and you can import or upload the files to a vCenter. However, starting such a virtual machine fails.

Workaround: None. Do not use the Restore to Directory option when restoring a virtual machine with a vRDM disk.

- Storage policy for VVol datastores not considered during the restore.

Workaround: You can edit the storage policy of the virtual machine after a restore. For more details, see the following URL:

<http://pubs.vmware.com/vsphere-60/index.jsp#com.vmware.vsphere.storage.doc/GUID-D6A099C5-8F80-474C-A79A-64F5EC4455DA.html>

- On Windows systems, when backing up VMs from a vCenter installed on a French-language version of the operating system, the backup fails with the following error:

```
[Major] From: VEPALIB_VMWARE@hostname "/Datacenter"  
Time: <Date> <Time>  
Virtual Machine 'vm_2': Could not backup disk scsi0:0 ...
```

```
[Major] From: VEPALIB_VMWARE@hostname "/Datacenter"  
Time: <Date> <Time>  
Virtual Machine 'vm_2': No disk backed up ...
```

If you back up the same VM from a vCenter installed on an English-language version of the operating system, the backup finishes successfully. After that, a backup from a French-language operating system host is also successful. The issue is caused by UTF-16 characters in the VDDK temporary directory path, which is by default set to `%TEMP%\vmware-%USERNAME%`.

Workaround: Set a different temporary directory in the file `Data_Protector_Program_data\Config\client\vepa_vddk.config`, for example: `tmpDirectory=c:\tmp`.

- Creating a virtual machine snapshot, either when the virtual machine is being migrated using VMotion migration or when such migration is not taking place, may fail with one of the following errors:

```
[Major] From: BSM@hostname "Barlist-BackupRestore-1054"
```

```
Time: <Date> <Time>
[61:3003] Lost connection to OB2BAR Backup DA named "ERROR" on host hostname
  Ipc subsystem reports: "IPC Read Error System error: [10054] Connection reset by
  peer"

[Normal] From: VEPALIB_VMWARE@hostname ""
Time: <Date> <Time>
Creating Virtual Machine 'jeos-e10x-001' ...
Datacenter: /ESX5.0_Name
Host/Cluster: cluster.company.com
Datastore: VMData_env39 name
2013-04-12T17:30:35.651+02:00 [04048 trivia 'ThreadPool'] PrepareToWait: Starting
new thread
2013-04g45280-12T17:30:35.806+02:00 [02928 trivia 'ThreadPool'] PrepareToWait:
Starting new thread
2013-04-12T17:30:35.899+02:00 [02928 trivia 'ThreadPool'] PrepareToWait: Starting
new thread
2013-04-12T17:30:35.902+02:00 [03184 trivia 'ThreadPool'] PrepareToWait: Starting
new thread
2013-04-12T17:30:35.902+02:00 [02852 trivia 'ThreadPool'] PrepareToWait: Starting
new thread
2013-04-12T17:30:35.902+02:00 [03184 trivia 'ThreadPool'] PrepareToWait: Starting
new thread
[Major] From: RSM@hostname ""
Time: <Date> <Time>
[61:3003] Lost connection to OB2BAR restore DA named "" on host hostname.
  Ipc subsystem reports: "IPC Read Error System error: [10054] Connection reset by
  peer"
Workaround: Restart the backup session.
```

Lotus Notes

- In case the Data Protector application specific pre-exec script is specified for the Lotus integration backup on Solaris, the backup may fail with the following error:
Data Protector Lotus Agent cannot back up object *objectname*.
Workaround:
Disable the Data Protector application specific pre-exec script and restart the backup.
- During a point-in-time restore of the Lotus Notes database, the following error may occur:
Lotus Notes C API 'NSFRecoverDatabases' returned error 5099: Recovery Manager:
Backup was later than recovery point in time.
The problem occurs if the specified point in time is earlier than the time when the source backup was performed.

Workaround:

Specify a different point in time or a different source backup, so that the point in time is later than the time when the source backup was performed.

Disk array integrations

- The configuration requirements for ZDB of Oracle or SAP R/3 databases have changed in the following cases:
 - if Oracle is used as a part of Oracle ZDB integration and you intend to perform instant recovery sessions,
 - if Oracle is used as a part of SAP R/3 ZDB integration and you intend to perform instant recovery sessions.

In these cases, the Oracle database needs to be reconfigured. For more information on configuration requirements, see description of the `ZDB_ORA_INCLUDE_CF_OLF omnirc` option in the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

- Database recovery after instant recovery for the Microsoft Exchange Server and Microsoft SQL Server integrations cannot be performed from the Command Line Interface (CLI).

Workaround: Perform the recovery using the GUI.

- The Data Protector session in the NetApp storage environment fails with an error similar to the following:

```
[Major] From: SMISA@hostname "SMISA" Time: <Date> <Time>  
The presentations have not been created for this storage volume:  
Storage volume:
```

Workaround: Delete the LUNs left after the backup session and restart the Data Protector ZDB session for the NetApp storage. The name of such LUN has the following syntax: `LUN_name.ID` (ID is generated by the NetApp storage), for example, `LUN_2_4.DP-2014.11.14-10-05465F54A`.

Granular Recovery Extension issues

VMware vSphere

- When performing a partial restore using the VMware Granular Recovery Extension for cached recovery, Data Protector reads all disks from tape device instead of only the selected ones, which can significantly slow down the restore process.

Workaround: Backup to either the Smart Cache device or 3PAR replica to take advantage of cached recovery.

- The granular recovery or browsing of swap partitions in Linux systems may result in the following error message:

```
Exception: The virtual disk does not have any partitions that the host system  
knows how to mount.
```

Workaround: Navigate to other partitions or volumes.

Microsoft SharePoint Server

- Microsoft SharePoint Server 2013: When you select Site Settings > HPE Data Protector Granular Recovery Extension > Granular Recovery, the Granular Recovery link does not work if Minimal Download Strategy (MDS) is enabled. If MDS is disabled, the link properly opens the Granular Recovery Extension page.

Workarounds:

- Open the link in a new tab or window.
- Alternatively, disable Minimal Download Strategy.

Disaster recovery issues

- An encrypted backup of the Internal Database (a prerequisite for Cell Manager disaster recovery) fails unless an active encryption key was created prior to the backup.

Workaround: Create an active encryption key prior to invoking an encrypted backup of the Internal Database. For details, see the *omnikeytool* man page or the *HPE Data Protector Command Line Interface Reference*.

- On Windows Server 2003 systems, when performing EADR on a ProLiant BL460c system, the DR OS cannot find the network card and a restore cannot be started.

Workaround:

Enable the safe boot mode:

- Edit the `drm.cfg` file *before* creating the ISO image:
 - i. Open the file `drm.cfg.tpl` in `\\OmniBack\bin\drim\config`
 - ii. Edit the variable `safe_boot`:

```
safe_boot = normal
```
 - iii. Save the file `drm.cfg.tpl` and rename it to `drm.cfg`.
 - iv. Create the ISO image.

The disaster recovery process should now start normally.

- Or, if you are already performing disaster recovery, edit the `boot.ini` file and restart the system.
 - i. Once the DR OS boots and the Disaster Recovery Wizard starts, abort the countdown.
 - ii. Start a command prompt and launch Notepad.
 - iii. Open the file `C:\boot.ini` and search for the string `/SAFEBOOT:NETWORK`.
 - iv. Remove the string from the `boot.ini` file and save it.
 - v. Restart the computer and leave boot sequence to start from the disk (do not boot from CD-ROM again).
 - vi. When the system logs on, proceed with the standard disaster recovery procedure.
- If the `omnidr` command is started with invalid parameters, a message to press the F8 key in the next 10 seconds is displayed instead of the command synopsis. After pressing any key, the command properly displays the command synopsis.

- When hidden dynamic volume is mounted to an empty NTFS folder and backed up, the disaster recovery process creates the volume but skips mounting it to the folder and restoring the actual data.

Workaround:

Once the disaster recovery completes, you can restore the data using the following procedure:

- Mount the volume to the same NTFS folder.
- From the Restore context of the Data Protector GUI, execute data restore.

Cluster related issues

Common issues

- When backup system is in a cluster environment and the backup session is performed using the name of the cluster node, instant recovery fails if you try to perform recovery using the other cluster node.

Workaround: To avoid this problem, use the name of the virtual host for configuration of the backup specification.

- If a backup session stops responding during a cluster failover, and all Backup Agents fail, a timeout will be reported but the session itself will not abort. The default session timeout occurs after 7200 seconds (two hours). As long as the session is not responding, another session using the same backup specification cannot be started.

Workaround: Manually abort the backup session and restart the session.

- If a cluster failover occurs during a Data Protector backup session in which an application database that resides on the cluster is being backed up with the appropriate application integration agent, particular problem may occur after the failover which prevents the session from succeeding.

Under such circumstances, in Monitoring context of the Data Protector GUI, two backup sessions are displayed: the backup session that was restarted after the failover, and another, unknown session. Output of the unknown session contains messages similar to the following:

```
[Critical] From: BSM@ClusterNode01Name "BackupSpecificationName" Time: <Date> <Time>
```

```
[12:1243] Device not found.
```

```
[Critical] From: OB2BAR_VSSBAR@ClusterNode02Name "MSVSSW" Time: <Date> <Time>  
Failed VSSBAR agent.
```

```
[Major] From: OB2BAR_VSSBAR@ClusterNode02Name "MSVSSW" Time: <Date> <Time>  
Aborting connection to BSM. Abort code -1.
```

```
[Critical] From: BSM@ClusterNode01Name "BackupSpecificationName" Time: <Date> <Time>
```

```
None of the Disk Agents completed successfully.
```

```
Session has failed.
```

The root cause of the problem is unsuccessful identification of the restarted backup session after the failover. The involved integration agent is not notified about the backup session restart. Depending on the particular situation, the integration agent either starts a new backup session or connects to

the restarted backup session manager (BSM) process. In both cases, such behavior of the integration agent is wrong.

Workaround: None.

Issues in HPE Serviceguard

- After failover on the secondary application system (application runs on the HPE Serviceguard cluster) instant recovery may fail with the following error message, if the **Check data configuration consistency** option is selected:

```
[Critical] From: SSEA@hostname"" Time: <Date> <Time>
```

```
Data consistency check failed!
```

```
Configuration of volume group /dev/vg_sap has changed since the last backup session!
```

Two workarounds are possible:

- Make sure that the vg configuration on the system is not changed, deselect the **Check data configuration consistency** option, and restart the instant recovery.
- When setting up the cluster, use the `ioinit` command to ensure that all disk device files are identical.
- If you export a physical node from an HPE Serviceguard cluster, you cannot import it back as the `cell_server` file is deleted. This file is shared among all nodes of a cluster, so you need to recreate it.

