



Hewlett Packard
Enterprise

HPE Data Protector

Software Version: 9.09

Administrator's Guide

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

Legal Notices

Warranty

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

The information contained herein is subject to change without notice.

Restricted Rights Legend

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

Copyright Notice

© Copyright 2017 Hewlett Packard Enterprise Development LP

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

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

UNIX® is a registered trademark of The Open Group.

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

Documentation Updates

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

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

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

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

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

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

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

Support

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

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

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

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

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

To register for an HPE Passport ID, go to:

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

To find more information about access levels, go to:

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

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

Contents

Chapter 1: Introduction	1
About Data Protector	1
Major Data Protector features	1
Data Protector Architecture	1
Cell Manager	1
Installation Server	2
Client Systems	2
Systems to be backed up	2
Systems with backup devices	2
Overview of Tasks to Set Up Data Protector	2
Steps	2
Data Protector Operation	4
Backup session	4
Restore session	4
Pre-exec and post-exec commands	4
Object copy, object consolidation and object verification sessions	4
Chapter 2: Configuration Tasks	6
Enabling Security	6
About Security Considerations	6
Cell Manager Security	6
Client Security	6
Trusted clients	7
The allow_hosts and deny_hosts files	7
Users Security	8
User rights	8
Start backup specification user right	8
Hiding the contents of backup specifications	8
Host trusts	8
User groups	9
User restrictions	9
User validation	9
Strict Hostname Checking	9
Limitations	9
Requirements	10
Hostname resolution	10
Security Logs	11
Client security events	11
Cell Manager security events	11
Securing the Entire Data Protector Cell	11
Steps	11

Securing a Client System	11
Steps	12
Unsecuring the Entire Data Protector Cell	12
Steps	12
Unsecuring a Client System	12
Steps	12
Configuring Host Trusts	13
Steps	13
Encryption	13
About Encryption	13
Enabling AES 256-bit Encryption	14
Prerequisite	14
Limitations	14
Enabling encryption in a filesystem backup specification	14
Steps	14
Enabling encryption in a disk image backup specification	14
Steps	14
Enabling encryption in an Internal Database backup specification	15
Steps	15
Enabling encryption in an application integration backup specification	15
Limitations	15
Steps	15
Exporting and Importing Media with Encrypted Backups	15
Cell Manager environment or MoM environment without CMMDB	16
Steps	16
MoM environment with CMMDB	16
Steps	16
Enabling Drive-Based Encryption	16
Prerequisite	17
Limitations	17
Recommendation	17
Enabling drive-based encryption in the drive configuration	17
Steps	17
Enabling drive-based encryption in a backup specification	17
Steps	17
Enabling drive-based encryption for an automated media operation	18
Steps	18
Encrypted Control Communication	18
About Encrypted Control Communication	18
Managing Encrypted Control Communication	18
Considerations	19
Enabling encrypted control communication	19
Encrypted control communication with user-created certificates	22
Selecting TLS version	23
Disabling encrypted control communication	24
Viewing certificate expiration date in Data Protector GUI	25

Adding a Client to the Security Exceptions List	26
Steps	26
Introduction to User Authentication and LDAP	26
Initializing and Configuring the LDAP Login Module	27
Initializing the LDAP Login Module	27
Configuring the LDAP Login Module	29
Granting Data Protector Permissions to LDAP Users or Groups	31
Adding LDAP Users to Data Protector User Groups	32
Adding LDAP Groups to Data Protector User Groups	32
Logging In using LDAP Credentials	32
Checking the LDAP Configuration	33
Certificate Generation Utility	33
Introduction to the Certificate Generation Utility	33
Directory Structure for the Certificate Generation Utility	34
Overwriting Existing Certificates	36
Overwriting Existing Certificates	36
Overwriting Certificates in Existing Keystore and Truststore Files	36
Replacing Existing Server and Client Store Files	36
Replacing the CA Certificate	37
Updating the Distinguished Name (DN) String	37
Overwriting Certificates by Creating New Keystore and Truststore Files	38
Replacing Existing Server and Client Store Files	38
Replacing the CA Certificate	39
Updating the Distinguished Name (DN) String	39
Updating the Configuration File with the Stores Password	39
Firewall Support	40
About Firewall Support	40
Communication in Data Protector	40
Configuration mechanism	41
From Data Protector 9.09 and later	42
Limiting a Port Range in Data Protector prior to 9.09	42
For all Data Protector processes	43
For a specific Data Protector agent	43
For Data Protector processes and a specific Data Protector agent together	44
Port Usage in Data Protector 9.09 and later	44
Limitations	46
Disk, Media and Application Agent in the DMZ	46
Configuration figure	47
Open ports	47
Limitations	48
Disk and Application Agents in DMZ	48
Configuration figure	49
Open ports	49
Limitations	50
Cell Manager, Disk Agent, and Media Agent in the DMZ	50
Configuration figure	51

Open ports	51
Limitations	52
Chapter 3: Users and User Groups	54
About User Management	54
Users	54
UNIX	54
Windows	54
Predefined users	54
User Groups	56
Predefined user groups	56
Available User Rights	57
Chapter 4: Internal Database	58
About the IDB	58
What is the IDB used for?	58
IDB size and growth consideration	58
Regular IDB backups	58
IDB Architecture	58
IDB parts	59
Media Management Database (MMDB)	59
MMDB records	59
MMDB size and growth	60
MMDB location	60
Catalog Database (CDB)	60
CDB records	60
CDB (objects and positions) size and growth	60
CDB location	60
Detail Catalog Binary Files (DCBF)	61
DCBF information	61
DCBF size and growth	61
DCBF location	61
Session Messages Binary Files (SMBF)	61
SMBF records	61
SMBF size and growth	62
SMBF location	62
Encryption keystore and catalog files	62
Keystore location	62
Catalog file location	62
IDB Operation	63
Backup	63
IDB backup and archived log files	63
Restore	63

Object copy and object consolidation	64
Object verification	64
Exporting media	64
Removing the Detail Catalog	64
Maintenance of DC Directories	64
Checking the IDB Size	65
Steps	65
Reducing the IDB Growth	66
Reducing logging level	66
Steps	66
Reducing catalog protection	66
Steps	67
Reducing the IDB Current Size	67
Changing catalog protection for a session	67
Steps	67
Changing catalog protection for an object	68
Steps	68
Extending the IDB Size	68
Reconfiguring DC directories for higher capacity	68
Steps	68
IDB Consistency Check	68
Moving the IDB to a Different Cell Manager	69
Steps	69
Steps	70
Customizing the Data Protector Global Options	71
Prerequisites	71
Setting the global options using GUI	71
Steps	71
Customizing Options By Editing The Global File	72
Steps	72
Configuration of IDB Reports	72
IDB reports	72
Configuration of IDB Notifications	72
IDB notifications	72
Restoring the IDB	73
Restoring the IDB	73
Prerequisites	73
Limitations	73
Steps	73
Preparing for IDB restore from an encrypted backup	74
Steps	75
About IDB Recovery	75
Complete recovery (restore and update the IDB beyond the last IDB backup)	75
Overview of IDB Recovery Methods	75
The most convenient complete recovery	76
Omitting (removing) corrupted IDB parts	76

More recovery methods	76
IDB Corruption Levels	77
Identifying the Level of IDB Corruption	77
Steps	77
Performing Guided Autorecovery (IDB Restore and Replay Archived Log Files)	78
Prerequisites	78
Steps	79
Handling Minor IDB Corruption in the DCBF Part	79
Recovery if DC binary files are missing	79
Steps	79
Recovery if DC binary files are corrupted	80
Steps	80
Restoring the IDB Using IDB Recovery File and Changed Device	80
Prerequisites	80
Steps	81
Restoring the IDB Without IDB Recovery File	81
Prerequisites	81
Steps	82
Restoring the IDB from a Specific IDB Session	83
Prerequisites	83
Steps	84
Restoring the IDB database on a different Cell Manager host	84
Updating IDB by Importing Media	86
Steps	86
 Chapter 5: Manager-of-Managers Environment	 88
About MoM Environment	88
About CMMDB	88
How media are shared	88
How media are initialized	89
MoM Environment Configuration Procedure	89
Prerequisites	89
MoM environment configuration procedure	89
Setting Up MoM Manager	89
Steps	90
Adding a MoM Administrator to Cells	90
Prerequisite	90
Steps	90
Importing Cells	90
Prerequisites	90
Steps	90
Restarting the Data Protector Services in MoM	91
Stopping the Data Protector services	91
Cell Manager in a non-cluster environment	91
Cell Manager on HPE Serviceguard	91

Cell Manager on Symantec Veritas Cluster Server	91
Cell Manager on Microsoft Cluster Server	91
Starting the Data Protector services	91
Cell Manager in a non-cluster environment	91
Cell Manager on HPE Serviceguard	91
Cell Manager on Symantec Veritas Cluster Server	91
Cell Manager on Microsoft Cluster Server	92
Configuring CMMDB	92
Consideration	92
Prerequisites	92
Configuring CMMDB on a client cell	92
Steps	92
Configuring CMMDB on the MoM Manager	93
Steps	93
About Centralized Licensing	93
Setting Up Centralized Licensing	94
Prerequisite	94
Steps	94
Deactivating Centralized Licensing	95
Steps	95
About MoM Environment Administration	96
Exporting Cells	96
Steps	96
Moving Client Systems Among Cells	96
Steps	97
Deactivating Centralized Licensing	97
Prerequisites	97
Steps	97
Configuring Data Protector Users	97
Steps	97
Adding a User to Other Cells	98
Steps	98
Removing a User from Cells	98
Steps	98
Managing Devices and Media for a Specific Cell	98
Steps	98
Managing Internal Database for a Specific Cell	99
Steps	99
 Chapter 6: Clustering	 100
About Clustering	100
About the Data Protector Microsoft Cluster Server Integration	100
Licensing and MSCS	100
Configuration	100
How to Manage Cluster-Aware Backups	101

Failover of Data Protector	101
Failover of application other than Data Protector	101
About Disaster Recovery of a Microsoft Cluster Server	102
Possible scenarios	102
About the Data Protector HPE Serviceguard Integration	102
Licensing and HPE Serviceguard	102
Configuration	103
About the Data Protector HACMP Cluster Integration	103
Nodes	104
Shared external disk interfaces	104
Networks	105
Clients	105
Tasks	105
 Chapter 7: Devices	 106
About Backup Devices	106
What is a backup device?	106
About Configuring backup devices	106
Types of Backup Devices	106
Standalone	107
Backup to Disk device	107
SCSI library	107
Stacker	108
Magazine device	109
Jukebox	109
Standalone file device	109
File library device	109
External control	109
ADIC/GRAU DAS library	109
StorageTek ACS library	111
About Cloud (Helion) Devices	112
Prerequisites	113
Limitations	113
Recommendations	114
Preparing for the Cloud (Helion) device	114
About Cloud (Azure) Devices	115
Prerequisites	115
Limitations	115
Recommendations	116
Preparing for the Cloud (Azure) device	116
Device Performance Tuning	117
Block size	117
Determining the optimal block size	117
Limitations	118

Changing the block size	118
Device Performance	118
Support of New Devices	119
Preparing Backup Devices	119
Prerequisite	119
Steps	120
In the SAN Environment	120
Steps	120
File devices	121
Steps	121
Magazine	121
Steps	121
SCSI library, Jukebox, External Control	121
Steps	121
Windows robotics drivers	121
Steps	121
Creating SCSI Addresses on Windows Systems	121
Magneto-optical device	122
Tape device	122
Windows without the native tape driver	122
Windows using the native tape driver	122
Steps	122
Finding Device Filenames on UNIX System	122
Finding Device Filenames on HP-UX	123
Prerequisite	123
Steps	123
Finding Device Filenames on Solaris	123
Steps	123
Creating Device Files on UNIX Systems	123
Creating Device Files on HP-UX Systems	124
Prerequisites	124
Steps	124
Creating Device Files on Solaris Systems	124
Prerequisites	124
Steps	124
Auto-Detecting Device Filenames and SCSI Addresses	125
For an existing Data Protector device definition	125
Steps	125
While creating a Data Protector device definition	125
Steps	125
Auto-Detecting Device Filenames and SCSI Addresses for Libraries	126
For an already configured library	126
Steps	126
While configuring a library	126
Steps	126
About Configuring Backup Devices	126

About Library Management Console	127
What is a library management console?	127
Library management console support in Data Protector	127
Limitation	127
Autoconfiguring a Backup Device	127
Prerequisite	128
Device autoconfiguration	128
Steps	128
Device autoconfiguration in a SAN environment	128
Limitations	129
Steps	129
Configuring a Standalone Device	130
Steps	130
Configuring Backup to Disk Devices	131
Multi-Interface Support	131
Steps	131
Configuring a Backup to Disk Device - StoreOnce	132
Steps	132
Refreshing Cache for Stores	134
Refreshing cache using the Data Protector GUI	134
Refreshing cache using the Data Protector CLI	134
Configuring a Backup to Disk Device - Smart Cache	135
Configuring Smart Cache	135
Prerequisites	135
Limitations	136
Steps	136
Configuring a Backup to Disk Device - Data Domain Boost	137
Prerequisites	137
Limitations	137
Steps	137
Configuring Data Domain Boost on AIX Systems	139
Steps	139
Configuring a Backup to Disk Device - StoreOnce Software	139
Configuring the root directory of the deduplication stores	139
Creating a store	140
Configuring Cloud (Helion) Devices	141
Obtaining the HPE Public Cloud Project Name	141
Steps	141
Obtaining the Authentication Service URL	142
Steps	142
Creating the access keys	143
Steps	143
Configuring a Backup to Disk Device - Cloud (Helion)	143
Steps	143
Configuring a Backup to Disk Device - Cloud (Azure)	144
Steps	144

Configuring a File Library Device	144
Prerequisites	145
Limitations	145
Steps	145
About Configuring Multiple Paths to Devices	146
Why use multiple paths	146
Path selection	146
Backward compatibility	148
Limitations	148
Setting Advanced Options for Devices and Media	148
Steps	148
Configuring a VTL Device	148
Steps	148
Configuring a Stacker Device	149
Steps	149
Stacker device media management	150
Configuring a Jukebox Device (Optical Library)	150
Configuring a jukebox device	150
Steps	150
Configuring a drive in the jukebox device	151
Steps	151
Configuring a SCSI Library or a Magazine Device	151
Configuring a SCSI library robotics	151
Steps	151
Configuring a drive in a library	152
Steps	152
Configuring Devices in a SAN Environment	153
Considerations	153
Configuration Methods	153
Automatic device configuration using the GUI	153
Limitations	154
Automatic device configuration using the CLI (the sanconf command)	154
Device locking	155
Limitations	155
Recommendation	155
Manual configuration on UNIX systems	155
Phases	156
Configuring Devices in a SAN Environment Manually	156
Prerequisite	156
Configuration phases	156
Configuring a library in the SAN environment	156
Steps	156
Configuring a drive in a library	157
Steps	157
Configuring the libtab File in the SAN Environment	158
Steps	158

Configuring an ADIC/GRAU DAS Library Device	159
Configuration phases	159
Connecting library drives	159
Steps	159
Preparing for installation of a Media Agent	160
Steps	160
Installing a Media Agent	161
Prerequisites	161
Steps	162
Configuring the ADIC/GRAU DAS library device	162
Steps	163
Configuring a drive in the ADIC/GRAU DAS library device	163
Steps	163
Configuring a StorageTek ACS Library Device	164
Configuration phases	164
Connecting library drives	164
Steps	164
Installing a Media Agent	165
Prerequisites	165
Steps	166
Configuring the StorageTek ACS library device	166
Steps	167
Configuring a drive in the StorageTek ACS library device	167
Steps	167
About Using Backup Devices	168
Devices & Media Advanced Options	168
Advanced options - Settings	168
Options	168
Advanced options - Sizes	169
Advanced options - Other	169
Mount request	169
Device lock name	169
Library with Several Drive Types	169
Same density setting	169
Different media pool for each drive type	169
Free pool support	170
About Scanning	170
When to use scanning	170
Limitations	170
Drive Cleaning	171
Limitations	171
Conditions for automatic cleaning	171
Scheduled Eject of Media	172
Device Locking	172
Disabling a Backup Device	173
Disabling a backup device manually	173

Steps	173
Disabling a backup device automatically	173
Renaming a Backup Device	174
Steps	174
Removing a Backup Device	174
Steps	174
Responding to Mount Requests	174
Prerequisites	175
Steps	175
About Storage Area Network (SAN)	175
What is SAN?	175
FC-AL and LIP	175
Device Locking in the SAN Environment	176
Locking devices used exclusively by Data Protector	177
Locking devices used by multiple applications	177
Indirect and Direct Library Access	177
Indirect library access	177
Direct library access	177
Configuring Devices in a SAN Environment	178
Considerations	178
Configuration Methods	178
Automatic device configuration using the GUI	178
Limitations	178
Automatic device configuration using the CLI (the sanconf command)	179
Device locking	180
Limitations	180
Recommendation	180
Manual configuration on UNIX systems	180
Phases	180
About Backup to Disk	181
What is a disk-based backup device?	181
How to configure disk-based devices?	181
About Backup to Disk Devices	181
About Deduplication	182
When to use deduplication	183
Advantages of deduplication	183
Deduplication technologies	183
StoreOnce software deduplication	183
HPE StoreOnce Backup system devices	183
Deduplication setup	184
Source-side deduplication	184
Server-side deduplication	184
Target-side deduplication	184
About File Library Devices	185
How to maintain disk-based devices?	185
File Depots	185

File depot creation	185
File depot name	185
File depot size	186
File depot space consumption	186
Disk full handling	186
Number of devices per disk	186
Setting File Library Device Properties	186
Initial property setup	186
Steps	186
Changing device properties	187
Steps	187
Deleting File Library Devices	187
Deletion phases	187
Checking data protection	187
Steps	187
Recycling file depots	187
Steps	188
Deleting the exported file depot icon	188
Steps	188
Deleting the file library device	188
Steps	188
About Jukebox Devices	188
Jukebox Physical Devices	188
Jukebox File Devices	189
Recommended slot sizes for Windows and UNIX	189
How to maintain file jukebox devices?	189
Configuring a File Jukebox Device	189
Configuring a file jukebox device	190
Prerequisites	190
Steps	190
Configuring a drive in the file jukebox device	190
Steps	190
Recycling a File Jukebox Slot	191
Steps	191
About Standalone Devices	191
Standalone Physical Devices	191
Standalone File Devices	191
Configuring a Standalone File Device	192
Prerequisites	192
Steps	192
 Chapter 8: Media	 194
About Media Management	194
Customizing the Devices and Media View	194
About Media Pools	194

Free pools	195
Default media pool	195
Free Pool Characteristics	195
Free Pool Properties	195
When Is a Free Pool Used?	195
Media Quality Calculation	195
Free Pool Limitations	196
Media Pool Properties	196
Media pool properties - General	196
Media pool properties - Allocation	196
Allocation	196
Media pool properties - Condition	197
Media condition factors	197
Media pool properties - Usage	197
Media Pool Quality	197
Device error and media quality	198
Creating a Media Pool	198
Steps	198
Modifying a Media Pool	198
Steps	199
Deleting a Media Pool	199
Steps	199
Media Life Cycle	199
Preparing media for backups	200
Using media for backups	200
Vaulting media to a safe place	200
Retiring media	200
Media Types	200
Supported media types	201
Media Quality	201
Device error and media quality	201
How Media Are Selected for Backup	201
Media allocation policy	202
Preallocating media	202
Media condition	202
Media usage	202
Limitation	202
Media selection factors	203
Use of Different Media Format Types	203
Limitations	203
WORM Media	204
How to use WORM media with Data Protector	204
Supported WORM media	204
About Formatting Media	204
Formatting with padding blocks	204
When to format media	204

Media label	205
Recognized Media Formats	205
Data Protector media format categories	205
Formatting a Medium	206
Steps	206
Formatting All Media in a Magazine	207
Prerequisite	207
Steps	207
Formatting a Single Medium in a Magazine	207
Prerequisite	207
Steps	207
Formatting Media in a Library Device	208
Steps	208
About Importing Media	208
Considerations	208
When to import media?	209
Importing a Medium	209
Steps	209
Importing All Media in a Magazine	209
Prerequisite	209
Steps	209
Importing a Single Medium in a Magazine	210
Prerequisite	210
Steps	210
Importing Media in a Library Device	210
Steps	210
Exporting and Importing Media with Encrypted Backups	211
Cell Manager environment or MoM environment without CMMDB	211
Steps	211
MoM environment with CMMDB	211
Steps	212
About Media Copying	212
Prerequisites	212
Limitations	212
When to copy media	213
Results of Copying Media	213
Restoring from a copy	213
Copying a Medium	213
Copying a medium in a standalone device	213
Steps	213
Copying a medium in a library device	214
Automated Media Copying	214
Limitations	214
Automated media copying	215
Types of automated media copying	215
Post-backup media copying	215

Scheduled media copying	215
Configuring Post-Backup Media Copying	215
Limitations	216
Steps	216
Configuring Scheduled Media Copying	216
Limitations	216
Steps	216
Scheduling Media Copying on Specific Dates	217
Steps	217
Scheduling Periodic Media Copying	217
Steps	217
Disabling and Enabling an AMC Schedule	218
Steps	218
Disabling and Enabling AMC on Holidays	218
Steps	218
Resetting an AMC Schedule	219
Steps	219
Scanning a Device	219
Steps	219
Scanning Media in a Library Device	219
Steps	219
Scanning a Drive in a Library Device	220
Steps	220
Activating Barcode Reader Support	220
Steps	220
Barcode Scanning of a Library Device	221
Prerequisite	221
Steps	221
Searching and Selecting Media	221
Searching and selecting media in a media pool	221
Steps	221
Searching and selecting media in a library device	222
Steps	222
Searching for media using the List of Media report	222
Steps	222
Pre-allocation List of Media for Backup	222
Preallocating Media for Backup	222
Steps	223
Recycling a Medium	223
Steps	223
Importing the Catalog from Media	224
Steps	224
Verifying a Medium	224
Verifying a medium in a standalone device	225
Steps	225
Verifying a medium in a library device	225

Steps	225
Moving a Medium	225
Steps	225
Exporting a Medium	226
Steps	226
Copying the Catalog Media Data to the MCF File	226
Limitations	226
Recommendations	227
Steps	227
Importing the Catalog Media Data from the MCF Files	227
Prerequisites	227
Limitations	227
Steps	228
Modifying Media Description	228
Steps	228
Modifying Media Location	228
Steps	229
Creating a List of Locations	229
Steps	229
Setting the Media Location Priority	229
Steps	230
Vaulting a Medium	230
Prerequisites	230
Steps	230
Erasing a Medium	230
Steps	230
Detection of Write-Protected Media	231
About Mount Requests	231
About Library-Specific Media Management	231
The use of library media by other applications	232
About the Data Protector Query Operation Used with ADIC/GRAU DAS or STK ACS	
Libraries	232
Adding a Slot	233
Steps	233
Deleting a Slot	233
Steps	233
Entering a Medium	234
Steps	234
Ejecting a Medium	234
Bulk eject of media	234
Predefined eject of media	235
Steps	235
Erasing Media in a Library Device	235
Steps	235
Adding Volsers Manually	235
Steps	236

Querying the ADIC/GRAU DAS and StorageTek ACSLM Hosts	236
Limitation	236
Steps	236
Chapter 9: Backup	237
About Backup	237
Setting the Backup View	237
Steps	237
Full and Incremental Backups	238
Conventional Incremental Backup	238
How conventional incremental backup works	238
Detection of changes	238
Enhanced Incremental Backup	239
Why use enhanced incremental backup	240
Impact on disk space consumption	240
Limitations	240
Incremental Backup Using Change Log Provider	240
Prerequisites	241
Performance and Disk Space Consumption	241
Considerations	242
Limitations	243
Synthetic Backup	243
How to perform synthetic backup	243
Virtual full backup	243
Standard Backup Procedure	243
Prerequisites	244
Filesystem backup	244
Creating a Backup Specification	244
Limitations	244
Steps	245
Modifying a Backup Specification	246
Steps	246
Previewing and Starting a Backup	246
Limitations	246
Steps	247
Aborting a Backup	247
Steps	247
Restarting Failed Backups	247
Prerequisite	247
Considerations	248
Limitations	248
Steps	248
Copying a Backup Specification	248
Steps	248
Deleting a Backup Specification	249

Steps	249
Advanced Backup Tasks	249
Prerequisites	249
What are advanced backup tasks?	249
Selecting Network Shared Disk for Backup	250
Prerequisite	250
Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012	250
Requirements	250
Limitations	250
Steps	251
Selecting Only Specific Files (Matching) for Backup	252
Steps	252
Skipping Files for Backup	252
Steps	252
Selecting the Location for the Shortcut for Starting a Backup	253
Limitations	253
Steps	253
Backing Up Using Multiple Disk Agents	253
Steps	254
Handling of Small Reoccurring Backups	254
Disk Image Backup	255
When to use a disk image backup?	255
How to specify a disk image section?	255
On UNIX systems	255
On Windows systems	256
Where to find a disk image section?	256
On UNIX systems	256
On Windows systems	256
Client Backup with Disk Discovery	257
When to use disk discovery	258
Backup specification	258
Web Server Backup	258
Enabling Wake ONLAN Support	258
Steps	259
About Backup Templates	259
Creating a New Backup Template	260
Steps	260
Modifying a Backup Template	260
Steps	261
Copying a Backup Template	261
Steps	261
Deleting a Backup Template	261
Steps	261
Applying a Backup Template to a Backup Specification	262
Steps	262

About Backup Options	262
Available backup options	263
Backup specification options	263
Filesystem options	263
Disk image options	263
Device options	264
Schedule options	264
Most Frequently Used Options	264
Interactive backups	264
Backups using a saved backup specification	264
Scheduled backups	265
Expired catalog protection	265
Catalog protection and backup	265
Catalog protection and restore	265
Logging level and backup speed	266
Logging level and browsing for restore	266
Logging level and restore speed	266
Who is a backup session owner?	267
Why change the backup owner?	268
Who can restore a private object?	268
Backup Specification Options	268
General backup specification options	268
Clustering backup specification options	269
Automatic session restart	269
Abort session and abort ID parameters	269
EMC Symmetrix backup specification options	269
Client systems	269
Mirror type	269
EMC Symmetrix split pre-exec and post-exec	269
EMC Symmetrix options	270
HPE P9000 XP Disk Array Family backup specification options	270
Client systems	270
Mirror type	270
Replica management options	270
At the start of the session	270
At the end of the session	270
Application system options	270
Backup system options	271
HPE P6000 EVA Disk Array Family backup specification options	271
Client systems	271
Replication mode	271
Replica handling during failover scenarios	271
Snapshot management options	271
Mirrorclone preparation / synchronization	271
Replica management options	271
Application system options	272

Backup system options	272
Filesystem Options	272
Filesystem options	272
Other filesystem options	272
WinFS filesystem options	273
Disk Image Options	274
Device Options	274
Device properties - General	274
Schedule Options	275
Session options	275
Split mirror/snapshot backup	275
Setting Backup Options	275
Steps	276
Specifying Data Protection	276
Specifying data protection on the backup specification level	276
Steps	276
Specifying data protection for individual backup objects	276
Steps	277
Specifying data protection for scheduled backups	277
Specifying data protection using the CLI	277
Steps	277
Changing Options for a Specific Object	277
Steps	278
Changing Backup Device Options	278
Steps	278
Setting Schedule Backup Options	279
Steps	279
About Pre- and Post-Exec Commands	280
What are pre- and post-exec commands?	280
Configuring pre- and post-exec commands for backup	280
Backup specification	280
Backup object	280
How are pre- and post-exec commands run?	280
Pre- and Post-Exec Commands for a Backup Specification	281
Pre- and Post-exec characteristics	281
Start-up and location of the commands	281
Windows systems	281
UNIX systems	281
Environment variables	282
SMEXIT values	282
Considerations for pre- and post-exec commands	282
Specifying Pre- and Post-Exec Commands for a Backup Specification	283
Pre- and Post-exec Commands for a Specific Backup Object	284
Start-up and location of the commands	284
Environment variable	285
Considerations for pre- and post-exec commands	285

Security considerations	286
Specifying Pre- and Post-Exec Commands for Backup Objects	286
Specifying pre- and post-exec commands for all objects	286
Specifying pre- and post-exec commands for individual objects	286
Specifying pre- and post-exec commands for integrations	287
About Backup Schedule	287
Scheduling and priority (Advanced Scheduler)	287
Scheduling options	289
Scheduling and different time zones	289
Scheduling tips	289
Backing up during holidays (basic scheduler only)	289
Handling scheduling conflicts (basic scheduler only)	290
Scheduling a Backup on a Specific Date and Time	290
Steps	290
Scheduling a Periodic Backup	291
Using a predefined backup schedule	291
Steps	291
Configuring a recurring backup	292
Steps	292
Running Consecutive Backups	292
Steps	292
Resetting a Backup Schedule	293
Steps	293
Disabling and Enabling a Backup Schedule	293
Steps (Basic Scheduler)	293
Steps (Advanced Scheduler)	294
Disabling and Enabling Backups on Holidays	294
Steps	294
Customizing the Schedule Calendar	294
Steps	294
About Backup Specification Groups	295
Example of backup specification groups	295
Viewing Backup Specification Groups	295
Steps	295
Creating a Backup Specification Group	296
Steps	296
Saving a Backup Specification into a Group	296
Steps	296
Moving a Backup Specification or Template Among Groups	296
Steps	297
Deleting a Backup Specification Group	297
Steps	297
About Windows Systems Backup	297
Limitation	297
What is backed up?	298
Windows Server 2012	298

Windows-specific information	298
What is not backed up?	298
Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012:	298
Windows Server 2012	298
Other Windows systems	299
NTFS 3.1 filesystem features	299
Reparse points	299
Sparse files	300
Warnings when backing up system disks	300
Configuration Backup (Windows)	300
Limitations	301
Windows configuration objects	301
Active Directory	301
DFS	301
DHCP and WINS	302
Profiles	302
Removable Storage Management Database	302
Terminal Service Database	302
Windows services	302
System State Data Backup	303
Remote Storage Service	303
Remote Storage Services:	304
Remote Storage databases:	304
Removable Storage Management Database	304
System File Protection	304
About UNIX Systems Backup	304
Limitations	304
What is backed up?	305
What should be excluded from a UNIX filesystem backup?	305
NFS Backup	305
When to use NFS backup?	305
Limitations	305
Prerequisites	306
Limitations	306
What is backed up?	307
About Novell Open Enterprise Server (OES) Backup	307
Prerequisites	307
Limitations	308
Backup and restore of compressed files	308
What is backed up?	308
Configuring Novell OES	308
Saving the username and password using the HPLOGIN utility	308
Steps	308
Loading Target Service Agent for File Systems (tsafs) in dual mode	308
Steps	308

Loading the Target Service Agent for Novell Directory Services (tsands)	309
Steps	309
Loading the GroupWise Target Service Agent for File Systems (tsafsgw)	309
Steps	310
About Backup Performance	310
Infrastructure	310
Object mirroring and backup performance	311
High Performance Hardware Other than Devices	311
Hardware Parallelism	311
Concurrency	312
Performance impact	312
Multiple data streams	312
Device Streaming	312
How to configure device streaming	312
Block Size	313
Segment Size	313
Number of Disk Agent Buffers	314
Software Compression	314
Hardware Compression	314
Disk Image Versus Filesystem Backup	315
Object Distribution to Media	316
Filesystem Scan	316
Miscellaneous Performance Hints	316
 Chapter 10: Object Consolidation	 318
About Object Consolidation	318
Types of object consolidation	318
Post-backup object consolidation	318
Scheduled object consolidation	318
How to Consolidate Objects	318
Selection of devices	318
Object consolidation options	319
Selection of the media set	319
Ownership of consolidated objects	319
Standard Object Consolidation Tasks	319
Prerequisites	319
Limitations	320
Consolidating Objects Interactively	320
Steps	320
Configuring Post-Backup Object Consolidation	321
Steps	321
Scheduling of Object Consolidation	322
Steps	322

Copying an Object Consolidation Specification	322
Steps	323
Chapter 11: Copy	324
About Duplicating Backed Up Data	324
About Object Copying	325
What is object copy?	325
Automated object copying	326
Post-backup object copying	326
Scheduled object copying	326
How to Copy Objects	326
Selection of devices	326
Object copy options	327
Selecting the media set to copy from	327
Object copy completion status	327
Copy objects	327
Source objects	328
Ownership of object copies	328
Standard Object Copy Tasks	328
Prerequisites	328
Limitations	328
Copying Objects Interactively	329
Steps	329
Configuring Post-Backup Object Copying	330
Steps	330
Scheduling of Object Copying	331
Steps	331
Restarting Failed Object Copy Sessions	332
Prerequisites	332
Limitations	332
Steps	332
Copying an Object Copy Specification	333
Steps	333
Advanced Object Copy Tasks	333
Freeing a Medium	333
Steps	334
Demultiplexing a Medium	334
Limitation	334
Steps	335
Consolidating a Restore Chain	335
Limitation	336
Steps	336
Migrating to Another Media Type	336
Steps	336
About Disk Staging	337

What is disk staging?	337
Why implement disk staging	337
Disk staging and small reoccurring backups	338
Troubleshooting Object Operations Sessions	338
Object copy problems	338
Fewer objects are copied than expected	338
Not all objects in the selected library are copied	338
Mount request for additional media is issued	338
When creating an object copy, the protection end time is prolonged	339
Replicating session with multiple objects stops responding	339
Replication session on Data Domain Boost devices is unable to respond to Abort operation during retry period	340
Object consolidation problems	340
Object consolidation of many points in time opens too many files	340
Object consolidation to B2D devices fails in the second attempt	341
About Replication	341
Automated replication	341
Post-backup replication	342
Scheduled replication	342
Limitations	342
Considerations	342
How to enable replication	343
Automated Replication Synchronization	343
Prerequisites	343
Considerations	343
Limitations	343
Importing the foreign Cell Manager	344
Performing an Object Copy session	344
About Object Mirroring	345
Benefits of object mirroring	345
Limitations	346
How to use object mirroring	346
Copying a Medium	346
Copying a medium in a standalone device	346
Steps	346
Copying a medium in a library device	347
Scheduling Media Copying on Specific Dates	347
Steps	347
Scheduling Periodic Media Copying	348
Steps	348
Customizing the Schedule Calendar	349
Steps	349
 Chapter 12: Object Verification	 350
About Object Verification	350

Data verification	350
Delivery to host	350
Types of object verification session	350
Post-backup object verification	350
Scheduled object verification	350
How to Verify Objects	351
Selection of backup objects	351
Automated operation	351
Interactive operation	351
Selection of a source device	351
Selection of target host	351
Scheduling	351
Standard Object Verification Tasks	351
Prerequisites	352
Limitations	352
Verifying Objects Interactively	352
Steps	352
Configuring Post-Backup Object Verification	353
Steps	354
Configuring Scheduled Object Verification	354
Steps	355
Customizing the Object Verification Environment	356
 Chapter 13: Restore	 357
About Restore	357
Standard Restore Procedure	357
Prerequisite	357
Selecting the Data to Restore	357
Prerequisite	358
Selecting the data from the list of the backed up objects	358
Steps	358
Selecting the data from the list of the backup sessions	358
Limitations	358
Steps	358
Selecting a Specific Backup Version	359
Selecting the backup version for each file or directory separately	359
Steps	359
Selecting the backup version for several files or directories simultaneously	359
Steps	359
Handling File Conflicts	360
Steps	360
Selecting a Device to Restore From	360
Steps	360

Finding Media Needed to Restore	361
Limitations	361
Steps	361
Previewing and Starting a Restore	362
Prerequisites	362
Limitations	362
Steps	362
Aborting a Restore	362
Steps	363
Restore Location Options	363
Selecting Restore Location	363
Steps	363
Specifying Restore Location for Individual Files and Directories	363
Restore into	364
Steps	364
Restore as	364
Steps	364
About Resuming Failed Sessions	365
Filesystem backup sessions	365
Limitations	366
Filesystem restore sessions	366
How the functionality works	366
Considerations	366
Limitations	367
Data Protector Oracle Server integration backup and restore sessions	367
Resuming Failed Sessions	367
Prerequisites	367
Steps	368
Advanced Restore Tasks	368
Prerequisites	368
Advanced restore tasks	368
Skipping Files for Restore	369
Steps	369
Selecting Only Specific Files (Matching) for Restore	369
Steps	369
Selecting Open Files for Restore	370
Steps	370
Denying Access to Files During Restore	370
Steps	370
Searching for a File to Restore	370
Steps	370
Selecting a Windows Shared Disk for Restore	371
Prerequisite	371
Steps	371
Restoring Objects in Parallel	372
Prerequisite	372

Limitation	372
Steps	372
Disk Image Restore	373
Prerequisites	373
Restore from Media in a Vault	373
Web Server Restore	373
Restore Without Browsing	374
Restoring the Entire Object and Extracting the Needed Parts	374
Prerequisite	374
Steps	374
Restoring Parts of the Backed Up Object Using Restore-Only Pattern Match	374
Prerequisites	375
Steps	375
Restoring the File or Directory Manually	375
Prerequisite	375
Steps	376
Restore Options	376
General restore options	376
Pre- and post-exec commands	378
Device selection	378
Handling file conflicts	379
Active Directory specific options	379
Replication mode	379
Setting Restore Options	379
Steps	380
About Windows Systems Restore	380
NTFS 3.1 filesystem features	380
Restoring objects backed as shared disks	381
Windows Filesystem Restore Limitations	381
Configuration Restore	382
Limitations	382
Windows configuration objects	382
Active Directory	383
DFS	383
Profiles	384
Registry	384
Removable Storage Manager Database	384
Server configuration objects	385
SysVol	385
Windows TCP/IP services	385
System State Data Restore	385
Remote Storage Service	386
System File Protection	386
About UNIX Systems Restore	387
UNIX systems specific information	387
About HP OpenVMS System Restore	387

Limitations	387
Filesystem information restored	388

Chapter 14: Monitoring, Reporting, Notifications, and Data Protector Event

Log	389
About Monitoring	389
Viewing Currently Running Sessions	389
Prerequisite	389
Steps	389
Viewing Finished Sessions	390
Prerequisite	390
Steps	390
Aborting Running Sessions	390
Prerequisite	390
Steps	390
About Reporting	391
Features	391
Reports Formats	392
Reports Types	392
Configuration reports	392
Cell Information	392
Client Backup	393
Clients not Configured for Data Protector	393
Configured Clients not Used by Data Protector	393
Configured Devices not Used by Data Protector	394
Licensing	394
Look up Schedule	394
IDB report	394
IDB Size	395
Pools and media reports	395
Extended List of Media	395
List of Media	395
List of Pools	396
Media Statistics	396
Session specification reports	396
Average Backup Object Sizes	397
Filesystems Not Configured for Backup	397
Object's Latest Backup	397
Objects Without Backup	398
Session Specification Information	398
Session Specification Schedule	399
Trees in Backup Specifications	399
Sessions in timeframe reports	399
Client Statistics	399
Device Flow	400
Extended Report on Used Media	400

List of Sessions	400
Object Copies	401
Report on Used Media	401
Session Errors	401
Session Flow	402
Session Statistics	402
Single session reports	402
Session Devices	402
Session Media	403
Session Object Copies	403
Session Objects	403
Session per Client	404
Single Session	404
Reports Send Methods	404
Broadcast message send method	405
E-mail send method	405
On Windows systems	405
On UNIX systems	405
E-mail (SMTP) send method	405
On Windows systems	406
On UNIX systems	406
External send method	406
Log to file send method	406
SNMP send method	406
On Windows systems	407
On UNIX systems	407
Configuring Report Groups Using the Data Protector GUI	407
Prerequisites	407
Configuration phases	407
Configuring a report group	407
Steps	407
Adding a report to a report group	408
Steps	408
Running Report Groups Using the Data Protector GUI	408
Prerequisites	408
Steps	408
Running Individual Reports Using the Data Protector GUI	409
Prerequisites	409
Steps	409
Running Reports and Report Groups Using the Data Protector CLI	409
Prerequisites	409
Steps	409
Creating a New Mail Profile	410
Steps	410
Configuring Windows SNMP traps	410
Prerequisites	410

Steps	410
About Notifications	411
Notification Types - Events that Trigger Notifications	412
Alarm	412
Expired Certificates	412
Csa Start Session Failed	412
Device Error	413
End of Session	413
File Library Disk Usage	413
Health Check Failed	414
IDB Backup Needed	414
IDB Corrupted	414
IDB Limits	415
IDB Reorganization Needed	415
IDB Space Low	415
License Warning	416
License Will Expire	416
Mail Slots Full	416
Mount Request	417
Not Enough Free Media	417
Session Error	417
Start of Session	417
Too Many Sessions	418
Unexpected Events	418
Check UNIX Media Agent	418
User Check Failed	419
Notifications Send Methods	419
Broadcast Message send method	419
E-mail send method	419
On Windows systems	420
On UNIX systems	420
E-mail (SMTP) send method	420
External send method	420
Log to File send method	420
Data Protector Event Log send method	421
SNMP send method	421
On Windows systems	421
On UNIX systems	421
Use report group send method	421
Configuring Notifications	421
Prerequisite	421
Steps	421
About Web Reporting and Notifications	422
Requirements	422
Limitations	422
Configuring and Launching Web Reporting and Notifications Interface	423

Prerequisite	423
Steps	423
Configuring a Password for Web Reporting	423
Steps	424
Configuring Report Groups Using the Web Reporting Interface	424
Prerequisites	424
Steps	424
Running Individual Reports Using the Web Reporting Interface	425
Prerequisite	425
Steps	425
Running Saved Reports Using the Web Reporting Interface	425
Prerequisite	425
Steps	425
Configuring Notifications Using the Web Reporting Interface	426
Prerequisite	426
Steps	426
About Data Protector Event Log	426
Process-triggered events	427
User-triggered events	427
Accessing Event Log Viewer	427
Prerequisite	427
Steps	428
Deleting Event Log Viewer Contents	428
Prerequisite	428
Steps	428
About Auditing	428
Generating an Audit Report	428
Steps	428
 Chapter 15: Telemetry	 430
How Telemetry works?	430
Telemetry Registration page	431
 Data Protector Documentation	 433
Documentation map	433
Abbreviations	433
Integrations	436
 Send Documentation Feedback	 438

Chapter 1: Introduction

About Data Protector

HPE Data Protector is a backup solution that provides reliable data protection and high accessibility for your fast-growing business data. Data Protector offers comprehensive backup and restore functionality specifically tailored for enterprise-wide and distributed environments.

Major Data Protector features

- Scalable and highly flexible architecture
- Mixed environment support
- Easy central administration
- High performance backup
- Easy restore
- Data and control communication security
- High availability support
- Automated or unattended operation
- Monitoring, reporting, and notification
- Service management
- Integration with online database applications
- Integration with other products

Data Protector Architecture

Data Protector can be used in environments ranging from a single system to thousands of systems on several sites. The basic management unit is the Data Protector cell.

The Data Protector cell is a network environment consisting of a Cell Manager system, one or more Installation Servers, client systems, and devices.

The Cell Manager and Installation Server can be on the same system, which is the default option, or on separate systems.

Cell Manager

The Cell Manager is the main system that controls the Data Protector cell from a central point, where the Data Protector core software with the IDB is installed. The Cell Manager runs Session Managers that control backup and restore sessions and write session information to the IDB. The IDB keeps track of the backed up files as well as of the configuration of the Data Protector cell.

Installation Server

The Installation Server is the computer where the Data Protector software repository is stored. You need at least one Installation Server for UNIX and one for the Windows environment so that you can perform remote installations through the network and distribute the software components to the client systems in the cell.

Client Systems

After installing Data Protector software on the Cell Manager system, you can install Data Protector components on every system in the cell. These systems become Data Protector clients. The role of a client depends on the Data Protector software you have installed on this system.

Systems to be backed up

Client systems you want to back up must have the Data Protector Disk Agent (DA also called backup agent) installed. The Disk Agent reads or writes data from a disk on the system and sends or receives data from a Media Agent. The Disk Agent is also installed on the Cell Manager, allowing you to back up data on the Cell Manager, the Data Protector configuration, and the IDB.

Systems with backup devices

Client systems with connected backup devices must have a Data Protector Media Agent (MA) installed. A Media Agent reads or writes data from media in the device and sends or receives data from the Disk Agent. A backup device can be connected to any system and not only to the Cell Manager. Client systems with backup devices are also called Drive Servers. A client system with several backup devices is called a multi-drive server.

Overview of Tasks to Set Up Data Protector

Although configuring Data Protector is easy, some advanced planning will help you configure the environment and optimize your backups. This section provides an overview of the global tasks to set up a backup environment.

Depending on the size and complexity of your environment, you may not need to go through all these steps.

Steps

1. Analyze your network and organizational structure. Decide which systems need to be backed up. For information, see the *HPE Data Protector Concepts Guide*.
2. Check whether there are any special applications and databases which you want to back up, such as Microsoft Exchange Server, Microsoft SQL Server, Oracle Server, SAP R/3, or others. Data Protector provides specific integrations with these products.

On how to configure the integrations, see the *HPE Data Protector Integration Guides*.

3. Decide on the configuration of your Data Protector cell, such as:
 - The system to be your Cell Manager
 - Systems on which you want to install the user interface
 - Local backup versus network backup
 - Systems to control backup devices and libraries
 - Type of connection, LAN and/or SAN
4. Purchase the required Data Protector licenses for your setup. This way you obtain the passwords you will need to install.

Alternatively, you can operate Data Protector using an instant-on password. However, this is valid only for 60 days from the date of installation. See the *HPE Data Protector Installation Guide*.

5. Consider security aspects:
 - Analyze security considerations. See the *HPE Data Protector Installation Guide*.
 - Consider which user groups you need to configure.
 - Enhance security by writing data to media in an encrypted format.
 - Help preventing unauthorized access by enabling encrypted control communication.
6. Decide how you want to structure your backups:
 - Which media pools would you like to have, and how will they be used?
 - Which devices will be used, and how?
 - How many copies of each backup do you want?
 - How many backup specifications do you want to have, and how should they be grouped?
 - If you are planning to back up to disk, consider advanced backup strategies such as synthetic backup and disk staging.
7. Install the Data Protector Cell Manager and Installation Server(s). Then use the Data Protector GUI to distribute Data Protector agents to other systems. For information, see the *HPE Data Protector Installation Guide*.
8. [Configure backup devices](#).
9. [Configure media pools](#) and prepare the media.
10. [Configure backup specifications](#), including backup of the IDB.
11. Configure reports, if required.
12. Prepare for disaster recovery. For more information on disaster recovery, see the *HPE Data Protector Disaster Recovery Guide*.
13. Become familiar with tasks such as:

- Handling failed backups
- Performing restores
- Duplicating backed up data and vaulting media
- Testing disaster recovery
- Maintaining the IDB

Data Protector Operation

Backup and restore tasks are completed within sessions. Several sessions can run at the same time. The maximum number of sessions is limited by resources in the cell, such as the configuration of the Cell Manager (processor speed, main memory size, disk space).

Backup session

A backup session is a process that backs up data from a client system to media. A backup session always runs on the Cell Manager system. A backup session is based on a backup specification and is started either interactively by an operator or unattended by the Data Protector Scheduler.

Restore session

A restore session is a process that restores data from previous backups to a disk. The restore session is interactive and started by an operator using the Data Protector user interface.

Pre-exec and post-exec commands

Pre-exec commands let you execute some actions before a backup or a restore session. Post-exec commands let you execute some actions after a backup or a restore session.

The pre-exec and post-exec commands can be set for a backup specification and, as such, executed on the Cell Manager system or they can be specified as a backup object option and be executed on the client system where the respective Disk Agent is running.

Pre-exec and post-exec script commands can be written as executables or batch files (on Windows systems) or shell scripts (on UNIX systems). These are not supplied by Data Protector and must be written separately (by the backup operator, for example).