Workaround: execute the command `/opt/omni/sbin/install/omniforsg.ksh -primary -upgrade`.

- Enabling the Encrypted Control Communication (ECC) on HP-UX cluster-aware cell manager installed on HPE Serviceguard cluster fails.

Workaround: Perform either of the following options:

- On all the clients, where the disabling fails, manually remove the `config/client/config` file.
- Enable ECC on the cell manager. Disable ECC on the clients either by selecting in the GUI or listing in `omnicc` file.
Disable the ECC on the cell manager and remove the `config/client/config` files on cell manager HPE Serviceguard nodes.
- Enable ECC back on the cell manager. Move the Data Protector cell manager package to another node. Disable ECC in the whole cell.

Issues in Microsoft Cluster Server

- When restoring the Cluster Database of Microsoft Cluster Server, you should stop the cluster service on all inactive nodes before starting the restore. If cluster service is active on any other node at the time of the restore, the restore API will fail and eventually cause a failover.
- When the Cell Manager is installed on Microsoft Cluster Server and you start a restore of the Cluster Database, the restore session will stop responding. This is because the cluster service is stopped by the restore API causing the Restore Session Manager to lose the connections to the IDB and the

MMD.

Workaround: Wait for the VRDA to complete and then abort the session. You then need to restart the GUI (or reconnect to the Cell Manager). Additionally, when starting a Cluster Database restore, make sure that this is the only item you are restoring and that no other sessions are running.

Reporting related issues

- On Linux systems, when sending a report using the e-mail send method, the e-mail does not have a subject and contains `root` in the **From** field. The correct **From** and **Subject** entries are inside the e-mail body.

Workaround: Use `sendmail` to send reports using the e-mail send method. For example, to use `sendmail` instead of `/usr/bin/mail`, create the following link:

```
ln -s /usr/sbin/sendmail /usr/bin/mail
```

Note that on some Linux distributions `/usr/bin/mail` already exists. HPE dissuades you from removing this existing path since some applications may rely on it.

Other known issues

- Data Protector is unable to detect duplicate IP addresses within a cell or MoM environment and cannot report them as such. If a duplicate IP address is assigned to a system with the Data Protector software installed, Data Protector reports such network misconfiguration with the message `Cell Manager host IP is not in IP range` in the GUI and CLI as follows:
 - In the About Data Protector Manager or About Data Protector MOM dialog box in the GUI, in the Remark column of the Passwords info property page
 - In the Remark lines of the Data Protector `omnicc -password_info` command output

Under these circumstances, Data Protector licensing is no longer covered. Until you resolve the problem, you can expect session and operation failures in the affected Data Protector cells.

Workaround: Reconfigure your network by assigning a unique IP address to the problematic system.

- If the clock on the client is not synchronized with the clock on the Cell Manager, the certificate may become invalid, thus resulting in failed authentication. For example, when the clock on the Cell Manager is ahead of the clock on the client, the certificate created during installation is not yet valid for the client attempting to connect to it.

Workaround: None. Make sure the clocks between the Cell Manager and the client are synchronized.

- The IDB restore session report displays the following message:

```
The OS reported error while accessing <new restore directory>:
```

```
[2] No such file or directory
```

Workaround: No action is required. In cluster Cell Manager Unix environment, `<DP_CONFIG_DIR>/server` is a link to the shared configuration folder between the cluster nodes. During restore, the Restore Session Manager tries to create `<DP_CONFIG_DIR>/server` as a link to the shared configuration folder between two nodes, which is not required.

- If you consolidate object versions that have already been consolidated, selecting the session in the

Restore context results in a message that the session contains no valid restore objects. This is because the session is treated as a copy and consequently cannot be selected for restore.

Workaround: Either select the session in which the objects were originally consolidated, or select the objects under **Restore Objects**.

- To prevent object consolidation sessions from using too much system resources, the number of object versions that can be consolidated in one session is limited to 500 by default. If more object versions match the selection criteria, the session is aborted.

Workaround: Either tighten the selection criteria, for example, by limiting the time frame, the number of backup specifications, and so on, or increase the value of the global option `ConsolidationAutomatedMaxObjects`.

- If you perform interactive object consolidation of objects that span more than one medium and the number of consolidation devices used is smaller than the number of objects being consolidated, the object consolidation session may become unresponsive.

Workaround: Either increase the number of consolidation devices, or select the object versions for consolidation in the order in which their full backups were performed.

- If full backups for multiple objects reside multiplexed on a device which is different than the file library hosting the corresponding incremental backups for these objects (for example, on a tape library), it may happen that some of the file writers (file library drives) needed as targets for the consolidation session get aborted because of a failure on the source Media Agent side (for example, in case of a media error, an incorrect block size, a canceled mount request, and similar). This may result in a hanging object consolidation session, in case there are not enough file writers remaining to complete the consolidation for other objects. Once all remaining objects are consolidated, all file writers will be freed up again at the end of the session.

Workaround: Ensure that the number of file library drives used as consolidation devices is equal or higher than the number of objects being consolidated. If the number of configured file library drives is smaller than the number of objects to be consolidated, HPE suggests to split the consolidation of multiple objects into more than one session.

- When replicating several backup sessions which backed up the same objects at the same time, for example, folders that reside on the same mountpoint on the same host, the session might fail with the following error:

```
[Major] From: CSM@hostname "new_0" <Date> <Time>  
[65:99] Import failed with possible cause:  
this media already has valid copy in DB.
```

Workaround: Group such objects into one session or explicitly use different labels for each object.

- When trying to restore to another organization, the original Vapp does not get deleted.

Workarounds:

You must manually delete the Vapp before performing a restore.

(Or)

As soon as the restore completes in the Data Protector GUI, the Administrator must manually delete the original Vapp and start the restored Vapp. If required, you must configure the network as well.

- When restoring to a new location from a medium that contains a consolidated object on HPUX 11.31 Itanium, the restore session fails with errors. For example:

```
[Minor] From: VRDA@hostname CONS02 Time: <Date> <Time> /tmp
```

Workaround: If the buffer size is less than 8, set it to 8 (the default value) and restart the restore session. For instructions, see the *Data Protector Help*, index: "setting, advanced options for devices and media".

- When performing a backup from HP-UX or AIX or Solaris system (Disk Agent) with deduplication functionality, the backup session fails with the following error, if the two omnirc variables (OB2ALIGN and/or OB2SAPALIGN) are set:

```
[Major] From: BMA@computer.com "D2D_NS_Baze_gw5 [GW 4560:4:1766155810]" Time: 7/15/2015 2:39:25 PM
```

```
[90:51] \\COFC-D2DNS\D2D_NS_Baze\df2ebb9f_55a6547a_2014_0233
```

```
Cannot write to device (JSONizer error: Invalid inputs)
```

Workaround: Ensure that the two omnirc variables (OB2ALIGN and/or OB2SAPALIGN) are not set. For more information, see the omnirc reference page in the HPE Data Protector Command Line Interface Reference Guide or the omnirc man page.

- If you have different logical devices for the same physical device and you use a different logical device for backup every day, the lock name concept prevents collisions between different logical devices assigned to the same physical device.

When trying to perform a restore, where several logical devices but only one physical device was used for different backups (Full, Incr1, Incr2, Incr3, ...), Data Protector does not check the lock name, and therefore does not recognize that the same physical device was used for all backups. An error message that the restore session is waiting for the next device to get free is displayed.

Workaround: Remap all logical devices to the same physical device by following the steps below:

- a. In the Context List, click **Restore**.
 - b. In the Scoping Pane, expand the appropriate data type and desired client system and object for restore.
 - c. When the Restore Properties window opens, select the files that you would like to restore.
 - d. In the Devices tab, select the original device and click **Change**.
 - e. When the Select New Device window opens, select the physical device name and click **OK**.
- Data Protector instant recovery fails when the filesystem is busy.

Workaround: List processes which occupy filesystem by using the `fuser` command. For example, if the filesystem `/oracle/P01` is busy, execute the command `fuser -kc /oracle/P01`.

- If a backup is performed on one node and then instant recovery attempted on another node with the **Check data configuration consistency** option selected, the following error message is displayed:

```
Volume group configuration has changed.
```

The message is displayed because the `vgdisplay` command detects that the LUN configuration on one client is different than that of the other client.

Workaround: If the `ext_bus` instance is the same, this message is not displayed. Alternatively, it is not displayed if the **Check data configuration consistency** option is not enabled.

- A backup may fail if the snapshot backup specification contains an invalid `rdsk` object in the first place.

Workaround: Change the order of the `rdsk` objects so that a valid `rdsk` is in the first place.

- Data Protector services may not be running after EADR or OBDR.

Workaround: In the **Control Panel > Administrative Tools > Services**, change the startup type for Data Protector services from Manual to Automatic. Start the services after you have changed the startup type.

- On HP OpenVMS systems, a restore session may become unresponsive and report errors due to an unusual delay while unloading a tape drive.

Workaround: Set the Cell Manager global parameter `SmPeerID` to 10 and restart all Data Protector services on your Cell Manager.

- When using SNMP traps on a Windows Cell Manager, Data Protector uses the default community name `public`. This applies to both the SNMP send method with Data Protector notifications or reporting and the SNMP traps for system and application management applications.

Workaround: In the registry key `HKEY_LOCAL_MACHINE\SOFTWARE\Hewlett-Packard\OpenView\OmniBackII\SNMPTrap` create a value named `Community` and set it to the community name you want to use. Note that all SNMP traps will be sent with the same community name and to the destinations associated with it in the Control Panel.

- Data Protector performance on Red Hat Enterprise Linux (RHEL) is negatively affected if the Name Server Caching (`nscd`) daemon is disabled.

Workaround: Enable Name Server Caching on RHEL, or switch to a local DNS, and then execute the `omnisv -start` command.

- The B2D copy session may allocate more source gateways than the maximum number of destination gateways.

Workaround: If you want to limit the number of source gateways, use the global variable `LimitInitGatewayExpension=0` (disabled). If you enable this global variable, the source gateways are expanded up to "x" number of devices, where "x" is maximum load balancing value specified in the target device settings.

- The command `omnistat -session [session ID] -detail` may incorrectly display a message `Restore started` or `Backup started`. This may result in both parameters appearing to be identical.
- When you collect the debug files matching a specific debug ID, using the Data Protector GUI or CLI, the relevant debug logs may not be collected.

Workaround: When you are collecting and saving the debug files, make sure you also specify all known source debug directory paths.

- If you use the Data Protector CLI, execute:

```
omnidlc -did debugID -debug_loc Dir1
```
- If you use the Data Protector GUI:
 - i. In the Context List, click **Clients**.
 - ii. In the Scoping Pane, expand **Data Protector Cell** and then **Clients**.
 - iii. Right-click the client and click **Collect debug files**.
 - iv. In the Debug File Collector – Directories page, enter the paths for all known non-default debug directories and click **Add**.
 - v. Click **Next** as many times as needed to reach the last page of the wizard.
 - vi. Click **Finish** to exit the wizard.
- When performing a point-in-time restore from a client system which is in a different time zone than the Cell Manager, the restore fails. If started from the Cell Manager, the restore succeeds. The issue

arises because the Cell Manager does not store the time zone of the client system during the backup.

Workaround: If you are performing a restore from a client system, account for the time zone difference when selecting the point in time to which to restore. For example, if the Cell Manager is in the time zone UTC+1 and the client is in UTC+5, and you performed the backup at 5:00 as seen on the client system, enter the time as it was on the Cell Manager at the time of the backup, that is 1:00.

The object version properties in the GUI and CLI only display information about the Data Protector software encryption. In case that only drive-based encryption was used, the object version is displayed as not encrypted.

Workaround: Drive-based encryption is not related to object version. To check for the object's drive-based encryption, use the `omnimm -media_info Medium -encryptioninfo` command.

When you run a restore to another location (using the `Restore into` option) of the multiple filesystem objects that are located in different subdirectories of the same top level directory, the following error may occur:

```
[Minor] From: VRDA@hostname "OBJECTNo1" Time: <Date> <Time> /tmp/RestoreDir Cannot create: ([17] File exists).
```

The restore finishes successfully.

Workaround: None. You can safely ignore the message.

Backup session hangs if the original and mirror device have their Max concurrency set to 32.

Workaround: Lower the concurrency on both the original and mirror device to 30.

When starting the Data Protector GUI, it will attempt to connect to the Cell server. The Cell server is identified by either its short `hostname` or `hostname.domain`. Just before authenticating, the certificate of the Certificate Authority (CA) is transferred from the Cell server and stored in the filesystem using the FQDN of the Cell in the path. When connecting to the Cell server for the second time, the already stored CA certificate is compared with the newly received one from the Cell server. When identifying the Cell server by the hostname only, the CA certificate is not found in the filesystem and an accept certificate dialog pops-up.

Workaround: The folder that stores the `cacert.pem` file on the client is located at the following location:
`\users\\AppData\Local\Hewlett-Packard\Data Protector\ca\\cacert.pem`. If you copy this folder to the following location then the issue does not occur again:
`\users\\AppData\Local\Hewlett-Packard\Data Protector\ca\\cacert.pem`.

A resumed session contains more than one object version per backed up object. But the `omnir` command expects each session to have one object version per object and does not restore the session completely with default parameters.

Workaround: Restore the resumed session by following the steps below:

1. Obtain the copy IDs of the session using the `omnidb -session <sessionId>` command.
2. Run `omnir` multiple times and add the `-copyid "UUID/seqnr"` one after the other in the `omnir` command line.

Note: The restore sequence has to be performed in a chronological order for a successful and complete restore.

In the restore context of Data Protector GUI, when you restore by session, a resumed session appears in the left pane repeatedly for the same object, given the number of object versions per object in the resumed session.

Workaround: For restore, select only the object coinciding with the last successful object version and one of the identical objects at the time.

When connecting the Data Protector Cell Manager on HP-UX or Linux using the Windows Data Protector Client GUI, the following error is displayed:

A server error has occurred. Reported error message: Unknown SSL protocol error in connection to <cell manager>: 7116.

The error occurs when a faulty Time-To-Live value of the SSL certificate is generated on the Cell Manager.

Workaround:

Manually generate a new SSL certificate with the `omnigencert.pl` script from the command line. To do so, perform the following:

1. Get the store password by performing the following:
 - Navigate to the `webservice.properties` file from the following location:
`/etc/opt/omni/client/components`
 - Search for the `keystorePassword` value, for instance, consider that the `keystorePassword = <store_password>`.

2. Generate a new SSL certificate by executing the following command:

```
/opt/omni/bin/perl /opt/omni/sbin/omnigencert.pl -server_id  
<CellManagerHostname.domain.com> -user_id hdpd -store_password <store_password>
```

3. Restart the Data Protector services using the following commands:

```
omnisv stop  
omnisv start
```

(Or)

Reload the jboss configuration using the following command:

```
/opt/omni/AppServer/bin/jboss-cli.sh -c --command="/:reload
```

Password encryption is done by AES keys with the entity name “Data Protector Passwords”. Do not delete or deactivate this key, as Data Protector will not be able to decrypt encrypted strings.

Deactivating the password key affects backup only; it does not affect restore, nor does it disable the key or act as a key revocation. For Cloud devices, deactivating or deleting the key renders the Cloud device unusable.

For details, see the `omnikeytool` man page or the *HPE Data Protector Command Line Interface Reference*.

Workaround: Do not delete or deactivate this password key.

The Cloud (Helion) and Cloud (Azure) devices encounter errors in while communicating with the Cloud object store. The Cloud (Helion) and Cloud (Azure) devices retries the operations if errors are encountered.

When communication errors occur, the following error displays:

Error in communication with cloud [ERROR], retrying

Workaround: The default retry count for the Cloud is 5. Set the omnirc option OB2_CLOUDDEV_MAXRETRIES on the Media Agent host to higher than 5.

The Cloud Azure device encounters problems during configuration when attempting to configure with Data Protector 9.08 or earlier.

Workaround: Cloud (Azure) devices are not supported by earlier versions of Data Protector.

Ensure that all Cell Managers, GUI servers, Installation Servers, and Media Agents are updated to the General Release Patch or later.

While performing Point-in-time recovery of PDB Oracle, displays the following error:

```
PLS-00306: wrong number or types of arguments in call to 'GETCNCTSTR'
```

Issue is present in Oracle Bundle patch when performing point-in-time recovery of PDB.

Workaround: To resolve this problem look for the newer Oracle Patch bundle or contact HPE/Oracle support.

In case where PDB has been dropped from CDB following error displays:

```
ORA-65011: Pluggable database does not exist.
```

Workaround: This bug is Oracle specific. Login to [Oracle Support](#) page and search for the Bug ID **18967466**.

```
Bug 18967466 : ALTER DATABASE BEGIN BACKUP" COMMAND FAILS DUE TO ORA-65011 IF PDB HAS BEEN DROPED
```

Known non-Data Protector issues and workarounds

Non-Data Protector issues related to installation or upgrade

- On Windows systems, the operating system might incorrectly report free disk space for an NTFS volume that is mounted to a directory on an NTFS filesystem: instead of the NTFS volume free space the amount of free space on the NTFS filesystem is reported. In such cases, the Data Protector Setup Wizard will not start the installation to the mounted NTFS volume if the amount of free space on the NTFS filesystem is smaller than the minimum disk space installation requirement.

Workaround: free disk space on the NTFS filesystem by removing unnecessary files until the installation requirement is met.

- On Windows systems which do not have the G5 root certificate issued by the VeriSign Class 3 Public Primary Certification Authority installed, verification of the binaries signed with the new Authenticode signing certificate may fail.

For more information on the root cause of this problem, see the related Verisign Knowledge Center article published on the web at

https://knowledge.verisign.com/support/ssl-certificates-support/index?page=content&id=AR1747&actp=search&viewlocale=en_US.

- On Windows systems, if you start local installation from a mapped drive through Remote Desktop Client, the installation may fail with the following error message:

Error 2755. Server returned unexpected error 3 attempting to install package *MappedDrive:\i386\DataProtector.msi*.

The Windows Installer service is running under a different user account than the user account under which the mappings were created, and therefore has different drive mappings. As a result, the installation fails.

Workarounds:

- Do not start the installation from a mapped drive. Use the UNC path specification instead (for example `\\hostname\shared_folder`).
 - For installation, use VNC instead of Remote Desktop Client.
 - Start the installation on the console.
- On Windows Server 2003 systems, the installation fails if the installation destination directory is a virtual drive, created for example with the `subst` command. The following error message is displayed:

Error: 1320. The specified Path is too long.

The Windows Installer service is running under a different user account than the `subst` command. As a result, the installation fails.

Workarounds:

- Use the UNC path specification (for example, `\\hostname\shared_folder`) instead of the virtual drive. This is the preferred solution.
 - Execute the `subst` command under the Local System user account.
- On Linux systems, the `rpm` utility does not correctly remove Data Protector components if you specify several installation packages in the same command line. For example, if you use `rpm -qa | grep OB2 | xargs rpm -e`, the `rpm` utility does not resolve dependencies in the correct order.

Workaround: Remove the Data Protector components one by one.

Non-Data Protector issues related to user interface

- When using CLI on UNIX systems, the characters may be displayed incorrectly.
Different encoding systems (Latin, EUC, SJIS, Unicode) cannot be used in the desktop environment and in the terminal emulator. For example, you start the desktop environment in EUC-JP, open a terminal emulator and change the locale to SJIS. Due to an operating system limitation, if you use any CLI command, the characters can be displayed incorrectly.
Workaround: To eliminate this problem, start the desktop in your desired locale.

Non-Data Protector issues related to Disk Agent

- If the LSI Logic 53C1010-66 card is used on an HPE Server rx2600 Itanium 2 client with Windows Server 2003 Enterprise Edition, restore may fail with an internal error.
- On Windows systems, due to filesystem limitations, files that were backed up on UNIX systems and whose names contain the backslash ("\") character may be restored to a wrong location and with the wrong file name. Windows operating system interprets the backslash in a file name as a directory separator. For example, if a file named `back\slash file` was backed up on a UNIX system and restored to a Windows system, it will be restored into the `back` directory with the file name `slash`.
- On Windows systems, when backing up several objects at once and with the option **Use shadow copy** selected, backup of some objects may fail with the following error message:

```
[Major] From: VBDA@ hostname"ObjectName" Time: SystemTime
```

```
It was not possible to create volume snapshot for 'Mountpoint'. System error:  
'VSS Snapshot creation failed'.
```

The issue can appear if there is not enough disk space available for storing multiple shadow copies or due to high disk I/O load. For details, see

<http://www.microsoft.com/technet/support/ee/transform.aspx?ProdName=Windows%20Operating%20System&ProdVer=5.2.3790.1830&EvtID=12298&EvtSrc=VSS&LCID=1033>.