Object copy, object consolidation and object verification sessions

An object copy session is based on an object copy specification. An object consolidation session is based on an object consolidation specification. Both sessions can be started interactively or

automatically.

An object verification session is based on an object verification specification. It checks the data integrity of objects created by backup, object copy or object consolidation sessions and the ability to deliver them to the required location. Sessions can be started interactively or automatically.

Chapter 2: Configuration Tasks

Enabling Security

This section describes the security elements of Data Protector. It describes the advanced settings that can be used to enhance the security of Data Protector with prerequisites and considerations that have to be taken into account.

Since enhancing security in an entire environment requires additional effort, many security features cannot be enabled by default.

The considerations described in this chapter apply not only when the security settings are changed, but must also be followed when configuring new users, adding clients, configuring application agents, or making any other changes these considerations apply to. Any changes in the security settings can have cell-wide implications and should be carefully planned.

About Security Considerations

For detailed information on security considerations with Data Protector cell components, see *HPE Data Protector Installation Guide*.

Cell Manager Security

The Cell Manager security is important because the Cell Manager has access to all clients and all data in the cell.

Security of the Cell Manager can be enhanced via the *Strict IP Checking* functionality. However, it is important that the Cell Manager is also secured as a client and that Data Protector users are configured carefully.

While it may not always be necessary to secure each and every client in the cell, it is important that the computers that other clients will trust are secured themselves. These are besides the Cell Manager also the Installation Server and Media Agent clients.

Security of a Cell Manager and subsequently all clients in the Data Protector cell can be additionally enhanced by enabling encrypted control communication.

Client Security

After you have installed the Data Protector clients and imported them to a cell, it is highly recommended to secure them.

Data Protector agents installed on the clients in the cell provide numerous powerful capabilities, like access to all the data on the system. It is important that these capabilities are available only to the processes running on cell authorities (Cell Manager and Installation Servers), and that all other requests are rejected.

Data Protector allows you to specify from which cell authorities a client will accept requests on the Data Protector port (default 5555). For activities such as backing up and restoring, starting pre- and post-exec commands, or importing and exporting clients, the client checks if the computer, which triggers one of these tasks via the Data Protector port, is allowed to do so. Other computers are not able to access such a client.

Trusted clients

Before securing clients, it is important to determine a list of trusted clients. This list must include:

- Cell Manager
- Relevant Installation Servers
- For some clients also a list of clients that will access the robotics remotely.

The list must contain all possible client names (or IP addresses) where connections can come from. Multiple client names may be needed if any of the above clients is multihomed (has multiple network adapters and/or multiple IP addresses) or is a cluster. The list should include:

- All additional client names (for all LAN cards) of the cell authority.
- All cluster nodes names where the Cell Manager might failover, as well as a cluster virtual server name.
- The target system name to which the cell authority will be moved in case of a total hardware failure of the cell authority. This target system has to be defined in the disaster recovery strategy.
- For clients that are allowed to access a client that controls the robotics of a library, all clients that use the drives of that library.

If the DNS configuration in the cell is not uniform, additional considerations may apply.

User interface clients do not need to be added to the list of trusted clients. Depending on the user rights, you can use the GUI to access either the complete Data Protector functionality or the specific contexts only.

Note: If an Installation Server residing on a system other than the Cell Manager is not added to the list of allowed clients, it will not have access to a secured client. In this case, the operations dependent on the Installation Server (such as checking installation, adding components and removing clients) will fail. If you want these operations to be available on the secured client, add the Installation Server to the list of allowed clients.

The `allow_hosts` and `deny_hosts` files

When you secure a client, the names of the systems allowed to access a client are written to the `allow_hosts` file. You can also explicitly deny access to a client from certain computers by adding their names to the `deny_hosts` file, located in the default Data Protector client configuration directory.

If you accidentally lock out a client, you can manually edit or delete the `allow_hosts` file on this client.

Specify each client name in a separate line.

On Windows systems, the files are in double-byte format (Unicode), whereas on UNIX systems the files are in single-byte format or multi-byte format (for example, Shift-JIS).

You can allow or deny access to all systems with Data Protector installed. For example, you can allow or deny the access of Cell Managers to clients, Cell Managers to Cell Managers, or clients to clients.

Users Security

Data Protector users is one of the security-critical layers of Data Protector. The configuration of users must be carefully planned and tested.

User rights

Some user rights are very powerful and therefore represent a security issue. For example, the user configuration and clients configuration user rights enable a user to change the security settings.

The **Restore to other clients** user right is also very powerful, especially if combined with either the **Back up as root** or **Restore as root** user rights.

Even less powerful user rights bear an inherent risk associated with them. Data Protector can be configured to restrict certain user rights to reduce these risks.

Start backup specification user right

The user is allowed to start backup sessions for a backup specification from the command line by using the `omnib` with the `-datalist` option.

By combining the **Start Backup Specification** with the **Start Backup** user rights, a user is allowed to see the configured backup specifications in the GUI and is able to start a backup session for a backup specification or an interactive backup.

Allowing users to perform interactive backups may not always be desired. To allow interactive backups only to users which also have the **Save backup specification** user right, set the `StrictSecurityFlags` global option to `0x0200`.

Hiding the contents of backup specifications

In a high security environment, the contents of saved backup specifications may be considered to be sensitive or even confident information.

Data Protector can be configured to hide the contents of backup specifications for all users, except for those who have the **Save backup specification** user right. To do so, set the `StrictSecurityFlags` global option to `0x0400`.

Host trusts

The host trusts functionality reduces the need to grant the **Restore to other clients** user right to users when they only need to restore the data from one client to another within a limited number of clients.

You can define groups of hosts that will trust each other with the data.

Host trusts are typically used in the following situations:

- For clients in a cluster (nodes and virtual server).
- If the hostname of a client is changed and the data from the old backup objects needs to be restored.

- If there is a mismatch between the client hostname and backup objects due to DNS issues.
- If a user owns several clients and needs to restore the data from one client to another.

User groups

Data Protector has by default only a few predefined user groups. It is recommended to define specific groups for each type of user in the Data Protector environment to minimize the set of rights assigned to them.

User restrictions

In addition to defining specific user groups, you can further restrict user actions to be performed only on specific systems of the cell. You can enforce such restrictions by configuring the `user_restrictions` file on the Cell Manager. The restrictions apply only to members of the Data Protector user groups other than admin and operator.

User validation

The configuration of users is connected with user validation. Enhanced validation can be worthless without careful user configuration and the other way round - even the most careful user configuration can be worked around without the enhanced validation.

It is important that there are no “weak” user specifications in the Data Protector user list. Note that the client part of a user specification is the strong part (especially with the enhanced validation), while user and group parts cannot be verified reliably.

Any user with powerful user rights should be configured for the specific client they will use for Data Protector administration. If multiple clients are used, an entry should be added for each client, rather than specifying such a user as `user, group, <Any>`. Non-trusted users should not be allowed to log on to any of those systems.

Strict Hostname Checking

By default, the Cell Manager uses a relatively simple method for validating users. It uses the hostname as known by the client where a user interface or an application agent is started. This method is the easier to configure and provides a reasonable level of security in environments where security is considered as “advisory” (that is, malicious attacks are not expected).

The strict hostname checking setting on the other hand, provides enhanced validation of users. The validation uses the hostname as it is resolved by the Cell Manager using the reverse DNS lookup from the IP obtained from the connection. To enable the strict hostname checking, set the `StrictSecurityFlags` global option to `0x0001`.

Limitations

- IP based validation of users can only be as strong as the anti-spoof protection in the network. The security designer must determine whether the existing network provides a sufficient degree of anti-spoof safety for the particular security requirements. Anti-spoof protection can be implemented by

segmenting the network with firewalls, routers, VPN, and such.

- The separation of users within a certain client is not as strong as the separation between clients. In a high security environment, regular and powerful users should not be mixed within the same client.
- Hosts that are used in user specifications cannot be configured to use DHCP, unless they are bound to a fixed IP and configured in the DNS.

Be aware of the limitations in order to correctly assess the degree of safety that can be achieved with this setting.

Requirements

The enhanced validation does not automatically grant access for certain internal connections. Therefore, when this validation is used, a new user must be added for each of the following:

- Any application agent (OB2BAR) on Windows clients. It is required that the user `SYSTEM`, `NT AUTHORITY, client` is added for each client where an application agent is installed. Note that if `Inet` on a certain client is configured to use a specific account, the account must have already been configured.
- If you are using Web Reporting, user `java, applet, hostname` must be added for every hostname from where Web Reporting will be used. Note that for full Web Reporting functionality, the users must be in the `admin` group. Therefore, these clients must be trusted. Also, before making any data or functionality of Web Reporting available to other users (for example, via a web server), consider the security implications of making such data generally available.

Hostname resolution

The hostname that Data Protector uses for validation may differ between the default user validation and strict hostname checking in the following situations:

- Reverse DNS lookup returns a different hostname. This can be either intentional or can indicate misconfiguration of either the client or the reverse DNS table.
- The client is multihomed (has multiple network adapters and/or multiple IP addresses). Whether or not this consideration applies to a specific multihomed client, depends on its role in the network and on the way it is configured in the DNS.
- The client is a cluster.

The nature of checks that are enabled with this setting may require reconfiguration of Data Protector users. Existing specifications of Data Protector users must be checked to see if they may be affected by any of the above reasons. Depending on the situation, existing specifications may need to be changed or new specifications added to account for all the possible IPs from which the connections can come.

Note that users have to be reconfigured also when reverting back to the default user validation, if you had to modify user specifications when you enabled the strict hostname checking. It is therefore recommended to decide which user validation you would like to use and keep using it.

A prerequisite for a reliable reverse DNS lookup is a secure DNS server. You must prevent physical access and log on to all unauthorized personnel.

By using IPs for validation (instead of using hostnames), you will resolve some potential DNS related validation problems, but it is more difficult to maintain.

Security Logs

If you encounter problems accessing the Data Protector functionality or clients, you can use the information in the log files to determine your problem. For example, logged events can help you to determine misconfigured users or clients.

Client security events

Client security events are logged to the `inet.log` file residing in the default Data Protector log files directory on every client in the cell.

It is useful to check the recent activity of Data Protector on the clients.

Cell Manager security events

Cell Manager security events are logged in the `security.log` file residing in the default Data Protector server log files directory.

The `security.log` file is created with the first security event.

Securing the Entire Data Protector Cell

You can secure all clients in the cell.

Steps

1. In the Context List, click **Clients**.
2. In the Scoping Pane, right-click **Clients** and click **Cell Secure**.
3. Type the names of the systems that will be allowed to access all clients in the cell or search for the systems using the **Network** (on Windows GUI only) or **Search** tabs. Click **Add** to add each system to the list.
4. Click **Finish** to add the selected systems to the `allow_hosts` file.

Clients will verify the source for each request and allow only those requests received from clients selected in the `Enable Security on selected client(s)` window. These clients are listed in the `allow_hosts` file. If the request is denied, the event is logged to the `inet.log` file residing in the default Data Protector log files directory.

When you secure an entire cell, all clients residing in this cell at the time are secured. When you add new clients to the cell, you should also secure them.

For more information on securing clients and security considerations, see the *HPE Data Protector Installation Guide*.

Securing a Client System

You can secure the selected clients in the cell.

Steps

1. In the Context List, click **Clients**.
2. In the Scoping Pane, expand **Clients**, right-click the client(s) you want to secure, and click **Secure**.
3. Type the names of the systems that will be allowed to access the selected client(s) or search for the systems using the **Network** (on Windows GUI only) or **Search** tabs. Click **Add** to add each system to the list.
4. Click **Finish** to add the selected systems to the `allow_hosts` file.

Clients will verify the source for each request and allow only those requests received from the clients selected in the `Enable Security on selected client(s)` window. These clients are listed in the `allow_hosts` file. If the request is denied, the event is logged to the `inet.log` file residing in the default Data Protector log files directory.

Tip: If you do not select any Cell Manager and you simply click **Finish**, your Cell Manager is automatically provided with access and (his primary client name) added to the `allow_hosts` file. You cannot exclude the Cell Manager from the list.

For more information on securing clients and security considerations, see the *HPE Data Protector Installation Guide*.

Unsecuring the Entire Data Protector Cell

You can remove security from all clients that are imported to the cell.

Steps

1. In the Context List, click **Clients**.
2. In the Scoping Pane, expand **Data Protector Cell**, right-click **Clients**, and then click **Cell Unsecure**.
3. Click **Yes** to confirm that you want to allow access to all the client(s) in your cell.

Unsecuring a Client System

You can remove security from the selected client systems.

Steps

1. In the Context List, click **Clients**.
2. In the Scoping Pane, expand **Clients**, right-click the client from which you want to remove security, and then click **Unsecure**.
3. Click **Yes** to confirm that you want to allow access to the selected client.

Configuring Host Trusts

You can define groups of hosts that will trust each other with the data.

Steps

1. On a Windows Cell Manager, create the *Data_Protector_program_data\Config\Server\cell\host_trusts* file.

On a UNIX Cell Manager, create the */etc/opt/omni/server/cell/host_trusts* file.

2. In the file, list the trusted hosts.

For example:

```
GROUP="cluster.domain.com"
{
  cluster.domain.com
  node1.domain.com
  node2.domain.com
}
GROUP="DFG"
{
  computer.domain.com
  anothercomputer.domain.com
}
```

3. Save the file.

Encryption

About Encryption

Data Protector lets you encrypt backup data so that it becomes protected from others. Two data encryption techniques are available: software-based and drive-based encryption.

Data Protector software encryption, referred to as AES 256-bit encryption, is based on the Advanced Encryption Standard (AES) cryptographic algorithm that uses the same key for both encryption and decryption. Data is encrypted before it is transferred over the network and written to media.

Data Protector drive-based encryption uses the encryption functionality of the drive. The actual implementation and encryption strength depend on the drive's firmware. Data Protector only turns on the feature and manages encryption keys.

After the encryption is turned on, no additional configuration is required. However, for AES 256-bit encryption, Data Protector offers you advanced manual management of encryption keys (such as expiring, reactivating, exporting, importing, and deleting keys) via the command-line interface (CLI).

Using the Data Protector GUI, or the CLI, it is possible to determine which backup objects are encrypted, or which backup media contain encrypted objects, and to obtain encryption details for those objects.

Enabling AES 256-bit Encryption

You can enable software-based AES 256-bit encryption while creating a new backup specification or modifying one that is already configured.

Prerequisite

- You must have an active encryption key prior to performing an encrypted IDB backup. For details, see the `omnikeytool` man page or the *HPE Data Protector Command Line Interface Reference*.

Limitations

- AES 256-bit encryption does not encrypt metadata, such as the file name and file size.
- Encryption is not applicable for ZDB to disk and the disk part of ZDB to disk+tape.
- Objects that are backed up using AES 256-bit encryption cannot be consolidated.

Enabling encryption in a filesystem backup specification

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Filesystem**. All saved backup specifications are displayed.
3. Click the backup specification that you want to modify.
4. In the Options property page, click the **Advanced** button for Filesystem Options.
5. In the Filesystem Options window, click the **Other** tab. In the **Data security** drop-down list, select the **AES 256-bit** option.
6. Click **OK** and then click **Apply** to save the changes.

Tip: To encrypt only selected backup objects, go to the **Backup Object Summary** tab and select the **AES 256-bit** option in the object's properties.

Enabling encryption in a disk image backup specification

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Filesystem**. All saved backup specifications are displayed.
3. Click the backup specification that you want to modify.
4. In the Backup Object Summary page, click the **Properties** button.
5. In the Object Properties window, click the **Other** tab. In the **Data security** drop-down list, select the **AES 256-bit** option.
6. Click **OK** and then click **Apply** to save the changes.

Enabling encryption in an Internal Database backup specification

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then **Internal Database**. All saved backup specifications are displayed.
3. Click the backup specification that you want to modify.
4. In the Options page, under Common Application Options, click **Advanced**.
5. In the Common Application Options window, click the **Other** tab. From the **Data security** drop-down list, select the **AES 256-bit** option.
6. Click **OK** and then click **Apply** to save the changes.

Enabling encryption in an application integration backup specification

Limitations

- For an up-to-date list of application integrations that support AES 256-bit encryption, see the latest support matrices at <https://softwaresupport.hpe.com/>.
- It is not possible to use a combination of the options **Fast direct mode** and **AES 256-bit** for the Microsoft SQL Server integration.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then the appropriate type of backup specification (for example, **MS SQL Server**). All saved backup specifications are displayed.
3. Click the backup specification that you want to modify.
4. In the Options property page, click the **Advanced** button for Common Application Options.
5. In the Common Application Options window, click the **Other** tab. In the **Data security** drop-down list, select the **AES 256-bit** option.
6. Click **OK** and then click **Apply** to save the changes.

Exporting and Importing Media with Encrypted Backups

To restore data from encrypted backup to a client in a different Data Protector cell, you need to import the media and the encryption keys to the destination Cell Manager, as described in the following sections.

Note: Data Protector also provides advanced manual management of encryption keys (such as expiring, reactivating, exporting, importing, and deleting keys) via the command-line interface (CLI). For details, see the *omnikeytool* man page or the *HPE Data Protector Command Line Interface Reference*.

Cell Manager environment or MoM environment without CMMDB

In a Cell Manager environment or in a MoM environment where local MMDs are used, perform the following steps to export and import a medium with encrypted backup:

Steps

1. On the original Cell Manager, export the medium from the IDB. This operation also exports the relevant encryption keys from the keystore into the file *mediumID.csv*, in the default exported encryption keys directory.
2. Transfer the *mediumID.csv* file to the destination Cell Manager and place it into the directory default imported encryption keys directory.
3. Insert the exported medium into the drive that will be used by the destination Cell Manager.
4. On the destination Cell Manager, import the medium. This operation also imports the keys from the *mediumID.csv* file.

Note: If the key file is not present, you can still import the medium, but the catalog import will abort because of missing decryption keys.

MoM environment with CMMDB

In a MoM environment where the CMMDB is used, all media information is stored on the MoM Manager, but encryption keys IDs used by these media as well as the CDB are stored in a local keystore on each respective Cell Manager. Note that all media management operations need to be done on the MoM Cell Manager.

To export and import a medium with encrypted backup if the CMMDB resides on the MoM Manager, perform the following steps:

Steps

1. Export the medium from the CMMDB. The key IDs are exported into the file *mediumID.csv*, in the default exported encryption keys directory.
2. Transfer the *mediumID.csv* file to the destination Cell Manager and place it into the default imported encryption keys directory.
3. From the MoM Manager, eject a medium from a library.
4. Move a medium from the original media pool to the destination media pool, which is associated with a drive in the destination cell. This operation also imports the catalogue.
5. Insert the exported medium into the drive that will be used by the destination Cell Manager.
6. On the destination Cell Manager, import the medium. This operation also imports the keys from the *mediumID.csv* file.

Enabling Drive-Based Encryption

For an up-to-date list of devices that support drive-based encryption, see the latest support matrices at <https://softwaresupport.hpe.com/>.

You can enable drive-based encryption:

- While configuring a drive or modifying an already configured one.
- While configuring a backup, object copy, or object consolidation specification or modifying an already configured one.
- While configuring a automated media operation or modifying an already configured one.

Prerequisite

- You must have an active encryption key prior to performing an encrypted IDB backup. For details, see the *omnikeytool* man page or the *HPE Data Protector Command Line Interface Reference*.

Limitations

- It is not possible to use drive-based encryption for NDMP Server controlled devices or for drives in a library with external encryption control (for example, an ESL library under HPE SKM control).

Recommendation

- For optimal performance, the block size used should be at least 256 kilobytes.

Note: When backing up to a medium that contains both encrypted and unencrypted backups, you might get the message *Drive-based decryption enabled*. This means that the last backup on the medium is an encrypted one and it was automatically decrypted so it could be checked by Data Protector before the new backup was added.

Enabling drive-based encryption in the drive configuration

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Devices**, expand the desired device and then its drives.
3. Right-click the desired drive and click **Properties**.
4. In the Settings property page, click the **Advanced** button.
5. In the Advanced Options window, in the **Settings** tab, select the **Drive-based encryption** option, and then click **OK**.
6. Click **Apply** to save the changes.

Enabling drive-based encryption in a backup specification

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then the appropriate type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Click the appropriate backup specification.
4. In the Destination page, right-click the device that is selected for the backup and click **Properties**.
5. In the Device Properties window, select the **Drive-based encryption** option, and then click **OK**.
6. Click **Apply** to save the changes.

Tip: To modify an object copy or object consolidation specification, open the specification in the **Object Operations** context and perform steps 4 to 6.

Enabling drive-based encryption for an automated media operation

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Automated Operations**. All configured automated operations are displayed.
3. Click the media operation for which you want to enable drive-based encryption.
4. In the Options page, select the **Drive-based encryption** option, and then click **Apply**.

Note: The **Drive-based encryption** option applies to all devices that are involved in the automated media operation.

Encrypted Control Communication

About Encrypted Control Communication

Data Protector encrypted control communication helps preventing unauthorized access to clients in Data Protector cell. It is based on Secure Socket Layer (SSL), a cryptographic protocol, which provides network connections and encapsulates existing Data Protector communication protocol.

Since SSL requires certificates to establish encrypted communication, Data Protector provides default certificates during the installation or upgrade.

Using the Data Protector GUI or the CLI, you can remotely enable encrypted control communication for all clients in the Data Protector cell. You must first enable encrypted control communication on a Cell Manager and then on the clients in the cell. Clients that are not supposed to communicate confidentially can be placed in a Cell Manager exception list, which allows those clients to communicate in non-encrypted mode.

Managing Encrypted Control Communication

Data Protector encrypted control communication helps in preventing unauthorized access to clients in Data Protector Cell Managers. Using the Data Protector GUI or the CLI you can enable or disable encrypted control communication for all clients in the Data Protector cell.

- [Enabling encrypted control communication](#)
- [Selecting TLS version](#)
- [Disabling encrypted control communication](#)
- [Viewing certificate expiration date in Data Protector GUI](#)
- [Upgrading an encrypted environment](#)

Considerations

- The Cell Manager has to be upgraded to the latest patch for using the new encrypted control communication with the Data Protector automatically generated certificates. If the Cell Manager had encrypted control communication enabled from a prior release, you have to disable the encrypted control communication before you proceed to use the new procedure.
- It is also recommended to upgrade the Data Protector clients. Data Protector clients that have not been upgraded will not be able to disable the earlier encrypted control communication.
- Hosts with General Media agent, acting as gateway clients, and hosts with StoreOnce Software Deduplication agent must be upgraded.
- The Installation Server cannot be shared between Cell Managers, if the Installation Server has enabled encrypted control communication. However, if the Installation Server has enabled encrypted control communication as part of the MoM environment, then the Installation Server can have encrypted control communication enabled and shared between the Cell Managers in the MoM environment.
- In a Windows environment, you can enable encrypted control communication from the GUI and from the CLI.
- In a UNIX environment, you can enable encrypted control communication only after installing the Cell Manager, using the CLI.
- StoreOnce Software may fail if the certificate key length is 512 bits or less when the encrypted control communication is enabled. Therefore, use a certificate that has a key length of more than 512 bits.
- After you enable encrypted control communication with Data Protector automatically generated certificates on the Cell Manager, the clients added will also have encrypted control communication as enabled.
- As the encryption used for the certificates generated by Data Protector for Encrypted Control Communication (ECC) is stronger, therefore it cannot be used with client versions earlier than 8.0.

Note: It is only possible to manage encryption locally on a Cell Manager or from a client that has enabled encrypted control communication.

Enabling encrypted control communication

You can enable encrypted control communication on the following:

In a cell: This includes the Cell Manager and individual clients. You do not need to enable encrypted control communication on all clients.

In a MoM environment: This includes all cells that are a part of the MoM environment.

Enabling encrypted control communication for all clients in the cell, using the CLI:

Execute the following command: `omnicc -encryption -enable -all`

If encrypted control communication has been disabled on the Cell Manager, then it is not possible to enable encrypted control communication for a client in a cell.

To enable encrypted communication only on the Cell Manager, run:


```
omnicc -encryption -enable <CellManager_name>
```

To enable encrypted communication on the Cell Manager (if it has not yet been enabled) and all clients in the cell, run:

```
omnicc -encryption -enable -all
```

Enabling encrypted control communication for all cells in a MoM environment, using the CLI:

It is recommended that you first disable encrypted control communication on all the Cell Managers (including the clients of the Cell Managers) before importing them to the MoM environment, otherwise the Cell Managers cannot communicate and the creation of the MoM environment will not complete. After creating the MoM environment, proceed to enable encrypted control communication in the MoM environment.

Enabling MoM encryption only works:

1. If all the Cell Managers are upgraded to the latest patch. Some clients in Cell Managers can be older, but disabling will not work in this case.
2. If the MoM server and the other Cell Managers can connect and communicate:
 - Encrypted control communication has not been enabled on the MoM server and all the other Cell Managers or
 - Encrypted control communication has been enabled with Data Protector generated certificates on the MoM server and on some or all of the Cell Managers, which are a part of the MoM environment. Additionally, trust has been established between the MoM server and member servers.

Establishing trust

To enable encrypted control communication without disabling the earlier encrypted control communication, the MoM server has to be able to communicate with the other (member) servers. Before MoM can be created, trust has to be established between the MoM server and the member servers.

Note: Save the initial state of the files so that you can revert the changes in case of an error.

To establish trust between the Cell Managers, do the following:

1. Get the CA certificate for MoM server.
 - a. On the MoM server, open the MoM server trusted certificates file `Data_Protector_program_data/config/client/config` and find the line `trusted_certificates_file=`
For example, `trusted_certificates_file='C:\ProgramData\OmniBack\config\client\certificates\<CMhostname>_cacert.pem'`;
 - b. Open the file `client\certificates\<CMhostname>_cacert.pem` file in a text editor (unless it has been modified, the standard file name format is `<CMhostname>_cacert.pem`) and copy its contents (MoM server CA certificate).
2. Get the CA certificate for server1.
 - a. On server 1 open the server 1 trusted certificates file `Data_Protector_program_data/config/client/config` and find the line `trusted_certificates_file=`

For example, trusted_certificates_
file='C:\ProgramData\OmniBack\config\client\\certificates\<CMhostname>_
cacert.pem'

- b. Open the file client\\certificates\<CMhostname>_cacert.pem file in a text editor (unless it has been modified, the standard file name format is <CMhostname>_cacert.pem) and copy its contents (server 1 CA certificate).
3. Edit both trusted certificate files '<CMhostname>_cacert.pem' to include all the certificates that exist on each server that needs to be trusted. In this example, the MoM Server and Server1 need to establish trust with each other.
 - a. On the MoM server, open the MoM server trusted certificates file and include the server 1 CA certificate to the file.
 - b. On server 1, open the server 1 trusted certificates file and include the MoM server CA certificate to the file.
4. If there are more servers (server 2) and so on. Repeat steps 2 and 3 for every server, to be added to the MoM environment.

The Cell Manager trusted certificate file is initially a copy of Data_Protector_program_data/config/server/certificates/<CMhostname>_cacert.pem

To enable encrypted control communication, in the MoM environment run `omnicc -encryption -enable_mom{CSHostname1 [CSHostName2...] |-all} [-recreate_cert]`

For more details, see the `omnicc` command in the *HPE Data Protector Command Line Interface Reference*.

Enabling encrypted control communication for all clients in the cell, using the GUI:

1. In the Context List, click **Clients**.
 2. In the Scoping Pane, expand **Data Protector Cell** and then **Clients**. All clients are displayed.
 3. Right-click the client that you want to modify and select **Enable encrypted control communication**. In case of multiple clients, select one or more clients for which you want to enable encrypted control communication.
- Note:** If you select enable encrypted communication option for a client whose Cell Manager is not yet encrypted, you are prompted with a message “You can change encrypted communication configuration only from a client with encrypted communication enabled or the Cell Manager” and the options on that page become unavailable.
4. In the Connection tab, the **Encrypted control communication** option is selected by default.
 5. Select **Use Existing certificates**, if you need to use the existing certificates on the Cell Manager.
 6. Click **Apply** to save the changes.

Note: You can also enable encrypted control communication in the following scenarios:
Adding or importing: clients to a cell.
Editing the Properties of a client or Cell Manager.

Encrypted control communication with user-created certificates

This section is applicable for users who want to generate the certificates themselves.

Encrypted control communication with certificates created manually

The earlier versions of Data Protector did not create certificates automatically, you had to create the certificates and point Data Protector to the certificate files.

If you generate the certificates manually, then you have to place the certificates in the following certificates directory on the Cell Manager :

Windows: Data_Protector_program_data\Omniback\Config\Server\certificates ;

UNIX: /etc/opt/omni/server/certificates directory.

In addition, the certificates have to comply with the following naming convention.

<computer.company.com>_cert.pem for the certificate

<computer.company.com >_key.pem for the private key

<CellManager.company.com>_cacert.pem for the trusted certificate

When you enable encryption (while adding / importing / editing properties of a client or a Cell Manager), these certificates are used by Data Protector. When encryption is enabled, ensure that you select the **Use existing certificates** option from the Data Protector GUI otherwise the existing certificates will get overwritten.

Note that you can also generate the certificates to be used for encrypted control communication, using the script `omnigencert.pl` and then select **Use existing certificates** option from the Data Protector GUI. This enables faster encryption of the clients.

To create the certificates for encrypted control communication use the script `omnigencert.pl`, and run:

```
omnigencert.pl -pem_client -user_id <computer.company.com> [-recreate]
```

The `-recreate` option overwrites the existing certificates, if they exist.

Note: The `omnigencert.pl` script can also be used for generating certificates for other purposes.

Encrypted control communication with certificates created automatically

If you need to generate certificates automatically, and as per your specification, then you can create a Perl script file `gencert.pl` and place it in the following location:

ws: %Data_Protector_home%\bin

UNIX: /opt/omni/lbin

Data Protector starts using the `gencert.pl` instead of the `omnigencert.pl` script after it is added to the specified folder. You can enable encryption using the Data Protector GUI or CLI. This `gencert.pl` script must comply with the following certificate naming conventions:

<computer.company.com>_cert.pem for the certificate

<computer.company.com>_key.pem for the private key

<CellManager.company.com>_cacert.pem for the trusted certificate

The gencert.pl script should be able to accept the following parameters:

```
gencert.pl-pem_client -user_id <computer.company.com> [-recreate]
```

Replacing CA certificates in an encrypted control communication environment

It is possible to replace certificates with the ones signed by a different CA. If you need to replace the CA and the certificates in the cell you must perform the following steps:

1. Concatenate the CA certificates:

Copy the new CA certificate to the following path:

- Windows - Data_Protector_program_data\Omniback\Config\Server\certificates and
- UNIX - /etc/opt/omni/server/certificates

To update all the clients in the cell to also trust this new CA, run the following command:

```
omnicc -encryption -update_trust -all -trust newCA.pem
```

2. Recreate the certificates:

You can recreate the certificates either manually or use Data Protector to trigger certificate generation. Data Protector triggers omnigencert.pl or gencert.pl (if it exists) for creating certificates when you run the following command:

```
omnicc -encryption -enable -all -recreate_cert
```

3. Update the clients to trust only the new CA:

```
omnicc -encryption -update_trust -all -trust newCA.pem -replace
```

Selecting TLS version

To configure the TLS versions, execute the following omnicc command:

```
omnicc -encryption -encr_param <hosts> -tls_min <min_ver> -tls_max <max_ver>
```

This command specifies both minimum and maximum versions of TLS. The default range after the installation is TLSv1 to TLSv1.1.

By default, Data Protector uses TLSv1.1 for Encrypted Control Communication. TLSv1 is the default minimum version supported to support communication with previous versions of Data Protector binaries. Binaries prior to version 9.07 supported only TLSv1.

When setting the range of minimum and maximum TLS versions, ensure that a common version is available for all the pairs of systems and Data Protector processes that communicate. If there is no overlap between the two clients, then the connection between them cannot be established.

The maximum version of TLS is TLS1.2. To enable TLS1.2 for a host, use the following command:

```
omnicc -encryption -encr_param <hostname> -tls_max TLS1.2
```

Note: The file hpdpcert.pem is not suitable for TLS1.2 version.

When using the `hdpdcert.pem` or a similar short certificate, update the encryption before setting TLS1.2. It is recommended to switch to the Data Protector generated certificates. This can be done by disabling the old encrypted control communication and enabling it again. This causes the certificates to be newly generated by Data Protector.

The `<ssl/>` element with protocol attribute defines the allowed versions of TLS protocol. The default value is comma-separated list of three versions:

protocol = **TLSv1,TLSv1.1,TLSv1.2**

For Windows:

`c:\ProgramData\OmniBack\Config\Server\AppServer\standalone.xml`

For Linux:

`/etc/opt/omni/server/AppServer/standalone.xml`

Disabling encrypted control communication

You can disable encrypted control communication:

- In a cell: This includes the Cell Manager and clients
- In a MoM environment: This includes all Cell Managers in a MoM environment.

Note: You can change encrypted communication configuration only from a client with encrypted communication enabled or from the Cell Manager.

Disabling encrypted control communication, using the CLI:

- In a cell, run: `omnicc -encryption -disable -all`
- In a MoM environment, run: `omnicc -encryption -disable_mom -all`
- On a specific client, run: `omnicc -encryption -disable <client_name>`
- On multiple clients, run: `omnicc -encryption -disable{Hostname1 [HostName2 ...] | -all}`

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

Disabling encrypted control communication for multiple clients, using the GUI:

1. In the Clients context, select a client or multiple clients.
2. Right-click the selection and select **Disable Encrypted Communication**.
The Disable encrypted control communication page appears. All the clients are selected.
3. Click **Finish** to disable encrypted control communication for the clients.

Disabling encrypted control communication for each client, using the GUI:

1. In the Clients context, select a client.
2. Right-click the selection and select **Properties**.
3. In the **Connection** tab deselect the **Encrypted control communication** option.
4. Click **Apply**.

Note: You can also disable encrypted control communication in the following scenarios:
Adding or importing: Clients to a cell.
Editing the Properties of a client and Cell Manager.

Viewing certificate expiration date in Data Protector GUI

To view the duration from when the certificates are valid using the Data Protector GUI, proceed as follows:

1. In the Context List, click **Clients**.
2. In the Scoping Pane, expand **Data Protector Cell** and then **Clients**. All clients are displayed.
3. Select a Cell Manager host.
You get to view the General tab details.
4. Select the **Certificates** tab.
You get to view the list of all certificates and their valid from and to dates.

Note: From Data Protector 9.09 onwards, the list of certificates does not contain private keys (*_key.pem) as they are no longer available on the Cell Manager.

Upgrading an encrypted environment

By default, after upgrading to the latest patch, changes made to the encrypted control communication functionality do not affect the existing environment. You can choose from the following options to maintain the existing encrypted environment:

Option 1

Remove the encryption from the entire cell and enable encryption in the cell in the new way (recommended). See [Enabling encrypted control communication](#).

Option 2

Keep the existing certificates on the clients and maintain the environment using the `omnicc` command:

```
omnicc -encryption -enable {Hostname1 [HostName2 ...] | -all} [-cert Cert [-key Key]] [-trust TrustedCerts]
```

In this method, it is not possible to configure encrypted control communication using the GUI. Also, the clients will not be encrypted automatically after import. You can encrypt the clients after importing them using the CLI.

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

Note: With the earlier method of enabling encrypted control communication, if certificates were not specified, then using the command line `omnicc -encryption -enable` defaulted to `hdpdcert.pem`. With the new approach, the default mechanism is for Data Protector to generate the certificates. To enable encrypted control communication with `hdpdcert.pem`, the certificate

```
has to be specified: omnicc -encryption -enable <host> -cert hdpdcert.pem -key  
hdpdcert.pem -trust hdpdcert.pem
```

Adding a Client to the Security Exceptions List

You can add a client to the **Security Exceptions** list on the Cell Manager while modifying the connection properties.

Adding security exceptions is available if encrypted control communication is enabled on the Cell Manager.

Remote disabling of encrypted control communication by using the Data Protector GUI or the CLI is for security reasons not supported.

Note: To simplify the import of a client with enabled encrypted control communication to another Data Protector cell, encrypted control communication is disabled during the export from the primary Data Protector cell.

Steps

1. In the Context List, click **Clients**.
2. In the Scoping Pane, expand **Data Protector Cell** and then **Clients**. All clients are displayed.
3. Click the Cell Manager that you want to modify.
4. Type the names of the systems that will be added to the Security Exceptions list in the cell or search for the systems using the **Network** (on Windows GUI only) or **Search** tabs.
5. Click **Add** to add systems to the list, then click **Apply** to save the changes.

The clients that accept communication in a plain text mode are written to the `config` file, located on the Cell Manager in the default Data Protector server configuration directory.

Tip: To remove a system from the **Security Exceptions** list, perform steps 1 to 4 and click **Remove**, then click **Apply** to save the changes.

Introduction to User Authentication and LDAP

Authentication and authorization of Data Protector as an enterprise system should be connected to the enterprise user management infrastructure. This connection allows users and groups configured in a corporate user directory to be granted access to Data Protector services.

User authentication is performed over secure connections, and Lightweight Directory Access Protocol (LDAP) is used as the underlying technology. Consequently, users can use their corporate credentials to access Data Protector services and are not required to maintain separate passwords. In addition, administrators or operators can be maintained as groups in the corporate directory, adhering to established authorization and approval processes.

LDAP integration is configured in a security domain of Data Protector's embedded application server (JBoss) using Java Authentication and Authorization Service (JAAS) login modules. An optional LDAP login module provides LDAP authentication and authorization services, which are mapped to Data Protector permissions by a mandatory Data Protector Login Module. If LDAP integration is not configured, then Data Protector works just as it did in previous releases.

Data Protector uses the login modules in an login module stack to authenticate users. When a user connects to the Cell Manager using the Data Protector GUI, user authentication is performed by the following login modules:

1. LDAP Login Module: Authenticates user credentials, such as username and password, against an existing LDAP server. See [Initializing and Configuring the LDAP Login Module](#).
2. Data Protector Login Module: Authenticates user credentials against the Data Protector user list and the Web access password. See [Granting Data Protector Permissions to LDAP Users or Groups](#).
3. After performing all the steps necessary to complete LDAP initialization and configuration, you can also check the configuration. See [Checking the LDAP Configuration](#).

Note: Whenever a user or client is configured in Data Protector to allow the CLI access in the classic way, the Data Protector GUI does not use the LDAP feature.

Initializing and Configuring the LDAP Login Module

The LDAP login module is located in the security domain of JBoss Application Server, which is installed with Data Protector. The LDAP login module must be initialized and configured prior to the first use of the LDAP security feature.

1. [Initializing the LDAP Login Module](#).
2. [Configuring the LDAP Login Module](#).

Initializing the LDAP Login Module

To initialize the LDAP login module, use the `jboss-cli` utility, which is also installed with Data Protector

1. The `jboss-cli` utility is located in: `%Data_Protector_home%/AppServer/bin`. Execute the following command:
 - Windows: `jboss-cli.bat --file=ldapinit.cli`
 - UNIX: `jboss-cli.sh --file=ldapinit.cli`

This command creates an LDAP login module in JBoss configuration and populates this new login module with default values. The default values generated by the command line within the `standalone.xml` configuration file:

```
<security-domain name="hdpd-domain">
  <authentication>
    <login-module code="LdapExtended" flag="optional">
      <module-option name="java.naming.factory.initial"
        value="com.sun.jndi.ldap.LdapCtxFactory"/>
      <module-option name="java.naming.security.authentication" value="simple"/>
      <module-option name="roleFilter" value="(member={1})"/>
      <module-option name="roleAttributeID" value="memberOf"/>
    </login-module>
  </authentication>
</security-domain>
```



```
<module-option name="roleNameAttributeID" value="distinguishedName"/>
<module-option name="roleAttributeIsDN" value="true"/>
<module-option name="searchScope" value="SUBTREE_SCOPE"/>
<module-option name="allowEmptyPasswords" value="true"/>
<module-option name="password-stacking" value="useFirstPass"/>
</login-module>
<login-module code="com.hp.im.dp.cell.auth.DpLoginModule" flag="required">
<module-option name="password-stacking" value="useFirstPass"/>
</login-module>
</authentication>
</security-domain>
```

Note: The default values generated by the command line within the standalone.xml configuration file changes, if the Cell Manager is installed on UNIX environment and uses LDAP authentication. The following are the changes:

```
<login-module code="LdapExtended" flag="optional">
  <module-option name="java.naming.factory.initial"
value="com.sun.jndi.ldap.LdapCtxFactory"/>
  <module-option name="java.naming.security.authentication" value="simple"/>
  <module-option name="roleFilter" value="(member={1})"/>
  <module-option name="roleAttributeID" value="memberOf"/>
  <module-option name="roleNameAttributeID" value="distinguishedName"/>
  <module-option name="roleAttributeIsDN" value="true"/>
  <module-option name="searchScope" value="SUBTREE_SCOPE"/>
  <module-option name="allowEmptyPasswords" value="false"/>
  <module-option name="password-stacking" value="useFirstPass"/>
  <module-option name="java.naming.provider.url" value="ldap://<IP_of_
Active_Directory_host>"/>
  <module-option name="baseCtxDN" value="OU=_Benutzer,DC=godyo,DC=int"/>
  <module-option name="rolesCtxDN" value="OU=_Gruppen,DC=godyo,DC=int"/>
  <module-option name="bindDN" value="CN=backup-service,OU=_Service_
Accounts,DC=godyo,DC=int"/>
  <module-option name="bindCredential" value="password"/>
  <module-option name="baseFilter" value="(userPrincipalName={0})"/>
```

```
</login-module>
```

The configuration parameters `baseCtxDN` and `rolesCtxDN` are the main ones. The Organization Unit (OU) parameter is used to authenticate the UNIX Cell Manager.

2. To access the JBoss admin console, located on the Cell Manager, from a remote client, enable the remote access to the JBoss admin console. To do this, use a text editor and change the bind address of the management interface from 127.0.0.1 to 0.0.0.0 in the interfaces section of the `standalone.xml` file:

```
<interfaces>
<interface name="management">
<inet-address value="{jboss.bind.address.management:0.0.0.0}"/>
</interface>
<interface name="public">
<inet-address value="0.0.0.0"/>
</interface>
<interface name="unsecure">
<inet-address value="{jboss.bind.address.unsecure:127.0.0.1}"/>
</interface>
</interfaces>
```

3. Restart the Data Protector services:

```
omnisv stop
omnisv start
```

Configuring the LDAP Login Module

To configure the LDAP login module, use the web-based admin console of JBoss Application Server, which gets installed with Data Protector. Proceed as follows:

1. To access the JBoss admin console, create a JBoss user. To create a JBoss user, run the `add-user` utility:
 - Windows: `add-user.bat` located in `%Data_Protector_home%/AppServer/bin`
 - UNIX: `add-user.sh` located in `/opt/omni/AppServer/bin`
2. Provide inputs for the following parameters:
 - **Type of user to add:** Select Management User.
 - **Realm:** Leave this field blank, as the default value `ManagementRealm` is selected by the utility.

- **Username:** Add a username.
 - **Password:** Add a password.
3. To access the JBoss admin console, use a browser and open the URL: <http://cell-manager-name:9990/console>
 4. In the Authentication screen, specify the **Username** and **Password** created using the add-user utility.
 5. Click **Log In**. JBoss Application Server admin console appears.
 6. In the JBoss admin console, select the **Profile** tab.
 7. In the **Profile** tab, expand the **Security** node and then click **Security Domains**.
 8. From the list of registered security domains, click **View** for hdpd-domain. The following login modules are defined for the security domain, hdpd-domain:
 - LdapExtended
 - Com.hp.im.dp.cell.auth.DpLoginModule
 9. Select the **LdapExtended** module.
 10. From the Details section, click the **Module Options** tab. All of the pre-configured module options are listed in the **Module Options** tab.
 11. To customize and use the LDAP login module, you need to add additional Module Options. Click **Add** and specify the **Name** and **Value** for each module option. See the following table for more information:

Module Options	Name	Value	Description
Provider URL	java.naming.provider.url	Specify the URL of the LDAP server in the following format: ldap://<server>:<port>	A standard property name
Base Context Distinguished Name (DN)	baseCtxDN	Specify the DN of the LDAP location that contains the users.	The fixed DN of the context from where you start the user search
Base Filter	baseFilter	Specify the attribute in the LDAP setup that matches the user's login name in the following format: (<user-login-name-attribute>={0}) where <user-login-name-attribute> needs to be replaced by the corresponding LDAP attribute name.	A search filter used to locate the context of the user to authenticate
Roles Context DN	rolesCtxDN	Specify the DN of the LDAP location that contains the user	The fixed DN of the context to

		groups.	search for user groups
Bind DN	bindDN	Specify the DN of an LDAP user that is used by the login module to perform the initial LDAP bind. You must have the required permission to search the LDAP location of the users and groups to obtain the users and their groups. These locations are defined in the <code>baseCtxDN</code> and <code>rolesCtxDN</code> module options.	The DN used to bind against the LDAP server for the user and roles queries. This is a DN with read/search permissions on the <code>baseCtxDN</code> and <code>rolesCtxDN</code> values
Bind Credential	bindCredential	Specify the password for the LDAP user provided in the <code>BindDN</code> module option.	The password for the <code>bindDN</code> .

For more information on other Module Options, visit the following URLs:

- <https://community.jboss.org/wiki/LdapExtLoginModule>
 - [http://technet.microsoft.com/en-us/library/cc773354\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/cc773354(v=ws.10).aspx)
12. The changes will take effect when you reload JBoss Application Server configuration. To reload the configuration, use the `jboss-cli` utility located in `%Data_Protector_home%/AppServer/bin`.
 13. Execute the following command:
 - Windows: `jboss-cli.bat -c :reload`
 - UNIX: `jboss-cli.sh -c :reload`

Note: When configuring the LDAP Login Module in MoM environments, ensure that you perform the steps described above on every Cell Manager. Every Cell Manager in the MoM environment needs to have the same configuration for the LDAP login module.

Granting Data Protector Permissions to LDAP Users or Groups

LDAP users can connect to a Cell Manager only if they are granted the Data Protector permissions. After configuring the LDAP login module, you can grant the LDAP users the required Data Protector permissions.

To grant the Data Protector permissions, proceed as follows:

1. Start the Data Protector GUI and grant Data Protector permissions to the LDAP users or groups.
 - [Add LDAP users to Data Protector user groups.](#)
 - [Add LDAP groups to Data Protector user groups.](#)

2. [Log In using LDAP credentials.](#)

Adding LDAP Users to Data Protector User Groups

To add LDAP users to Data Protector user groups, proceed as follows:

1. In the Context List, click **Users**.
2. In the Scoping Pane, expand **Users** and right-click the user group to which you want to add the LDAP user(s).
3. Click **Add/Delete Users** to open the wizard.
4. In the **Manual** tab of the Add/Delete Users dialog box, provide the following details:
 - **Type:** Select LDAP.
 - **Name:** Specify the LDAP user in the LDAP user principal name format.
 - **Entity:** Enter LDAP User.
 - **Description:** This is optional.
5. Click **Finish** to exit the wizard.

Adding LDAP Groups to Data Protector User Groups

To add LDAP groups to Data Protector user groups, proceed as follows:

1. In the Context List, click **Users**.
2. In the Scoping Pane, expand **Users** and right-click the user group to which you want to add the LDAP group.
3. Click **Add/Delete Users** to open the wizard.
4. In the **Manual** tab of the Add/Delete Users dialog box, provide the following details:
 - **Type:** Select LDAP.
 - **Name:** Specify the LDAP group name in the Distinguished name (DN) format.
 - **Entity:** Enter LDAP Group.
 - **Description:** This is optional.
5. Click **Finish** to exit the wizard.

Note: An LDAP user is automatically granted the same permission level as the LDAP group this user belongs to.