Workarounds:

- Run the backup session again.
- Provision enough dedicated storage space on a separate volume, reserved only for shadow copy storage.

You can allocate additional shadow storage space using the command `vssadmin add shadowstorage`.

For details, see [http://technet.microsoft.com/en-us/library/cc788051\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/cc788051(v=ws.10).aspx).

Recommendation:

Allocate separate storage space to store shadow copies. This storage space should not be on the volumes that are backed up.

- On Solaris 9 systems, filesystem backup may fail with error messages similar to the following:
Cannot open attribute directory /BC/fs/VxVM/UFS/Test6.doc: read-only filesystem!
Extended attributes not backed up.
Workaround: Set the `omnirc` option `OB2SOL9EXTATTR` to `0` to disable the backup of extended attributes.
- On HP-UX 11.31 systems, the NFS mount points become inactive when performing an NFS backup.
Workaround: Upgrade the Open Networking Computing component (ONCplus) to B.11.31.10 or a newer version.

Non-Data Protector issues related to Media Agent

- Erase operation on magneto-optical drive connected to an HP-UX system fails with the following error:

```
[Major] From: MMA@hostname "MO-lada" Time: <Date> <Time> [90:90] /dev/rdisk/c2t0d1  
Cannot erase disk surface ([22] Invalid argument) => aborting
```

- Breece Hill's Saguaro libraries use the stack mode for entering and ejecting cartridges. One mail slot has two SCSI addresses, one for the enter operation and the other for the eject operation. For Data Protector to function properly in this mode, the following `omnirc` command options must be configured as follows:

- `OB2LIB_STACKEXP` must contain the SCSI address of the export slot
- `OB2LIB_STACKIMP` must contain the SCSI address of the import slot

- Data Protector Media Agent cannot coexist with CA ArcServe installed on the same Windows client system. Such setup may lead to a data loss.
- When a DLT8000 library is used, media cannot be imported and the `omnimlist` command does not function properly. In this case, the following errors are reported:

```
[Major] From: MMA@hostname "Datacenter" Time: <Date> <Time>
```

```
[90:182] Cannot forward segment. ([5] I/O error)
```

```
[Major] From: MMA@hostname "Datacenter" Time: <Date> <Time>
```

```
[90:53] /dev/rmt/1m Cannot seek to requested position ([5] I/O error)
```

Quantum has confirmed a problem with the controller firmware. There is a cumulative slip occurring in the tach relative to the tape. When such a slip occurs and the drive detects the BOT marker, the drive reconstructs its internal directory. The problem occurs only when tape media containing large amounts of data are used.

Workaround: Consult your HPE support representative before you proceed. You need to upgrade the DLT8000 drive firmware to version V51. More details about the firmware changes can be found in Service Note A5597A-27.

- If another medium is loaded to a tape drive during a backup session that writes backup data to this drive's tape library, the Data Protector Media Agent running on an IBM AIX system may not appropriately handle the loaded medium. In this case, the backup sessions fails.

The root cause of the problem is an AIX operating system limitation in shared memory allocation functionality. The problem occurs more frequently when a relatively high Disk Agent concurrency is used.

Workaround: Enable the IBM AIX extended shared memory model by setting the Data Protector `omnirc` option `EXTSHM` to the value `ON`.

- If an LTO 4 device is connected to a SmartArray 6i controller, drive based encryption may fail due to an issue with the SmartArray 6i firmware.

Workaround: Check if a newer version of the firmware resolves the issue or use a different SCSI controller.

- When configuring a StoreOnce Backup system device, if the Client ID field contains any of the characters `"_"`, `"-"`, `"."`, or `"+"`, the following error is displayed:

Could not get information for host *Hostname* using gateway *GatewayName*...
Workaround: None. Do not use the above mentioned characters for client IDs.

Non-Data Protector issues related to integrations

Microsoft Exchange Server

- Due to MAPI behavior, if the subject line of a backed up message begins with a sequence of up to 4 non-space characters followed by a space, and any of these non-space characters is a colon (":"), the message, once restored, will have a wrong subject line. For example, a message with the original subject line ABC: ha1a will get the subject line ABC: ABC: ha1a. after the restore.

This does not apply to standard prefixes for e-mail subjects, such as Re:, Fwd:, and so on, if they are generated automatically by your e-mail client (for example, by pressing the **Reply** button in Microsoft Outlook).

Workaround: None.

Microsoft SQL Server

- Instant recovery of Microsoft SQL Server databases fails.

Workaround: Follow the instant recovery procedure in the *HPE Data Protector Zero Downtime Backup Integration Guide*. You need to restart the services of the SQL Server instance after the instant recovery completes. If this action does not automatically start a recovery of all system databases, perform the following:

- a. Start the SQL Server instance in the single-user mode.
- b. Manually run a recovery of the master database.
- c. Run a recovery of every other system database. SQL Server instance must still be running in the single-user mode.
- d. Restart the services of the SQL Server instance.

Microsoft Volume Shadow Copy Service

- The following MSDE Writer and Microsoft SQL 2005 writer components cannot be restored while the Microsoft SQL Server is online: master, model, and msdb.
- A snapshot backup of an Exchange Server 2003 database fails, and event ID 9607 is logged.

Workaround: For information on how to resolve this problem, see the Microsoft webpage <http://support.microsoft.com/kb/910250>.

- With HPE P6000 EVA Disk Array Family, a backup session may fail if there are more than 4 source volumes (original disks) in a snapshot set.

Workaround: None. Make sure that the number of source volumes in a backup specification does not exceed 4 and that the next snapshot creation starts no earlier than 30 minutes after the last snapshot was deleted.

Also ensure that you upgrade firmware and HPE Command View (CV) EVA to the latest version.

- With HPE P6000 EVA Disk Array Family, a backup session that uses software provider fails,

reporting that shadow copies could not be created.

Workaround: Install the latest HBAs firmware and start a new backup session.

- With HPE P9000 XP Disk Array Family with hardware providers configured, the client system fails abnormally every second or third backup. This may be caused by particular versions of HPE MPIO DSM for HPE P9000 XP Disk Array Family.

Workaround: Ensure that you are using a supported version of HPE MPIO.

- The full path of any virtual disk in the HPE Command View (CV) EVA Virtual Disks hierarchy should not exceed 650 characters in length.

Workaround: None. A future release of the hardware provider may remove this limitation.

- With the VSS P9000 XP Array hardware provider on Windows Server 2008 systems, warning messages are logged to the Application Event Log during each shadow copy import. The issue does not appear on Windows Server 2008 R2 systems.

Workaround: None. A future release of the VSS P9000 XP Array hardware provider may remove this issue.

- When using the P6000 EVA Array hardware provider on Windows Server 2008 systems, during a transportable backup when Data Protector tries to break the shadow copy set, the following errors are reported:

```
[Minor] From: OB2BAR_VSSBAR@hostname "MSVSSW" Time: <Date> <Time> Failed to break Shadow Copy Set of session '2011/01/11-4:tpc211'.
```

```
[Warning] From: OB2BAR_VSSBAR@hostname "MSVSSW" Time: <Date> <Time> [145:714] Rescanning system due to Break Shadow Copy Set failure.
```

```
[Minor] From: OB2BAR_VSSBAR@hostname "MSVSSW" Time: <Date> <Time> Failed to disable backup '2011/01/11-4:tpc211'
```

The issue appears if the P4000 SAN Solutions hardware provider is also installed on the same system. The issue does not appear on Windows Server 2008 R2 systems.

Workaround: Remove the P4000 SAN Solutions hardware provider or use a different client as the backup system. A future release of the P4000 SAN Solutions hardware provider may remove this issue.

- On Windows Server 2008 R2 systems with the VDS hardware provider installed, when performing an instant recovery with the Switch of disks method and using a P6000 EVA Array with a large number of LUNs, the operation may fail.

Workaround: Use the Copy of replica data method instead of the Switch of disks method.

Recommendation: To avoid the problem, remove the VDS hardware provider. Note that there are use cases which require the VDS hardware provider to be installed. For details, see the *HPE Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*.

- If two command devices are configured for a disk array of the HPE P9000 XP Disk Array Family which supports authorization verification, one operating in the user authentication mode and the other in the conventional mode, a problem may occur when you run a ZDB or IR session if no or wrong user credentials exist in the ZDB database (XPDB). In such circumstances, the problem occurs if the HPEP9000 XP Agent first connects to the command device with enabled authentication, and after failing to start the requested operation, it connects to the command device with disabled authentication. At this point the session fails unexpectedly.

Workaround: Do one of the following and restart the session afterwards:

- Using the `omnidbpx -user` command, add correct user credentials to the XPDB or update the existing ones appropriately.
For command syntax and usage examples, see the `omnidbpx` reference page in the *HPE Data Protector Command Line Interface Reference* or the `omnidbpx man` page.
- Disable the user authentication mode on the command device.
- Prevent the HPE P9000 XP Agent from connecting to the command device operating in the user authentication mode using either method:
 - Unpresent the command device from the application system and the backup system.
 - Follow the steps:
 - A. On the application system and the backup system, set the `SSEA_QUERY_STORED_CMDDEVS omnirc` option to 1.
 - B. Remove the data belonging to the command device from the XPDB using the `omnidbpx -cm -remove` command.

Microsoft SharePoint Server

- When the number of site collections of the backed up content database equals to the value of the parameter `Site Level of Warning`, during the restore the values of the `Site Level of Warning` and the `Maximum Number of Sites` parameters increases as follows:
`Site Level of Warning = number of site collections + 500`
`Maximum Number of Sites = number of site collections + 1000`
- After a restore of the configuration database, the data in the Microsoft SharePoint Server filesystem caches on front-end Web Server systems might not be consistent with the data in the newly-restored configuration database.
Workaround: Clear the Microsoft Office SharePoint Server file system cache on all server systems in the farm and retry the restore. For details, see the Microsoft web page <http://support.microsoft.com/kb/939308>.

SAP MaxDB

- Backup completes with errors if filenames contain spaces.
Workaround:
Windows systems:
 - a. Change the `RUNDIRECTORY` parameter to short (8+3) path names and edit filenames in the registry key `HKEY_LOCAL_MACHINE\SOFTWARE\SAP\SAP_DBTech\IndepData`.
 - b. Restart the database.**HP-UX and Linux systems:**
 - a. Create a symbolic link to the directory with a space in the name and adjust the `RUNDIRECTORY` parameter of the database to use the symbolic link.
 - b. Adjust the values of the `IndepData` parameter in the file `/var/spool/sql/ini/SAP_DBTech.ini` (on HP-UX) or `/usr/spool/sql/ini/SAP_DBTech.ini` (on Linux).
- On SUSE Linux Enterprise Server 10 x86-64 systems with SAP MaxDB 7.6 installed, you cannot

back up SAP MaxDB data with more than 19 streams. If you set the **Parallelism** option to a higher value, the session fails.

Workaround: Contact SAP MaxDB support.

SAP HANA Appliance

- After restore and recovery of an SAP HANA database using Data Protector, a problem with starting SAP HANA services occurs. The following issue is reported when the mouse pointer is paused on an instance name in the scoping pane of SAP HANA Studio:

```
InstanceID SystemName InstanceNumericalID - Some services not started
```

For example, the notification in a pop-up window reads:

```
H95 hanasys 95 - Some services not started
```

This symptom might indicate a failure to automatically start the SAP HANA name server. After changing the logon user account to the one that is specified in the corresponding SAP HANA backup specification (*SAPHANAUserAccount*), and attempting to manually start the name server, the system reports the following error:

```
su - SAPHANAUserAccount
```

```
cd InstallationPath/InstanceID/HDBInstanceNumericalID
```

```
./HDB start
```

```
Start service InstallationPath/sapservices : Permission denied.
```

Workaround: Add the missing access permissions to the SAP HANA name server binary file, and manually start the name server again. Follow the steps:

- a. Log on to the SAP HANA system as the root user and open a Terminal window.
- b. Execute the commands in sequence:

```
cd InstallationPath
```

```
chmod 777 sapservices
```

```
su - SAPHANAUserAccount
```

```
cd InstanceID/HDBInstanceNumericalID
```

```
./HDB start
```

Oracle Server

- In case the backup system is low on resources (CPU, memory, and so on), the following error is reported by the Oracle Server Manager in the Data Protector Monitor context for the Oracle HPE P9000 XP Disk Array Family integration:

```
ORA-12532: TNS: invalid argument
```

Workaround: Configure the backup system so that it has sufficient resources to simultaneously run the Oracle Server instance and execute a backup session.

- While performing a backup set ZDB session, the following warning is displayed for each database datafile:

```
RMAN-06554: WARNING: file n is in backup mode
```

Processing of each such message may take up to 20 seconds. It considerably slows down backups of databases with a large number of datafiles (200 or more).

- If you abort a running backup set ZDB session after the Oracle RMAN command has already been invoked by Data Protector, the ZDB session is terminated properly, but the RMAN-related Oracle Server processes continue to run on the backup system. These processes need to be manually terminated.

Workaround: To automatically terminate problematic processes properly, run another Data Protector session based on the same ZDB backup specification.

VMware vSphere

- While running multiple parallel VEPA backup sessions, a small fraction (1 or 2) of the total number of sessions may stall, causing the VEPA and BSM to not respond until the time-out period. This may occur because vCenter is loaded with multiple concurrent connection requests from the VEPA agent (from parallel VM backup sessions). The stalled VEPA and BSM processes for these sessions may finally time-out with the following message:

```
[Major] From: BSM@hostname "barlist7" Time: <DATE> <TIME>  
[61:1002] The OB2BAR Backup DA named "/Datacenter" on host machineName reached  
its inactivity timeout of xxxx seconds. The agent on host will be shutdown.
```

Workaround: After the time-out period:

- a. Stop the vepa_bar processes manually and wait for the bsm process associated with it to close. The bsm closes once the vepa_bar exits.
- b. Restart the backup specification that contains the failed VM objects.

Note: If the time-out period elapses before the backup starts (while resolving the objects), increase the SmWaitForFirstBackupClient parameter in Internal Database -> Global Options.

- The restore of thin disks using the SAN transportation mode is slower compared to the backup. The vCenter Server system receives a large number of requests to "Clear lazy zero" and "Allocate blocks" during the restore, which cause a slower restore and fill up the vCenter Server system database by logging the tasks.

This is caused by a known issue described at

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1035096.

Workaround: Do one of the following:

- Use an ESX Server as a Data Protector client and select the target datastore for restore.
- Use a vCenter Server system as a Data Protector client and set the OB2_VEAGENT_RESTORE_TRANSPORT_METHOD omnirc option on the backup host to NBD.
- The backup and restore of a virtual machine on VMware VVol datastores does not use the SAN transport mode. This limitation is documented by VMware in the VMware vSphere 6.0 Documentation Center (Table: Summary of restore snapshot requirements). See, <http://pubs.vmware.com/vsphere-60/topic/com.vmware.vddk.pg.doc/vddkBkupVadp.9.4.html>
- The backup and restore of a virtual machine with greater than 2TB disks on VMware VVol (Virtual Volumes) datastores, fall back to NBDSSL or NBD transport mode, when the selected transport

mode is HotAdd.

This is caused by a Known Issue (VDDK cannot HotAdd a > 2TB disk on VVol datastores) described in *VMware VDDK 6.0 Release Notes*. See, <https://www.vmware.com/support/developer/vddk/vddk-600-releasenotes.html>

- After an upgrade to Data Protector 9.05, the SAN transport mode falls back to NBDSSL on the vSphere versions 5.1 and 5.5.

This is a VMware VDDK 6.0 Update1 issue. For more details, see the following URL: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2135621.

Sybase Server

- On Solaris systems, aborting a Sybase backup session makes the system unresponsive.
Workaround: Abort the backup session by terminating the `$SYBASE_HOME_DIR/bin/sybmultbuf` process from the command-line interface.

Disk array integrations

- The Data Protector integration with HPE P6000 EVA Disk Array Family provides instant recovery by use of snapclones. The snapclone creation takes time and requires disk array resources. The actual performance impact depends on factors such as disk management, configuration, I/O load, and disk usage. Thus, HPE strongly recommends to perform performance benchmarking in sensitive environments before using snapclones.

Data Protector also provides built-in performance boosting functionality. For example:

- You can allocate snapclones to a different disk group than the one used for the original virtual disks, thus redirecting read and write operations on a replica from the original disk group to a replica disk group, or allocating low-performance disks for replicas.
- During a ZDB-to-disk+tape or ZDB-to-tape session, you can postpone the backup to tape until the snapclones are fully created, thus preventing performance degradation of the application during this phase.

For further assistance, contact HPE support.

- On Windows systems, if performing a snapshot backup on P6000 EVA Array, the following message may occur:

```
[Normal]Starting drive discovery routine.
```

```
[Major]Resolving of filesystem fsname has failed. Details unknown.
```

Workaround: Install Secure Path version 4.0B and patch v4.0B-3. The patch can be obtained from the HPE webpage <https://softwaresupport.hpe.com/patches>.

- When using the Secure Path 4.0C driver, unrecoverable error occurs occasionally on the backup system.
- On Windows Server 2008 systems without the Windows Server 2008 Service Pack 2 installed, it may occur that the Data Protector ZDB agent cannot dismount a volume during an ZDB or IR session, although no processes are running which could keep the volume locked and prevent the dismount operation.

Workaround:

- a. On the system where the problematic volume resides, perform one of the following:
 - o Update the operating system to Windows Server 2008 Service Pack 2.
 - o Install a specific Windows Server 2008 hotfix. The hotfix package can be obtained from the Microsoft website <http://support.microsoft.com/kb/952790>.
 - o Set the `omnirc` option `SMISA_FORCE_DISMOUNT` (in the case of the Data Protector HPE P6000 / HPE 3PAR SMI-S Agent) or `SSEA_FORCE_DISMOUNT` (in the case of Data Protector HPE P9000 XP Agent) to 1.
 - b. Restart the session that failed.
- In circumstances when several ZDB sessions that involve the HPE P6000 / HPE 3PAR SMI-S Agent and a Windows Server 2008 SP2 backup system are running simultaneously, the backup administrator logged on the backup system using the system default administrative account might be occasionally presented with a pop-up window, asking them to format a disk which is presented to the backup system. A message similar to the following is displayed in the pop-up window:

You need to format the disk in drive *DriveLetter*: before you can use it.

Do you want to format it?

This was recognized as a known issue by Microsoft, and was addressed by the hotfix available at <http://support.microsoft.com/kb/971254>. When installed, the hotfix significantly reduces the frequency of such occurrences, but it does not completely eliminate them. According to Microsoft, the problem might also occur on Windows Server 2008 R2.

Workaround: Click **Cancel** to close the pop-up window. To avoid such pop-up windows from reappearing, disable the system-default administrative account and use another user account. The workaround might not be useful on Windows Server 2008 R2 systems. For further assistance, contact HPE Customer Support Service or Microsoft Support directly.

- On Windows Server 2008 R2 systems, when using the Data Protector Microsoft Volume Shadow Copy Service integration or the Data Protector HPE P6000 / HPE 3PAR SMI-S Agent, you may encounter either of the following problems after several zero downtime backup sessions have been simultaneously and continuously running for several days:

- Although presented to the backup system, target volumes are not recognized by the operating system. As a result, the affected ZDB session ends abnormally. All consecutive ZDB sessions fail as well.

Although this problem turned out not to occur under ordinary conditions, it is not possible to exclude that it will occur in your environment.

Workaround: None. HPE is collaborating with external partners to find a solution.

- A critical system error occurs on the application system, resulting in a Stop error message (displayed in white text on a blue screen).

This was recognized by Microsoft as a known issue in the Microsoft Multipath I/O (MPIO) framework driver, and was addressed by the hotfixes available at <http://support.microsoft.com/kb/2511962> and <http://support.microsoft.com/kb/2549567>. The hotfixes resolve one aspect of the issue and significantly reduce the probability of a system failure.

Workaround: Install the hotfixes on the application system, and rerun the problematic sessions. If the problem persists, avoid running multiple ZDB sessions in parallel.

- On SUSE Linux Enterprise Server 10.3 and 11.1 and on Oracle Enterprise Linux 5.3, after multiple simultaneous zero downtime backup sessions that involve a P6000 EVA Array and the same backup system have been running continuously for a longer period, the virtual disks of the disk array are unexpectedly unrepresented from the backup system. Additionally, creation of the virtual disk device files on the backup system fails sporadically, even after a user-triggered disk rescan is complete.

Workaround: Restart the backup system, and rerun the problematic zero downtime backup sessions.

- If two command devices are configured for a disk array of the HPE P9000 XP Disk Array Family which supports authorization verification, one operating in the user authentication mode and the other in the conventional mode, a problem may occur when you run a ZDB or IR session if no or wrong user credentials exist in the ZDB database (XPDB). In such circumstances, the problem occurs if the HPEP9000 XP Agent first connects to the command device with enabled authentication, and after failing to start the requested operation, it connects to the command device with disabled authentication. At this point the session fails unexpectedly.