Logging In using LDAP Credentials

To log in using your LDAP credentials, proceed as follows:

1. Start the Data Protector GUI and connect to a Cell Manager.
2. In the LDAP Authentication screen, provide the LDAP credentials to access Data Protector. The LDAP user can belong to any available Data Protector user group.

Checking the LDAP Configuration

The following procedure explains how to check if the user rights are set correctly for a specific LDAP user or group by querying the Data Protector login provider service `getDpAc1` from a web browser.

To obtain the Data Protector Access Control List (ACL) for a specific user, proceed as follows:

1. Connect to the Data Protector login provider web service using a browser.
2. The browser may prompt you to accept the server certificate. Click **Accept** to confirm the request.
3. A dialog box appears, prompting you to provide login credentials. Provide a valid LDAP user name and password that was configured using Data Protector.
4. The browser returns the following Access Control List (ACL): `https://<server>:7116/dp-loginprovider/restws/dp-ac1`
5. Use the ACL to check if the assigned rights match the Data Protector user rights specified for the corresponding Data Protector user group.

Certificate Generation Utility

Introduction to the Certificate Generation Utility

The X.509 certificate generation utility—`omnigencert.pl`—generates the Certificate Authority (CA), server, and client certificates. It is responsible for the following tasks:

- Setting up a single-level root CA
- Generating CA, server, and client certificates
- Creating the necessary directory structure for storing keys, certificates, configuration, and keystore files
- Storing the generated certificates in predefined locations on the CM
- Generating the properties files of web service roles

Note: The `omnigencert.pl` utility can be run only by the Administrator user (Windows) or the root user (UNIX).

The `omnigencert.pl` utility is developed as a script and gets installed along with the Cell Manager (CM) installation kit. As part of the CM installation, the script is run for the first time, and the certificates are generated and stored in predefined locations.

If required, the Data Protector administrators can run this utility any time after the installation to regenerate certificates using the new keys pair or the new CA setup. However, it is not mandatory to use the certificates generated by this utility for the certificate-based authentication. Instead, you can use an existing CA setup for generating the necessary certificates.

Directory Structure for the Certificate Generation Utility

The following sections list the directories where certificates are stored.

Windows Directory	Unix Directory	Description
ProgramData\Omniback\Config\Server\certificates	/etc/opt/omni/server/certificates	Contains the CA certificate file, cacert.pem, which contains the CA public key.
ProgramData\Omniback\Config\Server\certificates\ca	/etc/opt/omni/server/certificates/ca	Contains the configuration, input, and other files necessary for the CA functioning.
ProgramData\Omniback\Config\Server\certificates\ca\keys	/etc/opt/omni/server/certificates/ca/keys	Contains the CA private key file, cakey.pem.
ProgramData\Omniback\Config\Server\certificates\server	/etc/opt/omni/server/certificates/server	Contains two kinds of stores: keystore and truststore. These stores are created by the Java utility, keytool, for protecting server certificates and its keys. These stores are protected by the store password. It contains the

		<p>following stores:</p> <p>ca.truststore</p> <p>server.keystore</p> <p>server.truststore</p>
ProgramData\Omniback\Config\Server\certificates\client	/etc/opt/omni/server/certificates/client	<p>Contains two kinds of stores: keystore and truststore. These stores are created by the Java utility, keytool, for protecting client certificates and its keys. These stores are protected by the store password. It contains the following stores:</p> <p>client.keystore</p> <p>client.truststore</p>
ProgramData\Omniback\Config\Server\App Server	/etc/opt/omni/server/AppServer	<p>Contains the properties files created by this utility. This directory contains other files apart from the following properties files:</p>

		jce- webservice- roles.prope rties dp- webservice- roles.prope rties
--	--	-----------------------------------------------------------------------------------------------

Overwriting Existing Certificates

Overwriting Existing Certificates

To overwrite existing certificates—generated by the utility as part of the CM installation—with the certificates generated by an existing CA setup, you can use one of the following options:

- [Overwriting certificates in existing keystore and truststore files](#)
- [Overwriting certificates by creating new keystore and truststore files](#)

Note: After regenerating certificates or using new certificates, you must restart the Data Protector services on the CM. You must do this before performing any operation that uses certificates, as restarting the services ensures that new certificates are in effect.

Overwriting Certificates in Existing Keystore and Truststore Files

To overwrite certificates in existing keystore and truststore files, complete the following tasks:

- [Replace existing server and client store files](#)
- [Replace the CA certificate](#)
- [Update the Distinguished Name\(DN\) string](#)

Replacing Existing Server and Client Store Files

To replace existing server and client store files, proceed as follows:

1. Retrieve the keystore and trustore files' store password from the `webservice.properties` and `standalone.xml` configuration files, which are available at the following locations:

Windows:

- `ProgramData\OmniBack\Config\client\components\webservice.properties`
- `ProgramData\OmniBack\Config\server\AppServer\standalone.xml`

UNIX:

- `/etc/opt/omni/client/components/webservice.properties`
 - `/etc/opt/omni/server/AppServer/standalone.xml`
2. Remove all entries from the existing server and client store files, `server.keystore`, `server.truststore`, `client.keystore`, and `client.truststore`, available at the following locations:

Server:

- Windows: `ProgramData\Omniback\Config\Server\certificates\server`
- Unix: `/etc/opt/omni/server/certificates/server`

Client:

- Windows: `ProgramData\Omniback\Config\Server\certificates\client`
- UNIX: `/etc/opt/omni/server/certificates/client`

To make these changes, you can use the Java keytool utility, located in

Windows: `Program Files\Omniback\jre\bin`

UNIX : `/opt/omni/jre/bin`

3. Import the generated certificates into the following stores using the Java keytool utility:
 - Server and CA certificates into `server.keystore`
 - CA and Client certificate into `server.truststore`
 - CA certificate into `ca.truststore`
 - Client and CA certificates into `client.keystore`
 - CA and Server certificate into `client.truststore`

Replacing the CA Certificate

To replace the existing CA certificate, proceed as follows:

1. Note the permissions of the existing CA certificate file `cacert.pem`, which is located in:
 - **Windows:** `ProgramData\Omniback\Config\Server\certificates`
 - **UNIX:** `/etc/opt/omni/server/certificates`
2. Replace the existing CA certificate `cacert.pem` file with the generated CA certificate.

Updating the Distinguished Name (DN) String

Replace the existing Distinguished Name (DN) string in the `jce-webservice-roles.properties` and `dp-webservice-roles.properties` files with the DN string used for the client certificate. These files are located in:

Windows: `ProgramData\Omniback\Config\Server\AppServer`

UNIX: /etc/opt/omni/server/AppServer

Note: In the DN string, precede spaces and "=" characters with the backslash (\) character.

Overwriting Certificates by Creating New Keystore and Truststore Files

To overwrite certificates by creating new keystore and truststore files, complete the following tasks:

- [Replace existing server and client store files](#)
- [Replace the CA certificate](#)
- [Update the Distinguished Name \(DN\) string](#)
- [Update the configuration file with the stores password](#)

Note: You must retain the password for server and client stores.

Replacing Existing Server and Client Store Files

To replace existing server and client store files, proceed as follows:

1. Note the permissions of the existing server and client store files, `server.keystore`, `server.truststore`, `client.keystore`, and `client.truststore`, located in:

Server:

- Windows: ProgramData\Omniback\Config\Server\certificates\server
- UNIX: /etc/opt/omni/server/certificates/server

Client:

- Windows: ProgramData\Omniback\Config\Server\certificates\client
- UNIX: /etc/opt/omni/server/certificates/client

2. Remove the server and client store files.
3. Create stores with the same file names and permissions.
4. Import the generated certificates into the following stores using the Java keytool utility:
 - Server and CA certificates into `server.keystore`
 - CA and Client certificate into `server.truststore`
 - CA certificate into `ca.truststore`
 - Client and CA certificates into `client.keystore`
 - CA and Server certificate into `client.truststore`

Note: The Java keytool utility is located at Windows: Program Files\Omniback\jre\bin and UNIX: /opt/omni/jre/bin.

Replacing the CA Certificate

To replace the existing CA certificate, proceed as follows:

1. Note the permissions of the existing CA certificate file, `cacert.pem`, which is located in:
Windows: `ProgramData\OmniBack\Config\Server\certificates`
UNIX: `/etc/opt/omni/server/certificates`
2. Replace the existing CA certificate file, `cacert.pem`, with the generated CA certificate.

Updating the Distinguished Name (DN) String

Replace the existing Distinguished Name (DN) string in the `jce-webservice-roles.properties` and `dp-webservice-roles.properties` files with the DN string used for the client certificate. These files are located in:

Windows: `ProgramData\OmniBack\Config\Server\AppServer`

UNIX: `/etc/opt/omni/server/AppServer`

Note: In the DN string, precede spaces and "=" characters with the backslash (\) character.

Updating the Configuration File with the Stores Password

To update the configuration file with the stores password, proceed as follows:

Note: This task is required only if new stores are created with a new password.

1. Update the `webservice.properties` and `standalone.xml` configuration files with the store password used while creating store files, such as `server.keystore`, `server.truststore`, `ca.truststore`, `client.keystore`, and `client.truststore`.

The configuration files are located in:

Windows:

- `ProgramData\OmniBack\Config\client\components\webservice.properties`
- `ProgramData\OmniBack\Config\server\AppServer\standalone.xml`

UNIX:

- `/etc/opt/omni/client/components/webservice.properties`
- `/etc/opt/omni/server/AppServer/standalone.xml`

2. In the `standalone.xml` file, update the stores password (highlighted in bold):

```
<ssl name="ssl" password="M6.p0ino06L3w" certificate-key-  
file="/etc/opt/omni/server/certificates/server/server.keystore" protocol="TLS"  
verify-client="want" ca-certificate-  
file="/etc/opt/omni/server/certificates/server/ca.truststore" ca-certificate-  
password="M6.p0ino06L3w"/>
```

3. In the `webservice.properties` file, update the password (highlighted in bold):

```
<jsse keystore-password="M6.p0ino06L3w" keystore-  
url="/etc/opt/omni/server/certificates/server/server.keystore" truststore-  
password="M6.p0ino06L3w" truststore-  
url="/etc/opt/omni/server/certificates/server/server.truststore"/>  
  
<jsse keystore-password="M6.p0ino06L3w" keystore-  
url="/etc/opt/omni/server/certificates/server/server.keystore" truststore-  
password="M6.p0ino06L3w" truststore-  
url="/etc/opt/omni/server/certificates/server/server.truststore"/>  
  
<ssl name="ssl" password="M6.p0ino06L3w" certificate-key-  
file="/etc/opt/omni/server/certificates/server/server.keystore" protocol="TLS"  
verify-client="want" ca-certificate-  
file="/etc/opt/omni/server/certificates/server/ca.truststore" ca-certificate-  
password="M6.p0ino06L3w"/>
```

Firewall Support

About Firewall Support

You can configure Data Protector in an environment where the Data Protector processes communicate across a firewall. Starting with the Data Protector 9.09, the number of ports that must be open in the firewall is reduced. The change happens only after the cell is upgraded, until then the old clients continue to operate in the legacy mode, and use the same open ports as in previous Data Protector versions.

Communication in Data Protector

Data Protector processes communicate using TCP/IP connections. Data Protector needs the following ports:

- Inet port (5555 by default) on every Data Protector system.

Note: Windows inet is multi-threaded.

- IDB service port (7112 by default) on the Cell manager system.
- Application Server port (7116 by default) on the Cell Manager system.
- Rule for `StoreOnceSoftware.exe` binary must be left in the inbound firewall exceptions. Since `StoreOnceSoftware.exe` does not support (third party code based) single port passing, but it does open inbound ports and accepts communication on them.

These ports must be open in the firewall (as accessible from remote hosts).

In addition, Data Protector opens a number of dynamic ports. They must remain open in the firewall until Data Protector cell is upgraded. Once the cell is upgraded to Data Protector 9.09 or later, these ports are used within processes (IPC) and they are not required to be open from firewall perspective.

Following changes are required to configure the range of dynamic ports:

- Limit the ports that still need to be open in a firewall until whole cell is upgraded.
- Prevent Data Protector from opening ports that may be needed by the third-party applications.
- Data Protector can communicate with the non-Data Protector software that may need to open own ports.

Configuration mechanism

You can configure the port allocation behavior using two `omnirc` options:

- **OB2PORTRANGE**
This option sets port range from which all Data Protector processes open the dynamic ports.
- **OB2PORTRANGESPEC**
This option sets per-binary port range from which matching Data Protector processes opens the dynamic listening ports. This option gives you more control over the ranges and helps to keep their sizes smaller.

Before Data Protector 9.09, port ranges served two purposes:

- **Security:** Limit the ports opened by Data Protector, therefore port user needs to open in firewall.
- **Conflict with other software:** Non-Data Protector software may need ports 1000-2000 for specific use, so by using `OB2PORTRANGE` prevents Data Protector from using that range. This is valid for `OB2PORTRANGE`.

Note:

- By default, dynamic ports are assigned by the operating system.
- These options do not affect the fixed ports for `Inet` (5555), `IDB` service port (7112) and `Application Server` (7116).
- Port range options limit the Data Protector port usage. They cannot prevent non-Data Protector applications from allocating ports from this range.

The number of open ports in firewall is reduced when agents participating in communication are upgraded. Until then, Data Protector uses the old method for communication. Old disk agents work with new media agent and vice-versa. Users should keep the ports open until all the cell hosts are upgraded.

On platforms where `inetd` supports `-p <proc_limit>` option, if possible avoid using this option otherwise it is recommended to use `proc_limit` value greater than 2200.

Before enabling real firewall, it is recommended to test the Data Protector with Data Protector firewall enabled environment.

To enable Data Protector firewall in the entire cell, run the following commands:

```
omnicc -firewall -all -enable_dp
```

To enable Data Protector firewall in one part of a cell, specify individual hosts instead of `-all`.

For example, to test if disk agent to media agent communication is possible through firewall, close the firewall by running the following command:

```
omnicc -firewall -host MAhost DAhost -enable_dp
```

Along with `-enable_dp` option, use `-enable_os` option to disable Data Protector rules in Windows firewall.

After testing with these options, user can proceed to close third-party firewalls (for example, routers).

From Data Protector 9.09 and later

OB2PORTRANGE and OB2PORTRANGESPEC variables are not required to be set for listening ports for Data Protector processes. These can still be set for Data Protector to use them for inter process communication within a host. The following examples explain the usage:

Example 1:

Data Protector 9.09 MA and earlier version of Data Protector 9.09 DA

MA opens port which is bound to all addresses. This is displayed in netstat output as MA listening on the 0.0.0.0:1234.

- Here, "1234" is an example, actual port depends on dynamic port range which is set using OB2PORTRANGE variables and Data Protector configuration.
- The "0.0.0.0" is for "all addresses" which is equivalent for IPv6 is `:::`. This means a client can connect through any route.

Other processes from the same host cannot open port 1234. DA connects directly to MA host 1234. You must leave port 1234 open in firewall.

Example 2:

Data Protector 9.09 MA and Data Protector 9.09 DA:

- Since MA does not know if the connecting DA is Data Protector version 9.09 or earlier, therefore MA still keeps the port 0.0.0.0:1234 open.
- When compared with case 1, you need not open the port 1234 in firewall if you are sure that only Data Protector 9.09 clients would connect. This depends on the order in which you upgrade the hosts.

Example 3:

Data Protector 9.09 MA, Data Protector 9.09 DA, and Data Protector firewall enabled on MA host

- MA opens the port which is bound to loopback interface only. This is displayed in netstat output as MA listening on the port 127.0.0.1:1234 , or `::1`:1234 for IPv6.
- Old DA would not be able to connect, as port 1234 is not accessible from the remote host, irrespective of Windows or some other firewall.

Note: Firewalls do not prevent processes from opening ports, only the connection to these ports from remote hosts.

For more information on port usage in Data Protector 9.09, refer to the [Port Usage in Data Protector 9.09 and later](#) table.

Limiting a Port Range in Data Protector prior to 9.09

You can limit the port range in Data Protector prior to 9.09:

- For all Data Protector processes
- For a specific Data Protector agent
- For Data Protector processes and a specific Data Protector agent together

For all Data Protector processes

To limit the port range for all Data Protector processes on a system, use the OB2PORTRANGE option in the `omnirc` file:

```
OB2PORTRANGE=start_port-end_port
```

Data Protector processes use dynamically allocated ports and select ports from the specified range. The port range is consumed by taking the first available port, starting with port `start_port`. If there is no available port within the specified range, the port allocation fails and the requested operation is not performed.

For a specific Data Protector agent

You can limit the port range on a system on which a specific agent runs by using the OB2PORTRANGESPEC option in the `omnirc` file:

```
OB2PORTRANGESPEC=AGENT:start_port-end_port;...
```

All agent processes check the OB2PORTRANGESPEC for range restrictions. If there is a range defined for agent processes, all dynamically allocated ports will be selected from this specified range. The port range is consumed by taking the first available port, starting with port `start_port`. If there is no available port within the specified range, the port allocation fails and the requested operation is not performed.

The table below lists all possible Data Protector agent identifiers that can be used in the OB2PORTRANGESPEC option. Note that agent processes that do not dynamically allocate listen ports are not listed in the following table.

Agent identifiers

Data Protector component	Agent identifier	Description	Port consumption
Cell Manager	BSM	Backup Session Manager	1 port per concurrently running BSM
	RSM	Restore Session Manager	1 port per concurrently running RSM
	DBSM	Database Session Manager	1 port per concurrently running DBSM
	xSM	Wildcard matching Session Managers	1 port per database operations (such as database purges or database upgrades) + 1 port per concurrently running

			Session Manager
	MMD	Media Management Daemon	1 port
	CRS	Cell Request Server Service	1 port
Media Agent	BMA-NET	Backup Media Agent ¹	1 port per concurrently running Media Agent
	RMA-NET	Restore Media Agent ¹	1 port per concurrently running Media Agent
	xMA-NET	Wildcard matching Media Agent ¹	1 port per concurrently running Media Agent

¹ BMA and RMA fork two processes, the main process and a NetIO process. The listen port is allocated by the BMA-NET / RMA-NET process.

For Data Protector processes and a specific Data Protector agent together

If both options are set, OB2PORTRANGESPEC overwrites the settings of OB2PORTRANGE. For example, the setting

```
OB2PORTRANGESPEC=BMA-NET:18000-18009
```

```
OB2PORTRANGE=22000-22499
```

limits the port range used by a Media Agent to port numbers 18000-18009, while all other Data Protector processes use port numbers from the range 22000-22499.

By using both options it is possible to force a specific agent to use only a dedicated port range (OB2PORTRANGESPEC) and, at the same time, prevent other Data Protector processes from selecting port numbers from this range.

Port Usage in Data Protector 9.09 and later

The following table provides information on the port requirements of the different Data Protector 9.09 components:

Data Protector hosts	Port requirement	Port requirement for Data Protector on application hosts	Third-party requirement for the application
As install target	Linux/Unix: <ul style="list-style-type: none"> • REXEC¹ (non-secure): 512 • RSH (non-secure): 514 	None	None

	<ul style="list-style-type: none"> SSH: 22 Windows: SMB service: 445		
As Cell Manager	<ul style="list-style-type: none"> Inet: 5555 hdpd-as: 7116 IDB service: 7112 	None	None
As Data Protector Disk Agent	Inet : 5555	None	None
As Data Protector Integration Agent(s) ⁴	Inet : 5555	None	None
As Data Protector Media Agent or NDMP Media Agent	Inet : 5555	None	Store Once Deduplication system ² <ul style="list-style-type: none"> Command port: 9387 Data port: 9388 NDMP Server ² <ul style="list-style-type: none"> Server: 10000 DDBoost ³ <ul style="list-style-type: none"> NFS: 2049 Managed file replication: 2051 NFS port mapper: 111

¹ Only one of REXEC/RSH/SSHD ports is needed, depending on install method.

² For exact information on other 3rd party ports that can be open please consult 3rd party software documentation.

³ For more information on DDBoost ports please consult EMC® Data Domain® Boost for OpenStorage Administration Guide.

⁴ On a Windows Hyper-V server, following ports need to open for Hyper-v backup and restores

- **WMI Instance: 135 (Initiation)**
- **Windows Remote Management (HTTPS): 5986**

When writing the firewall configuration rules, the process in the first column must be able to accept new TCP connections (SYN bit set) on the ports defined in the second column from the process listed in the third column.

In addition, the process listed in the first column must be able to reply to the process in the third column on the existing TCP connection (SYN bit not set).

For example, the Inet process on a Media Agent system must be able to accept new TCP connections from the Cell Manager on port 5555. A Media Agent must be able to reply to the Cell Manager using the existing TCP connection.

Limitations

- This feature is not available on OpenVMS and SCO hosts. If media agent (or any Data Protector component that opens ports) is running on these systems, user still needs to open these ports in firewall.
- When you upgrade your Cell Manager to Data Protector 9.09 version, you will be unable to connect to WebReporting page on that Cell Manager from non-Data Protector 9.09 version of the Cell Manager.

Disk, Media and Application Agent in the DMZ

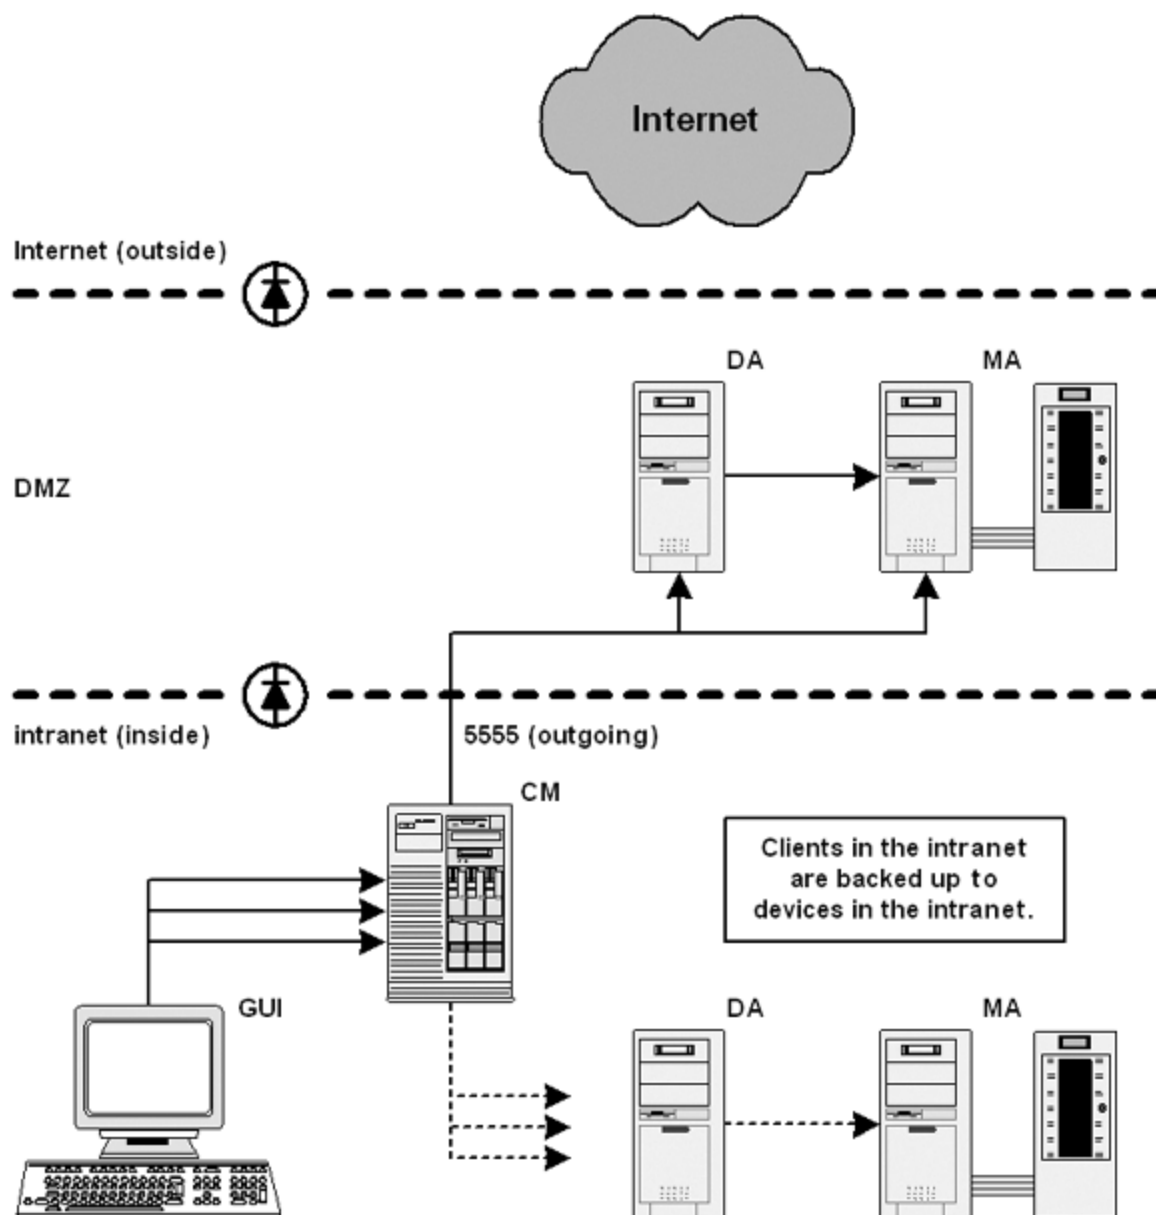
You can configure your backup environment so that the Cell Manager and GUI are in the intranet and some Disk, Application and Media Agents are in the DMZ.

[Configuration figures](#)

[Open ports](#)

[Limitations](#)

Configuration figure



Open ports

Data Protector opens the following ports for configuration:

1. The Disk Agent and a Media Agent need to accept connections from the Session Manager on port 5555:

- Allow connections from the CM system to port 5555 on the DA system
 - Allow connections from the CM system to port 5555 on the MA system
2. When **reconnect broken connections** is enabled, MA and DA connect to the Session Manager:
 - Allow connections from the MA and DA systems to port 5555 on the CM system
 3. The application agent needs to connect to the Session Manager and the CRS:
 - Allow connections from the application server system to port 5555 on the CM system

Note: Point 2 and 3 allow connections from the DMZ to the intranet, which is a potential security risk.

Limitations

- This cell can back up clients in the DMZ as well as clients in the intranet. However, each group of clients must be backed up to devices configured on clients that are on the same side of the firewall. If your firewall does not restrict connections from the intranet to the DMZ, it is possible to back up clients in the intranet to devices configured on clients in the DMZ. However, this is not recommended, as the data backed up in this way becomes more vulnerable.
- If a device in the DMZ has robotics configured on a separate client, this client must also be in the DMZ.

Disk and Application Agents in DMZ

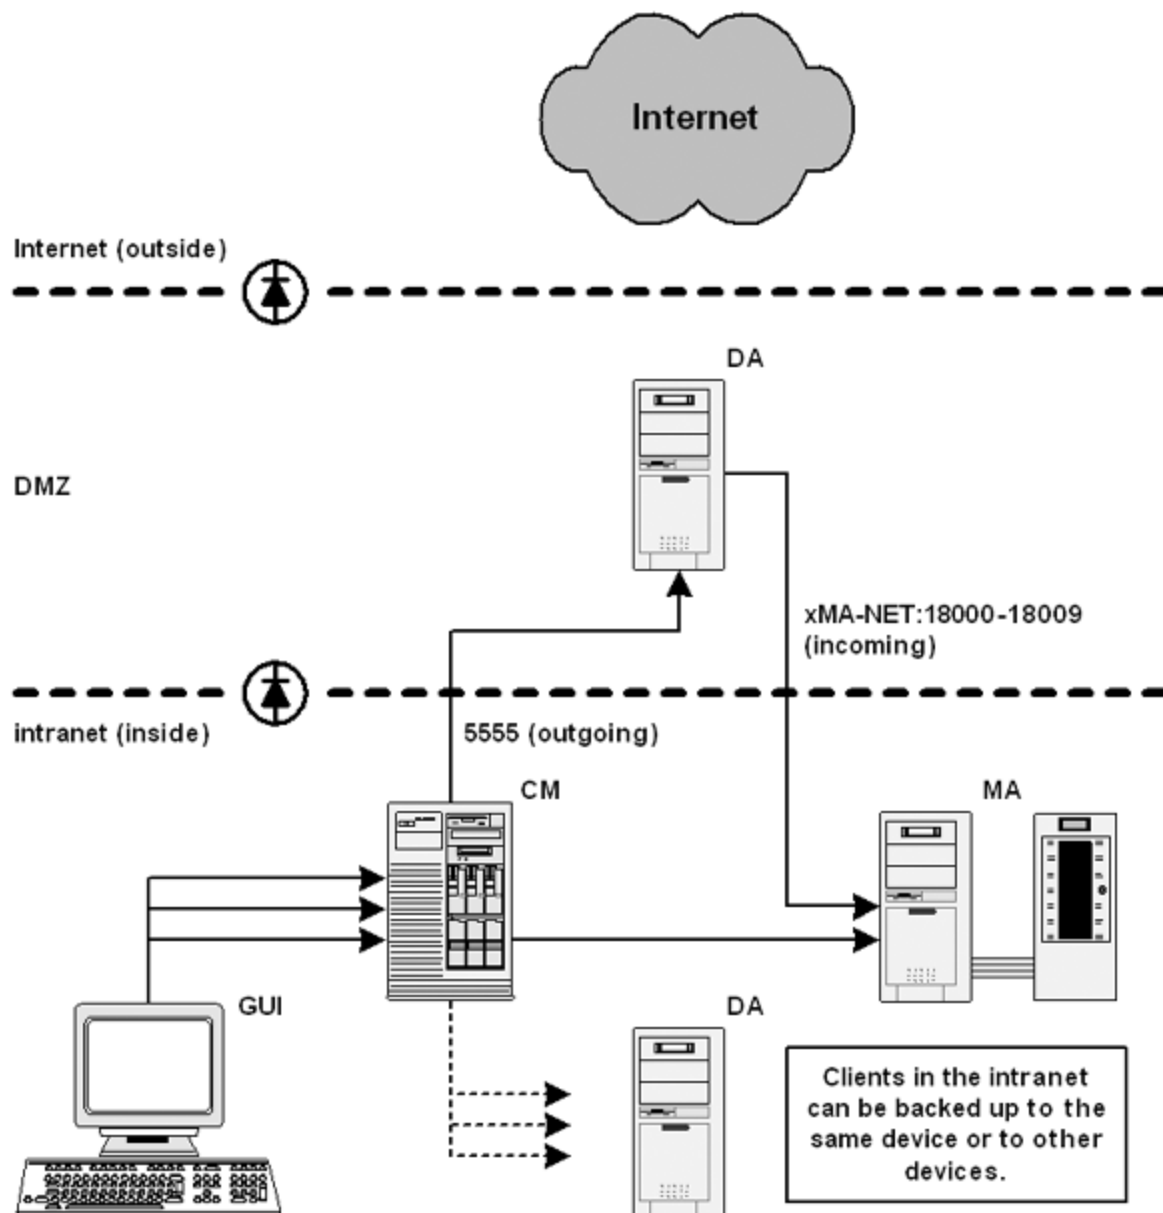
You can configure your backup environment so that the Cell Manager, Media Agents and GUI are in the intranet and some Disk and Application Agents are in the DMZ.

[Configuration figure](#)

[Open ports](#)

[Limitations](#)

Configuration figure



Open ports

Data Protector opens the following ports for configuration:

1. The Disk Agent needs to accept connections from the Session Manager on port 5555:
 - Allow connections from the CM system to port 5555 on the DA system
2. The Disk Agent needs to connect to the port 5555 on MA system:

- Allow connections from the DA system to port 5555 on the MA system

Note that this rule allows connections from the DMZ to the intranet, which is a potential security risk.

3. When **Reconnect broken connections** option is enabled, DA connects to the Session Manager:
 - Allow connections from DA system to port 5555 on the CM system.
4. The Application agent needs to connect to Session Managers and CRS:
 - Allow connections from the Application Agent system to port 5555 on the CM system.

Limitations

- Remote installation of clients across a firewall is not supported. You need to install clients locally in the DMZ.
- This setup does not allow the backup of databases or applications using Application Agents on clients in the DMZ.

Cell Manager, Disk Agent, and Media Agent in the DMZ

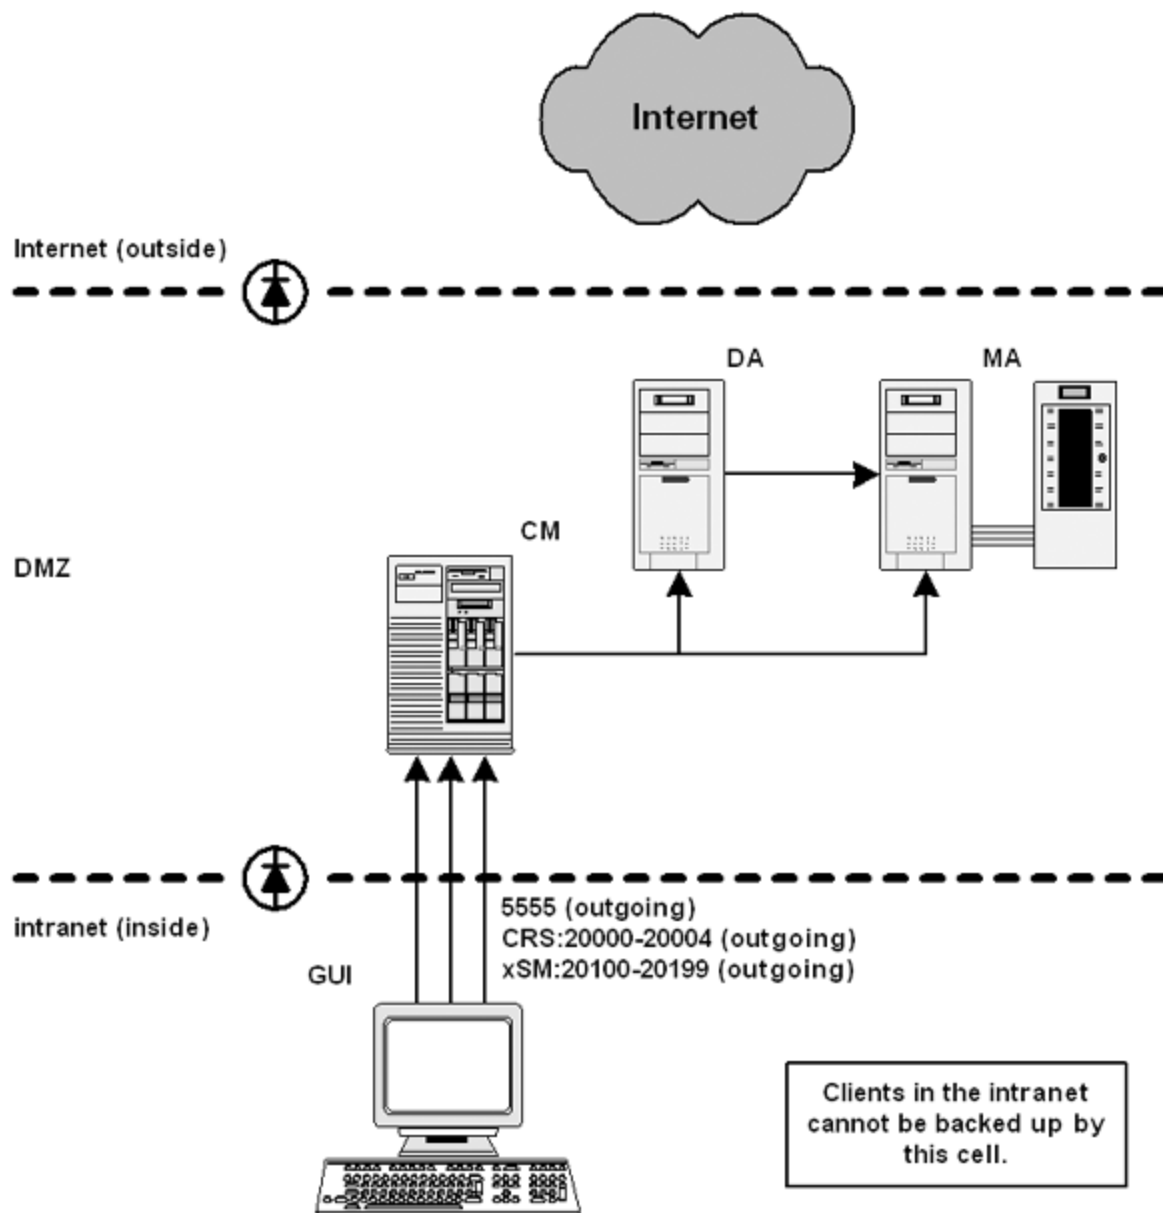
You can configure your backup environment so that the entire cell is in the DMZ and only the graphical user interface is in the intranet.

[Configuration figure](#)

[Open ports](#)

[Limitations](#)

Configuration figure



Open ports

Data Protector opens the following ports for configuration:

1. The GUI does not accept any connections. However, it needs to connect to the following processes on the Cell Manager:

Process	Port
Inet	5555
Application Server	7116

This leads to the following firewall rule:

- Allow connections from the GUI system to ports 5555 and 7116 on the CM system

Limitations

For this configuration, almost all Data Protector functionality is available, including remote installation and online backup of databases and applications.

- This cell cannot be a part of a MoM environment if centralized media management or centralized licensing is used and the MoM cell is inside.
- All backup clients must be in the DMZ. The GUI client cannot be backed up by a Media Agent from the DMZ. The GUI can also be run from a client that is a member of another cell located in the intranet, provided that both cells use the same Inet listen port.

Chapter 3: Users and User Groups

About User Management

The Data Protector user management functionality provides a security layer that prevents systems and data from being accessed by unauthorized personnel.

Security is based on a user-related security concept. Users that want to use Data Protector have to be configured as Data Protector users. User groups together with a rich set of user rights enable you to flexibly map your security requirements to your Data Protector user configuration.

By default, backed up data is hidden from other users, except the backup owner. Other users do not even see that data was backed up. If desired, data can be made visible to other user via appropriate user rights.

Users

To work with Data Protector, you have to be an authorized Data Protector user. For this, you need a Data Protector account, which restricts unauthorized access to Data Protector and to backed up data. In small environments, one person is sufficient for the backup tasks. The Data Protector administrators create this account specifying user logon name, the systems from which the user can log on, and a Data Protector user group membership. This is checked whenever the user starts the Data Protector user interface or performs specific tasks.

Each user belongs to one user group only. This defines the user's user rights.

You can configure both UNIX and Windows users:

UNIX

Users are defined by their logon name, UNIX user group, and a system from which they log on. A wildcard character can be used.

Windows

Users are defined by their logon name, Windows domain or workgroup, and a system from which they log on. A wildcard character can be used.

Predefined users

After the initial installation, all default user groups are empty, except for the `admin` group. Data Protector adds the following users to the `admin` group:

Cell Manager	User account	Remarks
--------------	--------------	---------

UNIX Cell Manager	The root user on the Cell Manager (root, any group, CeLL Manager host).	This user account should not be modified. It is required for proper operation of the CRS daemon and other processes on the Cell Manager. Only this user is initially allowed to administer the cell. To administer the cell from any other client, add a new user.
	The java user (java, applet, webreporting).	This user account enables Web Reporting. It needs to be modified when certain security settings are changed.
Windows Cell Manager	The CRS service account, as specified during the Data Protector installation (limited to the Cell Manager host).	The CRS service account should remain unchanged unless you modify the logon parameters of the CRS service. It is required for proper operation of the CRS daemon and other processes on the Cell Manager.
	The user who installed the Cell Manager (the initial cell administrator).	This user is configured as the initial cell administrator and can administer the cell from any client. It is recommended to modify this user account after the Data Protector installation is complete. Specify the client from which you will administer the cell instead of allowing access from any host. If you will be using another account, add this account and then remove the initial cell administrator or allow it only from the Cell Manager.
	The local system account on the Cell Manager (SYSTEM, NT AUTHORITY, CeLL Manager host).	This account is provided in case the CRS service is configured to log on as the local system account.
	The java user (java, applet, webreporting)	This user account enables Web Reporting. It needs to be modified when certain security settings are changed.

It is recommended to define specific groups for each type of users in an environment to minimize the set of rights assigned to them.

For more information on the java user, see the *HPE Data Protector Installation Guide*.

Admin group capabilities are very powerful. A member of the Data Protector admin user group has system administrator capabilities for the whole cell. For more information on security, see the *HPE Data Protector Installation Guide*.

User Groups

A user group is a collection of users who have the same rights. The administrator simplifies user configuration by grouping users according to their access needs. That is, the administrator puts users who need the same specific rights into the same group. Users might need the rights, for example, to monitor sessions in the cell, to configure backup, or to restore files.

Data Protector provides default user groups. You can use these groups as provided, modify them, or create new groups.

Predefined user groups

To simplify configuration, Data Protector provides three predefined user groups with the following user rights:

User right	Admin	Operator	User
Clients configuration	✓		
User configuration	✓		
Device configuration	✓		
Media configuration	✓	✓	
Reporting and notifications	✓		
Start backup	✓	✓	
Start backup specification	✓	✓	
Save backup specification	✓		
Back up as root	✓		
Switch session ownership	✓	✓	
Monitor	✓	✓	
Abort	✓	✓	
Mount request	✓	✓	
Start restore	✓	✓	✓
Restore to other clients	✓		
Restore from other users	✓	✓	
Restore as root	✓		
See private objects	✓	✓	

After initial installation, all predefined groups are empty, except the `admin` user group.

Admin capabilities are very powerful! A member of the Data Protector `admin` user group has system administrator rights on the whole cell.

The user rights you have set on the Cell Manager determine the availability of the Data Protector Cell Manager GUI or GUI contexts to the computer from which you connect to the Cell Manager. For example, if you have only the Start Restore user right set, then only the Restore context is available when you install the User Interface component.

Available User Rights

Data Protector provides a rich set of user rights to implement advanced security functionality. For more detailed information on user rights, see the HPE Data Protector Help.

Chapter 4: Internal Database

About the IDB

The Internal Database (IDB) is a database embedded in Data Protector, located on the Cell Manager, that keeps information regarding what data is backed up, on which media it resides, the result of backup, restore, object copy, object consolidation, object verification, and media management sessions, and which devices and libraries are configured.

What is the IDB used for?

The information stored in the IDB enables the following:

- Fast and convenient restore
You are able to browse the files and directories to be restored. You can quickly find the media required for a restore and therefore make the restore much faster.
- Backup management
You can verify the result of backup sessions.
- Media management
You can allocate media during backup, object copy, and object consolidation sessions, track media management operations and media attributes, group media in different media pools, and track media location in tape libraries.
- Encryption/decryption management: The information stored in the IDB enables Data Protector to allocate encryption keys for encrypted backup or copy sessions, and to supply the decryption key required for the restore of encrypted backup objects.

IDB size and growth consideration

The IDB can grow very big and can have a significant impact on backup performance and the Cell Manager system. The Data Protector administrator has to understand the IDB and decide which information to keep in the IDB and for how long. It is the administrator's task to balance restore time and functionality with the size and growth of the IDB. Data Protector offers two key parameters, logging level and catalog protection, that assist you in balancing your needs.

Regular IDB backups

HPE highly recommends to back up the IDB regularly. For more information, see [IDB Backup Configuration](#).

IDB Architecture


The Internal Database (IDB) consists of the following parts:

- Media Management Database (MMDB)
- Catalog Database (CDB)
- Detail Catalog Binary Files (DCBF)
- Session Messages Binary Files (SMBF)
- Encryption keystore and catalog files

Each of the IDB parts stores certain specific Data Protector information (records), influences the IDB size and growth in different ways, and is located in a separate directory on the Cell Manager.

IDB parts

Database architecture



The MMDB and CDB parts are implemented using an embedded database consisting of tablespaces. This database is controlled by the `hdp-idb`, `hdp-idb-cp`, and `hdp-as` processes. CDB (objects and positions) and MMDB present the core part of IDB.

The DCBF and SMBF parts of the IDB consist of binary files. Updates are direct (no transactions).

In the Manager-of-Managers (MoM) environment, the MMDB can be moved to a central system to create the Central Media Management Database (CMMDB).

Media Management Database (MMDB)

MMDB records

The Media Management Database stores information about the following:

- Configured devices, libraries, library drives, and slots
- Data Protector media
- Configured media pools and media magazines

MMDB size and growth

The MMDB does not grow very big in size. The largest part of the MMDB is typically occupied by information about the Data Protector media.

MMDB location

The MMDB is located in the following directory:

Windows systems: *Data_Protector_program_data\server\db80\idb*

UNIX systems: */var/opt/omni/server/db80/idb*

Catalog Database (CDB)

CDB records

The Catalog Database stores information about the following:

- Backup, restore, object copy, object consolidation, object verification, and media management sessions. This is the copy of the information sent to the Data Protector Monitor window.
- Backed up objects, their versions, and object copies. In the case of encrypted object versions, key identifiers (KeyID-StoreID) are also stored.
- Positions of backed up objects on media. For each backed up object, Data Protector stores information about the media and data segments used for the backup. The same is done for object copies and object mirrors.

CDB (objects and positions) size and growth

The CDB records occupy minor share of space in the IDB.

CDB location

The CDB is located in the following directory:

Windows systems: *Data_Protector_program_data\server\db80\idb*

UNIX systems: */var/opt/omni/server/db80/idb*

Detail Catalog Binary Files (DCBF)

DCBF information

The Detail Catalog Binary Files part stores information about the following:

- Pathnames of backed up files (filenames) together with client system names. Filenames of the files created between backups are added to the DCBF.
- File metadata. This is information about backed up file versions, their file sizes, modification times, attributes/protection, and positions of the backup copies on the backup media..

One DC (Detail Catalog) binary file is created for each Data Protector medium used for backup. When the medium is overwritten, the old binary file is removed and a new one is created.

DCBF size and growth

In an environment where filesystem backups using the Log All option are typical, the DCBF occupies the largest part of the IDB. Logging level and catalog protection can be used to specify what is actually stored in the IDB and for how long.

By default, five DC directories are configured for the DC binary files. If the number of backup media or DC binary files grows extremely big or you have disk space issues, you can create more of them, thus extending your IDB size.

The biggest and fastest growing part of the DCBF is the filenames part.

The growth of the filenames part is proportional to the growth and dynamics of the backup environment as well as to the number of backups.

A file or directory occupies approximately 100 bytes in the IDB.

DCBF location

By default, the DCBF is located in subdirectories `dcbf0` through `dcbf4` in the following directory:

Windows systems: `Data_Protector_program_data\server\db80\dcbf`

UNIX systems: `/var/opt/omni/server/db80/dcbf`

Consider the disk space on the Cell Manager and relocate the DC directory, if necessary. You can create more DC directories and locate them to different disks.

Session Messages Binary Files (SMBF)

SMBF records

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

SMBF size and growth

The SMBF size depends on the following:

- Number of performed sessions.
- Number of messages in a session. One session message occupies approximately 200 bytes. You can change the volume of messages displayed when backup, restore, and media management operations are performed by changing the Report level option. This influences the amount of messages stored in the IDB.

SMBF location

The SMBF is located in the following directory:

Windows systems: *Data_Protector_program_data\server\db80\msg*

UNIX systems: */var/opt/omni/server/db80/msg*

You can relocate the directory by editing the *SessionMessageDir* global option.

Encryption keystore and catalog files

All the keys created, either manually or automatically, during encrypted backups are stored in a keystore. The keys can also be used for object copy, object verification, and restore sessions. In the case of hardware encryption, they can also be used for object consolidation sessions.

In the case of software encryption, the key identifiers (each consisting of a KeyID and a StoreID) are mapped to the object versions encrypted. This mapping is stored in the Catalog Database. Different objects in a medium can have different (software) encryption keys.