Workaround: Do one of the following and restart the session afterwards:

- Using the `omnidbpx -user` command, add correct user credentials to the XPDB or update the existing ones appropriately.

For command syntax and usage examples, see the `omnidbpx` reference page in the *HPE Data Protector Command Line Interface Reference* or the `omnidbpx` man page.

- Disable the user authentication mode on the command device.
- Prevent the HPE P9000 XP Agent from connecting to the command device operating in the user authentication mode using either method:
 - Unpresent the command device from the application system and the backup system.
 - Follow the steps:
 - A. On the application system and the backup system, set the `SSEA_QUERY_STORED_CMDDEVS omnirc` option to 1.
 - B. Remove the data belonging to the command device from the XPDB using the `omnidbpx -cm -remove` command.
- If you remove the last HPE 3PAR StoreServ Storage snapshot presentation that belongs to a specific application system, the application system is also removed from its HPE 3PAR StoreServ Storage virtual domain.

The problem is caused by a known issue in handling the *Super* user account privilege level by the HPE 3PAR StoreServ Storage firmware.

Workaround: Configure the Data Protector HPE 3PAR StoreServ Storage integration agent by choosing an HPE 3PAR StoreServ Storage system user account that only has the *Edit* privilege level assigned.

Non-Data Protector issues related to Granular Recovery Extensions

VMware vSphere

- When you select a virtual machine and navigate to the **HPE Data Protector** plugin tab, the following error message is displayed:

```
HTTP Status 500 -
```

```
Exception javax.servlet.ServletException: Unable to read shell environment variables
```

```
org.apache.catalina.servlets.CGIServlet.init(CGIServlet.java:310)
```

```
org.apache.catalina.valves.ErrorReportValve.invoke(ErrorReportValve.java:102)
```

```
...
```

This error appears in Windows Server 2008 environments in combination with some versions of the Tomcat Server for vCenter.

Workaround: Change the `os.name` property in the Tomcat configuration:

- a. Run the Tomcat configuration utility: **Start > VMware > VMware Tomcat > Configure Tomcat.**
- b. Click the **Java** tab and add the following line to Java options: `-Dos.name="Windows 2008"`.
- c. Restart the Tomcat Server.

Microsoft Exchange Server

- Recovery of an entire mailbox or single mailbox items using the Data Protector Granular Recovery Extension for Microsoft Exchange Server in an Exchange Server 2013 environment with the Cumulative Update 1 for Exchange Server 2013 (Exchange 2013 CU1) installed fails with the following error message:

```
The call to 'net.tcp://serverName/Microsoft.Exchange.MailboxReplicationService serverName (15.0.620.29 caps:3F)' failed. Error details: must be logging in with GUIDs, not legDN
```

```
Parameter name: owner.
```

The problem is caused by a known Exchange 2013 CU1 issue, and does not appear in Exchange Server 2013 environments without the Exchange 2013 CU1 installed. For details, see [http://technet.microsoft.com/en-us/library/jj150489\(v=exchg.150\).aspx](http://technet.microsoft.com/en-us/library/jj150489(v=exchg.150).aspx).

Workaround: None. Installing a future Microsoft Exchange Server patch will resolve the issue.

Non-Data Protector issues related to disaster recovery

- During an Enhanced Automated Disaster recovery of an Red Hat Enterprise Linux 5.1, the restore session completes successfully, but the operating system is left in an inconsistent state after the disaster recovery and does not start successfully.

Workaround: Update the GRUB bootloader package to `grub-0.97-13.5.src.rpm` or a later version, as described in <http://rhn.redhat.com/errata/RHBA-2008-0440.html>.

Non-Data Protector issues related to reporting

- While using Microsoft Outlook, when you add a report to a report group specifying e-mail as the send method, and then try to start the report group, the CRS service stops responding and must be restarted. The same happens if you configure a notification and select the e-mail send method. The cause of the problem is that Outlook requires user interaction before sending an e-mail notification.

Workaround: To prevent this behavior, customize security settings so that you set the **When sending items via Simple MAPI** option to `Automatically approve`. For information on how to customize security settings for Microsoft Outlook XP, 2003, or 2007, see the respective Office Resource Kit.

Additionally, Outlook Express can be used as an alternative to Outlook, as it does not require any user intervention for sending e-mails. Data Protector is able to send reports in HTML format if used in combination with Outlook Express. Otherwise an HTML report is sent as an attachment. Outlook Express is installed by default on specific Windows operating systems and is the default MAPI handler on those systems. If you plan to use Outlook Express, do not install any other e-mail software (including Outlook) since it typically replaces the default MAPI handler. If you are using Microsoft Office, ensure that you do not select Microsoft Outlook during Microsoft Office installation. Outlook Express supports only the SMTP protocol as e-mail carrier. If you plan to use Outlook Express with Microsoft Exchange Server systems, the **SMTP Mail Connector** option must be enabled on the Microsoft Exchange Server. For details of how to configure SMTP on Microsoft Exchange Server system, see the Microsoft webpage <http://support.microsoft.com/kb/265293>.

- If a Data Protector Cell Manager and Microsoft Exchange Server 2003 or 2007 coexist on the same system, e-mail reporting using MAPI does not function. This is because Microsoft does not support installing Outlook on a system with Microsoft Exchange Server 2003 or 2007 installed.

Workaround: Use the e-mail SMTP send method for reports and notifications.

- On UNIX systems, due to the operating system limitations, international characters in localized e-mail notifications and reporting may be displayed incorrectly if they are transmitted between systems using a different locale.
- When viewing web reports using Netscape Navigator, after resizing the browser window the applet does not automatically adjust its size appropriately.

Workaround: Start Netscape Navigator manually, resize the window to the desired size and only then open the `WebReporting.html` file.

- In localized UNIX environments with SJIS or EUC Japanese locale set, the non-UTF-8 web reporting input data is converted into UTF-8 (Unicode) before it is written to the Data Protector configuration files. Such characters will not be displayed correctly when using web reporting.
- When you are backing up Data Protector clients not configured for Data Protector report, the report lists all clients from a specified network range. In case you specify a C-class network that is in another subnet, the report can take significant time to be created.
- If you use Data Protector reporting and the HTML output format, a Unicode file is generated. Some older web browsers do not support local viewing of Unicode files. However, the files may be displayed correctly if retrieved from a Web server.
- If you receive localized Data Protector e-mail notifications containing Japanese characters on the

host where Japanese is not the default locale, the output of the notifications may not be displayed correctly.

Workaround:

- a. If you have this problem with the Microsoft Outlook, save the message in the HTML format, then open it in a web browser and follow the next step.
 - b. If you use a web browser, select the Japanese locale, Shift-JIS, EUC, or UTF-8. For example, select **View > Character Encoding > More Encodings > East Asian-Japanese (Shift_JIS)**.
- Due to the Microsoft Office Word 2007 limitation which states that the maximum number of columns in a table is 63, the following issue can occur:

When using Microsoft Outlook 2007 and “email SMTP” send method, HTML format, for Device Flow Report and Session Flow Report, Outlook does not display properly the tables in the reports since these reports contain more than 63 columns. The same issue occurs if you log such a report to an HTML file and then try to open it with Microsoft Office Word. Also, in both cases, the tooltips are not displayed.

Workaround: Do not use Word to display such a report. Use a web browser supported by Data Protector. You can open the report with a web browser in one of the following ways:

- Open the mail. In the **Other Actions** menu, click **View in Browser**.
- Since the report is sent also as the HTML format attachment, you can open the attachment directly from Outlook, or you can save the attachment first and then open it with a supported browser.

Other known non-Data Protector issues

- When mounting a CIFS share on a UNIX system, the shared directory size is not calculated correctly and Data Protector backup statistics consequently report a wrong backup size at the end of the backup session. The reason are inter-operability problems between Windows and UNIX platforms.
- Backup on UNIX systems may fail because of the shared memory shortage with the following error:

```
Cannot allocate shared memory pool (IPC Cannot Create Shared Memory Segment  
System error: [22] Invalid argument ) => aborting
```

Workaround: The actions are different for different operating systems. After you have applied the changes, you need to restart the system.

HP-UX systems:

Set the OB2SHMEM_IPCGLOBAL option to 1 in the file `/opt/omni/.omnirc`.

Solaris systems:

Set the kernel parameters in the `/etc/system` file as follows:

```
set shmsys:shminfo_shmmax=4294967295  
set shmsys:shminfo_shmmmin=1  
set shmsys:shminfo_shmmni=100  
set shmsys:shminfo_shmseg=10  
set semsys:seminfo_semmni=100
```

```
set semsys:seminfo_semmsl=100  
set semsys:seminfo_semmns=256  
set semsys:seminfo_semopm=100  
set semsys:seminfo_sevmx=32767
```

If the problem persists, the parameter value needs to be increased.

- Data Protector uses host name resolution for communication between different systems. This is done either via DNS servers or via `/etc/hosts` or `/etc/lmhosts` file. On Windows clients, if the DNS service is not available or correctly configured, you can edit the `hosts` (`lmhosts`) file, which is located in the `%SystemRoot%\System32\drivers\etc` directory. Use the `hosts` file if you want to map IP addresses to host names and `lmhosts` file if you want to map IP addresses to computer (NetBIOS) names. Additional information on how you can edit these files can be found in the beginning of these two files. Restart Data Protector GUI for changes to take effect. You must ensure that the name resolution is consistent throughout the Data Protector cell.
- Secure path on HP-UX external device filename may change after restart. This changes the mapping to volume managers. Raw device backups may fail due to a different device file being specified in the backup specification.
- When creating a file system backup for a Windows Vista, Windows 7, Windows 8, Windows Server 2008, or Windows Server 2012 system, Data Protector GUI does not list `TerminalServiceDatabase` among Windows configuration objects available for backup.
Workaround: To enable backup of the `TerminalServiceDatabase` configuration object, install the Terminal Server Licensing service on the system which will be backed up.
- When creating a file system backup for a Windows Vista, Windows 7, Windows 8, Windows Server 2008, or Windows Server 2012 system, Data Protector GUI does not list `RemovableStorageManagementDatabase` among Windows configuration objects available for backup.
Workaround: To enable backup of the `RemovableStorageManagementDatabase` configuration object, install Removable Storage Manager on the system to be backed up.
- If a FAT32 boot partition exists on Windows Server 2003 system, you cannot use a Windows Vista client for creating an ISO image for this system, since the resulting CD-ROM cannot be used to start the system.
Workaround: Use the Windows Server 2003 system to create the ISO image.
- Data Protector clients without the Internet Protocol version 6 (IPv6) functionality are not able to connect to IPv6-only clients in the cell.
Workaround: For all clients running a newer version of Data Protector in such mixed environments, a dual-stack configuration (enabled both IPv4 and IPv6 protocols) is recommended.
- After installing the Quality Pack Patch Bundle 1103 or 1109 on HP-UX 11.31, the Data Protector backup session performance decreases significantly.
Workaround: To resolve this issue, install the kernel patch `PHKL_41967`.
After installing this patch and setting the parameter, the Data Protector backup performance is restored.

Chapter 5: Resolved issues and enhancements

List of enhancements implemented and resolved issues in Data Protector

This section lists the issues resolved and enhancements implemented in Data Protector 9.00. For issues related to the latest Data Protector patches, see <https://softwaresupport.hpe.com>.

Resolved issues

Global ID	Summary
QCCR2A49005	The IDB recovery results in a critical error: "Replacement of the Internal Database failed".
QCCR2A49469	The Data Protector 8.0 "omnigrate.pl" script fails if a space exists in the DPDATA path name.
QCCR2A49135	The Data Protector 6.21 Microsoft SQL Server Integration backup session takes a long time to start.
QCCR2A49740	Selection of the restore version is inconsistent in the Data Protector 8.0 Graphical User Interface (GUI).
QCCR2A49541	The Data Protector 7.03 Graphical User Interface (GUI) ends abnormally while performing the "List of Sessions" report.
QCCR2A49478	The Backup Media Agent (BMA) ends abnormally in SysWrite_SODA.
QCCR2A49334	The "Apply Template" procedure does not work for integrations.
QCCR2A49673	In the DP 8.0 Graphical User Interface (GUI), the session filter in the "Backup Specifications" field is too small.
QCCR2A49833	Size of the restored file does not match the size of the backed-up file.
QCCR2A49820	Instant recovery of SQL ZDB with P9500 ends unexpectedly.

Global ID	Summary
QCCR2A49412	The Virtual Environment Integration restore session fails with the following error: "Wrong datastore path format".
QCCR2A49875	The media replication session ends abnormally without unloading the media in Data Protector 8.0.
QCCR2A49840	The object copy session uses only one device.
QCCR2A50044	Data Protector 7.0 is unable to start the StoreOnce software deduplication datastore.
QCCR2A49623	The Data Protector 7.01 Graphical User Interface (GUI) ends abnormally in MoM during listing of objects.
QCCR2A49369	Licenses are not available after the system restart.
QCCR2A50119	Possible memory irregularities in the Volume Restore Disk Agent (VRDA).
QCCR2A50158	Data Protector 7.03 has notifications and schedule issues for replication sessions.
QCCR2A50147	The Data Protector 8.0 Database Session Manager (DBSM) on the Linux system fails with an error message.
QCCR2A49944	During the IDB SMBF check, a lot of session message files are missing.
QCCR2A50091	Data Protector 7.03 cannot back up files that have path names longer than 1024 characters.
QCCR2A49966	Data Protector Microsoft SharePoint Server 2007/2010/2013 integration backup session ends abnormally.
QCCR2A50162	Unable to remove items from the StoreOnce devices.
QCCR2A50232	Restore session fails with the following error: "Could not find datacenter".
QCCR2A50272	The restore session fails with the following error: "Error: sharepoint_bar: The farm is unavailable".
QCCR2A49715	After an upgrade to Data Protector 8.00, the Detail Catalog (DC) gets corrupted when using the non-migrated tapes.
QCCR2A50348	The Volume Backup Disk Agent (VBDA) displays

Global ID	Summary
	the following error: "Cannot read ACLs: (22 Invalid argument)".
QCCR2A50453	Data Protector HPE P9000 XP Agent (SSEA) displays the following error: "Mapping from character device file of disk to LDEV failed".
QCCR2A50227	The Backup Catalog on Data Protector 8.0 is not displayed after a media copy from VLS to LTO.
QCCR2A50449	After an upgrade to Data Protector 7.03, the session fails with the following error: "The creation of ISO image failed".
QCCR2A49655	Scheduled reports do not work.
QCCR2A46934	Query data is not found in the Data Protector Graphical User Interface (GUI).
QCCR2A50642	The IDB backup session stops responding on Data Protector 8.0.
QCCR2A50468	The LTO6 medium cannot be initialized though 2.5 TB is its native size.
QCCR2A50605	The TBL database unpacking in Data Protector 8.1 fails when the string size is 0.
QCCR2A50688	The "omnidbutil -remap_dkdir" command fails and thereafter the DCBF information is corrupted in the IDB.
QCCR2A50576	The incremental backup session fails on NSS/OES.
QCCR2A50769	Access to Data Protector 8.1 Graphical User Interface (GUI) is denied for users with special characters in the name.
QCCR2A50760	The Data Protector 8.1 IDB is not correctly upgraded.
QCCR2A50349	The catalog migration fails with the following error: "Cannot upgrade catalog of medium".
QCCR2A39081	Oracle 11g on RHEL 6.1 Veritas cluster 5.1 support.
QCCR2A48578	The "omnirpt -report obj_lastbackup" command does not report some of the VEagent objects.
QCCR2A49416	The CSM locks to many target devices during the Data Protector 8.0 object copy session.

Global ID	Summary
QCCR2A48841	The backup session fails due to a disabled primary device in the mirroring devices.
QCCR2A49827	When the first gateway server is down, the replication session does not switch to the second configured gateway server.
QCCR2A50476	The backup sessions stop responding.
QCCR2A49603	The output of the "omnidownload" command is not sorted correctly.
QCCR2A49284	The original virtual machine is overwritten when the "Restore as" option is used.
QCCR2A48734	The Data Protector Virtual Environment Integration component ends abnormally.
QCCR2A50149	After an upgrade to Data Protector 8.0, a warning message is displayed during the daily maintenance checks.
QCCR2A50809	After upgrading to Data Protector 8.0, the restore sessions are very slow.
QCCR2A50750	The Novell Open Enterprise Server 11 SP1 backup session in Data Protector 8.1 fails with a segmentation error.
QCCR2A50365	A restore session fails with the following error: "Cannot write: (5 Access is denied.) => not restored".
QCCR2A51024	After an upgrade from Data Protector 6.20, the folder restore is not possible.
QCCR2A50942	The script "Mount.bat" does not start in the Data Protector Windows Cell Manager.
QCCR2A50868	There are memory allocation issues during the object copy sessions.
QCCR2A48632	Data Protector 7.03 fails due to the SCSI Reserve/Release barcode scan issue.
QCCR2A41587	EMC SYMMETRIX VMAXE storage array support.
QCCR2A47238	Enhancement request for SYMA to support new device special files.
QCCR2A51124	The Data Protector 8.1 "omnib -rawdisk" command does not run.

Global ID	Summary
QCCR2A51127	vCenter 5.1 update 2 starts with a blank page.
QCCR2A46384	In the Data Protector Graphical User Interface (GUI), add filters for the virtual machine name.
QCCR2A50891	The Data Protector 8.0 "omnidbcheck -bf" command does not run correctly.
QCCR2A50817	A standalone Media Agent or Disk Agent stops responding.
QCCR2A51086	The Data Protector 8.1 Virtual Environment Integration backup session fails with the following error: "Exception Occurred while creating VM snapshot".
QCCR2A50745	After upgrading to Data Protector 7.03, VBDA ignores the exclude list from the backup specification.
QCCR2A50983	The Data Protector 8.0 restore session of German vowels fails on SUSE Linux Enterprise Server 11 with Novell Open Enterprise Server 11.
QCCR2A51125	The Data Protector 7.03 Lotus integration cannot restore a backup on which the objects appear as failed.
QCCR2A51295	The "Encrypted control communication" option does not function correctly on an HP-UX client.
QCCR2A51257	An enhancement request to disable the "encrypted" checkbox for the StoreOnceSoftware store.
QCCR2A51154	The restored files from deduplicated volumes to a different host are corrupted.
QCCR2A51033	The ZDB SSE Agent backup session fails if the "SSEA_ATOMIC_SPLIT=1" omnirc option is used.
QCCR2A49880	The configured Daily report is not executed correctly.
QCCR2A50399	Objects with the "Completed/Errors" status are excluded from the object copy session.
QCCR2A51225	The Data Protector "omniofflr" command reports the "Unknown option -dfmf." error.
QCCR2A50089	The total "Store Data Size" does not match the size of the StoreOnce root directory.

Global ID	Summary
QCCR2A51142	The Data Protector Virtual Environment Integration (VEagent) template creation fails with the "details unknown" error.
QCCR2A51060	The backup session fails with the following error: "The string was not recognized as a valid DateTime".
QCCR2A51465	Upgrade to Data Protector 8.11 for the VMware Granular Recovery Extension agent fails.
QCCR2A50798	The Data Protector 8.1 SAP HANA Appliance integration fails.
QCCR2A51498	Upgrade to Data Protector 8.0 displays rejected lines from media.txt.
QCCR2A51527	Data Protector 8.1 incorrectly handles empty spaces in data paths.
QCCR2A51398	Auto-registration of virtual machine fails.
QCCR2A50218	An error occurs while running the Data Protector "omnidl" command.
QCCR2A50550	The Data Protector Granular Recovery Extension for SharePoint Server installation fails on the German SharePoint environment.
QCCR2A51446	When restoring a vApp, if a machine is backed up on a Windows XP storage, the restore fails with the following message: "handleDiskMetaData-Unknown".
QCCR2A51324	Data Protector Virtual Environment Integration backup sessions fail with the following error: "Connection reset by peer".
QCCR2A51456	The backup session fails with the following error: "Backup failed because unable to utilize CBT as it was required".
QCCR2A50434	Upgrade to Data Protector 8.1 does not proceed.
QCCR2A51522	The IDB upgrade to Data Protector 8.0, 8.01, or 8.1 fails and is missing the catalog information.
QCCR2A51318	EADR on Linux systems fails with the update boot disk error.
QCCR2A51396	The "omnikeytool -import CSVFile" command causes the kms.exe file to end abnormally.