For hardware encryption, the key identifiers are mapped to medium ID and these mappings are stored in a catalog file. This file contains the information required to allow an encrypted medium to be exported to another cell.

Keystore location

The keystore is located in the following directory:

Windows systems: *Data_Protector_program_data\server\db80\keystore*

UNIX systems: */var/opt/omni/server/db80/keystore*

Catalog file location

The catalog files are located in the following directory:

Windows systems: *Data_Protector_program_data\server\db80\keystore\catalog*

UNIX systems: */var/opt/omni/server/db80/keystore/catalog*

IDB Operation

Find out about the IDB behavior during the following Data Protector operations:

- [Backup](#)
- [Restore](#)
- [Object copy and object consolidation](#)
- [Object verification](#)
- [Exporting media](#)
- [Removing the Detail Catalog](#)

Backup

When a backup session is started, a session record is created in the IDB. Also, for each object in the session, an object version record is created. Both records are stored in the CDB part and have several attributes. The Backup Session Manager updates media during a backup. All media records are stored in the MMDB part and are allocated for a backup depending on policies.

When a data segment (and a catalog segment after it) is written on the tape, a media position record is stored in the CDB for each object version that was part of this data segment. In addition, the catalog is stored in the Detail Catalog (DC) binary file. One DC binary file is maintained per Data Protector medium. The DC binary file is named *MediumID_TimeStamp.dat*. The name is not changed when backups append to the same medium. If a medium is overwritten during a backup, its old DC binary file is removed and a new DC binary file is created.

All session messages generated during backups are stored in session messages binary files (the SMBF part).

IDB backup and archived log files

Depending on configuration of your Internal Database backup specification, the IDB backup process can remove old archived log files and starts creating new ones that are necessary for IDB recovery.

Restore

When configuring restore, Data Protector performs a set of queries in the CDB and DCBF parts to enable users to browse virtual filesystems of backed up data. These browse queries are done in two steps. The first step is to select a specific object (filesystem or logical drive). If this object has many backup versions stored, this can take some time because Data Protector scans the DCBF to build a lookup cache for later browsing. The second step is browsing directories.

After specific versions of files are selected, Data Protector determines the required media and locates media position records used by the selected files. These media are read by the Media Agents and data is sent to the Disk Agents that restore the selected files.

Object copy and object consolidation

During an object copy or object consolidation session, the same processes run as during a backup and a restore session. Basically, data is read from source media as if it was restored and written to target media as if it was backed up. An object copy or object consolidation session has the same effect on the IDB operation as backup and restore. For details, see the preceding sections.

Object verification

During an object verification session, the same database processes run as during a restore session. Basically, data is read from the source media, as if it were being restored, and is sent to the host Disk Agent(s) where the verification is performed. An object verification session has the same effect on the IDB operation as a restore session. For details, see the Restore section above.

All session messages generated during verification sessions are stored in session messages binary files.

Exporting media

When a medium is exported, the following is removed:

- All the media position records from that medium are removed from the CDB part.
- All objects that now have no positions on any other media are removed from the CDB part.
- Obsolete sessions (whose media have been either overwritten or exported) are removed. Session messages of such sessions are also removed.
- The medium record is removed from the MMDB part and the DC binary file for that medium is removed from the DCBF part.

Removing the Detail Catalog

When the Detail Catalog is removed for a specific medium, its DC binary file is removed. The same result is achieved by removing the catalog protection for all object versions on that medium (the next daily maintenance of DC binary files removes the binary file). All other records stay in the CDB and MMDB parts and it is possible to run a restore from such media (however, browsing is not possible).

Maintenance of DC Directories

The IDB allows several directories to be registered where the Detail Catalog Binary Files (DCBF) part of the IDB is stored. This allows the DC binary files to be distributed over more disks or volumes. By default, there are five directories named `dcbf0` through `dcbf4`.

Each DCBF directory has several configuration parameters:

- Allocation sequence
- Path

- Maximum size
- Maximum files
- Low space

For more detailed information on configuration parameters, see the HPE Data Protector Help.

Whenever there is a need to create a new binary file, the "DCBF allocation procedure" is performed by Data Protector:

1. From the list of all possible DC directories, Data Protector eliminates all that are deactivated or missing. Note that in the case of a missing DC directory, an `IDBCorrupted` event is generated.

All full DC directories are not considered. A DC directory is full if at least one of the following conditions is true:

`Maximum size - Current size < Low space`

`Free disk space < Low space`

`Maximum files <= Current files`

2. A set of user selectable algorithms (the `DCDirAllocation` global option) selects the actual DC directory:

- Fill in sequence

Data Protector creates a new DC binary file in the first non-full DC directory according to the configured sequence.

- Balance size

Data Protector selects the DC directory that contains (proportionally to the effective limit on the total size) the least DCBF data. The minimum for the following value is selected:

$(\text{Maximum size} - \text{Current size} - \text{Low space}) / (\text{Maximum size} - \text{Low space})$

- Balance number

Data Protector selects the DC directory that contains (proportionally to the effective limit on the number of files) the fewest DC binary files. The minimum for the following value is selected:

$\text{Current files} / \text{Maximum files}$

See the `DCDirAllocation` and `MaxDCDirs` global options that influence the DCBF behavior.

Checking the IDB Size

You can check the current size of the Internal Database parts using the Data Protector GUI.

Also, if configured, the IDB Size Report as well as the IDB Space Low notification inform you about the IDB size.

Steps

1. In the Context List, click **Internal Database**.
2. In the Scoping Pane, expand the **Usage** item. The following IDB items are displayed: Catalog Database, Media Management Database, Detail Catalog Binary Files, Session Messages Binary

Files, and Serverless Integrations Binary Files.

The item Serverless Integrations Binary Files relates to the functionality that is no longer supported in the installed HPE Data Protector version.

3. Check the size of the IDB by viewing properties of IDB parts and their records:
 - Right-click an IDB item, for example, **Catalog Database** and click **Properties** to view Disk Usage of the part of the IDB. Disk Usage shows how much of a disk space is currently being occupied by specific part of the IDB. Click the **Records Statistic** tab to view statistics for all records in the specific part of the IDB.
 - To check Disk Usage of a DC directory, expand **Detail Catalog Binary Files**, double-click the DC directory, and then click the **Disk Usage** tab.

Reducing the IDB Growth

You can reduce the growth of the Internal Database by reducing the logging level and catalog protection settings of your backup, object copy, and object consolidation specifications. These actions do not influence the current size of the IDB but they do impact its future growth.

The effect of reducing the logging level is a reduction in browse comfort at restore time.

The effect of reducing the catalog protection is that browsing is not possible for some restores (namely of those backups that have exceeded the catalog protection).

The following procedures describe how to change these settings in a backup specification.

Reducing logging level

By reducing the logging level settings for a backup specification, you reduce the amount of data (files/directories) that will be stored in the IDB (**Log All** -> **Log Files** -> **Log Directories** -> **No Log**).

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the backup specification for which you want to change the logging level and click the **Options** tab.
4. In the Options property page, click the appropriate **Advanced** button (under **Filesystem Options**).
5. Click the **Other** tab and, under **Logging**, change the logging level.
6. Click **OK** to apply the changes.

Reducing catalog protection

By reducing the catalog protection, you reduce the protection for the restore browse information in the IDB only. The information is still stored on media.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the backup specification for which you want to change the catalog protection and click the **Options** tab.
4. In the Options property page, click the appropriate **Advanced** button (under **Filesystem Options**).
5. Click the **Options** tab and, under **Catalog protection**, change the catalog protection.
6. Click **OK** to apply the changes.

Reducing the IDB Current Size

You can reduce the Internal Database current size by changing the catalog protection settings for a complete backup, object copy, or object consolidation session (all objects in the session) or for specific objects only.

The effect of reducing the catalog protection is that browsing is not possible for some restores (namely of those backups that have exceeded the catalog protection).

This action does not influence the IDB growth in the future.

The change takes effect:

- If the catalog protection is removed from all objects on a medium.
- Once per day (by default, at noon) when Data Protector automatically removes obsolete data from the IDB. You can specify the time using the `DailyMaintenanceTime` global option. Use the twenty-four hour clock notation.

You can start the purge immediately by running the `omnidbutil -purge -dcbf` command. For information on removing other obsolete items from the IDB, see the `omnidbutil` man page or the *HPE Data Protector Command Line Interface Reference*.

By changing the catalog protection, you change the protection for the restore browse information in the IDB only. The information is still stored on media. Therefore, if you export a medium and import it back, Data Protector rereads information about catalog protection from the media.

Changing catalog protection for a session

Changing the protection for a backup session changes the protection of all objects backed up in the session.

Steps

1. In the Context List, click **Internal Database**.
2. In the Scoping Pane, expand the **Sessions** item.
3. Right-click the session(s) for which you want to change protection and click **Change Catalog**

Protection.

4. Specify the new catalog protection for the session(s), and then click **Finish** to apply the changes.

Changing catalog protection for an object

Changing the protection for a specific object changes the protection of this object regardless of the session it was backed up with.

Steps

1. In the Context List, click **Internal Database**.
2. In the Scoping Pane, expand the **Objects** item.
3. Right-click the object(s) for which you want to change protection and click **Change Catalog Protection**.
4. Specify the new catalog protection for the object(s), and then click **Finish** to apply the changes.

Extending the IDB Size

Due to free disk space shortage for the detail part of the IDB (names, versions, and metadata of the backed up objects), you may need to extend the Internal Database by creating new DC directories or reconfiguring existing ones for higher capacity.

Reconfiguring DC directories for higher capacity

You can reconfigure an existing DC directory by modifying its Allocation sequence, Maximum size, Maximum files, or Low space options. Note that the number and the current total size of files in the chosen DC directory may limit the adjustment range.

Steps

1. In the Context List, click **Internal Database**.
2. In the Scoping Pane, expand **Usage** and then **Detail Catalog Binary Files**.
3. Right-click the path of the chosen DC directory and click **Properties**.
4. In the Results Area, modify the available options as desired.
5. Click **Finish** to apply your changes.

IDB Consistency Check

The contents of the Internal Database must be logically correct, in other words, the IDB parts must be consistent and in order. You can manually perform consistency checks for specific parts and for the whole IDB.

Data Protector checks the consistency of the IDB by default before the IDB is backed up (quick check). This is extremely important for recovering the IDB and backed up data in case of a disaster on the Cell Manager.

IDB check type	What it checks	Command
Quick check of the IDB	The core (MMDB and CDB), the filenames, and the simple check of the DCBF parts.	omnidbcheck -quick
Simple check of the DCBF part	If the DC binary files exist and what their size is.	omnidbcheck -bf
Complete check of the DCBF part	The consistency of media positions and the DC binary files.	omnidbcheck -dc
Check of the SMBF part	Presence of session messages binary files.	omnidbcheck -smbf
Media consistency check	The consistency of media. It also lists inconsistent media names in case of a media consistency failure.	omnidbcheck -media_consistency
Schema consistency check	The consistency of IDB schema. Detects also all changes in the schema since its first creation during the Data Protector installation.	omnidbcheck -schema_consistency
Database consistency check	The consistency of the database. It also lists errors in case of a database consistency failure.	omnidbcheck -database_consistency
Extended check of the IDB	All checks with the exception of the SMBF are performed.	omnidbcheck -extended

Moving the IDB to a Different Cell Manager

You can move the Internal Database to a different Cell Manager that runs on the same operating system.

In a first scenario, where you perform a restore of the IDB from a backup device on a Data Protector client, proceed as follows:

Steps

1. Prepare a backup device *PreparedDevice* on the Data Protector client *client.company.com*.
2. Run the IDB backup using the backup device *PreparedDevice*.
3. Prepare the new Data Protector Cell Manager on the host *cmb.company.com* – clean installation.
4. Export the client *client.company.com* from the Cell Manager on the host *cma.company.com*.
5. Import the client *client.company.com* to the new Cell Manager on the host *cmb.company.com*.
6. Import the backup device *PreparedDevice* to the new Cell Manager.
7. Run the IDB restore from the *PreparedDevice* backup device.
8. Stop Data Protector services.

9. For all passwords (keystore-password, truststore-password, ssl password, and ca-certificate-password) located in the standalone.xml configuration file, use KeystorePassword from the webservice.properties configuration file.

These configuration files are available at the following locations:

Windows:

- ProgramData\OmniBack\Config\client\components\webservice.properties
- ProgramData\OmniBack\Config\server\AppServer\standalone.xml

UNIX:

- /etc/opt/omni/client/components/webservice.properties
- /etc/opt/omni/server/AppServer/standalone.xml

10. Start Data Protector services.
11. Import clients from the original Cell Manager to the new Cell Manager.
Each client has to be previously exported from the original Cell Manager.
12. Reconnect the GUI to the new Cell Manager.

In a second scenario, where you perform a restore of the IDB from a backup device on the original Cell Manager, proceed as follows:

Steps

1. Prepare a backup device *PreparedDevice* on the original Cell Manager.
2. Run the IDB backup using the backup device *PreparedDevice*.
3. Prepare the new Data Protector Cell Manager on the host *cmb.company.com* – clean installation.
4. Export the backup device *PreparedDevice* from the original Cell Manager.
5. Import the backup device *PreparedDevice* to the new Cell Manager on the host *cmb.company.com*.
6. Run the IDB restore from the *PreparedDevice* backup device.
7. Stop Data Protector services.
8. For all passwords (keystore-password, truststore-password, ssl password, and ca-certificate-password) located in the standalone.xml configuration file, use KeystorePassword from the webservice.properties configuration file.

These configuration files are available at the following locations:

Windows:

- ProgramData\OmniBack\Config\client\components\webservice.properties
- ProgramData\OmniBack\Config\server\AppServer\standalone.xml

UNIX:

- /etc/opt/omni/client/components/webservice.properties
- /etc/opt/omni/server/AppServer/standalone.xml

9. Start Data Protector services.
10. Import clients from the original Cell Manager to the new Cell Manager.
Each client has to be previously exported from the original Cell Manager.
11. Reconnect the GUI to the new Cell Manager.

Customizing the Data Protector Global Options

In the Data Protector global options file, you can modify values of global options or add new ones.




Prerequisites

- Your user account must be a member of a Data Protector `Admin` user group.

Setting the global options using GUI

Steps

To set global options using the GUI:

1. In the Context List, click **Internal Database**.
2. In the Scoping Pane, under **Internal Database**, click **Global Options**.
In Results area, the **Data Protector Global Options** table is displayed, consisting of six columns:
 - Group - represents the contextual section the option belongs to.
 - In use - indicates the status of an option. Selected options are active, while the empty check box indicates the inactive options that are commented out in the global options file.
 - Name
 - Origin - indicates the file which the option is loaded from.
 - Value - represents the value to which the option is currently set.
 - Description - informs you how to use the option.
3. To modify an option - in the Results Pane, in the Value column - click on the value you want to change, click the Edit icon  and enter a new one. Click **Save** to save the option.
To add an option, click the Add icon , fill in the dialog box with option parameters and click **Add**.
4. At the top of the Results Pane, click the Save icon .

You can also modify multiple rows before saving.

To change the table appearance, use the filters in the table headings.

In case anything goes wrong during the saving process, a copy of the original global options file named `global.old` is made in the global options folder.

Customizing Options By Editing The Global File

Besides using the GUI, you can edit the `global` file in a text editor to set the Data Protector global options.

Caution: HPE recommends using the GUI to set the global options, as it ensures validation of changes upon saving and reduces the chance of issues arising from the out-of-range or invalid settings, accidental deletions, typographical or spelling errors.

Steps

1. Open any text editor
2. In the text editor, open the `global` file, located in the default Data Protector server configuration directory, in the `options` subdirectory.
3. To activate an option, remove the `#` mark in front of its name and set it to the desired value.
4. Save the file in the Unicode format.

Configuration of IDB Reports

You can configure the Internal Database report so that you are informed when you need to perform some of the IDB maintenance tasks, such as extending the IDB size and reducing the IDB growth.

IDB reports

Report	Informs you ...
Internal Database Size Report	... about the size of the particular parts of the IDB.

Configuration of IDB Notifications

Configure the Internal Database notifications so that you are informed when you need to perform some of the IDB maintenance tasks, such as extending the IDB size, checking the IDB consistency, and so forth.

IDB notifications

Notification	Informs you ...
IDB Space Low	... if the IDB is running out of space.
IDB Limits	... if any of the MMDB or CDB parts has reached its

	limit.
IDB Backup Needed	... if an IDB backup does not occur frequently or there are too many successive incremental IDB backups.

Restoring the IDB

You can restore the Internal Database (IDB) from a backup image created in the standard IDB backup procedure. If the IDB is corrupted, you cannot use this restore procedure but you need to perform one of the IDB recovery methods.

To restore the Internal Database, perform the following procedure:

- [Restoring the IDB](#)

When restoring from an encrypted IDB backup, additional steps are required before the actual restore:

- [Preparing for IDB restore from an encrypted backup](#)

Restoring the IDB

During online Internal Database restore, the basic IDB parts (CDB, MMDB, SMBF) can only be restored to a location different from the original, while the Cell Manager configuration data and the Detail Catalog Binary Files (DCBF) part of the IDB can be either restored to their original or different locations.

Prerequisites

- Depending on the size of your Internal Database backup image, make sure there is enough free disk space available on the Cell Manager.

Limitations

Use of the restored IDB as a new IDB through the option "use the restored database as new internal database" is not supported on the SG cluster setup. You can set the omnirc variable `OB2SGENABLED`, which provides the procedure to use the restored IDB as a new IDB, in the session report. On setting the omnirc variable you will see the following message in the session report:

```
[Warning] From: OB2BAR_POSTGRES_BAR@<host name> "DPIDB" Time: <date time>
```

```
[175:316]Automatic replacement of the Internal Database on cluster environment not supported.
```

Click the error number in the message for details on the procedure that needs to be followed for using the restored IDB as a new IDB.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand **Restore Objects**, and then expand **Internal Database**.

3. Expand the Cell Manager from where the IDB was backed up from, and click **Internal Database**.
4. On the Internal Database property page, to restore the basic Internal Database parts, keep the **Restore Internal Database** option selected. The basic parts of the IDB are the Catalog Database (CDB), the Media Management Database (MMDB), and the Session Messages Binary Files (SMBF). Specify the temporary port to be used for the Internal Database Service during the restore, and the location to which the basic IDB parts should be restored to.

Additionally, decide whether to perform the Internal Database recovery using the archived log files, and if the restored IDB should be put into use as the new Internal Database of the cell.
5. Select **Restore catalog binary files** to restore the DCBF part of the IDB, and choose its restore location: original or custom.
6. Specify whether Data Protector should restore the IDB to a specific point in time which is not the time of the latest IDB backup image creation. In this case, the basic Internal Database part will be restored to the latest backed up state before the specified time.
7. On the **Configuration Files** property page, make your choice about the restore of the Cell Manager configuration data. If this data is selected for restore, you should also specify its backup object version, the restore location, and decide how Data Protector will handle the configuration files that still exist at their original location.
8. On the **Options** property page, specify the optional pre-exec and post-exec commands for the restore session.
9. On the **Devices** property page, make your choice about which devices to use in the session.
10. On the **Media** property page, review the backup media that will be used for restoring the IDB. Optionally, adjust their priorities Data Protector will consider during the session.
11. From the Actions menu, select **Start Restore**, or in the Results Pane, click **Restore**.
12. Click **Finish**.

After a point-in-time IDB restore session, copy specific files from the `auditing_IDBRestoreSessionID_NNNNNNNNNN` directory to the original `auditing` directory. This will make auditing information consistent with the state of the restored IDB. The following audit logs should be copied:

`YYYY_MM_DD.med`

`YYYY_MM_DD.obj`

`YYYY_MM_DD.ses`

In the above filenames, the `YYYY`, `MM`, and `DD` strings correspond the date specified with the **Restore until** option on the Internal Database property page.

Note: After the restore, you may want to check the consistency of the IDB.

Preparing for IDB restore from an encrypted backup

During an encrypted IDB backup, encryption keys are automatically exported to the `IDB-ClientName-keys.csv` file located in the default Data Protector exported encryption keys directory.

Before restoring the IDB, proceed as follows:

Steps

1. Transfer the `IDB-ClientName-keys.csv` file to the Cell Manager where you will perform the IDB restore.
2. Import the key by running:

```
omnikeytool -import CSVFile
```

The Cell Manager will use the key from the online KMS to decrypt the data on the medium containing the IDB backup.

About IDB Recovery

The Internal Database recovery is needed if all or some of the IDB files are not available or they are corrupted.

There are three levels of IDB issues, each having its own fix:

- Troubleshoot IDB problems that are caused by operating system configuration issues, such as filesystems not mounted, nameservice problems, and so on.
- Omit or remove non-core parts (binary files) of the IDB that contain problems. This is possible if the identified level of IDB corruption is minor (corruption is not in the core part of the IDB).
- Perform complete recovery consisting of IDB restore and updating the IDB beyond the last IDB backup. This is a must if the identified level of IDB corruption is critical (corruption is in the core part).

Complete recovery (restore and update the IDB beyond the last IDB backup)

Complete recovery consists of two phases:

1. IDB restore, which gets the IDB to the last (available) consistent state.
2. Updating the IDB from the last consistent state up to the last moment when the IDB was still operational.

Depending on how well you prepared for IDB recovery before problems occurred (availability of the IDB recovery file, the IDB backup images, the original backup device, and the archived logs files), the recovery procedure can differ. If all these are available, you can use a very convenient IDB recovery method, guided autorecovery.

Overview of IDB Recovery Methods

Several recovery methods are available for recovering the Internal Database. Depending on the identified level of corruption, your requirements, and the availability of the IDB recovery file, the original backup device, and the archived logs files, the recovery procedure can differ.

The most convenient complete recovery

This recovery method guides you through restoring the IDB and replaying archived log files. If the archived log files are not available, you can still update the IDB by importing all media since the last IDB backup.

Corruption level	Problem type	Current situation	Recovery procedure
Critical	The complete IDB is missing or the core part is corrupted.	The IDB recovery file and the original device used for the IDB backup are available.	Perform the Guided Autorecovery (IDB Restore and Replay Archived Log Files) if possible. Otherwise, follow one of the methods given under "More recovery methods".

Omitting (removing) corrupted IDB parts

If the identified level of corruption is minor (corruption is not in the core part), you can consider omitting (removing) the missing or corrupted parts of the IDB or perform the complete IDB recovery instead.

Corruption level	Problem type	Recovery procedure
Minor	DC binary files are missing or corrupted.	Handle Minor IDB Corruption in the DCBF Part

More recovery methods

These recovery procedures are adapted to specific situations. They assume that you want to recover the complete IDB, but for some reason you cannot perform the guided autorecovery method. The recovery consists of restoring the IDB and updating the IDB.

Restore

Current situation	Remark	Recovery procedure (restoring IDB)
The IDB recovery file is available but the original device used for the IDB backup has changed.	The method is essentially the same as the guided autorecovery method, but less guided, more complex, and time consuming.	Restore the IDB Using IDB Recovery File and Changed Device
The IDB recovery file is not available.	The method is essentially the same as the guided autorecovery method, but less guided, more complex, and time consuming.	Restore the IDB Without IDB Recovery File

You want to recover the IDB from a specific IDB backup (not the latest one).	This method does not provide the latest state of the IDB as a result.	Restore the IDB from a Specific IDB Session
------------------------------------------------------------------------------	-----------------------------------------------------------------------	-------------------------------------------------------------

Update the IDB since the last IDB backup

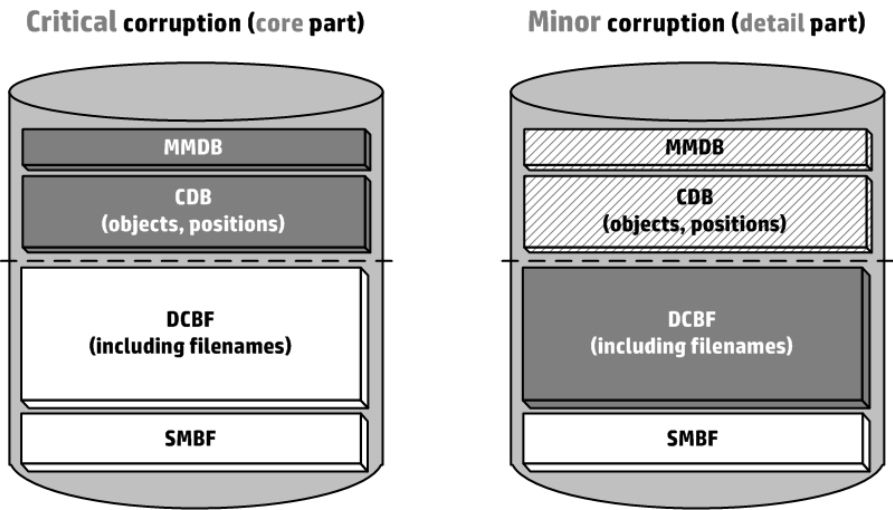
Current situation	Recovery procedure (updating the IDB)
The archived log files are not available.	Update IDB by Importing Media

IDB Corruption Levels

There are two levels of Internal Database corruption: critical and minor. The level depends on the part of the IDB where the corruption occurs.

You can use the IDB consistency check to determine which part of the IDB is corrupted.

Depending on the level of corruption, the IDB recovery procedure differs.



Identifying the Level of IDB Corruption

Identify the level of corruption in order to choose the appropriate Internal Database recovery method.

Steps

1. Identify the level of corruption using the `omnidbcheck -extended` command.

Note: The extended check may take a considerable amount of time. You can run parts of the `omnidbcheck` command instead. For example, run the `omnidbcheck -connection` to identify if the connection to the IDB is working.

After identifying the level of corruption, perform the appropriate recovery procedure.

Performing Guided Autorecovery (IDB Restore and Replay Archived Log Files)

Guided autorecovery is the most convenient Internal Database recovery method. You can perform it if the IDB recovery file and the original device used for the IDB backup together with the IDB backup medium are available.

This method guides you through restoring the IDB and replaying archived logs files since the last IDB backup. If the archived log files are not available, you can still update the IDB since the last IDB backup by importing media.

Transaction replay updates the core part of the IDB. Binary files are not updated and changes to binary files are lost. The following is not available for the backups that were running from the last IDB backup until the IDB corruption:

- Session messages
- Browsing of file versions (restores of complete objects are possible). Perform the import catalog on the media used by the backups to recover the changes.

Prerequisites

- Depending on the size of your Internal Database backup image, ensure there is enough free disk space available on the Cell Manager.
- Ensure the Cell Manager has twice as much total RAM as documented among Data Protector Cell Manager installation requirements in the *HPE Data Protector Product Announcements, Software Notes, and References*. If the Cell Manager is a UNIX system, ensure its kernel parameter `shmmax` is set to twice the required value documented in the same section.
- Mount a disk of the same size as before the disaster on the same directories as at the IDB backup time (on Windows systems, the same drive letters must be assigned). If this cannot be ensured, follow the procedure for recovering the IDB to a different disk/volume layout. You can use the `-preview` option of the `omniofflr` command to see where the files will be restored.
- Install Data Protector on the Cell Manager and the system where a device is attached (preferably, the device used for the IDB backup.)
- If the IDB is installed on HPE Serviceguard, the following commands have to be run on the active node before performing the guided autorecovery:
 - a. `cmhaltpkg PackageName`, where *PackageName* is the name of the Data Protector cluster package. This command stops the Data Protector package and dismounts the Data Protector shared volume group.
 - b. `vgchange -a e /dev/vg_name`, where *vg_name* is the name of Data Protector shared volume group. This command activates the Data Protector shared volume group. To list volume groups on your system, run `ll /dev/*/group`.
 - c. `mount /dev/vg_name/Lv_name/MountPoint`, where *MountPoint* is the name of the mount point for the Data Protector shared volume group. This command mounts the Data Protector shared volume group.

When the guided autorecovery has finished, run the `cmrunpkg PackageName` command on the active node to start the Data Protector package.

- If the IDB is installed on a Symantec Veritas Cluster Server, take the Data Protector application resource offline on the active node before performing the guided autorecovery.

When the guided autorecovery has finished, bring the Data Protector application resource online on the active node to start the Data Protector service.

- If the IDB is installed on Microsoft Cluster Server, take the `OOBVS_HPDP_AS`, `OBVS_HPDP_IDB`, and `OBVS_HPDP_IDB_CP` cluster groups offline using the Cluster Administrator utility and stop the `Inet` service on the active node before performing the guided autorecovery. When the guided autorecovery has finished, bring the `OBVS_HPDP_AS`, `OBVS_HPDP_IDB`, `OBVS_HPDP_IDB_CP`, and `OBVS_MCRS` cluster groups online using the Cluster Administrator utility and restart the `Inet` service.

Steps

1. Run the `omniofflr -idb -autorecover` command.

The command reads the IDB recovery file and if IDB backups are logged to the file, it stops the services and starts restore of the IDB back in place. All the options are generated automatically using data from the IDB recovery file.

Once the restore is complete, the `omniofflr` checks if archived log files are available to be replayed. If the log files are available, you are asked to confirm the replay of the logs. If this step is cancelled or archived log files are not available, output informs you how to update the IDB since the last IDB backup by:

- importing media
- finding the archived log files and replaying them later

Once you replay the log files or import media to update the IDB, the complete IDB should be successfully recovered.

Handling Minor IDB Corruption in the DCBF Part

If you detect that the Internal Database corruption is of minor severity, it means that some DC binary files are missing or corrupted. If this is the case, there is no need for complete IDB recovery. You can easily recreate the binary files by importing catalog from media. Choose the recovery procedure depending on the corruption type:

Recovery if DC binary files are missing

DC binary files are organized so that one binary file exists for each medium. If some DC binary files are missing, media positions of some media point to the non-existent files. An error message is displayed when browsing the relevant filesystems.

Steps

1. From the `omnidbcheck -bf` output, identify the Medium ID of the missing binary file. Run the `omnimmm -media_info medium-id` command to get other attributes of the medium, such as

medium label and media pool.

2. Run the `omnidbutil -fixmpos` command to establish consistency between media positions (mpos) and binary files.
3. Import catalog from media to recreate the binary files.

Recovery if DC binary files are corrupted

If some DC binary files are corrupted, you can remove the DC binary files and recreate them by importing the media with proper logging level. The only impact of removing the files is that some media positions point to the non-existent binary files, and thus an error message is displayed when browsing the relevant filesystems.

Steps

1. From the `omnidbcheck -dc` output, identify the Medium ID of the corrupted DC binary file. Run the `omnimm -media_info medium-id` command to get other attributes of the medium, such as medium label and media pool.
2. Identify the DC binary file for the affected medium. DC binary files are named: *MediumID_TimeStamp.dat* (in the MediumID, colons ":" are replaced with underscores "_").
3. Remove the corrupted DC binary files.
4. Run the `omnidbutil -fixmpos` command to establish consistency between media positions (mpos) and binary files.
5. Import catalog from media to recreate the binary files.

Restoring the IDB Using IDB Recovery File and Changed Device

Use this procedure to restore the Internal Database if the IDB recovery file is available but the original device used for the IDB backup is different from the one to be used for recovery or the medium is located in a different slot.

Prerequisites

- Mount a disk of the same size as before the disaster on the same directories as at the IDB backup time (on Windows systems, the same drive letters must be assigned). If this cannot be ensured, follow the procedure for recovering the IDB to a different disk/volume layout. You can use the `-preview` option of the `omniofflr` command to see where the files will be restored.
- If possible, move the `media.log` file from the previous installation to a safe place. It will provide you with the information about the media used since the last IDB backup. This is very helpful for updating the IDB if archived logs files are not available.
- Install Data Protector on the Cell Manager and a system where a device is attached (preferably the device used for the IDB backup.)
- If the IDB is installed on HPE Serviceguard, the following commands have to be run on the active node before performing the guided autorecovery:
 - a. `cmhaltpkg PackageName`, where *PackageName* is the name of the Data Protector cluster package. This command stops the Data Protector package and dismounts the Data Protector

shared volume group.

- b. `vgchange -a e /dev/vg_name`, where *vg_name* is the name of Data Protector shared volume group. This command activates the Data Protector shared volume group. To list volume groups on your system, run `ll /dev/*/group`.
- c. `mount /dev/vg_name/Lv_name/MountPoint`, where *MountPoint* is the name of the mount point for the Data Protector shared volume group. This command mounts the Data Protector shared volume group.

When the guided autorecovery has finished, run the `cmrunpkg PackageName` command on the active node to start the Data Protector package.

- If the IDB is installed on a Symantec Veritas Cluster Server, take the Data Protector application resource offline on the active node before performing the guided autorecovery.

When the guided autorecovery has finished, bring the Data Protector application resource online on the active node to start the Data Protector service.

- If the IDB is installed on Microsoft Cluster Server, take the `OBVS_HPDP_AS`, `OBVS_HPDP_IDB`, and `OBVS_HPDP_IDB_CP` cluster groups offline using the Cluster Administrator utility and stop the `Inet` service on the active node before performing the guided autorecovery. When the guided autorecovery has finished, bring the `OBVS_HPDP_AS`, `OBVS_HPDP_IDB`, `OBVS_HPDP_IDB_CP`, and `OBVS_MCRS` cluster groups online using the Cluster Administrator utility and restart the `Inet` service.

Steps

1. Run the following command to create a text file with the restore job options:

```
omniofflr -idb -autorecover -save C:\TEMP\restjob.txt -skiprestore -logview
```

The specified `-logview` command lists first archived log files next to the session IDs. Remember the first archived log file for the session you want to restore, because you need it in order to update the IDB after the restore. For example, from the output `2013/02/09-2 AAAAAAH`, you would remember the first archived log file `AAAAAAH` in order to restore the `2013/02/09-2` session.

The created `restjob.txt` file has information on original devices and on slots in which media were originally located (at IDB backup time).

2. Modify the `restjob.txt` file to specify the current device or the slot in which the media are currently located.
3. Run the restore with the `omniofflr -idb -read C:\TEMP\restjob.txt` command.

The command guides you through restoring the IDB and replaying archived log files beyond the last IDB backup. If the archived log files are not available, you can still update the IDB by importing all media used since the last IDB backup.

Restoring the IDB Without IDB Recovery File

Use this procedure to restore the Internal Database if the IDB recovery file is not available.

Prerequisites

- Mount a disk of the same size as before the disaster on the same directories as at the IDB backup time (on Windows systems, the same drive letters must be assigned). If this cannot be ensured,

follow the procedure for recovering the IDB to a different disk/volume layout. You can use the -preview option of the `omniofflr` command to see where the files will be restored.

- If possible, move the `media.log` file from the previous installation to a safe place. It will provide you with information about the media used since the last IDB backup. This is very helpful for updating the IDB if archived logs files are not available.
- Install Data Protector on the Cell Manager and a system where a device is attached (preferably, the device used for the IDB backup.)
- If the IDB is installed on HPE Serviceguard, the following commands have to be run on the active node before performing the guided autorecovery:
 - a. `cmhaltpkg PackageName`, where *PackageName* is the name of the Data Protector cluster package. This command stops the Data Protector package and dismounts the Data Protector shared volume group.
 - b. `vgchange -a e /dev/vg_name`, where *vg_name* is the name of Data Protector shared volume group. This command activates the Data Protector shared volume group. To list volume groups on your system, run `ll /dev/*/group`.
 - c. `mount /dev/vg_name/lv_name/MountPoint`, where *MountPoint* is the name of the mount point for the Data Protector shared volume group. This command mounts the Data Protector shared volume group.

When the guided autorecovery has finished, run the `cmrunpkg PackageName` command on the active node to start the Data Protector package.

- If the IDB is installed on a Symantec Veritas Cluster Server, take the Data Protector application resource offline on the active node before performing the guided autorecovery.

When the guided autorecovery has finished, bring the Data Protector application resource online on the active node to start the Data Protector service.

- If the IDB is installed on Microsoft Cluster Server, take the `OBVS_HPDP_AS`, `OBVS_HPDP_IDB`, and `OBVS_HPDP_IDB_CP` cluster groups offline using the Cluster Administrator utility and stop the `Inet` service on the active node before performing the guided autorecovery. When the guided autorecovery has finished, bring the `OBVS_HPDP_AS`, `OBVS_HPDP_IDB`, `OBVS_HPDP_IDB_CP`, and `OBVS_MCRS` cluster groups online using the Cluster Administrator utility and restart the `Inet` service.

Steps

1. Configure the device using the Data Protector GUI.
2. Find the medium with the latest IDB backup.
3. Insert the medium into the device and use the following command to display the contents of the medium:

```
omnimlist -dev device_name
```

For the IDB restore you need the Medium ID and Disk Agent ID for the backup session you want to restore.

4. Use the following command to display the information on the device configuration:

```
omnidownload -dev device_name
```

For the IDB restore you need the following information:

- Mahost (Media Agent host)
 - Policy (number): A policy number can be obtained using the following translation: 1 for Standalone devices, 3 for Stacker devices, 5 for Jukebox devices, 6 for External control devices, 8 for GRAU DAS library, 9 for StorageTek ACS library, and 10 for SCSI library.
 - Media type (number): Media type numbers are defined as media class in the `scsitab` file. For location, see the topic [Support of New Devices](#).
 - SCSI address
 - Robotics SCSI address (only if using Exchanger library devices)
5. Run the `omniofflr` command using the obtained information:

```
omniofflr -idb -policy PolicyNumber -type MediaTypeNumber [-ioctl  
RoboticsSCSIAddress] -dev SCSIAddress -mahost MAClientName -maid MediumID -daid  
DiskAgentID
```

For example, you would use the following command to restore the IDB from a backup session with the medium ID `0100007f:3a486bd7:0410:0001` and the Disk Agent ID `977824764`, performed using a standalone device of the type DLT, connected to the system `company.dot.com`, and with the SCSI address `scsi0:1:2:0`:

```
omniofflr -idb -policy 1 -type 10 -dev scsi0:1:2:0 -mahost company.dot.com -  
maid 0100007f:3a486bd7:0410:0001 -daid 977824764
```

The command guides you through restoring the IDB and replaying archived log files since the last IDB backup. If the log files are not available, you can still update the IDB by importing all media used since the last IDB backup.

Restoring the IDB from a Specific IDB Session

Use this procedure to restore the Internal Database from a backup other than the latest one if the IDB recovery file is available.

Prerequisites

- Mount a disk of the same size as before the disaster on the same directories as at the IDB backup time (on Windows systems, the same drive letters must be assigned). If this cannot be ensured, follow the procedure for recovering the IDB to a different disk/volume layout. You can use the `-preview` option of the `omniofflr` command to see where the files will be restored.
- If possible, store the `media.log` file from the previous installation to a safe place. It will provide you with information about the media used since the last IDB backup. This is very helpful for updating the IDB if archived logs files are not available.
- Install Data Protector on the Cell Manager and a system where a device is attached (preferably the device used for the IDB backup.)
- If the IDB is installed on HPE Serviceguard, the following commands have to be run on the active node before performing the guided autorecovery:
 - a. `cmhaltpkg PackageName`, where *PackageName* is the name of the Data Protector cluster package. This command stops the Data Protector package and dismounts the Data Protector

shared volume group.

- b. `vgchange -a e /dev/vg_name`, where *vg_name* is the name of Data Protector shared volume group. This command activates the Data Protector shared volume group. To list volume groups on your system, run `ll /dev/*/group`.
- c. `mount /dev/vg_name/Lv_name/MountPoint`, where *MountPoint* is the name of the mount point for the Data Protector shared volume group. This command mounts the Data Protector shared volume group.

When the guided autorecovery has finished, run the `cmrunpkg PackageName` command on the active node to start the Data Protector package.

- If the IDB is installed on a Symantec Veritas Cluster Server, take the Data Protector application resource offline on the active node before performing the guided autorecovery.

When the guided autorecovery has finished, bring the Data Protector application resource online on the active node to start the Data Protector service.

- If the IDB is installed on Microsoft Cluster Server, take the `OBVS_HPDP_AS`, `OBVS_HPDP_IDB`, and `OBVS_HPDP_IDB_CP` cluster groups offline using the Cluster Administrator utility and stop the `Inet` service on the active node before performing the guided autorecovery. When the guided autorecovery has finished, bring the `OBVS_HPDP_AS`, `OBVS_HPDP_IDB`, `OBVS_HPDP_IDB_CP`, and `OBVS_MCRC` cluster groups online using the Cluster Administrator utility, restart the `Inet` service, and run the `omnidbutil -fixmpos` command.

Steps

1. Check all backups using the following command:

```
omniofflr -idb -autorecover -logview -skiprestore
```

2. Choose the backup session you want to restore from and perform the IDB restore by running the command:

```
omniofflr -idb -autorecover -session SessionID
```

The command guides you through restoring the IDB and replaying archived log files since the last IDB backup. If the archived log files are not available, you can still update the IDB by importing all media used since the last IDB backup.

Restoring the IDB database on a different Cell Manager host

Use this procedure for Recovery of IDB database on a different Cell Manager host.

1. Install Data Protector on a new Cell Manager host and import the device containing the IDB backup of the old Cell Manager host.
2. Restore only the configuration files to a new location. For example, `/tmp/idb/config`.
3. Make a copy of the original file `/etc/opt/omni/server/cell/cell_info`.
4. Restore the complete IDB database to a new location. For example, `/tmp/idb/newidb`.
 - For restore of database files, select options **StartDatabaseServer** and **UseRestoredDatabaseAsNewDatabase**.
 - For catalog binary files as destination, select Restore to original location.

- For configuration files as destination, select **Restore to original location** and select conflict resolution **Overwrite**.
5. After restore is completed without any errors, create a copy of the following original files (as a precaution):
 - /etc/opt/omni/server/AppServer/standalone.xml
 - /etc/opt/omni/server/idb/idb.config
 - /etc/opt/omni/server/idb/ulist
 6. Stop Data Protector services by running the following command: `/opt/omni/sbin/omnisv stop`
 7. Overwrite the following file: `/etc/opt/omni/server/cell/cell_info` with a copy of the file made in step 3.
 8. Open the file `/etc/opt/omni/server/AppServer/standalone.xml` in a preferred editor, then find `keystore-password` and `truststore-password`, and make a note of it. They are usually the same.
 9. Open the file `/etc/opt/omni/client/components/webservice.properties` in a preferred editor, replace `keystore-password`, and `truststore-password` with values found in the `standalone.xml` file, save the changes and close the file.

NOTE: In a clustered environment, you must edit the `webservice.properties` file on all nodes of the cluster.

10. Regenerate the certificate by running the following command: `/opt/omni/bin/perl /opt/omni/sbin/omnigencert.pl -server_id <hostname> -user_id hdpd -store_password <your keystore password>`
11. Make sure the following files do not contain the hostname of the old Cell Manager:
 - /etc/opt/omni/client/components/dp-jobexecutionengine-backup\webservice.properties /etc/opt/omni/client/components/dp-jobexecutionengine-consolidation\webservice.properties
 - /etc/opt/omni/client/components/dp-jobexecutionengine-copy\webservice.properties
 - /etc/opt/omni/client/components/dp-jobexecutionengine-verification\webservice.properties
 - /etc/opt/omni/client/components/dp-loginprovider\webservice.properties
 - /etc/opt/omni/client/components/dp-scheduler-gui\webservice.properties
 - /etc/opt/omni/client/components/dp-webservice-server\webservice.properties
 - /etc/opt/omni/client/components/jce-dispatcher\webservice.properties
 - /etc/opt/omni/client/components/jce-serviceregistry\webservice.properties
 - /etc/opt/omni/client/components/webservice.properties

12. Add the following variable to the omnirc file - /opt/omni/.omnirc: OB2_CERT_VERIFYHOST=0. If the omnirc file does not exist, then create an empty text file and rename it to .omnirc or rename .omnirc.TMPL to .omnirc
13. Start the Data Protector services by running the following command: /opt/omni/sbin/omnisv start
14. Run the following command to change ownership of some Data Protector files:
/opt/omni/sbin/omnidbutil -change_cell_name <old_cm_hostname>
15. Run the following command to clear running sessions: /opt/omni/sbin/omnidbutil -clear
16. In the Windows GUI client, delete the folder with the old certificate. After you start the Data Protector services, the Data Protector GUI will import a new certificate from Cell Manager. You can find the old certificate in the following path:
C:\Users\<USERNAME>\AppData\Local\Hewlett-Packard\Data Protector\ca\<NEW_CM_HOSTNAME>"
17. Execute following additional, non-mandatory, steps:
 - a. Run the following command to confirm that IDB is using files from a new location (tablespace files and writeahead logs are in new location while DCBF's are in original folder):
/opt/omni/sbin/omnidbutil -show_db_files
 - b. Update files that contain the hostname of the old Cell Manager (usually in the userlist, barlists, and configuration files). You can find them by running the following command: grep -rnw /etc/opt/omni -e <OLD_CM_HOSTNAME>
 - c. Reconfigure the devices to use the new Cell Manager.

Updating IDB by Importing Media

If archived logs files are not available, update the Internal Database by importing all media used since the last IDB backup. Do this once the IDB restore has finished.

Steps

1. Start Data Protector processes and services.
2. Increase the session counter. When you initialized and restored the IDB, the counter was set to 0. Therefore, any new sessions would have the same session ID as a session already started that day.
The following command sets the session counter to 200, which suffices for most cases:
omnidbutil -set_session_counter 200
If necessary, you can now start with backups.
3. Export and import the media with the last IDB backup. This creates consistent information about the last IDB backup.
4. Import (export if already in IDB) the media used between the last IDB backup and the time of the IDB recovery. For a list of used media, see the media.log file residing in the default Data Protector server log files location.
5. Run the omnidbcheck command.

The complete IDB should be successfully recovered.

Note: If you are recovering an IDB that encompasses a CMMDB or a remote MMDB to a different disk layout, run the `omnidbutil -cdbsync` command after updating the IDB.

Chapter 5: Manager-of-Managers Environment

About MoM Environment

The Data Protector Manager-of-Managers concept allows administrators to manage a large environment, also known as enterprise backup environment, with multiple Data Protector cells centrally from a single point.

This way almost unlimited growth of the backup environment can be handled: new cells can be added or the existing ones can be split into several.

Note that each MoM client and the MoM Manager need to run the same version of Data Protector.

Manager-of-Managers provides the following features:

- Centralized management of all tasks
Data Protector enables configuration, management, and control of the enterprise backup environment from the single point. This includes configuring backup, restore, media management, monitoring, and reporting about the status of the whole backup environment.
- Centralized Media Management Database (CMMDB)
Optionally, all the cells in the environment can share a common, central database to manage devices and media within the enterprise. CMMDB enables you to share high-end devices and media across several cells in a MoM environment. This makes all devices of one cell (using CMMDB) accessible to other cells that use the CMMDB.
- Centralized license management
Data Protector enables you to configure centralized licensing for the whole MoM environment. All Data Protector licenses are installed and kept on the MoM Manager. You allocate licenses to specific cells to suit your needs.

About CMMDB

In large multicell environments with high-end backup devices, you may want to share the devices and media among several cells. This can be achieved by having one Centralized MMDB for all the cells and keeping an individual CDB for each cell. This allows media and device sharing while preserving the security capabilities of the multicell structure.

How media are shared

With the CMMDB, media can only be owned by the Data Protector cell that performed the first backup on those media. The media owner is displayed in the media view. While media are protected, only backups from that cell can be appended on the media. Each medium with protected data on it has information showing which cell currently owns the data. Once the protection expires, the media become available to other cells again.

How media are initialized

If a tape has been initialized by one cell, any other cell can use it as long as it does not have any protected data on it. If a tape is loaded in a library and not yet initialized, any cell can initialize it, assuming that there is a loose policy and no other tapes are available. The media allocation rules apply in exactly the same way to shared tapes, except that appendable media can only be appended by the cell that owns them.

Consider the following:

- The centralized MMDB has a significant effect on licensing. Immediately after the MMDB is changed from local to remote, all the licenses associated with libraries and devices are taken (validated) from the MoM Manager and should be removed from client cells.
- A cell in the enterprise environment must have access to the CMMDB to be able to run a backup. For example, this happens if a network failure occurs between the cell and the MoM cell. A reliable network connection is required between the MoM cell and the other Data Protector cells.

MoM Environment Configuration Procedure

Prerequisites

- You must choose a system for the MoM Manager. You must choose a highly reliable system that is a Data Protector Cell Manager with the software installed.
- Install the required licenses on the MoM cell and on every prospective MoM client cell.

MoM environment configuration procedure

The MoM environment configuration procedure consists of several phases. You need to:

1. Set up the MoM Manager.
2. Import Data Protector cells into the MoM environment.
3. Create a Data Protector user in the `admin` user group on every cell in the MoM environment who will act as MoM administrator.
4. Restart Data Protector services.

Optionally, you can also configure a Centralized Media Management Database, configure centralized licensing, and distribute the MoM configuration.

Setting Up MoM Manager

To set up an enterprise environment, configure one of your Cell Managers as an MoM Manager.

Steps

1. In the Context List, click **Clients**.
2. In the Actions menu, click **Configure CM as Data Protector Manager-of-Managers Server**.
3. Restart the Data Protector services.
4. Start the MoM User Interface by selecting **Data Protector Manager-of-Managers** in the Data Protector program group.

Alternatively, run the `mom` command from the `Data_Protector_home\bin` directory. For more information on the `mom` command, see the `omnigui` man page or the *HPE Data Protector Command Line Interface Reference*.

Adding a MoM Administrator to Cells

A MoM administrator can perform administration tasks in all cells in the enterprise environment.

Prerequisite

You need to have a certain user that is in the `admin` user group on every Cell Manager in the MoM environment. For example, you may have a user called `MoM_Admin`. This user will be the MoM administrator.

Steps

1. Using the Data Protector Manager, connect to each Cell Manager in the MoM environment as a member of the `admin` user group (the `User configuration` user right is required).
2. Add the user that will be the MoM Administrator to the Data Protector `admin` user group.

Importing Cells

Importing a cell into a MoM environment allows it to be centrally managed with the MoM Manager.

Cluster clients identify themselves to the MoM Manager with their virtual server names. If you import a cluster in a MoM environment, use only its virtual server name.

Prerequisites

- The active user must be a member of the `Admin` user group on the Cell Manager of the cell to be imported.

Steps

1. In the Data Protector Manager-of-Managers, click **Clients** in the Context List.
2. Right-click **Enterprise Clients** and then click **Import Cell Manager**.
3. Select a Cell Manager to import and click **Finish**.

Restarting the Data Protector Services in MoM

After you have configured the MoM environment, you are notified to restart the Data Protector services.

If you use the Windows `Service Control Manager` to stop and start services on the Cell Manager, only the current and previous copies of the database log are kept. Using the `omnisv` command will save all previous database logs.

Stopping the Data Protector services

Cell Manager in a non-cluster environment

Run the following command: `omnisv -stop`.

Cell Manager on HPE Serviceguard

Run the following command: `cmhaltpkg PackageName`, where *PackageName* is the name of the Data Protector cluster package.

This command stops the Data Protector package and dismounts the Data Protector shared volume group.

Cell Manager on Symantec Veritas Cluster Server

Take the Data Protector application resource offline.

Cell Manager on Microsoft Cluster Server

Take the `OBVS_HPDP_AS`, `OBVS_HPDP_IDB`, and `OBVS_HPDP_IDB_CP` cluster groups offline (using the `Cluster Administrator` utility on the active node).

Starting the Data Protector services

Cell Manager in a non-cluster environment

Run the following command: `omnisv -start`

Cell Manager on HPE Serviceguard

Restart the Data Protector package using the `cmrunpkg -n NodeName PackageName` command.

Cell Manager on Symantec Veritas Cluster Server

Bring the Data Protector application resource online.

Cell Manager on Microsoft Cluster Server

Bring the OBVS_HPDP_AS, OBVS_HPDP_IDB, OBVS_HPDP_IDB_CP, and OBVS_MCRS cluster groups online using the Cluster Administrator utility.

Configuring CMMDB

Set up CMMDB if you want to have central media management. If you do not set up a CMMDB, each cell will have its own IDB.

During the configuration, a local Media Management Database is merged into the CMMDB, if you select so. You can decide for each cell if it will use the CMMDB or its own local MMDB.

Once you have configured the CMMDB and start using it, you cannot split it back into local MMDBs. You should not try to recover the old state of an MMDB, but rather create a new MMDB from scratch.

Consideration

If you are configuring a new cell (and you do not yet have devices and media configured), there is no need to merge the database. You only want to merge cells with the CMMDB that already have devices and media configured.

Prerequisites

- Check that the Data Protector Cell Managers in all cells have the same version of Data Protector installed and running.
- Check that there are no backup, restore, or media management sessions running on any of the cells to be added to the CMMDB.

Configuring CMMDB on a client cell

Steps

1. Log on to the Cell Manager of the client cell as a member of the `admin` user group.
2. Create the file containing the name of the MMDB Server (fully qualified). On Windows systems, save the file in the Unicode format:
Windows systems: `Data_Protector_program_data\Config\server\cell\mmdb_server`
UNIX systems: `/etc/opt/omni/server/cell/mmdb_server`
3. Enable MoM Manager to establish connection to a cell, by modifying the `pg_hba.conf` file, located at the `pg` directory of the Internal Database location.
Open the file in text editor and add the line:
`host hdpidb hdpidb_app MoM_Server_IP_Address/32 trust`
after the following lines
`# IPv4 local connections:`

```
host all all 127.0.0.1/32 md5
```

Save the file.

Note: In case the Cell Manager on a MoM client is a part of a cluster environment, you need to specify either the IP address of all cluster nodes (one line item per node) or the subnet of the cluster in the `pg_hba.conf` file on the Cell Manager of the MoM client.

Open the file in text editor and add the line:

```
host hpdpidb hdpdpidb_app Cluster_Subnet trust
```

after the following lines

```
# IPv4 local connections:
```

```
host all all 127.0.0.1/32 md5
```

Save the file.

4. Restart the Data Protector services.
5. Update configuration files by running the following command:

```
omnicc -update_mom_server
```

Repeat the steps for all the client cells whose MMDB you want to merge into the CMMDB.

Configuring CMMDB on the MoM Manager

Steps

1. Log on to the Manager-of-Managers and copy the `idb tablespaces` directory to a temporary location for safety reasons.

The `idb` is a subdirectory at the Internal Database location.

2. Run the following command to merge the local MMDB into CMMDB:

```
omnidbutil -mergemmdb MoM_Client_Cell_Manager_Hostname
```

Make sure that the IDB service (`hdpd-idb`) port 7112 is opened on both, the MoM Manager and the client Cell Manager during the execution of the command. You can close the ports after the merge is done.

3. Run the following command to synchronize the local CDB:

```
omnidbutil -cdbsync MoM_Client_Cell_Manager_Hostname
```

4. Edit the duplicated names of media pools and devices. This duplication always happens to default pools if they exist on both cells. The duplicated names have a "_N" appended to their name, where N represents a number. In this case, manually change the backup specifications that use these devices to use the new device names.

Repeat the steps for all the client cells whose MMDB you want to merge into the CMMDB.

About Centralized Licensing

Centralized licensing means that all licenses are configured on the MoM Manager and can be allocated to specific cells as needed. Centralized licensing simplifies license management. Licensing

administration, including the distribution and moving of the licenses, is performed by the MoM administrator for all cells in the MoM environment.

Setting up centralized licensing is optional. Instead, individual licenses can be installed on each Cell Manager. These licenses are restricted to the cell on which they are installed and all licensing administration tasks have to be performed locally.

Setting Up Centralized Licensing

Set up centralized licensing to simplify license management in enterprise environments.

Prerequisite

If you are consolidating existing Data Protector cells into an MoM environment, send a request to the *HPE Password Delivery Center* to move the licenses from the existing Cell Managers to the new MoM Manager.

Steps

1. Log on to the MoM Manager and create the `licdistrib.dat` file:
Windows systems: `Data_Protector_program_data\Config\server\cell\licdistrib.dat`
UNIX systems: `/etc/opt/omni/server/cell/licdistrib.dat`
2. Log on to each Cell Manager in the MoM environment and create the `lic_server` file with the name of the MoM Manager:
Windows systems: `Data_Protector_program_data\Config\server\cell\lic_server`
UNIX systems: `/etc/opt/omni/server/cell/lic_server`
3. Stop and restart Data Protector services on each Cell Manager where you made the changes.
4. In the Data Protector Manager-of-Managers, click **Clients** in the Context List.
5. In the Scoping Pane, right-click the Cell Manager that has the licensing information you want to change and then click **Configure Licensing** to open the wizard. The types and numbers of licenses available to your selected Cell Manager are displayed.

Note: Identify a cluster client with its virtual hostname.

6. Click the **Remote** option to change the licensing from local to remote. The Used columns changes to Allocated.
7. Modify the license configuration. Only the Allocated column is available during the modification process.
 - To release (give up) a license type, thus increasing the number of available licenses, reduce its corresponding number in the Allocated column.
 - To assign a license type, increase its corresponding number in the Allocated column.
8. Click **Finish** to apply the configuration.
9. Repeat the steps for all Cell Managers for which you want to set up the centralized licensing.
10. Stop and restart the Data Protector processes using the `omnisv -stop` and `omnisv -start`

commands.

If the Cell Manager is configured on HPE Serviceguard, run the `cmhaltpkg PackageName` command to stop and the `cmrunpkg -n NodeName PackageName` to start the Data Protector package, where *PackageName* is the name of the Data Protector cluster package.

If the Cell Manager is configured on Symantec Veritas Cluster Server, take the Data Protector application resource offline and then bring the Data Protector application resource online.

The changes take effect after you stop and restart the Data Protector services on each Cell Manager where you made the changes.

Note: Data Protector checks the license configuration with the MoM Manager every hour. In case of communication problems or the MoM Manager being unavailable, the licensing status is kept for 72 hours. If the problems are not resolved within this 72 hour period, local licenses are used.

Deactivating Centralized Licensing

Centralized licensing can be deactivated and converted to local licensing.

Steps

1. In the Data Protector Manager-of-Managers, click **Clients** in the Context List.
2. In the Scoping Pane, right-click the Cell Manager for which you want to deactivate centralized licensing and then click **Configure Licensing** to open the wizard. The types and numbers of licenses available to the selected Cell Manager are displayed.

Note: Identify a cluster client with its virtual hostname.

3. Click the **Local** option to change licensing from remote to local.
4. Click **Finish** to apply the configuration.
5. Repeat the steps for all Cell Managers for which you want to deactivate centralized licensing.
6. Log on to the MoM Manager and mount the `cell` directory that resides in the default Data Protector server configuration directory.
7. Rename the `licdistrib.dat` file, for example, to `licdistrib.old`.

The changes take effect after you stop and restart the Data Protector services using the `omnisv -stop` and `omnisv -start` commands on the MoM Manager and each Cell Manager where you made the changes.

If the Cell Manager is configured on HPE Serviceguard, run the `cmhaltpkg PackageName` command to stop and the `cmrunpkg -n NodeName PackageName` to start the Data Protector package, where *PackageName* is the name of the Data Protector cluster package.

If the Cell Manager is configured on Symantec Veritas Cluster Server, take the Data Protector application resource offline and then bring the Data Protector application resource online.

About MoM Environment Administration

The MoM Manager lets you configure, manage, and control an enterprise backup environment from a single point.

In the MoM User Interface you can import and export cells, move clients among cells, and distribute the MoM configuration to other cells in the environment.

Other tasks are performed on the MoM Manager in the same way they would be if you were a local administrator. Follow the standard procedure to configure backup and restore, manage devices and media for a specific cell, configure Data Protector users and user groups, add clients, monitor running sessions and the status of the backup environment, and configure reporting and notifications.

Note: You can configure devices which are connected to clients in individual cells only from the respective Cell Managers, rather than from the MoM Manager. Only devices that are connected directly to Cell Managers can be configured from the MoM Manager.

Exporting Cells

Exporting a cell will remove it from the MoM environment.

Cluster clients identify themselves to the MoM Manager with their virtual server names. If you export a cluster in a MoM environment, use only its virtual server name.

Steps

1. In the Data Protector Manager-of-Managers, click **Clients** in the Context List.
2. In the Scoping Pane, right-click on the Cell Manager that you want to export, and then click **Export Cell Manager**.
3. Confirm your choice.

Moving Client Systems Among Cells

Data Protector allows you to move systems between cells. During the process, Data Protector does the following:

- Checks whether the client to be moved is configured in any backup specifications and removes all backup objects belonging to this client from backup specifications configured on the initial Cell Manager, while backup objects of other clients remain intact. Data Protector thus ensures no orphan backup objects remain in backup specifications after the client is moved to another cell.
- Checks if there are any devices configured on the system and leads you through the steps to move devices to another system.
- Checks if there are media used in the devices on this system and leads you through the steps to move media.

Steps

1. In the Data Protector Manager-of-Managers, click **Clients** in the Context List.
2. Expand the Cell Manager that has the client system you want to move to another cell.
3. Right-click that client system and then click **Move Client System to Other Cell** to open the wizard.
4. Select the target Cell Manager.
5. Click **Finish** to move the client.

Deactivating Centralized Licensing

Data Protector allows you to create common user class specification, Holidays file settings, global option settings, and vaulting on all Cell Managers in a MoM environment.

Prerequisites

Create the desired user class specification, holidays file settings, and global option settings on the MoM Manager.

Steps

1. In the Data Protector Manager-of-Managers, click **Clients** in the Context List.
2. Right-click **Enterprise Clients** and then click **Distribute Configuration**.
3. In the Distribute Configuration dialog box, select the type of configuration and the Cell Managers to which you want to distribute the selected configuration.
4. Click **Finish** to distribute the configuration.

Configuring Data Protector Users

You add users or user groups to a MoM environment as you would for a single Cell Manager. This procedure updates all Cell Managers with the new users.

Steps

1. In the Data Protector Manager-of-Managers, click **Users** in the Context List.
2. Select the Cell Manager to which you want to add users.
3. In the Edit menu, click **Add** and select **Users** if you want to add a new user or **User Group** if you want to add a new user group.
4. Enter the required information and click **Finish**.

Adding a User to Other Cells

You can add existing users to other cells in the MoM environment. The user is automatically added to the same user group on the target Cell Manager that he is in on the source Cell Manager.

Note: If the group the user is in on the source Cell Manager does not exist on the target Cell Manager, the user cannot be added to the cell.

Steps

1. In the Data Protector Manager-of-Managers, click **Users** in the Context List.
2. In the Scoping Pane, expand the Cell Manager and then the user group where the user is located.
3. Right-click the user and click **Add user to other cells** to open the wizard.
4. Select the target Cell Manager(s).
5. Click **Finish** to exit the wizard.

Removing a User from Cells

You can remove users from cells in the MoM environment.

Steps

1. In the Data Protector Manager-of-Managers, click **Users** in the Context List.
2. In the Scoping Pane, expand the Cell Manager and then the user group where the user is located.
3. Right-click the user and click **Remove user from cells** to open the wizard.
4. Select the Cell Manager(s) from which you want to remove the user.
5. Click **Finish** to exit the wizard.

Managing Devices and Media for a Specific Cell

You can configure devices and media for any cell within your enterprise environment.

Steps

1. In the Data Protector Manager-of-Managers, click **Clients** in the Context List.
2. Select the cell that has the devices or media that you want to manage.
3. In the **Tools** menu, click **Device & Media Administration**.
A Data Protector Manager opens with the Devices & Media context displayed.
4. Configure devices and media as if you were a local administrator.

Note: You can configure devices which are connected to clients in individual cells only from the

respective Cell Managers, rather than from the MoM Manager. Only devices that are connected directly to Cell Managers can be configured from the MoM Manager.

Managing Internal Database for a Specific Cell

You can manage the IDB for any cell in your enterprise environment.

Steps

1. In the Data Protector Manager-of-Managers, click **Clients** in the Context List.
2. Select the Cell Manager you want to manage.
3. In the **Tools** menu, click **Database Administration**. In the Internal Database context, perform database administration tasks as if you were a local administrator.

Chapter 6: Clustering

About Clustering

For more information on clustering concepts, architecture, and Data Protector in a cluster environment, see *HPE Data Protector Concepts Guide*.

For more information on installing Data Protector in a cluster environment, see *HPE Data Protector Installation Guide*.

About the Data Protector Microsoft Cluster Server Integration

For more information on clustering concepts, architecture, and Data Protector in a cluster environment, see *HPE Data Protector Concepts Guide*.

For more information on installing Data Protector in a cluster environment, see *HPE Data Protector Installation Guide*.

As a part of its high-availability, Data Protector provides an integration with Microsoft Cluster Server (MSCS), enabling you to back up a full cluster (local and shared disks) and applications running in a cluster environment. For details on supported operating system versions, level of cluster support and for supported configurations, see the *HPE Data Protector Product Announcements, Software Notes, and References*.

It is assumed that you are familiar with MSCS. If not, see the MSCS online documentation for more information.

Licensing and MSCS

When you purchase a license for the Data Protector Cell Manager, note that the license will be bound to the virtual server and will work regardless of which system inside an MSCS cluster runs the Data Protector Cell Manager.

Configuration

There are two possible ways to configure the integration:

- The Data Protector Cell Manager can be installed on the MSCS. This provides higher availability of the Data Protector Cell Manager and enables an automatic migration of Data Protector services from one cluster node to another in case of a failover, and thus an automatic restart of failed backup sessions.
- The Data Protector cluster-aware client can be installed on the MSCS, thus supporting filesystem backups and backups of cluster-aware applications.

To back up a cluster-aware application, use its virtual server name when configuring the backup specification.

Note: Cluster Service components (for example, Database Manager) maintain a coherent image of the

central cluster database, which stores information regarding changes in the status of a node, resource, or group. The cluster database must be stored on the cluster shared disk volume.

How to Manage Cluster-Aware Backups

In the Data Protector cluster Cell Manager, a backup session is cluster-aware. You can set options that define the backup behavior if a failover of Data Protector or other cluster-aware applications occurs.

Failover of Data Protector

If a failover of the cluster-aware Data Protector occurs during backup, all running and pending backup sessions fail. In the Data Protector GUI and in the backup specification, you can set one of three options that define automatic backup session restart at failover of Data Protector.

Failover of application other than Data Protector

As the cluster-aware Data Protector is a storage application within a cluster environment, it needs to be aware of other applications that might be running within the cluster. If they are running on a node other than Data Protector and if some application fails over to the node where Data Protector is running, this will result in a high load on this node. Thus, a node that had previously managed only backup operations must now handle critical application requests as well. Data Protector allows you to define what should happen in such a situation so that the critical application data are protected and the load is balanced again.

You can:

- **Abort all running backup sessions**
If the backup is less important than the application, Data Protector can automatically abort all running sessions to balance the load after the failover of the application.
To set this option, you need to create the appropriate script with the `omniclus` command.
- **Temporarily disable backup activities**
If the backup is less important than the application, Data Protector can also automatically disable the Cell Manager for a period of time to balance the load after the failover of the application. All running session continue but you cannot start new backups until the Cell Manager is enabled again.
To set this option, you need to create an appropriate script with the `omniclus` command.
- **Abort running sessions based on elapsed session time**
To balance the load after a failover of the application, you can abort backup sessions based on how long they have already been running. If a specific running backup session is just ending, Data Protector can continue the session. If the backup session has just started and if it is not important, Data Protector can abort the session.
To set one of these options, you need to create an appropriate script with the `omniclus` command and set the clustering backup options in the Data Protector GUI.
- **Abort running sessions based on a logical ID**
If a specific running backup session is more important than the application, Data Protector can continue this session. To balance the load after a failover, you can abort all backup sessions, except for an important one by using its abort ID.

To set this option, you need to create an appropriate script with the `omniclus` command and set the clustering backup options in the Data Protector GUI.

About Disaster Recovery of a Microsoft Cluster Server

Microsoft Cluster Server (MSCS) can be recovered using any disaster recovery method, except for Disk Delivery Disaster Recovery. All specifics, limitations and requirements pertaining a particular disaster recovery method also apply for the disaster recovery of the MSCS. Select the disaster recovery method that is appropriate for your cluster and include it in your disaster recovery plan. Consider the limitations and requirements of each disaster recovery method before making your decision. Perform tests from the test plan.

For details on supported operating systems, see the *HPE Data Protector Product Announcements, Software Notes, and References*.

All prerequisites for disaster recovery (for example, a consistent and up-to-date backup, an updated SRD file, replaced faulty hardware, and so on) must be met to recover the MSCS.

Possible scenarios

There are two possible scenarios for disaster recovery of an MSCS:

- a disaster occurred to a non-active(s) node
- all nodes in the cluster have experienced a disaster

About the Data Protector HPE Serviceguard Integration

For more information on clustering concepts, architecture, and Data Protector in a cluster environment, see *HPE Data Protector Concepts Guide*.

For more information on installing Data Protector in a cluster environment, see *HPE Data Protector Installation Guide*.

As part of its high-availability, Data Protector provides an integration with HPE Serviceguard (HPE SG) for HP-UX and Linux systems, enabling you to back up a full cluster (local and shared disks) and applications running in a cluster environment. For details on supported operating system versions, supported configurations, and level of cluster support, see the *HPE Data Protector Product Announcements, Software Notes, and References*.

It is assumed that you are familiar with HPE Serviceguard. If not, see the *Managing HPE Serviceguard* manual for more information.

Licensing and HPE Serviceguard

When you purchase a license for the Data Protector Cell Manager, note that the license will be bound to the virtual server and will work regardless of which physical node inside an HPE SG cluster runs the Data Protector cluster package, as long as the package is running on one of the nodes.

Configuration

There are two possible ways to configure the integration:

- The Data Protector Cell Manager can be installed in HPE SG. This enables an automatic migration of Data Protector services from one cluster node to another in case of a failover, and thus an automatic restart of failed backup sessions.

The inactive cluster node can also be used as an Installation Server.

- The Data Protector cluster-aware client can be installed in HPE SG, thus supporting filesystem backups and backups of cluster-aware applications.

About the Data Protector HACMP Cluster Integration

For more information on clustering concepts, architecture, and Data Protector in a cluster environment, see *HPE Data Protector Concepts Guide*.

For more information on installing Data Protector in a cluster environment, see *HPE Data Protector Installation Guide*.

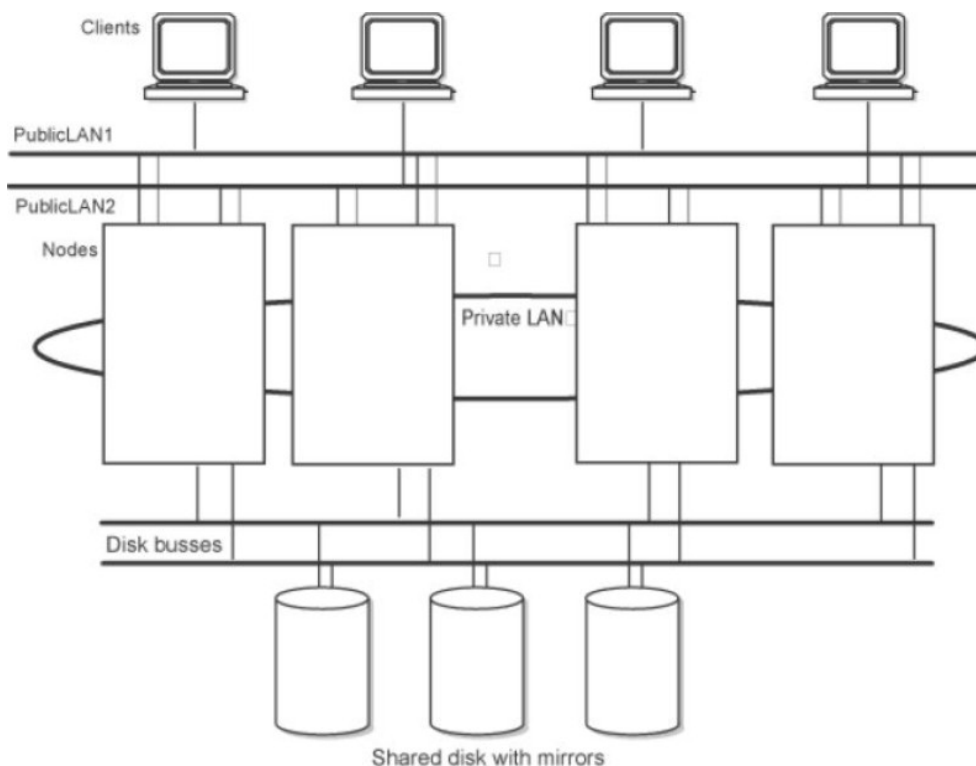
HACMP software is IBM's solution for building UNIX-based mission-critical computing environments, based on high availability (HA) and cluster multi-processing (CMP). It ensures that critical resources such as applications are available for processing.

The main reason for creating HACMP clusters is to provide a highly available environment for mission-critical applications. For example, an HACMP cluster could run a database server program which services client applications. The clients send queries to the server program which responds to their requests by accessing a database, stored on a shared external disk.

To ensure the availability of these applications in an HACMP cluster, they are put under HACMP control. HACMP ensures the applications remain available to client processes even if a component in a cluster fails. If a component fails, HACMP moves the application (along with resources that ensure access to the application) to another node in the cluster.

The entire cluster is accessed via a virtual server name (the Virtual Environment Domain Name), which represents the complete HACMP cluster over the network.

A typical HACMP cluster setup



As shown in the figure, an HACMP cluster is made up of the following physical components:

- Nodes
- Shared external disk interfaces
- Networks
- Network interfaces
- Clients

Nodes

Nodes form the core of an HACMP cluster. Each node is identified by a unique name, and contains a processor that runs an AIX operating system, the HACMP software, and the application software. A node may own a set of resources—disks, volume groups, filesystems, networks, network addresses, and applications.

Shared external disk interfaces

Each node has access to one or more shared external disk devices (disks physically connected to multiple nodes). Shared disks store mission-critical data, typically mirrored or RAID-configured for data redundancy. Note that nodes in an HACMP cluster also have internal disks storing the operating system and application binaries, but these disks are not shared.

Networks

As an independent, layered component of the AIX operating system, the HACMP software is designed to work with any TCP/IP-based network. Nodes use the network to:

- allow clients to access the cluster nodes,
- enable cluster nodes to exchange heartbeat messages,
- serialize access to data (in concurrent access environments).

The HACMP software defines two types of communication networks, depending on whether they use communication interfaces based on the TCP/IP subsystem (TCP/IP-based), or on non-TCP/IP subsystems (device-based).

Clients

A client is a processor that can access the nodes in a cluster over a LAN. Clients each run a "front end" or client application

that queries the server application running on the cluster node.

Tasks

How to Install and Configure the Data Protector IBM HACMP Cluster Integration

Chapter 7: Devices

About Backup Devices

Data Protector defines and models a physical device with Data Protector usage properties. It is possible to have several Data Protector device definitions referencing the same physical device. This device concept allows you to configure devices easily and flexibly and to use them in backup specification.

What is a backup device?

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

For a list of devices supported by Data Protector, see the *HPE Data Protector Device Support Matrix*. Unsupported devices can be configured using the `scsitab` file.

Some backup devices (such as tape drives) are subject to specific Data Protector licenses. See the *HPE Data Protector Installation Guide* for details.

About Configuring backup devices

After you have completed the preparation part, you can configure a backup device for use with Data Protector.

It is recommended that you let Data Protector configure backup devices automatically. Data Protector can automatically configure most common backup devices, including libraries. You still need to prepare the media for a backup session, but Data Protector determines the name, policy, media type, media policy, and the device file or SCSI address of the device, and also configures the drive and slots.

You can also configure a backup device manually. How you configure a backup device depends on the device type.

You can use devices that are not listed as supported in the *HPE Data Protector Product Announcements, Software Notes, and References*. Unsupported devices are configured using the `scsitab` file.

Note: External control is a means to control libraries not known to Data Protector. If Data Protector does not support a particular device, a user can write a script/program that will run the robotic control to load a medium from a particular slot into the specified drive. It is possible to configure a library as an external control by referring to a special script.

Types of Backup Devices

Data Protector supports the following device types that you can configure (depending on the components you have installed):

- Standalone
- Backup to Disk device
- SCSI library
- Stacker
- Magazine device
- Jukebox
- Standalone file device
- File library device
- External control
- ADIC/GRAU DAS library
- StorageTek ACS library

Standalone

A standalone device is a simple device with one drive that reads from or writes to one medium at a time, such as DDS or DLT. These devices are used for small-scale backups. As soon as the medium is full, an operator must manually replace it with a new medium for the backup to proceed. Standalone devices are, therefore, not appropriate for large, unattended backups.

Backup to Disk device

A Backup to Disk (B2D) device is a disk based storage device, which offers additional capabilities compared to a Data Protector Jukebox or file library device, such as access through multiple hosts (*gateways*) or, depending on the type of the device, deduplication.

SCSI library

SCSI library devices are large backup devices, also called autoloaders. They consist of a number of media cartridges in a device's repository and can have multiple drives handling multiple media at a time. Most library devices also allow you to provide automatic drive cleaning when the drive gets dirty.

A typical library device has a SCSI ID (Windows systems) or a device file (UNIX systems) for each drive in the device and one for the library's robotic mechanism, which moves media from slots to drives and back again. (For example, a library with four drives has five SCSI IDs, four for the drives and one for the robotic mechanism).

A medium is stored in a slot in the device's repository. Data Protector assigns a number to each slot, starting from one. When managing a library, you refer to the slots using their numbers.

The drive index identifies the mechanical position of the drive in the library. The index number is relevant for the robotics control. The library robotics is aware only of the drive index number and has no information about the SCSI address of the drive. The drive index is a sequential integer (starting at 1) that must be coupled with the SCSI address of the drive. Many Web interfaces to a SCSI Library, Commandview TL, or control panel of the SCSI library, will number drives starting at '0'. A drive '0' is not valid in Data Protector device configuration, the first drive must always be '1'.

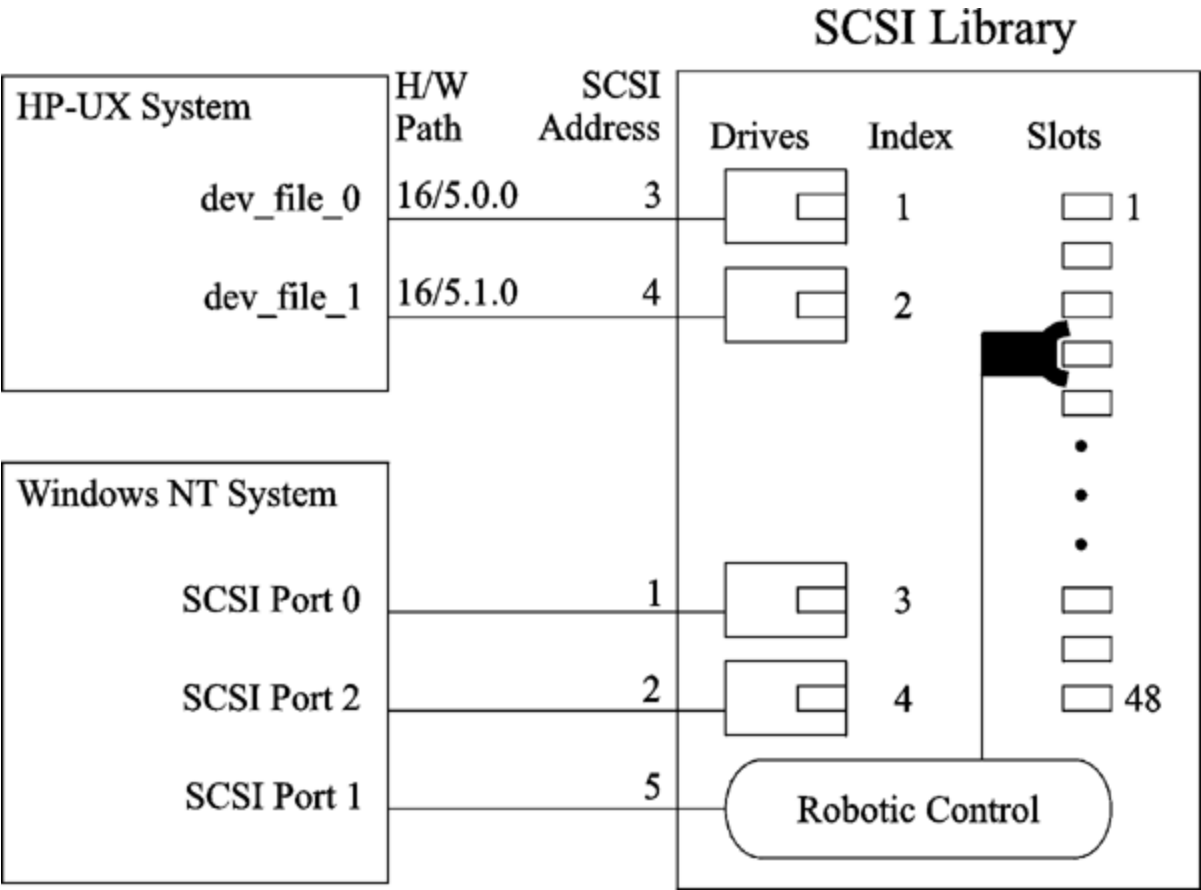
For example, for a four-drive library, the drive indexes are 1, 2, 3, and 4. If there is only one drive in the library, the drive index is 1.

The drive index must match the corresponding SCSI address. This means that you need to configure the pairs:

SCSI address_A for index one, SCSI address_B for index two, and so on.

Specify this type of device also when configuring a magazine device.

Drive index to SCSI address mapping



Stacker

A stacker is a single device that usually has only one drive. It loads media in a sequential rather than a random order, therefore a **Loose** media allocation policy is recommended. A stacker takes a medium from a "stack" (its repository) and inserts the medium into its drive. This exchange is always limited to ejecting the medium already in the drive and inserting the next medium from the stack. The load is done automatically, except the first medium has to be loaded manually. When a tape is full, it is ejecting and the next tape is loaded automatically. When all the tapes are used in a stacker magazine, the magazine has to be dismounted manually and the next one has to be inserted. Again the first tape has to be loaded manually into the drive.

A backup or restore session will not be aborted if media are not present, but a mount request will be issued instead. The whole session will not be aborted if you do not change stacker magazines within a time out period.

Magazine device

A magazine device groups a number of media into a single unit called a magazine. A magazine allows you to handle large amounts of data easily than using many individual media. The operations on each medium in the magazine are completely controlled by Data Protector. The HPE XP DAT 24x6 can be configured as a magazine device.

Jukebox

A jukebox is a library device. It can contain either optical or file media. If the device is used to contain file media it is known as a file jukebox device. The type of media the device will contain is defined during initial configuration.

If you are running a jukebox optical library on UNIX you need to have a UNIX device file configured for each exchanger slot or side of the platter.

Standalone file device

A standalone file device is a file in a specified directory to which you back up data instead of writing to a tape.

File library device

A file library device consists of a set of directories to which you back up data instead of writing to a tape.

External control

External control is a means to control libraries not known to Data Protector. If Data Protector does not support a particular device, a user can write a script/program, that will run the robotic control to load a medium from a particular slot into the specified drive. It is possible to configure a library as an external control by referring to a special script.

ADIC/GRAU DAS library

An ADIC/GRAU DAS library is a controlled, very large library (silo). It is used in complex environments where the amount of backed up data is exceptionally large, and so is the amount of media needed to store the data. It can handle from a hundred to several thousand tapes. Typically, an ADIC/GRAU DAS library can house many types of backup drives and thousands of media slots, all served by an internal robotic mechanism and controlled through special library control units. You can assign a dedicated set

of media in the library to an application so that the library can be shared between Data Protector and other applications.

All media operations can be executed from the Data Protector user interface. For media in a recognizable format, Data Protector displays the format as the media type, such as `tar`. For media in a non-recognizable format, the media type is `foreign`.

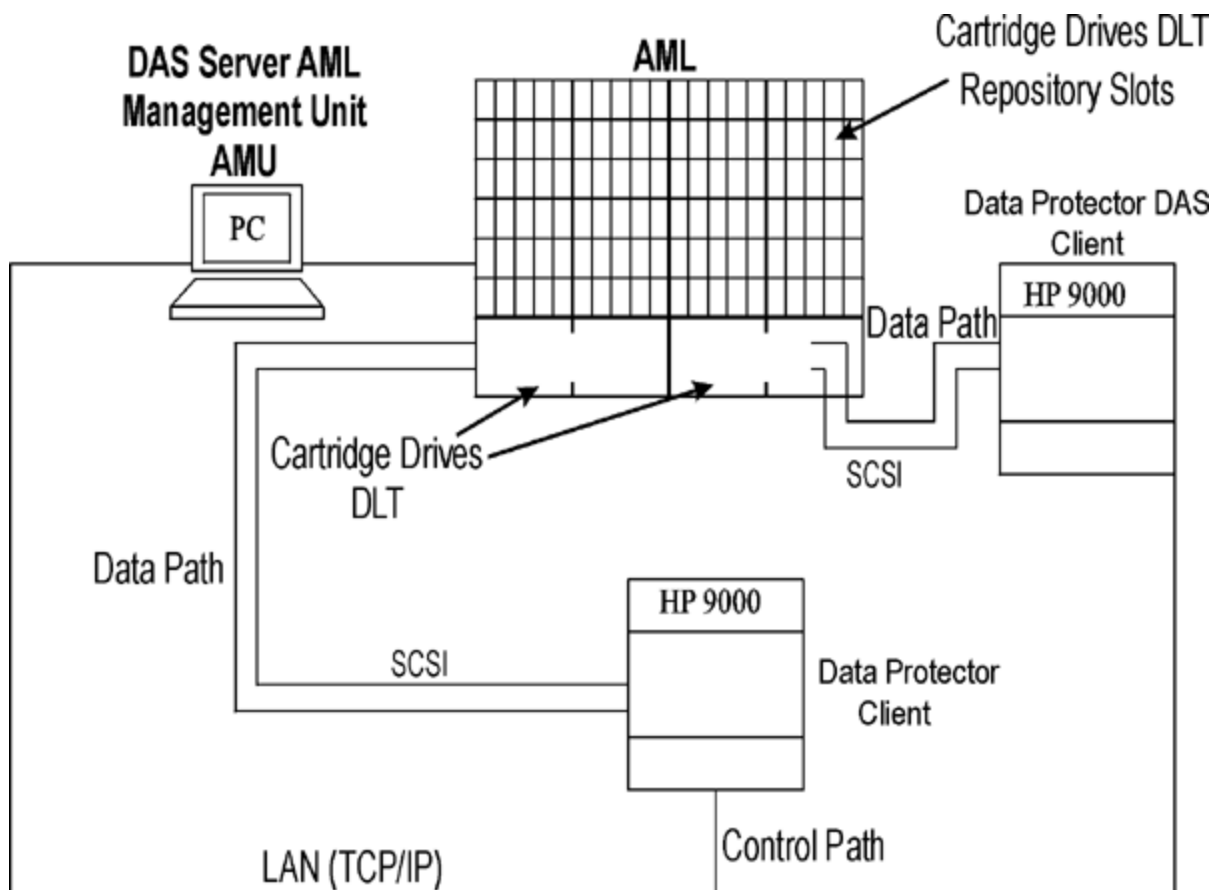
The Media Management Database tracks all Data Protector and non-Data Protector media, whether resident (media in the device's repository) or non-resident (media outside the device's repository), providing sophisticated overwrite protection. Data Protector will not overwrite media containing data in a recognizable format. However, it is not guaranteed that the Data Protector data on tapes will not be overwritten by some other application using the same media. It is recommended that media used by Data Protector are not used by any other application, and the other way round.

The actual location of a medium is maintained by the DAS Server, which tracks the location using its `volser`. When a medium is moved around the repository, it is not assigned to the same physical slot each time. Therefore, you cannot rely on the slot number when handling the media, but on the barcode (`volser`).

The ADIC/GRAU DAS library can automatically clean its drives after the drive has been used a set number of times. However, this is not recommended, as the drive cleaning interrupts the session running at that moment, causing it to fail. If you want to use the library's cleaning functionality, you have to ensure that drive cleaning is performed when no Data Protector sessions are running.

You have to create a logical Data Protector library for every media type. While the ADIC/GRAU or STK ACS system can store many physically different types of media, Data Protector can only recognize a library with a single type of media in it.

Data Protector and ADIC/GRAU DAS library systems integration



StorageTek ACS library

A StorageTek Automated Cartridge System (ACS) library is a robotic library (silo). It is used in complex environments where the amount of backed up data is exceptionally large, and so is the amount of media needed to store the data. It can handle hundreds of tapes. You can assign a dedicated set of media in the device to an application so that the library can be shared between Data Protector and other applications.

Typically, such a device has many types of backup drives and thousands of media slots, all served by an internal robotic mechanism and controlled through ACS Library Server (ACSL) software. Media- and device-related actions initiated by Data Protector are passed through the user interface to the ACSLS, which then directly controls the robotics and executes the moving and loading of media.

When the library is properly installed and configured, Data Protector provides easy handling of the media during a backup and restore session. All media operations can be executed from the Data Protector user interface. For media in a recognizable format, Data Protector displays the format as the media type, such as `tar`. For media in a non-recognizable format, the media type is `foreign`.

The Media Management Database tracks all Data Protector and non-Data Protector media, whether resident (media in the device's repository) or non-resident (media outside the device's repository), providing sophisticated overwrite protection. Data Protector will not overwrite media containing data in a recognizable format. However, it is not guaranteed that the Data Protector data on tapes will not be

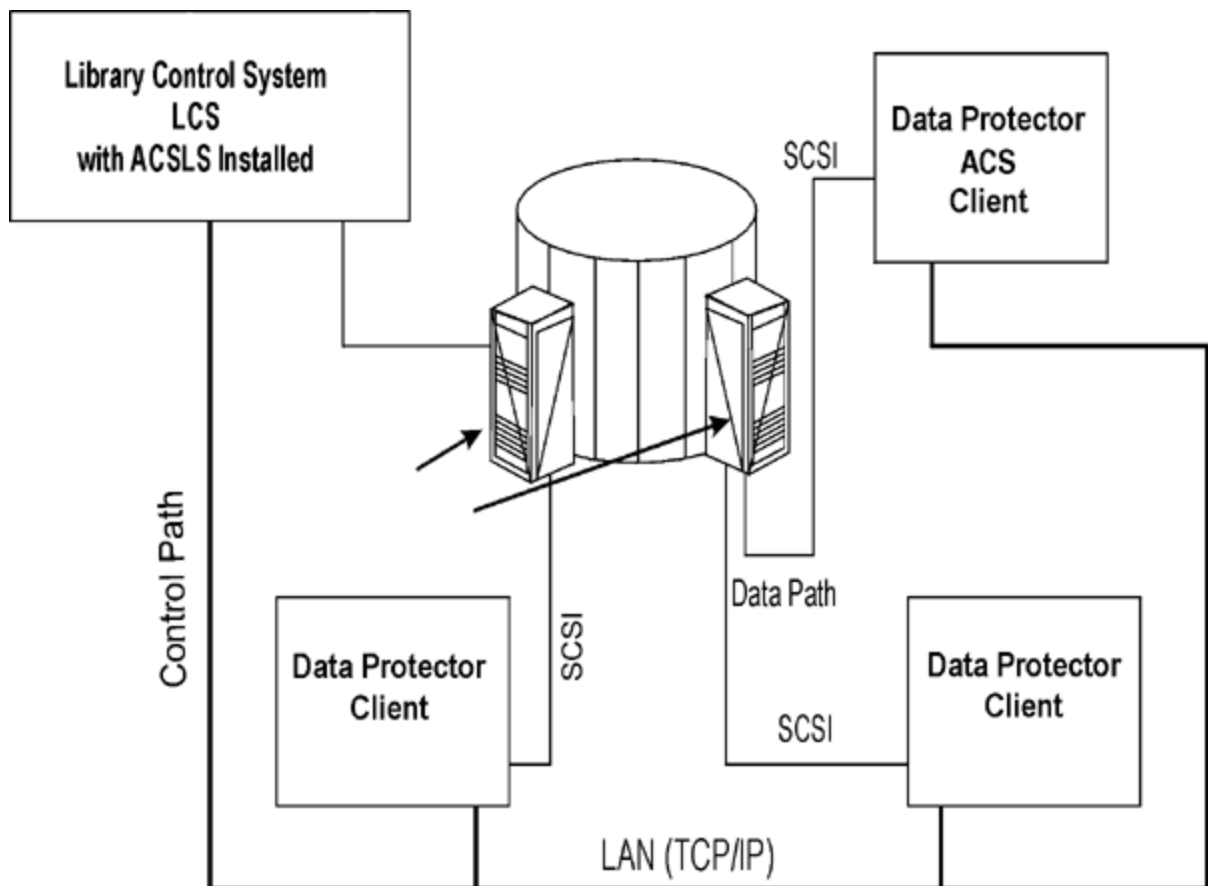
overwritten by some other application using the same media. It is recommended that media used by Data Protector are not used by any other application, and the other way round.

The actual location of a medium is maintained by the ACS Server, which tracks the location using its volser. When a medium is moved around the repository, it is not assigned to the same physical slot each time. Therefore, you cannot rely on the slot number when handling the media, but on the barcode (volser).

The StorageTek ACS library can automatically clean its drives after the drive has been used a set number of times. However, this is not recommended, as library the drive cleaning interrupts the session running at that moment, causing it to fail. If you want to use the library's cleaning functionality, you have to ensure that drive cleaning is performed when no Data Protector sessions are running.

You have to create a logical Data Protector library for every media type. While the ADIC/GRAU or STK ACS system can store many physically different types of media, Data Protector can only recognize a library with a single type of media in it.

Data Protector and StorageTek ACS library integration



About Cloud (Helion) Devices

The Cloud (Helion) device is a device configured with Cloud (Helion) credentials and supports HPE's Public Cloud. The Media Agent has been enhanced to act as a Cloud gateway to transmit data to the Cloud (Helion) device. It behaves similarly to a Backup to Disk (B2D) device.

Prerequisites

Prerequisites in HPE Public Cloud:

- You must have an HPE Public Cloud account and credentials. For more information, see <https://horizon.hpcloud.com>.
- You must have a subscription to the Object Store in the HPE Public Cloud.
- For your project in the HPE Public Cloud, you must take note of the Project name.
- Authentication Service URL for the geographic region closest to your datacenter.
- If you decide to use the access keys for authentication instead of username and password credentials, create your Access keys in the HPE Public Cloud.

Prerequisites in Data Protector:

- Ensure that the Data Protector latest Cell Manager, User Interface client, and Installation Server are installed on supported systems along with the latest 9.04 General Patch Release bundle.

For details, see the latest HPE Data Protector support matrices at <https://softwaresupport.hpe.com/>. See the *HPE Data Protector Installation Guide* on how to install Data Protector in various architectures.

- Install the Data Protector Media Agent or the NDMP Media Agent components on Windows and Linux systems that will become Cloud gateways, including the clients on which the Cloud Helion device will be enabled. For instructions, see the *HPE Data Protector Installation Guide*.

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

Limitations

- Cloud (Helion) device object copy has been tested and is supported with the following:
 - Source devices: file library devices and StoreOnce devices.
 - VMware backup specifications.
- When selecting or creating a container in the HPE Object Store, the following restrictions apply:
 - Each device can have only one container assigned.
 - Different devices cannot use the same container.
 - Once a container is assigned to a device, it cannot be changed.

- When configuring the Cloud (Helion) device, ensure that it has the same or large block size as the on-premises source device.

If you will be performing object copies back and forth between the on-premises device and the Cloud (Helion) device, the block sizes on both devices must match. Block size can be set in the gateway properties.