Global ID	Summary
QCCR2A47390	Connection to the application agent (OB2BAR) backup disk agent (DA) is lost.
QCCR2A50673	The Data Protector 8.0 Graphical User Interface (GUI) stops responding.
QCCR2A51239	The Data Protector 7.03 Microsoft SQL Server integration backup session fails with the following error: "Session has failed".
QCCR2A51198	The IDB backup session on Data Protector 8.1 fails if the NTFS mountpoint is used.
QCCR2A51234	The Data Protector Session Manager fails to load or unload.
QCCR2A51214	The "omnidbutil -mergemmdb" command fails if the hostname contains a special character.
QCCR2A51567	The Data Protector 8.1 object copy session fails with the following error: "Specified device for the object not found".
QCCR2A51589	The Data Protector 8.1 VEPA backup session reports a false warning message.
QCCR2A51726	A mailbox cannot be enumerated during the restore session.
QCCR2A48803	For verification sessions, the Session Flow report shows a Red bar under the Success column.
QCCR2A51243	Data Protector 8.1 backup sessions do not start during the removal of a virtual machine snapshot.
QCCR2A51956	Snapshot is not removed after the backup session completes.
QCCR2A50933	Data Protector 7.03 backup session fails with the following error: "ASR information could not be collected".
QCCR2A51056	Replication session stops responding.
QCCR2A52056	The backup specification preview progress is missing in the Monitor context.
QCCR2A51415	After upgrading to Data Protector 8.1 on a MOM environment, the IDB records related to tape are not migrated correctly.
QCCR2A52103	The Virtual Environment Integration backup session

Global ID	Summary
	fails after upgrading to the Data Protector 8.11
QCCR2A51518	Data Protector 8.01 client statistics reports display the integer overflow error.
QCCR2A52139	Enabling user restrictions causes the RDS to end abnormally.
QCCR2A51776	The EADR of SAP HANA appliance does not complete.
QCCR2A52338	The Data Protector 8.11 virtual machine restore session denies filter settings when a host or cluster is selected.
QCCR2A52438	The Data Protector 8.01 "devbra" command does not list the EML drives under robotics.
QCCR2A52218	The Data Protector 7.03 "omnic" command reports incorrect number of licenses.QCCR2A52308
QCCR2A52308	The virtual machine name is not displayed in the Object Latest Backup report.
QCCR2A52582	The Data Protector 8.11 Graphical User Interface (GUI) does not list all virtual machines in the Restore context.
QCCR2A52585	The Data Protector 8.11 Virtual Environment Integration backup session fails with the following error: "Exception Occurred while creating VM snapshot".
QCCR2A52943	Cannot open device (Semaphores).
QCCR2A52941	Logsense overflow on Solaris hosts.
QCCR2A52939	SCSI reserve/release failing on Solaris.
QCCR2A52410	Upgrading the catalog in Data Protector 8.1 and 8.11 fails.
QCCR2A52416	Problems with the "omnimm -delete_unprotected_media" command on Data Protector 7.03.

Chapter 6: Data Protector documentation

Note: The documentation set available at the HPE support website at <https://softwaresupport.hpe.com/> contains the latest updates and corrections.

You can access the Data Protector documentation set from the following locations:

- Data Protector installation directory.
Windows systems: `Data_Protector_home\docs`
UNIX systems: `/opt/omni/doc/C`
- **Help** menu of the Data Protector GUI.
- HPE Support website at <https://softwaresupport.hpe.com/>

Documentation map

The following table shows where to find information of different kinds. Squares shaded in gray are a good place to look first.

	Admin	Help	Getting Started	Concepts	Install	Troubleshooting	DR	CLI	PA	Integration VSS	Integration Guide				ZDB Guides		GRE Guide			
											MSFT	Oracle/SAP	IBM	Sybase/NDMP	Virtual Env	ZDB Admin	ZDB IG	Exchange	SharePoint	VMware
Admin tasks	X	X																		
Backup		X	X	X						X	X	X	X	X	X	X	X			
CLI								X												
Concepts, techniques		X	X							X	X	X	X	X	X	X	X	X	X	X
Disaster recovery				X			X													
Installation, upgrade			X		X				X											
Instant recovery				X	X											X	X			
Licensing					X				X											
Limitations		X			X	X			X	X	X	X	X	X	X		X			
New features		X							X											
Planning strategy		X		X																
Procedures, tasks	X	X			X	X	X			X	X	X	X	X	X	X	X	X	X	X
Recommendations				X					X											
Requirements					X				X	X	X	X	X	X	X					
Restore	X	X	X	X						X	X	X	X	X	X	X	X	X	X	X
Supported configurations				X																
Troubleshooting		X			X	X				X	X	X	X	X	X	X	X	X	X	X

Abbreviations

Abbreviations in the documentation map above are explained below. The documentation item titles are all preceded by the words “HPE Data Protector.”

Abbreviation	Documentation item	
Admin	Administrator's Guide	This guide describes administrative tasks in Data Protector.
CLI	Command Line Interface Reference	This guide describes the Data Protector command-line interface, command options, and their usage as well as provides some basic command-line examples.
Concepts	Concepts Guide	This guide describes Data Protector concepts, zero downtime backup (ZDB) concepts and provides background information on how Data Protector works. It is intended to be used with the task-oriented Help.
DR	Disaster Recovery Guide	This guide describes how to plan, prepare for, test, and perform a disaster recovery.
Getting Started	Getting Started Guide	This guide contains information to get you started with using Data Protector. It lists installation prerequisites, provides instructions on installing and configuring a basic backup environment and procedures for performing backup and restore. It also lists resources for further information.
GRE Guide	Granular Recovery Extension User Guide for Microsoft SharePoint Server, Exchange and VMware	This guide describes how to configure and use the Data Protector Granular Recovery Extension for: <ul style="list-style-type: none"> • Microsoft SharePoint Server • Exchange Server • VMware vSphere
Help	Help	
Install	Installation Guide	This guide describes how to install the Data Protector software, taking into account the operating system and architecture of your environment. This guide details how to upgrade Data Protector, as well as how to obtain the proper licenses for your environment.

Abbreviation	Documentation item	
Integration Guide	Integration Guide	This guide describes the integrations of Data Protector with the following applications: <ul style="list-style-type: none"> • MSFT: Microsoft SQL Server, Microsoft SharePoint Server, and Microsoft Exchange Server. • IBM: Informix Server, IBM DB2 UDB, and Lotus Notes/Domino Server. • Oracle/SAP: Oracle Server, MySQL, SAP R3, SAP MaxDB, and SAP HANA Appliance. • Sybase/NDMP: Sybase and Network Data Management Protocol Server. • Virtual Env: Virtualization environments integration with VMware vSphere, VMware vCloud Director, Microsoft Hyper-V, and Citrix XenServer.
Integration VSS	Integration Guide for Microsoft Volume Shadow Copy Service	This guide describes the integrations of Data Protector with Microsoft Volume Shadow Copy Service (VSS).
IG IDOL	Integration with Autonomy IDOL Server	This technical white paper describes all aspects of integrating Data Protector with Autonomy IDOL Server: integration concepts, installation and configuration, Data Protector backup image indexing, full content search-based restore, and troubleshooting.
PA	Product Announcements, Software Notes, and References	This guide gives a description of new features of the latest release. It also provides information on installation requirements, required patches, and limitations, as well as known issues and workarounds.
Troubleshooting	Troubleshooting Guide	This guide describes how to troubleshoot problems you may encounter when using Data Protector.
ZDB Admin	ZDB Administrator's Guide	This guide describes how to configure and use the integration of Data Protector

Abbreviation	Documentation item	
		with HPE P4000 SAN Solutions, HPE P6000 EVA Disk Array Family, HPE P9000 XP Disk Array Family, HPE 3PAR StoreServ Storage, NetApp Storage, EMC VNX Storage Family, EMC VMAX Storage Family, and EMC Symmetrix Remote Data Facility and TimeFinder. It is intended for backup administrators or operators. It covers the zero downtime backup, instant recovery, and the restore of filesystems and disk images.
ZDB IG	ZDB Integration Guide	This guide describes how to configure and use Data Protector to perform zero downtime backup, instant recovery, and standard restore of Oracle Server, SAP R/3, Microsoft Exchange Server, Microsoft SQL Server databases, and Virtual Environment for VMware .

Integrations

Software Application Integrations

Software application	Guides
Autonomy IDOL Server	IG IDOL
IBM DB2 UDB	Integration Guide
Informix Server	Integration Guide
Lotus Notes/Domino Server	Integration Guide
Microsoft Exchange Server	Integration Guide, ZDB IG, GRE Guide
Microsoft Hyper-V	Integration Guide
Microsoft SharePoint Server	Integration Guide, ZDB IG, GRE Guide
Microsoft SQL Server	Integration Guide, ZDB IG
Microsoft Volume Shadow Copy Service (VSS)	Integration VSS

Software application	Guides
Network Data Management Protocol (NDMP) Server	Integration Guide
Oracle Server	Integration Guide, ZDB IG
MySQL	Integration Guide
SAP HANA Appliance	Integration Guide
SAP MaxDB	Integration Guide
SAP R/3	Integration Guide, ZDB IG
Sybase Server	Integration Guide
VMware vCloud Director	Integration Guide
VMware vSphere	Integration Guide, ZDB IG, GRE Guide

Disk Array System Integrations

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

Disk array family	Guides
EMC Symmetrix	all ZDB
HPE P4000 SAN Solutions	Concepts, ZDB Admin, Integration Guide
HPE P6000 EVA Disk Array Family	all ZDB, Integration Guide
HPE P9000 XP Disk Array Family	all ZDB, Integration Guide
HPE 3PAR StoreServ Storage	Concepts, ZDB Admin, Integration Guide
NetApp Storage	Concepts, ZDB Admin, ZDB IG
EMC VNX Storage	Concepts, ZDB Admin, ZDB IG
EMC VMAX Storage	Concepts, ZDB Admin, ZDB IG

Localized documentation

Data Protector is localized into French, Japanese, and Simplified Chinese. The documentation for Data Protector 9.08 is available in English only. The localized documentation set is available in French, Japanese, and Simplified Chinese for Data Protector 9.06.

The following end-user documents are localized to French, Japanese, and Simplified Chinese:

- *HPE Data Protector Getting Started Guide*
- *HPE Data Protector Concepts Guide*
- *HPE Data Protector Administrator's Guide*
- *HPE Data Protector Disaster Recovery Guide*
- *HPE Data Protector Granular Recovery Extension for VMware Help*
- *HPE Data Protector Installation Guide*
- HPE Data Protector Product Announcements, Software Notes, and References
- *HPE Data Protector Troubleshooting Guide*
- *HPE Data Protector Help*

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 Product Announcements, Software Notes, and References (Data Protector 9.08)

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!