Recommendations

HPE recommends the following for Cloud (Helion) devices:

- When backing up VMware specifications, use a Cloud gateway local to the data source, as this will reduce the network load during object copy operations.
- Use **Access Keys** as the authentication mode where available. It provides more restricted overall access to the Cloud (Helion) device and more secure as its system generated.
- Break up large data sets into multiple backup specifications when copying to HPE Cloud.
This allows many copy sessions in parallel, increases the overall bandwidth utilization and enables more efficient data copies to HPE Cloud.
- Consolidation on the Cloud (Helion) device is not recommended because of the large bandwidth requirements and associated HPE Cloud costs.

Preparing for the Cloud (Helion) device

The following tasks must be performed to configure object copy operations to the Cloud (Helion) device.

1. Configure a backup specification to back up your data to a local backup device. For more information, see [Creating a Backup Specification](#).
2. In the HPE Public Cloud, obtain your user account credentials or the access keys required to authenticate, a subscription to the Object Store, the authentication service URL, and other HPE Public Cloud prerequisites. These will be used to configure the Cloud (Helion) device.
3. In Data Protector, configure a Cloud (Helion) device. For more information, see [Configuring Cloud Devices](#).
4. Configure object copy sessions using the local backup device as the source device and the Cloud (Helion) device as the destination backup device.
Creating a copy to Cloud (Helion) device object operation enables the data stored on the local backup device to replicate data to the HPE Public Cloud. Data sent to the Cloud (Helion) device is compressed and encrypted by default.
5. To restore data from the Cloud (Helion) device, you can either:
 - Create an object copy from the Cloud (Helion) device to your local backup device, and restore to your client from the local backup device.
 - Recycle and export the local media and restore directly from the Cloud (Helion) device to your client.
 - Restore directly from the Cloud (Helion) device even when there are local versions, by specifying the Cloud (Helion) device to be used for restore.
 - Set the media location priority to the Cloud (Helion) media instead of the local media. See [Setting Media Location Priority](#).

About Cloud (Azure) Devices

A new cloud device is used to enable backup and object copy to Microsoft Azure object store from Data Protector. The Cloud (Azure) device is configured with Azure credentials and sends data to the cloud.

Prerequisites

Prerequisites in Cloud (Azure) device portal:

- You must have Microsoft Azure account. For more information, see [Microsoft Azure](#) portal.
- You must have the two storage access keys that Microsoft Azure generates while creating the Microsoft Azure storage account.

There are two access keys generated for the account. These keys are needed while creating the related Data Protector device, as a part of the process of providing credentials.

- The system time must be set accurately to ensure proper synchronization between the gateway host and Microsoft Azure.

Prerequisites in Data Protector:

- Ensure that the Data Protector's Cell Manager, User Interface client, and Installation Server are installed on supported systems along with the latest 9.09 General Patch Release bundle.

For more information on how to install Data Protector in various architectures see *HPE Data Protector Installation Guide*.

- Install the Data Protector Media Agent or the NDMP Media Agent component on Windows and Linux systems that will become Cloud (Azure) gateways, including the clients on which the Cloud (Azure) device will be enabled. For instructions, see the *HPE Data Protector Installation Guide*.

For a detailed list of supported operating system versions, see the latest support matrices at <https://softwaresupport.hpe.com/manuals>.

Note:

- In case of media agent systems that require a proxy server to be configured to connect to the web, the `omnirc` variable `OB2_CLOUD_DEVICE_PROXY=<proxy_server:port_number>` needs to be set in the `omnirc` file.
- Media Agents have a built-in retry mechanism to handle different error conditions. As a result, user operations may occasionally take time to complete. Such issues are not observed under normal circumstances.

Limitations

Following are the limitations with Cloud (Azure) devices:

- Cloud (Azure) device object copy has been is supported with the following:
 - Source devices: File library devices and StoreOnce devices.
 - File system backup specifications.
- When selecting or creating a container in the Cloud (Azure) device, the following restrictions apply:
 - Each device can have only one container assigned.
 - Multiple devices cannot use the same container.
 - Once a container is assigned to a device, it cannot be changed.

Cloud (Azure) device blob size limit

A Data Protector medium is uploaded to Cloud (Azure) device as one or more blobs of data and metadata. Cloud (Azure) has a blob size limit of 195 GB, whereas Data Protector media has no upper size limit. However, to comply with this restrictions, one Data Protector medium can span multiple blobs, the maximum size of each being 75 GB. The exact amount of data stored per blob is a function of how compressible is the data.

Recommendations

Following are the recommendations for configuring Cloud (Azure) devices:

- When backing up file system specifications, use a Cloud (Azure) device gateway local to the data source, as this will reduce the network load during object copy operations.
- Because object copy jobs to Cloud (Azure) are encrypted by default, encryption should be turned off in the initial backup specification that generates data for the copy operation. If encryption is turned ON, data is encrypted twice, additional CPU resources are consumed, and object copy data becomes incompressible. As a result, the amount of data transferred to Cloud (Azure) increases, which extends the copy times.
- Break large data sets into multiple backup specifications when copying data to a Cloud (Azure) device, such that multiple copy sessions can be enabled in parallel. Thereby, increasing the overall bandwidth.
- Consolidation on the Cloud (Azure) device is not recommended because of the large bandwidth requirements and associated costs.

Preparing for the Cloud (Azure) device

The following tasks must be performed to configure object copy operations to a Cloud (Azure) device.

1. Configure a backup specification to back up your data to a local backup device. For more information, see ["Creating a Backup Specification " on page 244](#)
2. Login to the [Microsoft Azure](#) portal and obtain the access keys required to use the Microsoft Azure storage account.
3. Configure the Cloud (Azure) device in Data Protector. For more information, see ["Configuring a Backup to Disk Device - Cloud \(Azure\)" on page 144.](#)

4. Configure object copy sessions using the local backup device as the source device, and the Cloud (Azure) device as the destination backup device.

Creating a copy to a Cloud (Azure) device object operation enables the data stored on the local backup device to replicate data to the Cloud (Azure) device. Data sent to the Cloud (Azure) is compressed and encrypted by default.

5. Restore data from the Cloud (Azure) device. To restore you can either:
 - Create an object copy from the Cloud (Azure) device to your local backup device, and restore to your client from the local backup device.
 - Recycle and export the local media and restore directly from the Cloud (Azure) device to your client.
 - Restore directly from the Cloud (Azure) even when there are local versions, by specifying the Cloud (Azure) device to be used for restore.
 - Set the media location priority to the Cloud (Azure) media instead of the local media. See ["Setting the Media Location Priority " on page 229](#).

Device Performance Tuning

Block size

Every logical device can be configured to process data in units of a specific size (**block size**). Different devices have different default block sizes, which can be used (all sessions are completed successfully), but may not be optimal. By adjusting the block size, you can enhance the performance of Data Protector sessions.

The optimal block size value depends on your environment:

- Hardware (devices, bridges, switches, ...)
- Firmware
- Software (operating system, drivers, firewall, ...)

To achieve the best results, first optimize your environment by installing the latest drivers and firmware, optimize your network, and so on.

Determining the optimal block size

To determine the optimal block size, perform different tests by running usual Data Protector tasks (backup, restore, copy, and so on) with different block size values and measure the performance.

Note: Once you have changed the device block size, you cannot restore old backups (with the old block size) with this device anymore. .

Therefore, keep your old logical devices and media pools intact to be able to restore the data from the old media and create new logical devices and media pools with different block size values for

testing purposes. Alternately, know how to change the block size when performing a restore. The restore dialog prompts you for Block Size.

Limitations

- Disaster recovery: To be able to perform an offline EADR/OBDR recovery (Enhanced Automated Disaster Recovery, One Button Disaster Recovery), back up your data using the default block size.
- Library: If you are using several drive types of similar technology in the same library, the drives must have the same block size.
- SCSI adapters: Check if the selected block size is supported by the host SCSI adapter the device is connected to.
- Object copy functionality: The destination devices must have the same or larger block size than the source devices.
- Object consolidation functionality: The destination devices must have the same or larger block size than the source devices.
- Mirroring: Block size of devices must not decrease within a mirror chain. The devices used for writing mirror 1 must have the same or a larger block size than the devices used for backup; the devices used for writing mirror 2 must have the same or a larger block size than the devices used for writing mirror 1, and so on.

For other limitations, see the *HPE Data Protector Product Announcements, Software Notes, and References*.

Changing the block size

You can set the block size in the Sizes tab of the Advanced Options dialog box for a specific device. For more information, see [Setting Advanced Options for Devices and Media](#).

Device Performance

Device types and models impact performance because of the sustained speed at which devices can write data to a tape (or read data from it).

- DDS/DAT devices typically have a sustained data transfer rate of 510 kB/s to 3 MB/s, without hardware compression, depending on the model.
- DLT devices typically have a sustained data transfer rate of 1.5 MB/s to 6 MB/s, without hardware compression, depending on the model.
- LTO devices typically have a sustained data transfer rate of 10 MB/s to 20MB/s, without compression, depending on the model.

Data transfer rates also depend on the use of hardware compression. The achievable compression ratio depends on the nature of the data being backed up. In most cases, using high speed devices with hardware compression improves performance. This is true, however, only if the devices stream.

At the start and at the end of a backup session backup devices require some time for operations such as rewinding media and mount or dismount media.

Libraries offer the advantage of automation: new or reusable media must be loaded at backup time and media must be accessed quickly at restore time, but because library access is automated, the process is faster.

Disk-based devices are quicker to use than conventional devices. When using a disk-based device there is no need to mount and dismount media and the data in disk based devices is accessed faster, thus reducing the amount of time spent for backup and restore.

Support of New Devices

To use a device that is not listed as supported in the *HPE Data Protector Product Announcements, Software Notes, and References*, use the `scsitab` file.

The `scsitab` file is a machine-readable form of the Data Protector Support Matrix and includes information about all supported devices. The `scsitab` file is used by the Data Protector Media Agent to determine whether a given device or library is supported or not. It also provides information about the device and its specific parameters.

Modifying the `scsitab` file is not supported.

To use a device that is not listed as supported in the *HPE Data Protector Product Announcements, Software Notes, and References*, download the latest software package for the `scsitab` file from the HPE Data Protector website at <http://www.hpe.com/software/dataprotector>.

After you have downloaded the `scsitab` software package, follow the installation procedure provided with the software package.

The `scsitab` file is located on the system to which the device is connected, on the following location:

Windows systems: `Data_Protector_home\scsitab`

HP-UX, Solaris, and Linux systems: `/opt/omni/scsitab`

Other UNIX systems: `/usr/omni/scsitab`

If you still receive the same error while configuring your device, contact HPE Support to get the information when the device will be supported.

Preparing Backup Devices

Preparation of a backup device consists of connecting the device to the system or in an SAN environment to the SAN and knowing which of the (working) associated device files (SCSI address) is to be used.

Prerequisite

A Media Agent (the General Media Agent or the NDMP Media Agent) must be installed on each system that has a backup device connected or in an SAN environment on systems controlling backup devices on the SAN.

Steps

1. Connect the backup device to a computer system or in an SAN environment to the SAN.
2. Continue the preparation:

Windows systems:

[Specify SCSI address syntax](#) for a device connected to the Windows system.

UNIX systems:

[Find or create the device filename](#) for a device connected to the UNIX system.

3. In case you plan that several devices will be using the same media, you have to ensure that the writing density and the **block size** settings are identical.
4. Boot the system to have the device become known to the system.
5. For some backup devices, additional steps have to be performed.

After preparing the backup device, configure it for use with Data Protector. Prepare the media that you want to use with your backups.

- [In the SAN environment](#)
- [File devices](#)
- [Magazine](#)
- [SCSI library, Jukebox, External control](#)
- [Windows robotics drivers](#)

In the SAN Environment

Steps

1. Check that the same robotics device filename exists on all systems that need to access the shared library. Ignore this if you plan to use indirect library access.

HP-UX and Solaris systems:

The requirement of the device file identity is accomplished via hard or soft links, if necessary.

Windows systems:

Use the `libtab` text file to override default SCSI device identification and reassign robotics control devices to the logical drives defined on another host.

The `libtab` file should be created on the Media Agent client in the `Data_Protector_home` directory, as a text file with the following syntax (spaces in logical drive name are allowed):

`hostnamecontrol_device_fileddevice_name`

for example

`computer.company.com scsi2:0:4:0 DLT_1`

File devices

Disable the Windows compression option for a file you want to use as a device. You can do this using Windows Explorer:

Steps

1. Right-click the file, click **Properties**, and clear the **Compress** option under **Attributes**.

Magazine

Steps

1. Create a media pool with magazine support before you configure a magazine device. The device must have support for magazines (for example, HPE 12000e).

SCSI library, Jukebox, External Control

Steps

1. Decide which slots in the library you want to use with Data Protector. You will need to specify them when configuring a library.

Windows robotics drivers

On Windows systems, robotics drivers are automatically loaded for enabled tape libraries. To use the library robotics with Data Protector on a Windows system, disable the respective Windows driver.

Steps

1. In the Control Panel, double-click **Administrative Tools**.
2. Double-click **Computer Management** and then click **Device Manager**.
3. Expand **Medium Changers**.
4. Right-click the medium changer and select **Disable**.
5. Restart the system to apply the changes. The robotics is now ready to be configured with Data Protector.

Creating SCSI Addresses on Windows Systems

The SCSI address syntax depends on the type of physical device (magneto-optical or tape) connected to your Windows system. The device must have been connected to the system (and powered on) before the system is booted.

Tip: You can auto-detect SCSI addresses using Data Protector.

Magneto-optical device

If a magneto-optical device is connected to your system, the SCSI address syntax is N:B:T:P:L (N=mountpoint of the removable drive, B=Bus number, T=SCSI Target IDs, P=path, L=LUN).

Open **SCSI Adapters** in the **Control Panel** and double-click the name of the target device. Then click **Settings** to open the device property page. All the necessary information is displayed.

Tape device

If a tape device is connected to your system, the SCSI address syntax depends on whether the native tape driver is loaded or unloaded. The address syntax also depends on the system. See the following sections for instructions on creating a target SCSI address on:

[Windows without the native tape driver](#)

[Windows using the native tape driver](#)

Windows without the native tape driver

If the Native Tape Driver is unloaded, the SCSI address syntax is P:B:T:L (P=SCSI Port, B=Bus number, T=SCSI Target IDs, L=LUN). Look up the properties of the connected tape drives to gather this information.

Open **SCSI Adapters** in the **Control Panel** and double-click the name of the target device. Then click **Settings** to open the device property page. All the necessary information is displayed.

Windows using the native tape driver

If the Native Tape Driver is loaded, the SCSI address syntax is `tapeN` (N=drive instance number). The tape drive file can only be created using the drive's instance number, for example, `tape0` if N equals 0.

Steps

1. In the Windows Control Panel, double-click **Administrative Tools**.
2. In the Administrative Tools window, double-click **Computer Management**. Expand **Removable Storage** and then **Physical Locations**.
3. Right-click the tape drive and select **Properties**.

If the native tape driver is loaded, the device file name is displayed in the General property page. Otherwise, you can find the relevant information on the Device Information property page.

Finding Device Filenames on UNIX System

You need to know the device filenames for configuring devices connected to a UNIX system.

Device file creation depends on the specific UNIX operating system vendor. For devices on the HP-UX and Solaris platforms, see the following sections. For devices on other UNIX platforms, consult the respective vendor's information.

Finding Device Filenames on HP-UX

Prerequisite

Check if the device is properly connected using the `/usr/sbin/ioscan -f` command.

Steps

1. On your HP-UX system, start the **System Administration Manager (SAM)** application.
2. Click **Peripheral Devices** and then **Tape Drives**.
3. Click the target device.
4. In the Actions menu, click **Show Device Files**. The device filenames are displayed. Use the one with the syntax `*BEST`. For a no-rewind device, use the one with the syntax `'BESTn'`.

If there are no device filenames displayed, you need to create them.

Finding Device Filenames on Solaris

Steps

1. Press **Stop** and **A** to stop the client system.
2. At the ok prompt, use the `probe-scsi-all` command to check if the device is properly connected.
This will provide information on the attached SCSI devices, which should include the device id string(s) for the attached backup device.
3. At the ok prompt, enter `go` to return to normal running.
4. List the contents of the `/drv/rmt` and, if using a multi-drive library, `/drv` directories:
 - The `/drv/rmt` directory should contain the device filename(s) for the drive(s) of the backup device.
 - The `/drv` directory should contain the device filename for the robotics, if using a multidrive library device.

If there are no device filenames displayed, you need to create them.

For further information on device files, see the *HPE Data Protector Installation Guide*.

Creating Device Files on UNIX Systems

If the device files that correspond to a particular backup device have not been created during the system initialization (boot process), you have to create them manually. This is the case with the device files required to manage the library control device (library robotics).

Device file creation depends on the specific UNIX operating system vendor. For devices on the HP-UX and Solaris platforms, see the following sections. For devices on other UNIX platforms, consult the respective vendor's information.

Creating Device Files on HP-UX Systems

Prerequisites

- Check if the device is properly connected using the `/usr/sbin/ioscan -f` command.

Steps

1. On your HP-UX system, start the **System Administration Manager (SAM)** application.
2. Click **Peripheral Devices** and then **Tape Drives**.
3. Click the target device.
4. In the Actions menu, click **Create Device Files** and then **Create Default Device Files**.

Creating Device Files on Solaris Systems

Prerequisites

- Before you can use a new backup device on a Solaris client, you must first update the device and driver configuration files for the client, install another driver if using a library device, and create new device files on the client.

Steps

1. Press **Stop** and **A** to stop the client system.
2. At the ok prompt, run the `probe-scsi-all` command to check the available SCSI addresses on the client system, and choose an address for the device you want to attach (for a single drive device). In the case of a multi-drive device, you will need to choose a SCSI address for each drive and one for the robotic mechanism.
3. At the ok prompt, enter `go` to return to normal running.
4. Shut down and power down the client system.
5. Set the chosen SCSI address(es) on the backup device.
6. If necessary when connecting SCSI devices to the client system concerned, shut down and power down the system.
7. Attach the backup device to the client system
8. Power up first the backup device and then the client system (if powered down earlier).
9. Press **Stop** and **A** to stop the system again.
10. At the ok prompt, run the `probe-scsi-all` command.
This will provide information on the attached SCSI devices, including the correct device id string (s) for the newly attached backup device.

11. At the ok prompt, enter `go` to return to normal running.
12. Edit the configuration file `st.conf` and add the required device information and SCSI addresses for the drives.
For further information on how to do this, see the *HPE Data Protector Installation Guide*.
13. If you are attaching a multi-drive device with a robotics mechanism, also perform the steps below. For detailed information, see the *HPE Data Protector Installation Guide*.
 - a. Copy an `sst` driver onto the client and install it.
 - b. Copy the configuration file `sst.conf` (Solaris 8 or 9) or `sgen.conf` (Solaris 10) onto the client system concerned and edit it, adding an entry for the robotic mechanism.
 - c. Edit the `/etc/devlink.tab` file and add an entry for the robotic mechanism device file.
14. When you have updated the drivers and configuration files as required, create new device files for the client system:
 - a. Remove all existing device files from the `/drv/mnt/` directory.
 - b. Run the command `shutdown -i0 -g0` to shut down the system.
 - c. Run the command `boot -rv` to restart the system.
 - d. When the reboot has completed, list the contents of the `/dev` directory to check the device files created. Device files for robotic mechanisms should be in the `/dev` directory, and those for drives in the `/dev/rmt` directory.

Auto-Detecting Device Filenames and SCSI Addresses

You can auto-detect the device filenames (SCSI addresses) for most devices connected to Windows, HP-UX, or Solaris platforms.

For an existing Data Protector device definition

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**. The list of configured devices appears in the Results Area.
3. In the Results Area, right-click the device, and then click **Properties**.
4. Click the **Drives** tab.
5. Use the drop-down list to auto-detect the SCSI addresses (device filenames) for the device.

While creating a Data Protector device definition

Steps

1. Follow the procedure for configuring a device.
2. In the wizard, when prompted to specify the device filename (SCSI address), use the drop-down list to get a choice of available devices.

Auto-Detecting Device Filenames and SCSI Addresses for Libraries

You can auto-detect the device filenames (SCSI addresses) for the library robotics connected to Windows, HP-UX, or Solaris platforms.

For an already configured library

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**. The list of configured devices appears in the Results Area.
3. In the Results Area, right-click the library and then click **Properties**.
4. Click the **Control** tab.
5. In the Library's robotic SCSI address area, use the drop-down list to get a choice of available filenames (SCSI addresses) for the library robotics.

While configuring a library

Steps

1. Follow the procedure for configuring the library robotics.
2. In the wizard, when prompted to specify the SCSI address (filename), use the drop-down list to get a choice of available filenames (SCSI addresses) for the library robotics.

About Configuring Backup Devices

After you have completed the preparation part, you can configure a backup device for use with Data Protector.

It is recommended that you let Data Protector configure backup devices automatically. Data Protector can automatically configure most common backup devices, including libraries. You still need to prepare the media for a backup session, but Data Protector determines the name, policy, media type, media policy, and the device file or SCSI address of the device, and also configures the drive and slots.

You can also configure a backup device manually. How you configure a backup device depends on the device type.

You can use devices that are not listed as supported in the *HPE Data Protector Product Announcements, Software Notes, and References*. Unsupported devices are configured using the `scsitab` file.

About Library Management Console

What is a library management console?

Many modern tape libraries have integrated management consoles that provide you the possibility of performing remote library configuration, management, and monitoring tasks. A library management console is a web interface to the library, which is presented in the web browser like an ordinary webpage. A tape library that is equipped with such a web console enables you to perform various tasks from an arbitrary remote system. For example, you can set library configuration parameters, load tapes into library drives, and check the current library status. The scope of tasks that can be performed remotely depends on the management console implementation, which is independent of Data Protector.

Every library management console has its own URL (web address), which is the entry point to the management console interface. Type this URL in a web browser's address bar to access the console interface.

Library management console support in Data Protector

The library configuration contains a parameter representing the URL of the library management console. The **management console URL** can be specified during the library configuration or reconfiguration process.

Access to the management console interface is simplified by the extended Data Protector GUI functionality. You can invoke a web browser and load the console interface from the Data Protector GUI. Depending on the operating system, the system default web browser (on Windows systems) or the web browser specified in the Data Protector configuration (on UNIX systems) is used.

Before using the library management console, consider that some operations which you can perform through the console may interfere with your media management operations and/or your backup and restore sessions

Limitation

Entering spaces and double quotes as part of the management console URL is not supported; you should enter safe URL codes instead. Unsupported characters and their safe URL code equivalents are shown in the table below.

Character	Safe URL code
Space	%20
Double quote (")	%22

Autoconfiguring a Backup Device

After you have connected the backup device to the systems you want to configure and working device files (SCSI address) exist, you can configure it for use with Data Protector. Autoconfiguration implies

that Data Protector will create a device definition for you.

Data Protector can detect and automatically configure most common backup devices that are connected to a system or several systems in a SAN. You can modify the properties of the automatically configured device afterwards to adapt it to your specific needs.

Autoconfiguration is possible on the following operating systems:

- Windows
- HP-UX
- Solaris
- Linux

Note: When autoconfiguring libraries while the Removable Storage service is running, drives and robotics (exchangers) will not be combined correctly.

Prerequisite

Each client system you want to autoconfigure must have a Media Agent installed.

Device autoconfiguration

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and then click **Autoconfigure Devices** to open the wizard.
3. Select the client system with devices to be configured, and click **Next**.
4. Select the backup devices to be configured on your system. Click **Next**.
5. To enable automatic discovery of changed SCSI addresses select **Automatically discover changed SCSI address** and click **Finish**. For magazine devices change the media pool to one with magazine support after autoconfiguration.

The name of the device is displayed in the list of configured devices. You can scan the device to verify the configuration.

Device autoconfiguration in a SAN environment

Data Protector provides device autoconfiguration in SAN environment, where different clients use tape drives in one library. The Data Protector autoconfiguration functionality provides automated device and library configuration on multiple client systems.

Data Protector determines the name, lock name, policy, media type, media policy, and the device file or SCSI address of the device, and configures the drive and slots.

Note: When you introduce a new host into a SAN environment, the configured libraries and devices will not be updated automatically.

- If you want to use an existing library on a new host, delete this library and autoconfigure a new library with the same name on the new host.
- If you want to add devices to an existing library, you can delete the library, and autoconfigure a library with the same name and new drives on a new host, or you can manually add the drives to the library.

Limitations

Autoconfiguration cannot be used to configure the following devices in a SAN environment:

- mixed media libraries
- DAS or ACSLS libraries
- NDMP devices

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and then click **Autoconfigure Devices** to open the wizard.
3. Select the client systems you want to configure. In Microsoft Cluster Server environment, select the virtual server.
Click **Next**.
4. Select the devices and libraries you want to be configured on your system.
5. In case of configuring a library, select the Control Host - a client that will control library robotics when the library is visible by several clients. If there is a Cell Manager among the systems that see the library, it is selected by default. You can switch between the following two views:
 - **Group by Devices**
Displays a list of all devices and libraries. Expand the library or device and select the client system on which you want it to be configured.
 - **Group by Hosts**
Displays a list of clients that have devices attached. Expand the client on which you want devices or libraries to be configured.
6. Optionally, to enable multipath devices, select **Automatically configure MultiPath devices**. Click **Next**.
7. To enable automatic discovery of changed SCSI addresses select **Automatically discover changed SCSI address**.
8. Click **Finish**. The list of configured devices is displayed.

You can scan the device to verify the configuration.

Configuring a Standalone Device

After you have connected the backup device to the system and a working device file (SCSI address) exists, you can configure it for use with Data Protector.

It is recommended that you let Data Protector configure backup devices automatically.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and then click **Add Device** to open the wizard.
3. In the Device Name text box, enter the name of the device.
4. In the Description text box, enter a description (optional).
5. Optionally, select **MultiPath device**.
6. If the **MultiPath device** option is *not* selected, select the name of the client (backup system) from the Client drop-down list.
7. In the Device Type list, select the **Standalone** device type and then click **Next**.
8. Enter the SCSI address of the physical device (Windows systems) or a device filename (UNIX systems) and click **Add**.

For multipath devices, select the client from the drop-down list and enter the device filename for the device. Click **Add** to add the path to the list of configured paths.

Tip: You can enter multiple addresses to create a device chain.

The order in which the devices are added to the device chain determines the order in which Data Protector uses them.

When all of the media in a device chain are full, Data Protector issues a mount request. Replace the medium in the first device with a new medium, format it, and then confirm the mount request. Data Protector can immediately use media that are recognized and unprotected. Also blank media can be used.

9. Select **Automatically discover changed SCSI address** if you want to enable automatic discovery of changed SCSI addresses. Click **Next**.
10. In the Media Type list, select a media type for the device that you are configuring.
11. Specify a media pool for the selected media type. You can either select an existing pool from the Media Pool drop-down list or enter a new pool name. In this case, the pool will be created automatically.
12. Click **Finish** to exit the wizard.

The name of the device is displayed in the list of configured devices. You can scan the device to verify the configuration. If the device is configured correctly, Data Protector will be able to load, read, and unload media in the slots.

Configuring Backup to Disk Devices

Before performing a backup using a Backup to Disk (B2D) device, you need to configure the device for use with Data Protector. The available Backup to Disk devices are: StoreOnce Backup system, StoreOnce Software, Cloud (Helion), Cloud (Azure), Data Domain Boost, and Smart Cache.

Multi-Interface Support

Data Protector provides multi-interface support. 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*, and *HPE Data Protector Command Line Interface Reference*.

Steps

To add a B2D device (which targets an existing store), proceed as follows:

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and then click **Add Device** to open the wizard.
3. Specify a device name and its description (optional).
4. Select the **Backup to Disk** device type, and then select the **Interface Type: StoreOnce Backup system, Data Domain Boost, StoreOnce Software, Cloud (Helion), Cloud (Azure), or Smart Cache**.
5. The steps to configure the device vary based on the selected interface type.
 - [Configuring StoreOnce](#)
 - [Configuring Data Domain Boost](#)
 - [Configuring Smart Cache](#)
 - [Configuring StoreOnce Software](#)
 - [Configuring Cloud \(Helion\) Devices](#)
 - [Configuring Cloud \(Azure\) Devices](#)

The procedure for adding a B2D device is similar to the procedure for adding device types. Additionally, for StoreOnce Software deduplication devices, you must first configure a root directory and then create a store (see [Configuring a Backup to Disk Device - StoreOnce Software](#)).

Configuring a Backup to Disk Device - StoreOnce

Before performing a backup using a Backup to Disk (B2D) device, you need to configure the device for use with Data Protector.

If you are configuring a StoreOnce Software deduplication device, some additional steps are necessary. See [Configuring a Backup to Disk Device - StoreOnce Software](#).

Note: Data Protector supports federated stores of up to eight members. The number of members in a store can be changed in StoreOnce. To reflect this change, you can manually refresh the Data Protector cache using the Data Protector GUI or CLI. For more information, see [Refreshing Cache for Stores](#). All federation members must be online for a federated store to function.

Steps

To add a StoreOnce Backup System or StoreOnce Software B2D device (which targets an existing store), proceed as follows:

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and then click **Add Device** to open the wizard.
3. Specify a device name and its description (optional).
4. Select the **Backup to Disk** device type, and then select the Interface Type: **StoreOnce Backup system** or **StoreOnce Software**.
5. Optionally, enter a valid URL of the device management console in the **Management Console URL** text box. Click **Next**.
6. For StoreOnce Backup system devices, enter **Client ID** and optionally the password for accessing the store. You can use the following characters for the password: [a-z][A-Z][0-9][_ . + () {} : # \$ % ' = ? @ [] ^ | ~] ?
7. In the **Deduplication System** box, enter the IP address, hostname, fully qualified domain name (FQDN), or Fiber Channel (FC) address of the deduplication system (the hosting machine where the deduplication store is located).

Or click **Select Service Set** to query and retrieve the address of the deduplication system.

Note: For the StoreOnce Software interface, an IPv4 or IPv6 address, or FQDN is supported. However, for the StoreOnce Backup system interface, an IPv4 or IPv6 address, FQDN, or an FC global identifier is supported, provided you are using the latest StoreOnce Catalyst version.

If you are connecting to the StoreOnce Backup system device using FC, specify the FC address of the device. Ensure that you use Media Agents or Gateways that are connected to the FC device and that they are in the same zone as the StoreOnce Backup system device.

It is recommended that you use the IP address or the FQDN to take advantage of the multi-interface feature. To understand what the feature is about, see [Multi-Interface Support](#).

8. Click the **Select / Create Store** button to select an existing federated or non-federated store or to create a non-federated store. Select the store name from the list.

To create an encrypted store, select the option **Encrypted store**. Click **OK**.

Note: Encryption can only be enabled at the time of the store creation. After the store is created, it is not possible to convert it from encrypted to non-encrypted state, or vice versa. The StoreOnce Software deduplication devices do not support encryption of stores. You cannot create federated stores using the Data Protector GUI. You need to create them using the StoreOnce management console.

9. Optionally, select **Source-side deduplication** to enable source-side deduplication. The Source-side deduplication properties window opens. Review and if necessary modify the properties. By default, the source-side gateway will be named `DeviceName_Source_side`. Note that you can create only one source-side gateway per device. This (virtual) gateway will then be automatically expanded on the backed up system if source-side deduplication is enabled in the backup specification.

Note: For federated stores, all writing operations are performed in the low bandwidth mode (server-side deduplication). Even if a gateway is configured as target-side deduplication (high bandwidth mode), it automatically switches to the low bandwidth mode.

10. Select a gateway and click **Add** to display the properties dialog. If necessary, change any gateway properties and then click **OK** to add the gateway. If you are connecting to the StoreOnce Backup system device using FC, ensure that you use Media Agents or Gateways that are connected to the FC device and that they are in the same zone as the StoreOnce Backup system device.

Note: The federation member connected to the Data Protector gateway must be a member of the federated store. If the federation member is contracted away using StoreOnce, adjust the Data Protector gateway to attach to a different federation member using the steps mentioned in [Refreshing Cache for Stores](#).

To view gateway properties, select the desired gateway and click **Properties**. To set additional gateway options, click the **Settings** tab and then click **Advanced** to open the Advanced properties window.

In the Advanced Properties window, to limit the number of streams on each gateway, select **Max. Number of Parallel Streams per Gateway**. You can specify up to a maximum of 100 streams. If this option is not selected, the number of streams is not limited. Note that you can also set up this option when creating a backup specification. In this case, the value specified during the creation of a B2D device will be overwritten.

To limit the network bandwidth used by the gateway, select **Limit Gateway Network Bandwidth (Kbps)** and enter the limit in kilobits per second (kbps).

To enable server-side deduplication, select **Server-side deduplication**.

If you have configured an IP address or FQDN as your deduplication target, then the **Use FC** and **Fallback to IP** options are available and they are selected by default.

11. To verify the connection, click **Check**.
12. Click **Next** to proceed to the Settings window, where you can specify the following options:
 - **Max. Number of Connections per Store**
 - **Backup Size Soft Quota (GB)**

- **Store Size Soft Quota (GB)**
 - **Catalyst Item Size Threshold (GB):** Defines the threshold size of the catalyst item for StoreOnce Software Deduplication and StoreOnce Backup system devices. When this size is exceeded, the objects will no longer be appended to the current catalyst item. By default, the catalyst item size is unlimited.
 - **Single Object per Catalyst Item:** Select to enable one object per catalyst item for StoreOnce Software deduplication and StoreOnce Backup system devices.
13. Click **Next** to display the Summary window, which includes details of the configured B2D store. In addition, for a federated store, it includes a list of all federation members and their status (Online or Offline).
 14. Review the settings and click **Finish**. The newly configured B2D device is shown in the Scoping pane.

Refreshing Cache for Stores

With StoreOnce 3.12, you can add or remove federation members from federated stores. To reflect this change, you can manually refresh the Data Protector cache using the Data Protector GUI or CLI.

Refreshing cache using the Data Protector GUI

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Devices**.
3. Right-click the desired StoreOnce device, and click **Properties**.
4. Click the **Store and Gateways** tab, and click **Select/Create Store**. If necessary, change the directory path to include the address of a currently active federation member.
5. Select the same store, which is associated with this StoreOnce device, and click **OK**.
6. Click **Apply**.

Refreshing cache using the Data Protector CLI

1. Execute the following command:
`omnidownload -library <DPDeviceName> -file <DPDeviceOutputFile>`

2. Edit `DPDeviceOutputFile`.

If the device is not federated, remove the following lines:

```
B2DTEAMEDSTORE 1
B2DTEAMEDMEMBERS
"<teamed.device.one>"
"<teamed.device.two>"
...
```

If the device is federated, add these lines to `DPDeviceOutputFile` after substituting the appropriate teamed device IP addresses. If necessary, change the directory path to include the address of a currently active federation member.

Note: The addresses and format should exactly match the ones in the StoreOnce teaming policy file. For example, if the teaming policy file includes an IPv6 address, you must add the same address in this file too.

3. Save the modified file using the following command:

```
omniupload -modify_library <DPDeviceName> -file <DPDeviceOutputFile>
```

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

Configuring a Backup to Disk Device - Smart Cache

Before performing a backup using a Backup to Disk (B2D) device, you need to configure the device for use with Data Protector.

Configuring Smart Cache

Prerequisites

- You must have user credentials of the Media Agent host in which you want to create the Smart Cache device. The VMware plug-in uses these credentials to access the network share during the non-staged recovery.

Note: In a single Media Agent host, only one operating system user credential must be used to create a Smart Cache device. If multiple users simultaneously create Smart Cache devices on the same Media Agent host, then VMware Granular Recovery requests may encounter "Access Denied" errors.

- For the Linux operating system, you must install and run the Samba server on the Smart Cache client, as Data Protector uses the Samba server to create shares during recovery. To verify that the Samba server is running, execute the following command: `ps -ef | grep smbd`. The default mode of security for the Samba server is *user-level*. If the default mode is changed, you must update it to *user-level* using the following command: `[global] security = user`.
- Ensure that the Samba shares have read-write permissions. If the Security-Enhanced Linux (SELinux) kernel security module is deployed in your Linux system, then execute the command `# setsebool -P samba_export_all_rw on` to enable read-write permissions for the Samba shares.
- On the Samba server, you must add the user of the Media Agent host to the samba password database using the following command: `smbpasswd -a <user>`. You can verify if the user has been added to the password database using the following command: `pdbedit -w -L`.
- You must perform a regular cleanup of the Samba configuration file, (`smb.conf`). This ensures that the previous Samba share configuration information is removed.
- You must deploy the VMware non-staged recovery agent and the Media Agent module on the same host if the Smart Cache storage is a Windows ReFS file system, a CIFS, or an NFS share.
- If the Smart Cache storage is a local fixed disk or a SAN Storage LUN, the VMware non-staged recovery Agent host and Media Agent module can be different.
- You must dedicate the entire file system to one Smart Cache device. This file system should not be used by other applications, and should not be shared by other Smart Cache/Backup to Disk devices.
- Only a single media pool can be associated with one smart cache device.

Limitations

- Smart Cache is available only on Windows x64 and Linux x64 platforms.
- For a Windows Smart Cache device located on a network share, non-staged GRE is supported only for Windows Server 2008 and later systems.
- Smart Cache is available as a target for VMware backups only.
- On Linux operating systems, backup to Smart Cache is not supported if the NDMP Media Agent package is installed.
- Encoded or AES 256-bit encrypted VMware backups to a Smart Cache device is not supported.
- Encoded or AES 256-bit encrypted object copy of a source to a Smart Cache device is not supported. However, objects copies to and from tape devices with hardware encryption are supported.
- Only one mount point per Smart Cache device is supported.
- Backup to Smart Cache device might fail if there is insufficient space. Ensure that there is excess disk space available in the Smart Cache device.
- Export and import of media is not supported by the Smart Cache device.
- If you create a Smart Cache device on a Resilient File System (ReFS) volume or a network share (CIFS/NFS), install the mount-proxy component (used for recovery) on the same host, else non-staged recoveries will fail.
- CIFS is not supported with Smart Cache device configuration on StoreOnce 4500.

Steps

1. Create a directory for the Smart Cache device in the required location on the disk, for example, `c:\SmartCache`.
You can create a Smart Cache device on a local or network drive (or an NFS mounted filesystem for Linux systems). To specify a network drive, use the following format: `\\hostname\share_name`.
Hostnames and their share names and network drives do not appear in the Browse Drives dialog. You must enter the path to UNC names.
2. On the **Windows** operating system, to obtain permissions for accessing the shared disk containing a Smart Cache device, change the Data Protector Inet account on the Media Agent. You can do this by providing access permissions for both the local client system and remote shared disks. In addition, ensure that it is a specific user account, and not the system account. After you set up the Inet account, configure and use Smart Cache devices on shared disks.
3. In Data Protector, in the Context List, click **Devices & Media**.
4. In the Scoping Pane, right-click **Devices** and then click **Add Device** to open the wizard.
5. Specify a device name and its description (optional).
6. Select the **Backup to Disk** device type, and select the **Smart Cache** interface type.
7. In the Client drop-down list, select the system where the device will reside. Click **Next**.
8. Enter the User Name and Password of the user who needs access to the share created during the non-staged recovery.
9. Specify a directory for the Smart Cache device. Click **Add**.

10. To change the default properties of a directory, select the directory and click **Properties**.
11. Click **Next** to display the Summary window. Review the settings and click **Finish**. The newly configured B2D device is shown in the Scoping pane.

Configuring a Backup to Disk Device - Data Domain Boost

Before performing a backup using a Backup to Disk (B2D) device, you need to configure the device for use with Data Protector.

Prerequisites

- To support replication between Data Domain devices, virtual synthetics must be enabled on the Data Domain devices.
 - Using `ssh`, connect to the Data Domain devices and run the following command:

```
ddboost option set virtual-synthetics enabled
```
- To support replication, the same Data Domain Boost user must be configured on both source and target devices with the same administrative role. For more information, see your Data Domain documentation.

Limitations

- When performing an interactive replication, only one session at the time can be selected for replication.
- Data Protector operations are not supported when the Encryption Strength is modified from its default value.

When referring to Data Domain Boost devices, the term “storage unit” is used instead of the term “store”.

Steps

To add a DDBoost B2D device (which targets an existing store), proceed as follows:

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and then click **Add Device** to open the wizard.
3. Specify a device name and its description (optional).
4. Select the **Backup to Disk** device type, and then select the **Interface Type: Data Domain Boost**.
5. Optionally, enter a valid URL of the device management console in the **Management Console URL** text box. Click **Next**.
6. Enter **Username** and **Password**. You can use the following characters for the password: `[a-z][A-Z][0-9][_-.+(){}:#$*;=?@[]^~]?`
7. Enter the storage unit name (this assumes that the storage unit already exists).
8. In the **Deduplication System** text box, enter the hostname, IP address, or FC address of the

deduplication system (the hosting machine where the deduplication storage unit is located).

Note: It is recommended that you use the IP address or the FQDN to take advantage of the multi-interface feature. To understand what the feature is about, see [Multi-Interface Support](#).

9. Optionally, select **Source-side deduplication** to enable source-side deduplication. The Source-side deduplication properties window opens. Review and if necessary modify the properties. By default, the source-side gateway will be named `DeviceName_Source_side`. Note that you can create only one source-side gateway per device. This (virtual) gateway will then be automatically expanded on the backed up system if source-side deduplication is enabled in the backup specification.

10. Select a gateway and click **Add** to display the properties dialog. If necessary, change any gateway properties and then click **OK** to add the gateway.

To view gateway properties, select the desired gateway and click **Properties**. To set additional gateway options, click the **Settings** tab and then click **Advanced** to open the Advanced properties window.

To limit the number of streams on each gateway, select **Max. Number of Parallel Streams per Gateway**. You can specify up to a maximum of 100 streams. If this option is not selected, the number of streams is not limited. Note that you can also set up this option when creating a backup specification. In this case, the value specified during the creation of a B2D device will be overwritten.

To limit the network bandwidth used by the gateway, select **Limit Gateway Network Bandwidth (Kbps)** and enter the limit in kilobits per second (kbps).

If you have configured an IP address or FQDN as your deduplication target, then the **Use FC** and **Fallback to IP** options are available and they are selected by default.

To enable server-side deduplication, select **Server-side deduplication**.

11. To verify the connection, click **Check**.
12. Click **Next** to proceed to the Settings window, where you can specify the following options:
 - **Max. Number of Connections per Storage Unit:** Defines the median of maximum write and read streams limits the physical connection.
 - **Backup Size Soft Quota (GB):** Enter the backup size soft quota (in GB)
 - **Store Size Soft Quota (GB):** Supported if one storage unit is created, or if quotas are manually enabled for the entire Data Domain Operating System (DD OS) and specified when the storage unit is created.
 - **Store Media Item Size Threshold (GB):** Defines the threshold size of the store item for Data Domain Boost devices. When this size is exceeded, the objects will no longer be appended to the current store item. By default, the store item size is unlimited.
 - **Single Object per Store Media Item:** Select to enable one object per store item for Data Domain Boost devices.
13. Click **Next** to display the Summary window, which includes details of the configured B2D storage unit.
14. Review the settings and click **Finish**. The newly configured B2D device is shown in the Scoping

pane.

Configuring Data Domain Boost on AIX Systems

To configure the Data Domain Boost over Fibre Channel (FC) protocol on AIX systems, you must install the AIX DDdfc device driver. The driver file name is DDdfc.1.0.0.x.bff, where x is the version number.

Steps

1. Log in to the AIX client as a root user.
2. Enter the `# smitty install` command.
3. Select **Install and Update Software**.
4. Select **Install Software**.
5. Enter the path `/usr/omni/drv` to install the DDdfc.1.0.0.x.bff file, where x is the version number.
6. Press **F4** to select the DDdfc.1.0.0.x version that you want to install.
7. Press **Tab** to toggle the value on the Preview only? Line to No.
8. Press **Enter** to accept the information and install the driver.

Configuring a Backup to Disk Device - StoreOnce Software

If you are configuring a StoreOnce Software deduplication device, additional steps are necessary.

- [Configuring the root directory of the deduplication stores](#)
- [Creating a store](#)

Configuring the root directory of the deduplication stores

This section describes how to configure the root directory of the stores. This must be done after installing the software and before creating the first deduplication store.

One StoreOnce Software deduplication system can host multiple deduplication stores providing the stores share the same root directory. Each store operates independently of the other, that is, deduplication only occurs within one store and each store has its own index table. Although all stores run under the same process, they can be started / stopped individually (this does not mean to physically start / stop a store, see the *Deduplication, White Paper - Appendix A: StoreOnceSoftware utility* for details). Operations cannot be done on a store if it is stopped (offline).

Stores sharing the same root directory cannot be separated physically. This design guarantees uniform loading on all disks and provides better performance.

Following successful installation, the StoreOnceSoftware utility starts in a mode where it is running but waiting for the root directory of the stores to be configured. A B2D device cannot be added and a store cannot be created until the root directory is configured.

The root directory of the stores can be configured from:

- The GUI: Follow the procedure for adding a device and when prompted, specify the root directory (see below for details).

- The CLI: Use the command `StoreOnceSoftware --configure_store_root` (see *Deduplication White Paper - Appendix A: StoreOnceSoftware utility* for details).

Note: The root directory must already exist (on the server) and you must have write permissions before it can be configured. This is because the (GUI) configuration process asks you to specify its location.

The procedure for configuring the root directory using the GUI is similar to creating a store but includes a few additional steps. Once the root directory has been configured, these additional steps are no longer necessary. To configure the root directory (and create a store at the same time), proceed as follows:

1. Follow the procedure for adding a device:
 - a. In the Devices & Media context, right-click **Devices > Add Device**.
 - b. Specify a device name, add a description, select the device type **Backup To Disk**, and select the interface **StoreOnce software deduplication**.
 - c. Optionally, enter a valid URL of the device management console in the **Management Console URL** text box.
 - d. Click **Next** to display the screen where you specify a store and a list of gateways.
 - e. For StoreOnce Backup system devices, enter **Client ID** and optionally the **password** for accessing the store.
2. In the Deduplication System box, enter the hostname, IP address, or fully qualified domain name (FQDN) of the hosting machine where the deduplication store is located.
3. Select a gateway, click **Add** to display the properties dialog, then click **OK** to add the gateway.
4. Click **Check**. The message Root directory not configured is displayed.
5. In the dialog, specify the root directory path (for example, `C:\Volumes\StoreOnceRoot`) where all the stores are to reside and click **OK**. (Note: Browsing to the valid root directory is not possible).
6. If the root directory exists, the dialog closes and device configuration continues. The `StoreOnceSoftware` utility creates a subdirectory (the store) in the specified root directory. If the root directory does not exist, an error message is displayed.
7. Continue with the procedure for [adding a device](#).

Note the following points when configuring the root directory and creating stores:

- Do not use the same disk where the operating system (OS) is installed.
- Use dedicated (exclusive) storage disks.
- Data Protector supports a maximum of 32 stores per volume.

Note: On Windows systems, to improve the performance, apply the following options to the NTFS volume where the stores root will be located:

Disable creation of short (DOS-like) file names on the volume with the command: `fsutil behavior set Disable8dot3 Volume 1`
Increase NTFS internal LogFile size with the command: `Chkdsk Volume /L:131072`

Creating a store

Before creating a store, make sure the root directory of the stores has been configured and the physical storage disks (LUN devices) are formatted and mounted on the StoreOnce Software deduplication

system. The LUN devices may be on local disks or a disk array (SCSI or Fiber Channel interface) or on a NAS device in the same LAN (iSCSI interface). When using iSCSI interface, the reliable network connection must provide a latency of at most 2 ms and a throughput of at least 1 Gbit/s.

A store can be created from:

- The GUI: Follow the procedure for adding a device and when prompted, specify the name of the store (see below for details).
- The CLI: Use the command `StoreOnceSoftware --create_store` (see the *Deduplication White Paper - Appendix A: StoreOnceSoftware utility* for details).

The procedure for creating a store is similar to adding a device but includes a few additional steps. To create a store, proceed as follows:

1. Follow the procedure for adding a device:
 - a. In the Devices & Media context, right-click **Devices** > **Add Device**.
 - b. Specify a device name, add a description, select the device type **Backup To Disk**, and select the interface **StoreOnce software deduplication**.
 - c. Click **Next** to display the screen where you specify a store and a list of gateways.
2. Select the Deduplication System and specify a name for the store. The maximum length of the store name is 80 characters (alphanumeric characters only).
 - a. Select a gateway, click **Add** to display the properties dialog, then click **OK** to add the gateway.
 - b. Click **Check** to verify the connection. If the store does not exist, it is created. (Note: Click **Next** will also verify the connection.)
 - c. Continue with the procedure for [adding a device](#).

If you specify the store name incorrectly, you cannot change it through the GUI. Run through the procedure again and create the store with the correct name. Use the CLI to delete the incorrectly-named store (assuming data has not been written to it).

Configuring Cloud (Helion) Devices

Configure a [Cloud Helion device](#) in preparation for performing object copies to the Cloud object store.

In preparation, the following steps must be completed:

- [Obtaining the HPE Public Cloud Project Name](#)
- [Obtaining the Authentication Service URL](#)
- [Creating the Access Keys](#)

Next, you can configure a Cloud (Helion) device as a backup to disk device, in Data Protector.

[Configuring a Backup to Disk Device - Cloud \(Helion\)](#)

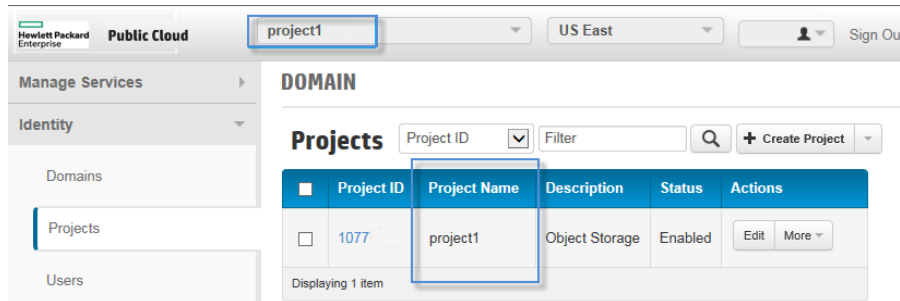
Obtaining the HPE Public Cloud Project Name

Steps

1. Log in to the HPE Public Cloud Console (<https://horizon.hpcloud.com>) with your HPE Public Cloud credentials.
2. Select the appropriate Project from the Project list.

- Take note of the Project name for later use in the Data Protector GUI. It will be specified in the Tenant / Project field during device creation.

Project in HPE Public Cloud



Obtaining the Authentication Service URL

Steps

- From the User menu, select **Roles and API Endpoints**. The User Roles and API Endpoints page opens.
- Click the **Service API Endpoints** tab. A list of Service API endpoints is displayed.
- For the geographic region closest to your datacenter, take note of the Service API Endpoint URL of the Service Type **identity**.

It will be specified later in the Authentication Service field during Cloud (Helion) device creation in the Data Protector GUI.

If you decide to use the access keys for authentication, take note of the Authentication Service URL that ends in the **/v3/** suffix.

For example:

<https://region-b.geo-1.identity.hpcloudsvc.com:35357/v3/>

Service API Endpoints in HPE Public Cloud

Current Roles

Service API Endpoints

Service API Endpoints

Service Name	URL(s)	Region	Service Type
Identity	Public URL: https://region-a.geo-1.identity.hpcloudsvc.com:35357/v2.0/	US West	identity
Identity	Public URL: https://region-a.geo-1.identity.hpcloudsvc.com:35357/v3/	US West	identity
Identity	Public URL: https://region-b.geo-1.identity.hpcloudsvc.com:35357/v2.0/	US East	identity
Identity	Public URL: https://region-b.geo-1.identity.hpcloudsvc.com:35357/v3/	US East	identity

Creating the access keys

Steps

1. From the User menu, select **Manage Access Keys**. The Manage Access Keys page opens.
2. To create a new key, specify a **Start Date** and **End Date** for the new key and click **Create Key**. The new key is created.

Create access keys in HPE Public Cloud

Manage Keys for:

Keys

Show Secret Keys

ID	Valid From	Valid To	Created On	Status	Actions
AAA123P09BOZ123	2013-07-01T15:30:07.000Z	2023-06-29T15:30:07.000Z	2013-07-01T15:30:07.273Z	active	Deactivate More ▾
B08DK123SDFAZ45	2014-03-25T00:00:00.000Z	2024-03-24T00:00:00.000Z	2014-03-25T15:51:36.465Z	active	Deactivate More ▾

Displaying 2 items

Create new key

Start Date *

End Date *

2014-04-02

2024-04-01

Create Key

3. Click **Show Secret Keys** to display the ID and secret keys for the new key.

Secret Keys in HPE Public Cloud

Manage Keys for:

Keys

ID	Valid From	Valid To	Created On	Secret Key	Status	Actions
ABC1DEF242CCCC	2013-07-01T15:30:07.000Z	2023-06-29T15:30:07.000Z	2013-07-01T15:30:07.273Z	1e43z2ABC09C1DWQasb30odL42ABC09C1D	active	Deactivate
1432ABC09CDWQ4	2014-03-25T00:00:00.000Z	2024-03-24T00:00:00.000Z	2014-03-25T15:51:36.465Z	A1bC1D6Fdh2ABC09C1D242Cm1C1C9dpaz2C	active	Deactivate

Displaying 2 items

4. Copy the Key ID and Secret Key information for later use. They will be specified during Cloud (Helion) device creation in the Data Protector GUI.

Configuring a Backup to Disk Device - Cloud (Helion)

In Data Protector, configure a backup to disk device with the interface type Cloud (Helion) device.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and then click **Add Device** to open the wizard.
3. Specify a device name and its description (optional).
4. Select the **Backup to Disk** device type, and then select the Interface Type: **Cloud (Helion)**. Click **Next**.
5. Specify the **Authentication Service URL**. This is the Service API Endpoint URL in [Obtaining the Authentication Service URL](#).
6. In the **Authentication mode** list, select a mode of authentication.
 - a. To use username and password authentication, select **Username and password** and input your HPE Public Cloud credentials.

- b. To use access keys for authentication, select **Access Keys** and input the **Access Key ID** and **Secret Key**. These are the keys noted in [Creating the access keys](#).

Note: To use the access keys for authentication, the Authentication Service URL must contain the **/v3/** suffix. For example:

`https://region-b.geo-1.identity.hpcloudsvc.com:35357/v3/`

7. Specify the **Tenant / Project**. This is the Project name from [Obtaining the Project Name](#).
8. Click **Select/Create Container** to select the containers from a list of already existing containers or create a new container.
9. Specify a gateway local to the data source.
 - a. Select a gateway and click **Add** to display the properties dialog. If necessary, change any gateway properties and then click **OK** to add the gateway.
10. Click **Next** to display the Summary window. Review the settings and click **Finish**. The newly configured device is shown in the Scoping pane.

Configuring a Backup to Disk Device - Cloud (Azure)

In Data Protector, configure a Backup to Disk device with the **Interface Type**: Cloud (Azure).

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and then click **Add Device** to open the wizard.
3. Specify a device name in **Device Name** field and its **Description** is optional.
4. Select the **Backup to Disk** device type, and then select the **Interface Type**: Cloud (Azure). Click **Next**.

The **Management Console URL** is entered by default.

5. Enter the **Storage Account Name**, **Secret key**, **Secret key 2** information in the fields. Click **Add** to add a gateway to send the data to Cloud (Azure). The Select Container window appears.
6. Select an existing container or create a new container to upload the data. Gateway can be added with default values.

There is a limitation for block size to object copy. If you are doing object copy from the local device to cloud and back to the same device for restore, the block size for the local device and cloud device must match.

7. Click **Check** to verify if the gateway is connected to Cloud (Azure). If the connection is successful, then the Status is displayed as **OK**. The device is created and is ready for use.

Configuring a File Library Device

Note that the disk on which the file library device resides must be local to the Media Agent. If it is not, device performance could be slow.

Prerequisites

- The disk on which the file library device will reside must be visible in the filesystem in which the file library device resides.
- The directory in which the contents of the file library device are to be created must exist on the disk where the file library device will reside.
- If you are creating a file library device on a Windows system, disable the Windows compression option for a file that you want to use as the file library device.

Limitations

- The file library device can include one or several directories. Only one directory can be located on a filesystem.
- The length of the pathnames of the directories that can be used for configuring devices of the file library type cannot exceed 46 characters.

Steps

1. Create a directory for the file library device on the disk where you want the device to be located, for example: `c:\FileLibrary`.

A file library device can be created on a local or network drive (or NFS mounted filesystem on UNIX systems). The network drive can be specified in the form `\\hostname\share_name` or can be mapped to a drive letter (`S:\datastore\My_FileLibrary`).

Hostnames along with the share names and network drives do not appear in the Browse Drives dialog where you enter the path. You need to enter the path to UNC names or network drives yourself.

On a **Windows** operating system, to get the right permissions for accessing the shared disk on which a file library device resides, change the Data Protector Inet account on the Media Agent (by giving it the permission to access both the local client system and remote shared disks). Also, make sure it is a specific user account, not the system account. Once you set up the Inet account, you can configure and use file library devices on shared disks.

It is critical that the directory created for the file library is not deleted from the disk. If it is deleted, any data within the file library device will be lost.

2. In the Data Protector Manager Context List, click **Devices & Media**.
3. In the Scoping Pane, right-click **Devices** and then click **Add Device** to open the wizard.
4. In the Device Name text box, type a name for the file library device.
5. In the Description text box, type a description of the library (optional).
6. In the Device Type drop-down list, select **File Library**.
7. In the Client drop-down list, select the system where the device will reside. Click **Next**.
8. Specify a directory or a set of directories where you would like the file library to reside. Click **Add**.
9. To change the default properties of a directory select the directory and click **Properties**.
10. Enter the number of writers to the file library. This defaults to the number of directories you added. If you add more writers than the number of directories in the device it is possible that you will

improve device performance. This depends on the hardware configuration you have. You will need to test this in your environment. Click **Next**.

11. The Media type of the file library device is File. To enable virtual full backup within this file library, select **Use distributed file media format**. Click **Next**.
12. Review the summary of the file library device configuration. Click **Finish** to exit the wizard.

The name of the device is displayed in the list of configured devices. The device name also appears in the media pool to which the device was assigned.

File depots will not appear in the device until it has been used for the first time.

You can scan the device to verify the configuration after the device has been used for the first time.

By default, media usage policy of the media pool used by the file library is non-appendable. The use of this policy is recommended, as this gives you the benefits of the file library, such as automatic reuse of expired media. Furthermore, to perform object copying or object consolidation using the file library, non-appendable media usage policy is required.

About Configuring Multiple Paths to Devices

A device in a SAN environment is usually connected to several clients and can thus be accessed through several paths, that is client names and SCSI addresses (device files on UNIX systems). Data Protector can use any of these paths. You can configure all paths to a physical device as a single logical device - *multipath device*.

For example, a tape device is connected to client1 and configured as `/dev/rs1` and `/dev/rs2`, on client2 as `/dev/r1s1` and on client3 as `scsi1:0:1:1`. Thus, it can be accessed through four different paths: `client1:/dev/rs1`, `client1:/dev/rs2`, `client2:/dev/r1s1` and `client3:scsi1:0:1:1`. A multipath device therefore contains all four paths to this tape device.

Why use multiple paths

With previous versions of Data Protector, a device could be accessed from only one client. To overcome this problem, several logical devices had to be configured for a physical device using a lock name. Thus, if you were using lock names for configuring access from different systems to a single physical device, you had to configure all devices on every system. For example, if there were 10 clients which were connected to a single device, you had to configure 10 devices with the same lock name. With this version of Data Protector, you can simplify the configuration by configuring a single multipath device for all paths.

Multipath devices increase system resilience. Data Protector will try to use the first defined path. If all paths on a client are inaccessible, Data Protector will try to use paths on the next client. Only when none of the listed paths is available, the session aborts.

Path selection

During a backup session, device paths are selected in the order defined during the configuration of that device, except when the preferred client is selected in the backup specification. In this case, the preferred client is used first.

During a restore session, the paths are selected in the following order:

1. Paths that are on the client to which the objects are restored, if *all* objects are restored to the same target client
2. Paths that were used for backup
3. Other available paths

For devices with multiple configured paths, the local paths are preferred. If no local path is available, any available path in the predefined order is used.

If direct library access is enabled, local paths (paths on the destination client) are used for library control first, regardless of the configured order.

The Data Protector Backup Session Manager (BSM) uses local devices as much as possible in multipath SAN environments. You can tune this behavior using the LANfree global option.

The LANfree global option has two possible values:

- 0 – Is the default value. No changes are required for earlier Data Protector versions below 8.11.
- 1 – Is applicable for multipath environment where Data Protector selects the host from where the object comes (if such a path is available), instead of selecting the preferred host or the first host from the multipath list.

The following describes the actual multipath device assignment improvements when the LANfree global option is set to 1:

- Data Protector prefers the host from where the data originates for a device that has a configured path to that host.
- Data Protector starts a new Media Agent (MA) on the host where the data originates for a device that has a configured path to that host. This is done even if a remote MA has already been started for the target device with a free concurrency slot.

Data Protector may still not use local paths for devices in the following scenarios:

- If a user has specified load balancing (MIN or MAX parameters), the BSM may choose and lock devices that are not local to any of the hosts from where the data originates.
- If a MA controlling a multipath device executes on one host, and an object comes from another host that has a path to the device, Data Protector will not migrate the MA to the local host but stream data over the LAN to the already started MA. This happens when the MAX value of load balancing has already been met.
- The LANfree setting is disabled when the IgnoreObjectLocalityForDeviceSelection global option is set. By default, the IgnoreObjectLocalityForDeviceSelection is not set.

In the following cases, the user may need to add additional device paths for achieving LAN-free backups:

- When a backup client has multiple network interfaces and hostnames. In this case, depending on the DNS configuration, Data Protector backups could go through multiple interfaces. Adding local paths for each interface would then be advisable.
- When performing a filesystem backup of a Windows file server which is a Windows cluster resource. In such a setup, each Windows cluster resource has its own hostname for which a separate device path entry should be created.

Backward compatibility

Devices configured with previous versions of Data Protector are not reconfigured during the upgrade and can be used as in previous releases of Data Protector without any changes. To utilize the new multipath functionality, reconfigure devices as multipath devices.

Limitations

The following limitations apply:

- Multiple paths are not supported for NDMP devices and Jukebox libraries.
- Device chains are not supported for multipath devices.

Setting Advanced Options for Devices and Media

You can set advanced options for devices and media when configuring a new device, or when changing device properties. The availability of these options depends on the device type.

Some of these options can also be set when configuring a backup. Device options set in a backup specification override options set for the device in general.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Devices**.
3. Right-click the device (the drive in the case of library devices) for which you want to change the options, and click **Properties**.
4. Click the **Settings** tab, and then click the **Advanced** button to open the Advanced Options pages: **Settings**, **Sizes**, and **Other**.
5. Specify the desired option(s), and then click **OK** to apply the changes.

Configuring a VTL Device

Before performing a backup to the Virtual Library System (VLS), you need to configure a Virtual Tape Library (VTL) device for use with Data Protector.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Environment**, then right-click **Devices** and click **Add Device** to open the wizard.
3. In the Device Name text box, enter the name of the VTL.
4. In the Description text box, enter a description (optional).
5. Optionally, select **MultiPath device**.

6. In the Device Type list, select **SCSI Library**. **SCSI** is then automatically selected in the Interface Type list.
7. If the **MultiPath device** option is *not* selected, select the name of the client in the Client list.
8. Optionally, enter a valid URL of the library management console in the **Management Console URL** text box. Click **Next**.
9. Specify the required information about the library SCSI address and drive handling, and click **Next**.
10. Specify the slots that you want to use with Data Protector, and click **Next**.
11. Select the media type that will be used with the device.
12. Click **Finish** to exit the wizard.

Note: If you are using the VTL device on RedHat Linux (RHEL) 7.1 systems, you must manually load the generic SCSI driver. You can do this by executing the command `modprobe -vs sg`. It is also recommended that you add this command to the RHEL `init scripts` or `cron job` to ensure that this command is initiated when the system starts.

Configuring a Stacker Device

After you have connected the backup device to the system and a working device file (SCSI address) exists, you can configure it for use with Data Protector.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and then click **Add Device** to open the wizard.
3. In the Device Name text box, enter the name of the device.
4. In the Description text box, enter a description (optional).
5. Optionally, select **MultiPath device**.
6. If the **MultiPath device** option is not selected, select the name of the client.
7. Click **Next**.
8. In the Device Type list, select the **Stacker** device type and then click **Next**.
9. In the Data Device text box, enter the SCSI address of the physical device (Windows systems), enter a device filename (UNIX systems), or use the drop-down arrow to auto-detect the drive addresses or filenames.

For multipath devices, select also the client name and click **Add** to add the path to the list of configured paths.
10. Select **Automatically discover changed SCSI address** to enable automatic discovery of changed SCSI addresses.
11. Click **Next**.
12. In the Media Type drop-down list, select a media type for the device that you are configuring.
13. Specify a media pool for the selected media type. You can either select an existing pool from the Media Pool drop-down list or enter a new pool name. In this case, the pool will be created

automatically.

14. Click **Finish** to exit the wizard.

The name of the device is displayed in the list of configured devices. You can scan the device to verify the configuration. If the device is configured correctly, Data Protector will be able to load, read, and unload media in the slots.

Stacker device media management

After configuring a stacker device, consider that managing media in such device has some specifics. For example, the operations scan, verify, or format have to be run separately on each medium in a stacker device. You should properly load a medium to be able to run Data Protector sessions.

Configuring a Jukebox Device (Optical Library)

After you have connected the backup device to the system and a working device file (SCSI address) exists, you can configure it for use with Data Protector.

Configuring a jukebox device

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and then click **Add Device** to open the wizard.
3. In the Device Name text box, enter the name of the device.
4. In the Description text box, enter a description (optional).
5. In the Device Type list, select the **Jukebox** device type.
6. In the Client list, select the name of the client.
7. Optionally, enter a valid URL of the library management console in the **Management Console URL** text box.
8. Click **Next**.
9. Specify a set of files/disks for the jukebox. Use a dash to enter multiple files or disks at a time, for example, /tmp/FILE 1-3, and then click **Add**. For magneto-optical jukeboxes, the disk names have to end on A/a or B/b. Click **Next**.
10. In the Media Type list, select a media type for the device that you are configuring.
11. Click **Finish** to exit this wizard. You are prompted to configure a library drive. Click **Yes** and the drive configuration wizard displays.

Configuring a drive in the jukebox device

Steps

1. In the Device Name text box, enter the name of the device.
2. In the Description text box, optionally enter a description.
3. Specify a media pool for the selected media type. You can either select an existing pool from the Media Pool list or enter a new pool name. In this case, the pool will be automatically created. You can configure one media pool for all drives or have an independent media pool for each drive. Click **Next**.
4. Optionally, select **Device may be used for restore** and/or **Device may be used as source device for object copy** and specify a **Device Tag**.
5. Click **Finish** to exit the wizard.

The name of the drive is displayed in the list of configured drives. You can scan the drives to verify the configuration.

Configuring a SCSI Library or a Magazine Device

After you have connected the backup device to the system and a working device file (SCSI address) exists, you can configure it for use with Data Protector.

The configuration procedure for a library and a magazine device is the same, except that you have to specify the media pool with the **Magazine support** option set when configuring a magazine device.

It is recommended that you let Data Protector configure backup devices automatically.

Configuring a SCSI library robotics

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and then click **Add Device** to open the wizard.
3. In the Device Name text box, enter the name of the device.
4. In the Description text box, enter a description (optional).
5. Optionally, select **MultiPath device**.
6. In the Device Type list, select the **SCSI Library** device type.
7. In the Interface Type list, select the **SCSI** interface type.
8. If the **MultiPath device** option is *not* selected, select the name of the client in the Client list.
9. Optionally, enter a valid URL of the library management console in the **Management Console URL** text box.
10. Click **Next**.
11. Enter the SCSI address of the library robotics or use the drop-down arrow to auto-detect the drive

addresses or filenames

For multipath devices, select also the client name and click **Add** to add the path to the list of configured paths.

12. In the **Busy Drive Handling** list, select the action Data Protector should take if the drive is busy.
13. Select **Automatically discover changed SCSI address** to enable automatic discovery of changed SCSI addresses.
14. Optionally, select **SCSI Reserve/Release (robotic control)**. Click **Next**.
15. Specify the slots for the device. Use a dash to enter slot ranges and then click **Add**. For example, enter 1-3 and click **Add** to add slot 1, 2, and 3 at once. Do not use letters or leading zeros. Click **Next**.
16. In the Media Type drop-down list, select a media type for the device that you are configuring.
17. Click **Finish** to exit this wizard. You are prompted to configure a library drive. Click **Yes** and the drive configuration wizard appears.

Configuring a drive in a library

Steps

1. In the Device Name text box, enter the name of the device.
2. In the Description text box, optionally enter a description.
3. Optionally, select **MultiPath device**.
4. If the **MultiPath device** option is *not* selected, select the name of the client in the Client list.

Tip: You can configure a library so that each drive receives data from a different system running a Data Protector Media Agent. This improves performance on high-end environments. From the Client drop-down list, select the client system that you want to use with each drive.

Click **Next**.

5. In the Data Drive text box, enter the SCSI address or filename of the data drive.
For multipath devices, select also the client name and click **Add** to add the path to the list of configured paths.
6. Select **Automatically discover changed SCSI address** to enable automatic discovery of changed SCSI addresses.
7. In the Drive Index text box, enter the index of the drive in the library. Click **Next**.
8. Specify a media pool for the selected media type. You can either select an existing pool from the Media Pool drop-down list or enter a new pool name. In this case, the pool will be created automatically. Using the default media pool is recommended.

Note: It is not necessary to configure all drives for use with Data Protector. You can configure one media pool for all drives or have an independent media pool for each drive.

When specifying the media pool for a magazine device, select one with the **Magazine Support** option set.

Click **Next**.

9. Optionally, select **Device may be used for restore** and/or **Device may be used as source device for object copy** and specify a **Device Tag**.
10. Click **Finish** to exit the wizard.

The name of the drive is displayed in the list of configured drives. You can scan the drives to verify the configuration. If the device is configured correctly, Data Protector will be able to load, read, and unload media in the slots.

Configuring Devices in a SAN Environment

A SAN environment can vary from one client using a library to several clients using several libraries. The clients can have different operating systems. The goals of the SAN Environment configuration from a Data Protector perspective are:

- On each host that is to share the library robotics, create a library robotics definition for each host. If there is only one host that is controlling the robotics, the library definition is created only for the default robotics control host.
- On each host that is to participate in sharing the same (tape) drives in the library:
 - Create a device definition for each device to be used.
 - Use a lock name if the (physical) device will be used by another host as well (shared device).
 - Optionally, select direct access if you want to use this functionality. If you use it, ensure that the `libtab` file is set up on that host.

Considerations

- Microsoft Cluster Server: Ensure that the drive hardware path is the same on both cluster nodes: once the device is configured, perform a failover to check it out.

Configuration Methods

There are three configuration methods that depend on the platforms that participate in the SAN configuration:

Automatic device configuration using the GUI

You can use Data Protector autoconfiguration functionality to automatically configure devices and libraries on multiple hosts in a SAN environment. Automatic configuration is provided on the following operating systems:

- Windows
- HP-UX
- Solaris
- Linux
- AIX

Limitations

Autoconfiguration cannot be used to configure the following devices in a SAN environment:

- mixed media libraries
- DAS or ACSLS libraries
- NDMP devices

Data Protector discovers the backup devices connected to your environment. For library devices, Data Protector determines the number of slots, the media type and the drives that belong to the library. Data Protector then configures the device by setting up a logical name, a lock name, the media type, and the device file or SCSI address of the device, as well as the drive and slots.

Note: When you introduce a new host into a SAN environment, the configured libraries and devices will not be updated automatically.

- To use an existing library on a new host, delete this library and autoconfigure a new library with the same name on the new host.
- To add devices to an existing library, either delete the library and then autoconfigure a library with the same name and new drives on a new host, or manually add the drives to the library.

Automatic device configuration using the CLI (the `sanconf` command)

You can configure devices and libraries in a SAN environment using the `sanconf` command. The `sanconf` command is a utility that provides easier configuration of libraries in SAN environments in single Data Protector cells as well as in MoM environments with Centralized Media Management Database (CMMDB). It can automatically configure a library within a SAN environment by gathering information on drives from multiple clients and configuring them into a single library. In MoM environments, `sanconf` can also configure any library in any Data Protector cell that uses CMMDB, provided that the cell in which `sanconf` is run uses CMMDB as well. `sanconf` is available on the following operating systems:

- Windows
- HP-UX
- Solaris

`sanconf` can detect and configure supported devices that are connected to clients running on the following operating systems:

- Windows
- HP-UX
- Solaris
- Linux
- AIX

Using this command you can:

- Scan the specified Data Protector, gathering the information on SCSI addresses of drives and robotic controls connected to the clients in the SAN environment.
- Configure or modify settings of a library or drive for given clients using the information gathered during the scan of Data Protector clients.
- Remove drives on all or the specified clients from a library.

Device locking

The `sanconf` command automatically creates lock names for drives that it is configuring. A lock name consists of the drive vendor ID string, the product ID string and the product serial number.

For example, the lock name for the HPE DLT 8000 drive with vendor ID "HP", product ID "DLT8000", and serial number "A1B2C3D4E5" will be HP:DLT8000:A1B2C3D4E5.

Lock names can also be added manually. Lock names are unique for each logical device.

You must not change the lock names that were created by the `sanconf` command. All other logical drives that are created manually and represent physical drives that have been configured by `sanconf` must also use lock names created by `sanconf`.

Limitations

- For a full list of libraries that are supported with `sanconf`, see the latest support matrices at <https://softwaresupport.hpe.com/>.
- `sanconf` does not provide the following features:
 - Placing spare drives in drive slots.
 - Mixing drive types; for example, combinations of DLT, 9840, and LTO drives.
 - Configuring clients that are currently unavailable. Configuration of such clients is possible only when the configuration of the library is performed using a configuration file that includes information gathered by scanning the clients.

Recommendation

Configure only one driver for a specific device on a system.

For information on how to use the `sanconf` command, see the `sanconf` man page or the *HPE Data Protector Command Line Interface Reference*.

Manual configuration on UNIX systems

When configuring shared devices connected to UNIX systems in a SAN environment manually, you have to:

- Create a device definition for each device to be used.
- Use a lock name.
- Optionally select direct access if you want to use this functionality. If you do so, you have to ensure that the `libtab` file on that host is properly configured.

Phases

1. [Manually configure devices](#)
2. [Manually configure the libtab file](#)

Configuring Devices in a SAN Environment Manually

The following procedure implies that the drive and the robotic are used by several systems, that the drive is used by several applications (not only Data Protector), and that all the systems send robotic control commands (direct library access). The following tasks also provide alternative steps to use if your environment differs.

For robotics control, you can use any client within the SAN. You need to configure the library robotics control first on a client which acts as the default robotics control system. This client will be used to manage media movements regardless of which client requests the media move. This is done in order to prevent conflicts in the robotics if several hosts request a media move at the same time. Only if the hosts fail, and direct access is enabled, is the robotics control performed by the local host requesting the media move.

Prerequisite

A Data Protector Media Agent (the General Media Agent or the NDMP Media Agent) must be installed on each client that needs to communicate with the shared library.

Configuration phases

[Configuring a library in the SAN environment](#)

[Configuring a drive in a library](#)

Configuring a library in the SAN environment

Note: If you want the robotic control to be managed by a cluster, you need to make sure that:

- The robotics control exists on each cluster node.
- The virtual cluster name is used in the library robotics configuration.
- The common robotics and device filenames are installed either using the `mksf` command or using the `libtab` file.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and then click **Add Device** to open the wizard.
3. In the Device Name text box, enter the name of the device.
4. Optionally, in the Description text box, enter a description.
5. Optionally, select **MultiPath device**.

6. In the Device Type drop-down list, select the **SCSI Library** device type.
7. In the Interface Type drop-down list, select the **SCSI** interface type.
8. If **Multipath device** is not selected, select the name of the client from the Client drop-down list.
9. Optionally, enter a valid URL of the library management console in the **Management Console URL** text box.
10. Click **Next**.
11. Enter the SCSI address of the library robotics or use the drop-down arrow to auto-detect the drive addresses or filenames.

For multipath devices, select also the client name from the client drop-down list. Click **Add** to add the path to the list of configured paths.
12. In the **Busy Drive Handling** list, select **Eject Medium**.
13. Select **Automatically discover changed SCSI address** if you want to enable automatic discovery of changed SCSI addresses. Click **Next**.
14. Specify the slots for the device. Use a dash to enter multiple slots at a time, and then click **Add**. For example, enter 1-3 and click **Add** to add slots 1, 2, and 3 all at once. Click **Next**.
15. In the Media Type drop-down list, select a media type for the device that you are configuring.
16. Click **Finish** to exit this wizard. You are prompted to configure a library drive. Click **Yes** and the drive configuration wizard appears. Follow the wizard as described in the task below.

Configuring a drive in a library

Configure each drive on each client from which you want to use it.

Steps

1. In the Device Name text box, enter the name of the drive.
It is recommended to use the following naming convention:
 - `LibraryLogicalName_DriveIndex_Hostname`, for example `SAN_LIB_2_hotdog` (for non-multipath devices)
 - `LibraryLogicalName_DriveIndex`, for example `SAN_LIB_2` (for multipath devices)
2. Optionally, in the Description text box, enter a description.
3. Optionally, select **Multipath device**.
4. If **Multipath device** is not selected, select the name of the client from the Client drop-down list.
5. Click **Next**.
6. In the Data Drive text box, enter the SCSI address or filename of the data drive

For multipath devices, select also the client name from the Client drop-down list. Click **Add** to add the path to the list of configured paths.
7. In the Drive Index text box, enter the index of the drive in the library.
8. Select **Automatically discover changed SCSI address** if you want to enable automatic discovery of changed SCSI addresses. Click **Next**.
9. Specify a media pool for the selected media type. You can either select an existing pool from the

Media Pool drop-down list or enter a new pool name. In this case, the pool will be created automatically.

You can configure one media pool for all drives or have an independent media pool for each drive.

10. Click the **Advanced** button. In the **Settings** tab, select the **Use direct library access** option.
Do NOT select the **Use direct library access** option if you want only one system to send robotic control commands that initiate Data Protector. The client system that you selected when configuring the library/drives with Data Protector will control the library robotic.
11. This step is *not* required for multipath drives. Click **Next**.
 - If Data Protector is the only application accessing the drive, click the **Other** tab, select the **Use Lock Name** option, and enter a name. Remember the name, since you will need it when configuring the same drive on another client. It is recommended to use the following naming convention:
LibraryLogicalName_DriveIndex, for example SAN_LIB_D2
 - If Data Protector is not the only application accessing the drive, select the **Use Lock Name** option, and ensure that operational rules provide exclusive access to all devices from only one application at a time.
 - If the drive is used by only one system, do NOT select the **Use Lock Name** option.
12. Optionally, select **Device may be used for restore** and/or **Device may be used as source device for object copy** and specify a **Device Tag**.
13. Click **Finish** to exit the wizard.

The drive is used by several systems and several applications (not only by Data Protector) Use device locking (define a Lock Name) and ensure that operational rules provide exclusive access to all devices from only one application at a time

The name of the drive is displayed in the list of configured drives. You can scan the drives to verify the configuration.

Configuring the libtab File in the SAN Environment

The purpose of the `libtab` files is to map the library robotic control access to work also on the "direct access requesting system", since here the local control path is likely to be different from the one used on the default library robotic control system.

You need one `libtab` file located on every Windows and UNIX client which needs "direct access" to the library robotic and is not equal to the system configured as the default library robotics control system.

Steps

1. Create the `libtab` file in plain text format on each system requesting direct access in the following directory:
 - Windows systems:** `Data_Protector_home\libtab`
 - HP-UX and Solaris systems:** `/opt/omni/.libtab`
 - Other UNIX systems:** `/usr/omni/.libtab`

2. Provide the following information in the `libtab` file:

FullyQualifiedHostname DeviceFile | SCSIPath DeviceName

- The *FullyQualifiedHostname* is the name of the client requesting direct access control for the library robotics. If the client is part of a cluster, the node name should be used.
- The *DeviceFile | SCSIPath* is the control path to the library robotic driver on this client.
- The *DeviceName* is the name of the device definition used on this client.

You need one line per device for which you request direct access.

If the system is part of a cluster, the *FullyQualifiedHostname* must be the virtual server name and the *DeviceFile | SCSIPath* must refer to the cluster node (physical system).

Configuring an ADIC/GRAU DAS Library Device

Data Protector provides a dedicated ADIC/GRAU library policy for configuring an ADIC/GRAU library as a Data Protector backup device.

Each system on which you install a Media Agent software and it accesses the library robotics through the DAS Server is called a DAS Client.

The following may provide additional information:

- The ADIC/GRAU functionality is subject to specific Data Protector licenses. For details, see the *HPE Data Protector Installation Guide*.
- Since this library manages media used by different applications, you have to configure which media and drives you want to use with Data Protector, and which media you want to track.
- Data Protector maintains its own independent media allocation policy and does not make use of scratch pools.

Configuration phases

1. [Connecting library drives](#)
2. [Preparing for installation of a Media Agent](#)
3. [Installing a Media Agent](#)
4. [Configuring the ADIC/GRAU DAS library device](#)
5. [Configuring a drive in the ADIC/GRAU DAS library device](#)

Connecting library drives

Steps

1. Physically connect the library drives and robotics to the systems where you intend to install a Media Agent software.

For information on how to physically attach a backup device to UNIX and Windows systems, see the *HPE Data Protector Installation Guide*.

2. Configure the ADIC/GRAU library. See the documentation that comes with the ADIC/GRAU library for instructions.

For details about supported ADIC/GRAU libraries, see <https://softwaresupport.hpe.com/>.

Preparing for installation of a Media Agent

Steps

1. If the DAS server is based on OS/2, before you configure a Data Protector ADIC/GRAU backup device, create or update the C:\DAS\ETC\CONFIG file on the DAS server computer.

In this file, a list of all DAS clients has to be defined. For Data Protector, this means that each Data Protector client with a Media Agent installed must be defined.

Each DAS client is identified with a unique client name (no spaces), for example OMNIBACK_C1. In this example, the contents of the C:\DAS\ETC\CONFIG file should look like this:

```
client client_name = OMNIBACK_C1,  
# hostname = AMU,"client1"  
ip_address = 19.18.17.15,  
requests = complete,  
options = (avc,dismount),  
volumes = ((ALL)),  
drives = ((ALL)),  
inserts = ((ALL)),  
ejects = ((ALL)),  
scratchHPools = ((ALL))
```

These names have to be configured on each Data Protector Media Agent client as the `omnirc` option `DAS_CLIENT`. The `omnirc` file is either the file `omnirc` in the *Data_Protector_home* directory (Windows systems) or the file `.omnirc` (UNIX systems). For example, on the system with the IP address 19.18.17.15, the appropriate line in the `omnirc` file is `DAS_CLIENT=OMNIBACK_C1`.

2. Find out how your ADIC/GRAU library slot allocation policy has been configured, either statically or dynamically. See the *AMU Reference Manual* for information on how to check what type of allocation policy you use.

The static policy has a designated slot for each volser while the dynamic allocation policy assigns the slots randomly. Depending on the policy that has been set configure Data Protector accordingly.

If the static allocation policy has been configured add the following `omnirc` option to your system controlling the robotics of the library:

```
OB2_ACIEJECTTOTAL = 0
```

Note that this applies to HP-UX and Windows.

Contact ADIC/GRAU support or review ADIC/GRAU documentation for further questions on the configuration of your ADIC/GRAU library.

Installing a Media Agent

You can either install the General Media Agent or the NDMP Media Agent on systems that will be physically connected to a backup drive in a ADIC/GRAU library and on the system that will access the library robotics through the DAS Server.

Note: You need special licenses, depending on the size of the repository with media or the number of drives and slots used in the ADIC/GRAU library. For more information see the *HPE Data Protector Installation Guide*.

Prerequisites

- The ADIC/GRAU library has to be configured and running. On how to configure an ADIC/GRAU library, see the documentation that comes with the ADIC/GRAU library.

- The DAS server has to be up and running and the DAS clients have to be properly configured.

You require the DAS software to control the ADIC/GRAU library. It consists of a DAS server and multiple DAS clients. For more information on DAS software, see the documentation that comes with the ADIC/GRAU library.

- Obtain the following information before you install the Media Agent:

- The hostname of the DAS Server.

- A list of available drives with the corresponding DAS name of the drive.

If you have defined the DAS Client for your ADIC/GRAU system, run the following commands to get this list:

```
dasadmin listd2 [client] or
```

```
dasadmin listd [client], where [client] is the DAS Client for which the reserved drives are to be displayed.
```

The `dasadmin` command is located in the `C:\DAS\BIN` directory on the OS/2 host, or in the directory where the DAS client has been installed:

Windows systems: `%SystemRoot%\system32`

UNIX systems: `/usr/local/aci/bin`

- A list of available Insert/Eject Areas with corresponding format specifications.

You can get this list in Graphical configuration of AMS (AML Management Software) on OS/2 host:

In the Admin menu, click **Configuration** to start the configuration. Double-click **I/O** to open the EIF-Configuration window, and then click the **Logical Ranges**. In the text box, the available Insert/Eject Areas are listed.

Note that one Data Protector library device can handle only one media type. It is important to remember which media type belongs to each one of the specified Insert and Eject Areas, because you will need this data later for configuring Insert/Eject Areas for the Data Protector library.

- **Windows systems:** A list of SCSI addresses for the drives, for example, `scsi4:0:1:0`.

- **UNIX systems:** A list of UNIX device files for the drives.

Run the `ioscan -fn` system command on your system to display the required information.

Steps

1. Distribute a Media Agent component to clients using the Data Protector graphical user interface and Installation Server.
2. Install the ADIC/GRAU library for client interface.

Windows systems:

- Copy the `aci.dll`, `winrpc32.dll` and `ezrpc32.dll` libraries to the `Data_Protector_home\bindirectory`. (These three libraries are part of the DAS client software shipped with the ADIC/GRAU library. They can be found either on the installation media or in the `C:\DAS\AMU\` directory on the AMU-PC.)
- Copy these three libraries to the `%SystemRoot%\system32` directory as well.
- Copy `Portinst` and `Portmapper` service to the DAS client. (These requirements are part of the DAS client software shipped with the ADIC/GRAU library. They can be found on the installation media.)
- In the Control Panel, go to **Administrative Tools, Services**, and start `portinst` to install `portmapper`.
- Restart the DAS client to start the `portmapper` service.
- In the Control Panel, go to **Administrative Tools, Services**, to check if `portmapper` and both `rpc` services are running.

HP-UX, Linux, and AIX systems:

Copy the shared library `libaci.sl` (HP-UX systems), `libaci.so` (Linux systems), or `libaci.o` (AIX systems) into the directory `/opt/omni/lib` (HP-UX and Linux systems) or `/usr/omni/lib` (AIX systems). You must have permissions to access this directory. Make sure that the shared library has read and executed permissions for everyone (root, group and others). (The `libaci.sl` and `libaci.o` shared libraries are part of the DAS client software shipped with the ADIC/GRAU library. They can be found on the installation media.)

3. After you have DAS software properly installed, execute the `devbra -dev` command to check whether or not the library drives are properly connected to your system. The command resides in the default Data Protector administrative commands directory.

A list of the library drives with the corresponding device files/SCSI addresses will be displayed.

Configuring the ADIC/GRAU DAS library device

When the ADIC/GRAU library is physically connected to the system and a Media Agent is installed, you can configure the ADIC/GRAU library device from Data Protector GUI. The DAS client will then access the ADIC/GRAU robotics during specific media management operations (Query, Enter, Eject).

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and click **Add Device**.
3. In the Device Name text box, type the name of the device.
4. In the Description text box, optionally type a description.
5. Optionally, select **MultiPath device**.
6. In the Device Type list, select **GRAU DAS Library**.
7. If the **MultiPath device** option is not selected, select the name of the Media Agent client that will access ADIC/GRAU robotics.
8. Optionally, enter a valid URL of the library management console in the **Management Console URL** text box.
9. Click **Next**.
10. In the DAS Server text box, type the hostname of the DAS Server.
For multipath devices, select also the client name and click **Add** to add the path to the list of configured paths.
11. In the **Busy drive handling** list, select the action Data Protector should take if the drive is busy and then click **Next**.
12. Specify the import and export areas for the library and then click **Add**. Click **Next**.
13. In the Media Type list, select the appropriate media type for the device.
14. Click **Finish** to exit the wizard. You are prompted to configure a library drive. Click **Yes** and the drive configuration wizard displays.

Configuring a drive in the ADIC/GRAU DAS library device

Steps

1. In the Device Name text box, type the name of the drive.
2. In the Description text box, optionally type a description.
3. Optionally, select **MultiPath device**.
4. If the **MultiPath device** option is not selected, select the name of the Media Agent client that will access ADIC/GRAU robotics.
5. Click **Next**.
6. In the Data Drive text box, specify the SCSI address of the device.
For multipath devices, select also the name of the Media Agent client that will access ADIC/GRAU robotics and click **Add** to add the path to the list of configured paths.
7. Select **Automatically discover changed SCSI address** to enable automatic discovery of changed SCSI addresses.
8. In the Drive Name text box, specify the ADIC/GRAU Drive name you obtained during the installation of a Media Agent. Click **Next**.

9. Select the **Default Media Pool** for the drive.
10. Click **Advanced** to set advanced options for the drive, such as **Concurrency**. Click **OK**. Click **Next**.
11. Optionally, select **Device may be used for restore** and/or **Device may be used as source device for object copy** and specify a **Device Tag**.
12. Click **Finish** to exit the wizard.

Configuring a StorageTek ACS Library Device

Data Protector provides a dedicated StorageTek ACS library policy for configuring a StorageTek ACS library as a Data Protector backup device.

Each system on which you install a Media Agent software and it accesses the library robotics through the ACSLS is called an ACS Client.

The following may provide additional information:

- The STK functionality is subject to specific Data Protector licenses. See the *HPE Data Protector Installation Guide* for details.
- Since this library manages media used by different applications, you have to configure which media and drives you want to use with Data Protector, and which media you want to track.
- Data Protector maintains its own independent media allocation policy and does not make use of scratch pools.

Configuration phases

1. [Connecting library drives](#)
2. [Installing a Media Agent](#)
3. [Configuring the StorageTek ACS library device](#)
4. [Configuring a drive in the StorageTek ACS library device](#)

Connecting library drives

Steps

1. Physically connect the library drives and robotics to the systems where you intend to install a Media Agent software.
For information on how to physically connect a backup device to UNIX and Windows systems, see the *HPE Data Protector Installation Guide*.
2. Configure the StorageTek ACS library. See the documentation that comes with the STK ACS library for instructions.
For details about supported StorageTek libraries, see <https://softwaresupport.hpe.com/>.

Installing a Media Agent

You can either install the General Media Agent or the NDMP Media Agent on systems that will be physically connected to a backup drive in a StorageTek library and on the system that will access the library robotics through the ACSLS.

Note: You need special licenses, depending on the size of the repository with media or the number of drives and slots used in the StorageTek library. For more information, see the *HPE Data Protector Installation Guide*.

Prerequisites

- The StorageTek library has to be configured and running. On how to configure a StorageTek library, see the documentation that comes with the StorageTek library.
- The following information has to be obtained before you start installing the Media Agent software:
 - The *hostname* of the host where ACSLS is running.

- A list of ACS drive IDs that you want to use with Data Protector. Log in on the host where ACSLS is running and execute the following command to display the list:

```
rlogin "ACSLS hostname" -l acssa
```

Enter the terminal type and wait for the command prompt. At the ACSSA prompt, enter the following command:

```
ACSSA> query drive all
```

The format specification of an ACS drive has to be the following:

```
ACS DRIVE: ID:#,#,#,# - (ACS num, LSM num, PANEL, DRIVE)
```

- Make sure that the drives that will be used for Data Protector are in the state `online`. If a drive is not in the `online` state, change the state with the following command on the ACSLS host:

```
vary drive drive_id online
```

- A list of available ACS CAP IDs and ACS CAP format specification. Log in on the host where ACSLS is running and execute the following command to display the list:

```
rlogin "ACSLS hostname" -l acssa
```

Enter the terminal type and wait for the command prompt. At the ACSSA prompt, enter the following command:

```
ACSSA> query cap all
```

The format specification of an ACS CAP has to be the following:

```
ACS CAP: ID:#,#,# (ACS num, LSM num, CAP num)
```

- Make sure that the CAPs that will be used for Data Protector are in the state `online` and in the `manual` operating mode.

If a CAP is not in the state `online`, change the state using the following command:

```
vary cap cap_id online
```

If a CAP is not in the manual operating mode, change the mode using the following command:

```
set cap manual cap_id
```

- **Windows systems:** A list of SCSI addresses for the drives, for example, `scsi4:0:1:0`.
- **UNIX systems:** A list of UNIX device files for the drives.

Run the `ioscan -fn` system command on your system to display the required information.

Steps

1. Distribute a Media Agent component to clients using the Data Protector GUI and Installation Server for Windows.
2. Start the ACS `ssi` daemon on all library hosts (Media Agent clients) with access to the robotics on the library.

Windows systems:

Install the LibAttach service. See the ACS documentation for details. Make sure that during the configuration of the LibAttach service the appropriate ACSLS hostname is entered. After the successful configuration, the LibAttach services are started automatically and will be started automatically after every system restart as well.

Note: After you have installed the LibAttach service, check if the `libattach\bin` directory has been added to the system path automatically. If not, add it manually.

For more information on the service, see the documentation that comes with the StorageTek library.

HP-UX and Solaris systems:

Execute the following command:

```
/opt/omni/acs/ssi.sh start ACS_LS_hostname
```

AIX systems:

Execute the following command:

```
/usr/omni/acs/ssi.sh start ACS_LS_hostname
```

3. From the default Data Protector administrative commands directory, execute the `devbra -dev` command to check whether or not the library drives are properly connected to your Media Agent clients.

A list of the library drives with the corresponding device files/SCSI addresses will be displayed.

Configuring the StorageTek ACS library device

When the StorageTek library is physically connected to the system and a Media Agent is installed, you can configure the StorageTek library device from Data Protector GUI. The ACS client will then access the StorageTek robotics during specific media management operations (Query, Enter, Eject).

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and click **Add Device**.
3. In the Device Name text box, type the name of the device.
4. In the Description text box, optionally type a description.
5. Optionally, select **MultiPath device**.
6. In the Device Type list, select **StorageTek ACS Library**.
7. If the **MultiPath device** option is not selected, select the Media Agent client that will access the StorageTek robotics.
8. Optionally, enter a valid URL of the library management console in the **Management Console URL** text box.
9. Click **Next**.
10. In the ACSLM Hostname text box, type the hostname of the ACS Server.
For MultiPath devices, select also the client name and add the path to the list of configured paths.
11. In the **Busy drive handling** list, select the action Data Protector should take if the drive is busy and then click **Next**.
12. Specify the **CAPs** for the library and then click **Add**. Click **Next**.
13. In the Media Type list, select the appropriate media type for the device.
14. Click **Finish** to exit the wizard. You are prompted to configure a library drive. Click **Yes** and the drive configuration wizard displays.

Configuring a drive in the StorageTek ACS library device

Steps

1. In the Device Name text box, type the name of the drive.
2. In the Description text box, optionally type a description.
3. Optionally, select **MultiPath device**.
4. If the **MultiPath device** option is not selected, select the Media Agent client that will access the StorageTek robotics.
5. Click **Next**.
6. In the Data Drive text box, specify the SCSI address of the device.
For multipath devices, select also the Media Agent client that will access the StorageTek robotics and click **Add** to add the path to the list of configured paths.
7. In the Drive Index text box, specify the StorageTek **drive index** you obtained during the installation of a Media Agent. Drive Index is a combination of four numbers separated by a comma. Click **Next**.
8. Select the **Default Media Pool** for the drive.
9. Click **Advanced** to set advanced options for the drive, such as **Concurrency**. Click **OK**. Click

Next.

10. Optionally, select **Device may be used for restore** and/or **Device may be used as source device for object copy** and specify a **Device Tag**.
11. Click **Finish** to exit the wizard.

About Using Backup Devices

Using backup devices applies to tasks such as scanning a device to identify the media in the device, locking a device by specifying a virtual lock name, performing a scheduled eject of media, automatic or manual cleaning of dirty drives, renaming a backup device and responding to a mount request to confirm that the needed medium is in a device.

Data Protector also provides a set of advanced options for devices and media, available according to the device type, which are beneficial to your device and media management.

Additionally, you can use several drive types in the same library, but you have to be aware of the media characteristics used.

When a device is for whatever reason inoperative, you can disable it for the backup and automatically use another device available from the list of devices. In case you don't want to use a device any more, you can remove it from the Data Protector configuration.

Devices & Media Advanced Options

Data Protector offers a set of advanced options for devices and media. The availability of these options depends on the device type. For example, more options are available for the configuration of a library than that of a standalone device.

You can set these options while configuring a new device or when changing the device properties. These options apply for the respective device in general. You can also tune a subset of the options listed to suit a specific backup specification. These options override options set for the device in general. You can access them while configuring or changing your backup specification.

For detailed information on advanced options, see the HPE Data Protector Help.

Advanced options - Settings

Concurrency

Options

- **CRC check**
- **Detect dirty drive**
- **Drive-based encryption**
- **Eject media after session**
- **Rescan**
- **Use direct library access** (SAN-specific option)

Advanced options - Sizes

- **Block size (KB)**
- **Disk agent buffers**
- **Segment size (MB)**

Advanced options - Other

Mount request

- **Delay (minutes)**
- **Script**

Device lock name

- **Use lock name**

Library with Several Drive Types

You can use several drive types of similar technology like DLT 4000/7000/8000 (the same is true within the DDS family) in the same library. This can lead into issues if you want to use the media in any drive, but do not ensure a common format on all media. For example, a DLT-4000 at restore time cannot read a tape which was written with a DLT-8000 (highest density). Compressed and non-compressed media also cannot be used interchangeably.

You can avoid these kind of problems by setting same density or creating different media pool for each drive type.

Same density setting

This method uses a common format on all media which allows to use all media interchangeably in any drive. For devices used on Windows systems, you need to consult the documentation of the drive on how to use a specific write density. On UNIX systems, you can set the density for drives when creating the device filename or by selecting the related device filenames and using them in the device definitions. The density must be set to the same value. For example, in case of DLT 4000 and DLT 7000, the DLT 4000 density should be set. You also have to ensure that the block size setting of the devices used is the same. You must use this setting in the device definition at the time the media get formatted. When all media have the same density setting, you can also use the free pool as desired. During restore, any drive can be used with any media.

Different media pool for each drive type

This method clearly separates the media used by one group of drives from the media used by another group of drives, allowing you to better optimize the drive and media usage. You can configure separate media pools for the different groups of drives. This allows you to use different density settings for different drive types. For example, create a DLT-4k-pool and a DLT-8k-pool. You must use this settings

in the device definition at the time the media get formatted. For example, the media in the pool for the DLT-8000-highest-density must be formatted by a DLT-8000 in highest density setting.

Free pool support

You cannot use one free pool "across" such pools. This would not identify media from the "other" pool to the devices correctly, they would be seen as foreign media. The free pool concept can be used only *with one pool* (like the DLT-8k-pool) *for each drive type*, in case the same media type (DLT) is written in an incompatible way. During restore you must be aware that media from a certain pool can only be used with related devices.

About Scanning

Scanning checks the format of media inserted in a drive, displays the content of the device's repository, and updates this information in the IDB.

- In a standalone device, you scan a medium in the drive.
- In a library device, you scan media in the selected slots.
- In a library device with barcode support, you scan media using barcodes.
- In a file library device, you update the information in the IDB about the file depots.
- With ADIC/GRAU DAS or STK ACS libraries, Data Protector queries an ADIC/GRAU DAS or an STK ACSLM Server and then synchronizes the information in the IDB with information returned from the Server.

When to use scanning

You scan a device whenever you want to update the Data Protector information about the media in the device. You must scan the device if you change the location of the media manually. Changing the location (slot, drive) manually creates inconsistencies with the information in the IDB because Data Protector is not aware of the manual change. Scanning synchronizes the MMDB with the media that are actually present at the selected locations (for example, slots in a library).

Ensure that all media in your cell have unique barcode labels. If an existing barcode is detected during a scan, then the medium that is already in the IDB is logically moved.

Perform scanning in a file library device, if you have moved one of the file depots to another location.

Limitations

Volsers scan may not complete successfully if the ADIC/GRAU library is configured with more than 3970 volsers in a repository. A workaround for this problem is to configure multiple logical ADIC/GRAU libraries in order to separate the slots from the large repository into several smaller repositories.

With ADIC/GRAU DAS and STK ACS libraries, when several logical libraries are configured for the same physical library, it is not recommended to query the DAS or STK ACSLM Server. Add volsers manually. With ADIC/GRAU DAS libraries, however, when logical libraries are not configured using Data Protector, but using the ADIC/GRAU DAS utilities, the Data Protector query operation can safely be used on such libraries.

Drive Cleaning

Data Protector provides several methods for cleaning dirty drives:

- Library built-in cleaning mechanism
Some tape libraries have a functionality for cleaning drives automatically when a drive requests head cleaning. When the library detects a dirty drive, it automatically loads a cleaning tape, and Data Protector is not notified of this action. This interrupts any active session, causing it to fail. This specific hardware-managed cleaning procedure is not recommended, since it is not compatible with Data Protector. Use automatic drive cleaning managed by Data Protector instead.
- Automatic drive cleaning managed by Data Protector
Data Protector provides automatic cleaning for most devices using cleaning tapes. For SCSI libraries and magazine devices, you can define which slots contain cleaning tapes. A dirty drive sends the cleaning request, and Data Protector uses the cleaning tape to clean the drive. This method prevents failed sessions due to dirty drives, provided that suitable media are available for backup. Automatic drive cleaning is supported for libraries with barcode support as well as for libraries without barcode support.
- Manual cleaning
If automatic drive cleaning is not configured, you need to clean the dirty drive manually. If Data Protector detects a dirty drive, a cleaning request appears in the session monitor window. You then have to manually insert a cleaning tape into the drive.

A special tape-cleaning cartridge with slightly abrasive tape is used to clean the head. Once loaded, the drive recognizes this special tape cartridge and starts cleaning the head.

Limitations

- Data Protector does not support the diagnostic vendor unique SCSI command for performing drive cleaning with cleaning-tapes stored in one of the special cleaning tape storage slots. These special cleaning tape storage slots are not accessible using the normal SCSI commands, and therefore cannot be used with automatic drive cleaning managed by Data Protector. Configure the standard slot(s) to store cleaning tape(s).
- Detection and use of cleaning tapes depends on the system platform where a Media Agent is running. For further information, see the *HPE Data Protector Product Announcements, Software Notes, and References*.
- You should not use another kind of device management application if you configure automatic drive cleaning managed by Data Protector, as this may cause unexpected results. This is due to the `cleanme` request being cleared as it is read, depending on the specific device type and vendor.
- Automatic drive cleaning for logical libraries with a shared cleaning tape is not supported. Each logical library needs to have its specific cleaning tape configured.

Conditions for automatic cleaning

- In libraries without barcode support, a cleaning-tape slot has been configured in the Data Protector device definition and contains a cleaning-tape cartridge. The cleaning-tape slot must be configured together with the other library slots.

- In libraries with barcode support, the barcode support must be activated to enable automatic drive cleaning. Cleaning tapes have a barcode label with CLN as its prefix, which enables Data Protector to recognize cleaning tape barcodes automatically.
- The configured drive has the **Detect dirty drive** option enabled.

When Data Protector receives notification that the drive needs cleaning, it automatically loads the cleaning tape, cleans the drive and then resumes the session. All cleaning activities are logged in the `cleaning.log` file residing in the Data Protector server log files directory.

Scheduled Eject of Media

Data Protector lets you perform a scheduled eject of media using the reporting functionality together with a script.

A program or script must be created on the Cell Manager to perform the ejection, and any applicable interpreters must also be installed on the Cell Manager.

You can set up and schedule a report group so that it creates a report and sends it as an input to a script. Such a report group should include a report that lists only the media you want to eject (you could, for example, use the List of Media Report). When the report group is started (as a result of a schedule or a notification such as the End of Session notification), Data Protector starts the script using the report result as input. The script parses the report and performs the eject of the specified media by using the Data Protector `omnim` CLI command.

You are notified in the Event Log Viewer, by default, if you need to remove media from mail slots so that the eject operation can continue (if, for example, there are more media to be ejected than there are empty mail slots in a library). If media are not removed from the mail slots after a default time and there are still media to be ejected, the `omnim` command aborts the operation. You can change the default time span in the `omnirc` file.

Device Locking

You can configure the same physical device many times with different characteristics simply by configuring the device with different device names. Thus, one physical device can be configured into several Data Protector backup devices and can be used for several backup sessions. The internal locking of logical devices prevents two Data Protector sessions from accessing the same physical device at the same time. For example, if one backup session is using a particular device, all other backup/restore sessions must wait for this device to become free before starting to use it. When a backup or restore session starts, the Data Protector locks the device, the drive, and the slot used for that session.

Media sessions performing media operations, such as initialize, scan, verify, copy, or import also lock devices. During that time, no other operations can lock and use the device. If a media session cannot obtain a lock, the operation fails, and you have to retry the operation at a later time.

When a backup or restore session issues a mount request, the lock is released, allowing you to perform media management operations only. The device will still be reserved so that no other backup or restore session can use the device. In addition, other media management operations are not allowed on the same drive during the first media operation. When the mount request is confirmed, the backup or restore session locks the device again and continues with the session.

Since the internal locking operates on logical devices rather than on physical devices, a collision can occur if you specify one device name in one backup specification and another device name for the same physical device in another backup specification. Depending on the backup schedule, Data Protector may try to use the same physical device in several backup sessions at the same time, which can cause collision. This can also happen when two device names are used in other operations, such as backup and restore, backup and scan, and so on. To prevent a collision when Data Protector might try to use the same physical device in several backup sessions at the same time, specify a virtual lock name in the device configurations. Data Protector then uses this lock name to check if the device is available, thus preventing collisions. You have to use the same lock name in all backup device configurations for the same physical device.

Note: The information about a physical device in the Device Flow report is taken from the currently configured device and it may not be the same as it was at the time when the device was actually used (for example, the device logical name was recently changed, but some sessions in the Internal Database still contain the former device name).

The Device Flow report always displays the current information - the current physical representation with the current logical device name.

Disabling a Backup Device

Disabling a backup device manually

If you disable a backup device, all subsequent backups skip the device. The next available device defined in the list of devices for the backup specification is used instead, provided that load balancing has been selected. All devices using the same lock name as the disabled device are also disabled.

This lets you avoid backups that fail because a device is damaged or in maintenance mode, while keeping other devices available (and configured) for backup.

Disabling a backup device is useful if a device is damaged or in maintenance mode.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**. The list of configured devices is displayed in the Results Area.
3. Right-click the device that you want to disable and then click **Properties**.
4. Click the **Settings** tab and then select the **Disable device** option.
5. Click **Apply**.

The device is disabled. To enable the device for backup, deselect the **Disable device** option.

Disabling a backup device automatically

You can configure Data Protector to automatically disable devices on which a certain number of unknown errors has occurred. You determine the threshold value by setting the `SmDeviceErrorThreshold` global option to `SmDeviceErrorThreshold=MaxNumberOfUnknownErrors`.

To enable the device for backup after it is fixed, right-click the device and click **Enable Device**.

Renaming a Backup Device

When you rename a backup device, the device is no longer used under its old name for backup or restore.

Make sure that you remove the device's old name from all backup specifications that used the device. Otherwise, Data Protector tries to back up to or restore from a device that does not exist, and the session fails.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**. The list of the configured devices is displayed in the Results Area.
3. Right-click the name of the device and then click **Properties**.
4. In the General property page, modify the name in the Device Name text box.
5. Click **Apply**.

The device is displayed in the list of the configured devices under the new name.

Removing a Backup Device

When you remove a backup device from the Data Protector configuration, the device is no longer used for backup or restore.

Make sure that you remove the device's old name from all backup specifications that used the device. Otherwise, Data Protector tries to back up to or restore from a device that does not exist, and the session fails.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**. The list of configured devices appears in the Results Area.
3. Right-click the device you want to remove and then click **Delete**. Confirm the action.

The device is removed from the list of the configured devices.

Tip: If you are not using a certain backup device with Data Protector anymore, you may want to remove the Media Agent software component from the system. This can be done using the Client context.

Responding to Mount Requests

You respond to a mount request to confirm that the needed medium is in a device. You have to be aware of how media are selected for backup.

Prerequisites

You either have to be added in the Admin user group or granted Monitor user rights.

Steps

1. In the Context List, select **Monitor**.
2. Insert the needed medium into the device. If you have a library device, it is not necessary to use the slot requested by mount request.
3. In the Results Area, double-click the session with the mount request status to display details about the session.
4. Select the device with the Mount Request status.
5. In the **Actions** menu, select **Confirm Mount Request** or right-click the device with the mount request status and select **Confirm Mount Request**.

The status of the session and device changes to Running.

About Storage Area Network (SAN)

What is SAN?

Storage Area Network (SAN), a network dedicated to data storage, is based on high-speed Fibre Channel technology. SAN provides off-loading storage operations from application servers to a separate network. Data Protector supports this technology by enabling multiple hosts to share storage devices connected over a SAN, which allows multiple-system to multiple-device connectivity. This is done by defining the same physical device multiple times, for example, once on every system that needs access to the device.

When using Data Protector in the SAN environment, you have to consider the following:

- Each system can have its (pseudo) local device, although the devices are typically shared among several systems. This applies to individual drives as well as the robotics in libraries.
- You have to take care to prevent several systems from writing to the same device at the same time. The access to the devices needs to be synchronized between all systems. This is done using locking mechanisms.
- SAN technology provides an excellent way to manage library robotics from multiple systems. This creates the ability to manage the robotics directly, as long as the requests sent to the robotics are synchronized among all the systems involved.

FC-AL and LIP

Using tape devices in Fibre Channel Arbitrated Loops (FC-AL) can cause certain anomalies that could abort a backup session. The problem appears because the FC-AL performs a Loop Initialization Protocol (LIP) whenever a new FC link is connected/disconnected, and whenever a system connected to the FC-AL is rebooted. This re-initialization of the FC-AL causes running backups to be aborted. Such terminated jobs should be restarted.


When a LIP occurs on the FC-AL Loop, any utility with an active I/O process gets an I/O error. For backup utilities attempting to use a shared tape, an I/O error causes failure of the current backup session:

- Tapes are rewound and unloaded
- The backup session aborts

The following is recommended:

- Do not add new devices or remove devices from the Arbitrated-Loop while backup sessions are running.
- Do not touch the FC components while the backup sessions are running. The static charge can cause a LIP.
- Do not use discovery on Windows or `ioscan` on HP-UX system since these also cause a LIP.

Example of multiple system to multiple device connectivity in SAN



Device Locking in the SAN Environment

Data Protector supports the SAN concept by enabling multiple systems to share backup devices in the SAN environment. The same device can be shared by multiple applications. It can also be shared by multiple systems in the Data Protector environment. The purpose of locking is to ensure that only one system at a time communicates with a device that is shared between several systems.

Locking devices used exclusively by Data Protector

If Data Protector is the only application that uses a drive, but that same drive needs to be used from several systems, you can use the device locking mechanism.

If Data Protector is the only application that uses a robotics control from several systems, Data Protector handles this internally assuming the library control is in the same cell as all the systems that need to control it. In such a case, all the synchronization of access to the device is managed by Data Protector internal control.

Locking devices used by multiple applications

If several systems are using Data Protector to access the same physical device, the device locking mechanism has to be used.

If Data Protector and at least one other application want to use the same device from several systems, the same (generic) device locking mechanism has to be used by every application. This mechanism has to work across several applications. This mode is currently not supported by Data Protector. In case this is required, operational rules must ensure exclusive access to all devices from one application only at a time.

Indirect and Direct Library Access

When configuring Data Protector with a SCSI Library device or silo libraries (ADIC/GRAU and StorageTek) there are two ways for client systems to access library robotics:

Indirect library access

With indirect library access, only one system (the default robotics control system), sends robotic control commands that are initiated from Data Protector. Any other system that requests a robotics function forwards the request to the robotics control system, which then sends the actual command to the robotics. This is done transparently within Data Protector for all requests from Data Protector.

Direct library access

With direct library access, every system sends control commands directly to the library robotics. Therefore, each system does not depend on any other system in order to function.

With direct library access and multiple systems sending commands to the same library, the sequence of this communication has to be coordinated.

In Data Protector every library definition is associated with a host controlling the library robotics (by default). If another host requests a medium to be moved, Data Protector will first access the system specified in the library definition to perform the media move. If the system is not available, a direct access from the local host to the library robotics can be used, if the `libtab` file is set. All of this is done transparently within Data Protector.

If direct library access is enabled for multipath devices, local paths (paths on the destination client) are used for library control first, regardless of the configured order. The `libtab` file is ignored with multipath devices.

Configuring Devices in a SAN Environment

A SAN environment can vary from one client using a library to several clients using several libraries. The clients can have different operating systems. The goals of the SAN Environment configuration from a Data Protector perspective are:

- On each host that is to share the library robotics, create a library robotics definition for each host. If there is only one host that is controlling the robotics, the library definition is created only for the default robotics control host.
- On each host that is to participate in sharing the same (tape) drives in the library:
 - Create a device definition for each device to be used.
 - Use a lock name if the (physical) device will be used by another host as well (shared device).
 - Optionally, select direct access if you want to use this functionality. If you use it, ensure that the `libtab` file is set up on that host.

Considerations

- Microsoft Cluster Server: Ensure that the drive hardware path is the same on both cluster nodes: once the device is configured, perform a failover to check it out.

Configuration Methods

There are three configuration methods that depend on the platforms that participate in the SAN configuration:

Automatic device configuration using the GUI

You can use Data Protector autoconfiguration functionality to automatically configure devices and libraries on multiple hosts in a SAN environment. Automatic configuration is provided on the following operating systems:

- Windows
- HP-UX
- Solaris
- Linux
- AIX

Limitations

Autoconfiguration cannot be used to configure the following devices in a SAN environment:

- mixed media libraries
- DAS or ACSLS libraries
- NDMP devices

Data Protector discovers the backup devices connected to your environment. For library devices, Data Protector determines the number of slots, the media type and the drives that belong to the library. Data Protector then configures the device by setting up a logical name, a lock name, the media type, and the device file or SCSI address of the device, as well as the drive and slots.

Note: When you introduce a new host into a SAN environment, the configured libraries and devices will not be updated automatically.

- To use an existing library on a new host, delete this library and autoconfigure a new library with the same name on the new host.
- To add devices to an existing library, either delete the library and then autoconfigure a library with the same name and new drives on a new host, or manually add the drives to the library.

Automatic device configuration using the CLI (the `sanconf` command)

You can configure devices and libraries in a SAN environment using the `sanconf` command. The `sanconf` command is a utility that provides easier configuration of libraries in SAN environments in single Data Protector cells as well as in MoM environments with Centralized Media Management Database (CMMDB). It can automatically configure a library within a SAN environment by gathering information on drives from multiple clients and configuring them into a single library. In MoM environments, `sanconf` can also configure any library in any Data Protector cell that uses CMMDB, provided that the cell in which `sanconf` is run uses CMMDB as well. `sanconf` is available on the following operating systems:

- Windows
- HP-UX
- Solaris

`sanconf` can detect and configure supported devices that are connected to clients running on the following operating systems:

- Windows
- HP-UX
- Solaris
- Linux
- AIX

Using this command you can:

- Scan the specified Data Protector, gathering the information on SCSI addresses of drives and robotic controls connected to the clients in the SAN environment.
- Configure or modify settings of a library or drive for given clients using the information gathered during the scan of Data Protector clients.
- Remove drives on all or the specified clients from a library.

Device locking

The `sanconf` command automatically creates lock names for drives that it is configuring. A lock name consists of the drive vendor ID string, the product ID string and the product serial number.

For example, the lock name for the HPE DLT 8000 drive with vendor ID "HP", product ID "DLT8000", and serial number "A1B2C3D4E5" will be HP:DLT8000:A1B2C3D4E5.

Lock names can also be added manually. Lock names are unique for each logical device.

You must not change the lock names that were created by the `sanconf` command. All other logical drives that are created manually and represent physical drives that have been configured by `sanconf` must also use lock names created by `sanconf`.

Limitations

- For a full list of libraries that are supported with `sanconf`, see the latest support matrices at <https://softwaresupport.hpe.com/>.
- `sanconf` does not provide the following features:
 - Placing spare drives in drive slots.
 - Mixing drive types; for example, combinations of DLT, 9840, and LTO drives.
 - Configuring clients that are currently unavailable. Configuration of such clients is possible only when the configuration of the library is performed using a configuration file that includes information gathered by scanning the clients.

Recommendation

Configure only one driver for a specific device on a system.

For information on how to use the `sanconf` command, see the `sanconf` man page or the *HPE Data Protector Command Line Interface Reference*.

Manual configuration on UNIX systems

When configuring shared devices connected to UNIX systems in a SAN environment manually, you have to:

- Create a device definition for each device to be used.
- Use a lock name.
- Optionally select direct access if you want to use this functionality. If you do so, you have to ensure that the `libtab` file on that host is properly configured.

Phases

1. [Manually configure devices](#)
2. [Manually configure the libtab file](#)

About Backup to Disk

Data Protector backup to disk saves data to disks rather than to tape. Data Protector writes to directories residing on one or many disks. The data is written to files residing in directories on the disk.

Disk backup is faster than backup to tape since there are no mechanical processes to carry out before the backup can be made, such as loading the tape for example. In addition, disk storage is becoming increasingly cheaper.

Many applications processing business critical data need to have each transaction backed up as soon as it is made. Disk-based backup means that disk can be written continuously to disk throughout the working day.

What is a disk-based backup device?

Conceptually, a disk-based backup device is similar to a tape drive or tape stack. The device has one or many directories which are the equivalent of a repository in a tape drive. When a backup is being made, a disk-based backup device writes data to file depots as if they were writing files to a tape. Since disk-based backup devices write data to files residing on disk, they are also referred to as 'file devices'.

How to configure disk-based devices?


Disk-based backup devices are configured using the Data Protector GUI. They use all of the Data Protector media management and backup and restore facilities.

About Backup to Disk Devices

A Backup to Disk (B2D) device is a device that backs up data to physical disk storage. The B2D device supports multi-host configurations. This means that a single physical storage can be accessed through multiple hosts called gateways. Each gateway represents a Data Protector client with the Media Agent component installed. The physical storage can also be partitioned into individual stores representing specific storage sections (this is similar to partitioning a hard disk). Each individual store on the physical storage can be accessed by one B2D device only. However, several B2D devices can access different stores on the same physical storage.

While similar to other library-based devices, B2D devices behave differently as gateways allow more flexibility. Unlike library drives, each gateway represents a host on which multiple Media Agents can be started simultaneously, either in single or multiple sessions.

B2D device (logical view)



The number of Media Agents that can be started on a specific gateway is defined by:

- Gateway limits. Each B2D gateway is limited to a maximum number of parallel streams.
- Connection limits on the store. This limit is specified in the GUI when configuring a B2D device. If the value is left unchecked, Data Protector uses the maximum available.
- The physical connection limitations of the physical storage unit. This value is retrieved from the physical store.
- Depending on the current operation, each Session Manager attempts to balance the number of Media Agents on a gateway with regards to the following input parameters:
 - The number of objects being backed up
 - Object location
 - Physical connection limitations.

B2D devices use a special data format for fast read/write access, which is incompatible with the traditional Data Protector tape format. The data format is automatically set when you select a B2D device.

About Deduplication

Data deduplication is a data compression technology which reduces the size of the backed up data by not backing up duplicate data. The deduplication process splits the data stream into manageable chunks (or blocks) of data. The contents of these data chunks are then compared to each other. If identical chunks are found, they are replaced by a pointer to a unique chunk. In other words, if 20 identical chunks are found, only one unique chunk is retained (and backed up) and the other 19 are replaced by pointers. The backed up data is written to a disk-based destination device called a deduplication store. When a restore operation is done, the unique chunk is duplicated and inserted in the correct position as identified by the pointer. With deduplication-type restore operations, the restore process is sometimes referred to as rehydration of the backed up data.

When to use deduplication

Typically, you would use data deduplication when backing up an e-mail system which may contain 100 instances of the same 1 MB graphic file attachment. If the system is backed up using a conventional backup technique, all 100 instances of the attachment are backed up. This requires approximately 100 MB of storage space. However, with data deduplication, only one instance of the attachment is actually stored. All other instances are referenced to the unique stored copy. In this example, the *deduplication ratio* is approximately 100 to 1. Although this example is referred to as file-level deduplication, it serves to demonstrate the benefits of using Backup to Disk devices and deduplication.

Advantages of deduplication

Generally, data deduplication increases the speed of the backup service as a whole and reduces overall storage costs. Data deduplication significantly reduces the amount of required disk storage space. Because data deduplication is a disk-based system, restore service levels are significantly higher and tape (or other media) handling errors are reduced.

Deduplication technologies

There are several deduplication technologies available in the marketplace. They are generally grouped into hardware-based and software-based solutions. These solutions can be further sub-grouped, for example, into file-level (single-instancing) or block-level deduplication.

Data Protector the following deduplication backends:

StoreOnce software deduplication

Data Protector's StoreOnce software duplication offers a software-based, block-level deduplication solution.

When using StoreOnce software deduplication, note the following:

- Deduplication backs up to disk-based devices only. It cannot be used with removable media such as tape drives or libraries.
- Because Data Protector uses a software-only approach to deduplication (that is, when using StoreOnce software deduplication), no specific hardware is required other than standard hard disks to store the backed up data.
- StoreOnce software deduplication uses hash-based chunking technology to split the data stream into sizeable chunks of data.
- In the deduplication process, duplicate data is deleted, leaving only one copy of the data to be stored, along with reference links to the unique copy. Deduplication is able to reduce the required storage capacity since only unique data is stored.
- Specifying a Backup to Disk target device in the backup specification, tells Data Protector to perform a deduplication-type backup.

HPE StoreOnce Backup system devices


HPE StoreOnce Backup system devices are disk to disk (D2D) backup devices which support deduplication.

Deduplication setup

Data Protector supports various deduplication setups:

- Source-side deduplication (1)—data is deduplicated at the source (the backed up system).
- Server-side deduplication (2)—data is deduplicated on the Media Agent system (the gateway).
- Target-side deduplication (3)—data is deduplicated on the target device (StoreOnce Backup system or StoreOnce Software system).

Deduplication setups



Source-side deduplication

With source-side deduplication (1), a Media Agent is installed together with the Disk Agent on the client that is backed up and thus the client becomes a gateway (a source-side gateway). The deduplication is performed by the Media Agent on the client itself so only deduplicated data is sent to the target device, thereby reducing the overall network traffic. The number of concurrent streams is limited by load balancing settings. Once a Media Agent finishes the backup of local objects, a new Media Agent is started on the next client system. Note that the backed up system must support deduplication.

Server-side deduplication

With server-side deduplication, deduplication is performed on a separate Media Agent client (a gateway) by the Media Agent. This reduces the load on the backed up system and on the target device, but does not reduce the amount of network traffic between the Disk Agent and Media Agent.

Note that the Media Agent client must support deduplication. Server-side deduplication enables you to deduplicate data from clients on which deduplication is not supported locally.

Target-side deduplication

The deduplication process takes place on the target device. It receives data to be backed from Media Agents installed on clients (gateways). Target-side deduplication does not reduce the amount of network traffic between Media Agent and deduplication system.

About File Library Devices

A file library device is a device which resides in a directory on an internal or external hard disk drive defined by you. A file library device consists of a set of directories. When a backup is made to the device files are automatically created in these directories. The files contained in the file library directories are called file depots.

There is no maximum capacity for the file library device that is set by Data Protector. The only limit on the size of the device is determined by the maximum size of the file system where the directory is located. For example, the maximum size of the file library device running on Linux would be the maximum size you can save on the file system.

You specify the capacity of each file depot in the file library device when you first configure the device. It is possible to re-set the sizing properties of the file depots at any time during use of the device using the file library properties.

The file library device can be located on a local or external hard drive, as long as Data Protector knows its path. You specify the path when configuring the file library device.

How to maintain disk-based devices?

If all the disk-based device you are using becomes full, you will need to do one of the following before continuing to make backups with it:

- Start moving data to tape, freeing up the file device or one or more file slots.
- Recycle file depots.
- Add a new file depot to the file device.

File Depots

File depots are the files containing the data from a backup to a file library device.

File depot creation

When the first backup is started using the file library device, file depots are automatically created in the device by Data Protector. Data Protector creates one file depot for each data backup session made using the device. If the amount of data being backed up is larger than the default maximum file depot size Data Protector creates more than a single file depot for a backup session.

File depot name

The name of each file depot is a unique identifier which is automatically generated by the system.

Data Protector also adds a media identifier to the file depot. This identifies the file depot as a media in the media pool. The identifier added to media helps to identify a particular backup session when performing a restore. The identifier can be seen when the file depot properties are viewed.

Note that if the file depot has been recycled, the file depot name may disappear from the GUI although the file depot icon is still visible in the GUI.

File depot size

The size of file depots is defined when you initially create the file library device. During this process you specify all sizing properties for the device, including the maximum size of the file depots. The sizing properties of the file depots, although only entered once, are globally applied to each file depot. If the size of data to be backed up within one session is larger than the originally specified file depot size, Data Protector automatically creates more file depots until the allocated disk space for the file library device has been consumed.

The default file depot size is 5 GB. You can increase this value (up to 2 TB) but some performance degradation is possible.

File depot space consumption

Data Protector automatically creates file depots until there is no more disk space available for the device. The amount of space which must stay free for the file library device is defined in the device properties when the device is initially being set up.

Disk full handling

If the total disk space available to the file library device goes below a user specified level, a notification is issued.

Number of devices per disk

The file library device can include one or several directories. Only one directory can be located on a filesystem.

In situations where the file depots are located on a variety of disks, it is not recommended to put file depots from two different file library devices on a single disk. This is owing to the fact that if the properties are different, it can cause a conflict in Data Protector (for example, if on one file library device the remaining disk space for the file depot is specified as 20 MB and on the other file library device 10 MB).

Setting File Library Device Properties

The file library device properties can be set during initial file library device configuration or they can be changed after the device is in operation.

Initial property setup

Steps

1. During the file library device configuration, select the file library device directory and click **Properties**.
2. Specify the sizing properties of the device. Click **OK**.
3. Click **Next** and continue to configure the file library device.

Changing device properties

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Devices** and then click the name of the file library device you want to alter.
3. Right-click on the file library device name and click **Properties**.
4. Click the **Repository** tab. Select the file library path in the list.
5. Click **Properties**. Specify all the sizing properties of the device, click **OK**.

Data Protector applies the properties specified in the Properties dialog to each file depot created in the file library device after the device properties were changed. The properties of any file depots created before the subsequent device properties changes will not be affected.

Deleting File Library Devices

Before you can delete a file library device it must not contain any protected data. This means that before you can delete the file library you have to change the data protection level on each file depot contained in the device.

Deletion phases

1. [Checking data protection](#)
2. [Recycling file depots](#)
3. [Deleting the exported file depot icon](#)
4. [Deleting the file library device](#)

Checking data protection

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, select the name of the file library device you would like to delete and open the Directories folder in the file library.
3. In the Results Area, locate the Protection column. Check which file depots have a protection level **Permanent**.

Recycling file depots

Disk space can be freed by recycling and deleting file depots or entire file library devices.

You can recycle either individual file depots or all of the file depots in a file library. This means that the disk space occupied by the recycled item can be recovered and used in the next backup. This is done by deleting the unprotected file depot(s) and creating new ones.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand the file library device file depots.
3. In the Results Area, select the file depots you want to recycle by clicking an individual depot.
4. Right-click the selected depot and click **Export**.
Exporting a file depot removes information about the file depot from the IDB. Data Protector no longer recognizes that the file depot exists. The depot information is however still retained, and can later be imported if it is necessary to recover the file depot.
5. Right-click the selected depot and click **Recycle**.
6. Repeat this exercise for every file depot in the file library with a data protection level of 'Full'.

Once a file depot has been marked for recycling the name which is automatically generated for it by Data Protector disappears and only the file depot icon is visible in the Data Protector GUI. It is possible to delete the exported file depot icon.

Deleting the exported file depot icon

Once a file depot has been exported its name disappears, and only the depot icon is visible in the Data Protector Manager.

Steps

1. In the Results Area, select the icon you want to delete.
2. Right-click the selected icon and click **Delete**.
3. Repeat this exercise for every exported file depot icon you would like to remove.

This removes the icon from the GUI but does not physically delete the file from the IDB.

Deleting the file library device

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, select the name of the file library device you would like to delete.
3. Right-click on the file library device and click **Delete**.

The file library device will now be deleted from the IDB.

About Jukebox Devices

Jukebox Physical Devices

A jukebox is a library device. It can contain either optical or file media. If the device is used to contain file media it is known as a file jukebox device. The type of media the device will contain is defined

during initial configuration. If you are running a jukebox optical library on UNIX you need to have a UNIX device file configured for each exchanger slot or side of the platter.

Jukebox File Devices

The file jukebox device is a logical equivalent of a tape stack. It contains slots whose size is defined by the user during the initial device configuration. This device is configured manually. The file jukebox properties can be altered while it is being used. If used to contain file media the device writes to disk instead of tape. The file jukebox device saves data in the form of files; each of these files is the equivalent of a slot in a tape device.

The recommended maximum data storage capacity of this device is limited only by the amount of data that can be stored in a filesystem by the operating system on which the file jukebox is running. Each slot in the file jukebox device has a maximum capacity of 2 TB. However, it is normally recommended that you keep the slot size between 100 MB and 50 GB (on Windows systems) or 100 MB to 2 TB (on UNIX systems). If, for instance, you have 1 TB of data to back up, the following device configuration is possible:

Windows systems: 1 File jukebox device with 100 file slots of 10 GB each

UNIX systems: 1 File jukebox device with 250 file slots of 4 GB each

To improve the jukebox file device performance, it is recommended to have only one device per disk and only one drive per device. You should also avoid other applications transferring large amounts of data from/to the disk when a Data Protector backup/restore is in progress.

Recommended slot sizes for Windows and UNIX

Available disk space	Number of slots	Slot size
1 TB	100	10
5 TB	250	20
10 TB	250	40

How to maintain file jukebox devices?

If all the file jukebox device you are using becomes full, you will need to do one of the following before continuing to make backups with it:

- Start moving data to tape, freeing up the file device or one or more file slots.
- Recycle or jukebox slots.
- Add a new jukebox slot to the file device.

Configuring a File Jukebox Device

It is recommended that the device you are creating is located on a disk other than the one on which the IDB is located. This ensures that there will be an adequate amount of disk space available for the database. Putting the device and the IDB on separate disks also improves performance.

Configuring a file jukebox device

Consider the following:

- Do not use the name of an existing device for configuring these devices, because the existing device will be overwritten.
- Do not use the same device name for configuring several devices, because every time the device is accessed it will be overwritten.

Prerequisites

- On Windows systems, for a file that you want to use as a device, disable the Windows compression option.
- The directory in which the device will reside must have been created on disk before you create the device.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and then click **Add Device** to open the wizard.
3. In the Device Name text box, enter the name of the device.
4. In the Description text box, enter a description (optional).
5. In the Device Type list, select the **Jukebox** device type.
6. In the Client list, select the name of the client.
7. In the Management Console URL text box, enter a valid URL address of the library management console (optional).
8. Click **Next**.
9. Specify a set of files/disks for the jukebox. Use a dash to enter multiple files or disks at a time, for example, /tmp/FILE 1-3, and then click **Add**. For magneto-optical jukeboxes, the disk names have to end on A/a or B/b. Click **Next**.
10. In the Media Type list, select **File** for the device that you are configuring.
11. Click **Finish** to exit this wizard. You are prompted to configure a library drive. Click **Yes** and the drive configuration wizard displays.

Configuring a drive in the file jukebox device

Steps

1. In the Device Name text box, enter the name of the device.
2. In the Description text box, optionally enter a description.
3. Specify a media pool for the selected media type. You can either select an existing pool from the Media Pool list or enter a new pool name. In this case, the pool will be automatically created. You can configure one media pool for all drives or have an independent media pool for each drive. Click **Next**.

4. Optionally, select **Device may be used for restore** and/or **Device may be used as source device for object copy** and specify a **Device Tag**.
5. Click **Finish** to exit the wizard.

The name of the drive is displayed in the list of configured drives. You can scan the drives to verify the configuration.

Recycling a File Jukebox Slot

Data protection is set for each individual file slot in a file jukebox, so it is possible to recycle a single slot by setting its **Protection** to **None**. Therefore, having multiple small slots can increase flexibility and enable more efficient data protection and space retention management. Recycling a slot in a file jukebox device removes its data protection, so that the slot can be reused for backup. The data in the slot will be overwritten in a subsequent backup session.

If this method is used, the existing data on the media will be overwritten and lost.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand the file jukebox device slots.
3. In the Results Area, select the slots you want to recycle.
4. Right-click the selected slot and click **Recycle**.

About Standalone Devices

Standalone Physical Devices

A standalone device is a simple device with one drive that reads from or writes to one medium at a time, such as DDS or DLT. These devices are used for small-scale backups. As soon as the medium is full, an operator must manually replace it with a new medium for the backup to proceed. Standalone devices are, therefore, not appropriate for large, unattended backups.

Standalone File Devices

A standalone file device is a file in a specified directory to which you back up data instead of writing to a tape. This device saves data in the form of files; each of these files is the equivalent of a slot in a tape device. The standalone file device is useful for smaller backups.

The maximum capacity of a file device is 2 TB. However, it is normally recommended that you keep the standalone file device size between 100 MB and 50 GB on Windows systems, or 100 MB to 2 TB on UNIX systems. Data Protector never measures the amount of free space on the filesystem; it takes either the default or the specified capacity for a file size limit. You cannot use compressed files for file devices. You can change the default file size by setting the `FileMediumCapacity` global option.

The default maximum size of a Standalone File device is 100 MB. If you wish to back up more than this, change the default file size by setting the `FileMediumCapacity` global option. For more information on setting global options, see [Customizing the Data Protector Global Options](#) or [Customizing Global Options by Editing the Global File](#).

For example, for a 20GB maximum (20Gb = 20000 MB), set the FileMediumCapacity global option:

```
# FileMediumCapacity=MaxSizeInMBytes  
  
FileMediumCapacity=20000
```

You specify the capacity of a file device when you first format the medium. When you reformat the medium, you can specify a new size; however, the originally specified size will be used. You can change the capacity of a file device only by deleting the file from the system.

The size specified should be at least 1 MB less than the maximum free space in the filesystem. When a file device reaches its size limit, Data Protector issues a mount request.

To improve the standalone file device performance, it is recommended to have only one device per disk and only one drive per device. You should also avoid other applications transferring large amounts of data from/to the disk when a Data Protector backup/restore is in progress.

The file can be located on a local or external hard drive, as long as Data Protector knows its path. You specify the path when configuring the file device.

Configuring a Standalone File Device

It is recommended that the device you are creating is located on a disk other than the one on which the IDB is located. This ensures that there will be an adequate amount of disk space available for the database. Putting the device and the IDB on separate disks also improves performance.

Consider the following:

- Do not use the name of an existing device for configuring these devices, because the existing device will be overwritten.
- Do not use the same device name for configuring several devices, because every time the device is accessed it will be overwritten.

Prerequisites

- On Windows systems, for a file that you want to use as a device disable the Windows compression option.
- The directory in which the device will reside must have been created on disk before you create the device.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Devices** and then click **Add Device** to open the wizard.
3. In the Device Name text box, enter a name for the device.
4. In the Description text box, enter a description (optional).
5. In the Client list, select the name of the client.
6. In the Device Type list, select the **Standalone** device type and then click **Next**.
7. In the text box, enter the pathname to the file device and a filename, for example, c:\My_Backup\file_device.bin.

8. Click **Add** and then click **Next**.
9. In the Media Type list, select a **File** media type.
10. Specify a media pool for the selected media type. You can either select an existing pool from the Media Pool drop-down list or enter a new pool name. In this case, the pool will be created automatically.
11. Click **Finish** to exit the wizard.

The name of the device is displayed in the list of configured devices. You can scan the device to verify the configuration.

At this point, the device has been specified to Data Protector, but it does not yet actually exist on disk. Before it can be used for backup, you have to format it.

Chapter 8: Media

About Media Management

Data Protector provides a powerful media management functionality that enables simple and efficient management of a large number of media. The system uses the IDB to store information about backup, restore, and media management events.

Advanced features of Data Protector media management are:

- Protection against accidental overwrites.
- Media pools that enable you to think about large sets of media without having to worry about each individual medium.
- The capability to transfer all media-related catalog data from one Data Protector Cell Manager to another one without physically accessing the media.
- The free pool functionality that enables you to avoid failed backups due to missing (free) media.
- Tracking of all media, the status of each medium, and the sharing this information among several Data Protector cells: data protection expiration time, availability of media for backups, and a catalog of what has been backed up to each medium.
- The ability to explicitly define which media and which devices you want to use for a certain backup.
- Automatic recognition of Data Protector media and other popular tape formats.
- Recognition and support of barcodes on large library and silo devices with a barcode support.
- Centralized media information that can be shared among several Data Protector cells.
- Support for media vaulting, also known as archiving or off-site storage.
- Interactive or automated creation of additional copies of the data on the media.
- Detailed filtering and paging settings.

Customizing the Devices and Media View

You can customize the default view of the Devices and Media context by configuring the `MediaView`, `MagazineView`, `SCSIView`, `ExternalView`, `JukeboxView`, `ACSView`, and `DASView` global options. Customize the attributes that will be displayed in the library or media management context by specifying the corresponding token strings. For more information, see [Customizing the Data Protector Global Options](#).

About Media Pools

A media pool represents a set of media of the same type that you use for backups. You may have one media pool for a regular backup, one for an archive backup, one for each department, and so on. Each media pool defines media usage and allocation policy, and the media condition factors.

Free pools

A free pool is an auxiliary source of media of the same type that can be used if all the media in a pool are in use. Having a free pool helps you to avoid failed backups due to missing free media.

Protected media belong to a specific pool, for example to SAP pool, while free media can be moved automatically to a free pool that is used by several other pools. This common free pool is used for allocation of free media for all pools that use this free pool. You can decide for each media pool whether you want to link it with a free pool or not.

Default media pool

A default media pool is a pool provided by Data Protector as part of the device definition. This pool is used if no media pool is specified in the backup specification.

Free Pool Characteristics

A free pool is a media pool that you can configure to allow the sharing of free media across media pools, which may reduce operator intervention due to mount requests. The usage of a free pool is optional.

A free pool has some characteristics you should consider before using it.

Free Pool Properties

A free pool:

- cannot be deleted if it is linked with a media pool or if it is not empty
- is different from a regular pool as it cannot be used for allocation because it cannot hold protected media. Consequently, allocation policy options (Strict / Loose, Appendable/Non-Appendable) are not available.
- contains only free Data Protector media (no unknown or blank media).

When Is a Free Pool Used?

Media are moved between regular pool and free pool on two occasions:

- If there is no free media in the regular pool anymore, then Data Protector allocates media from the free pool. This automatically moves the media to the regular pool.
- When all the data on the media expires (and the media is in a regular pool), media can be moved to the free pool automatically.

Media Quality Calculation

Media quality is calculated equally between "linked" pools. Media condition factors are configurable for a free pool only and are inherited by all pools using the free pool. Pools that do not use the free pool have their own separate calculation base.

Free Pool Limitations

- You cannot move protected media to a free pool.
- You cannot use some operations on media, such as Import, Copy, and Recycle, because they may operate on protected media.
- Pools with the Magazine support option selected cannot use a free pool.
- You may experience some temporary inconsistencies (1 day) in pools when using free pools (when there is an unprotected medium in a regular pool waiting for de-allocation to the free pool, for example).
- If a free pool contains media with different data format types, Data Protector automatically reformats allocated media if necessary. For example, NDMP media may be reformatted to normal media.

Media Pool Properties

You specify media pool attributes when configuring a media pool. Some of the properties can subsequently be modified.

For detailed information on media pool properties, see the HPE Data Protector Help.

Media pool properties - General

- Description
- Pool name
- Media type

Media pool properties - Allocation

Allocation

The media allocation policy defines the order in which media are accessed within a media pool so that media wear out evenly. It can be:

- Strict
- Loose
- Allocate unformatted media first
- Use free pool
- Move free media to free pool
- Magazine support

Media pool properties - Condition

Media condition factors

Media condition factors define the state of media, thus determining how long media can be reliably used for backup. For example, a backup to old or worn media is more likely to have read/write errors. Based on these factors, Data Protector changes the condition of media from good to fair or poor. The condition factors are set for the entire media pool, not for each medium.

For Data Protector to accurately calculate the condition of the media, use new media when adding media to the media pool.

Note: If a pool uses the free pool option, the media condition factors are inherited from the free pool.

The two media condition factors you can select are:

- Maximum overwrites
- Valid for (Months)

Media pool properties - Usage

The media usage policy controls how new backups are added to the already used media. It can be:

- Appendable
- Non-Appendable
- Appendable on incrementals only

Media Pool Quality

The media with the lowest quality in a pool determines the quality of the media pool. For example, as soon as one medium in a pool is poor, the whole media pool is marked as poor.

The quality of media influences how media are selected for a backup, as it affects the ability to write to the medium and read the data contained on it. Media in good condition are selected before media in fair condition. Media in poor condition are not selected for a backup.

Media status is based on one of the following media condition factors:

- Good
- Fair
- Poor

You can change the media condition factors that are used to calculate the condition of a medium in the Condition property page of the media pool properties. The new media condition factors are used to calculate the condition of all media in the media pool.

Device error and media quality

If a device fails during backup, the media used for backup in this device are marked as poor. This prevents future errors if the problem was caused by the bad media.

If this error was due to a dirty drive, clean the drive and verify the medium to reset its condition.

It is recommended that you investigate if media marked poor appear in a pool. You can use Verify to get more information on each medium's condition. It is not recommended to simply recycle the medium.

Creating a Media Pool

Data Protector provides default media pools, but you can create your own media pool to suit your needs.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media**, right-click **Pools**, and click **Add Media Pool** to open the wizard.
3. In the Pool Name text box, enter the name of the media pool, enter a description in the Description text box (optional) and in the Media Type drop-down list, select the type of media that you will use with your backup device. Click **Next**.
4. Set the following options:
 - Change the defaults for Media Usage Policy and Media Allocation Policy (optional).
 - To use a free pool, first select the **Use free pool** option and then select the free pool from the drop-down list.
 - To disable the automatic de-allocation of free media to a free pool, select the **Move free media to free pool** option.
 - Select the **Magazine Support** option if you are configuring a media pool for a device with magazine support. This option cannot be used together with free pools.

Click **Next**.

5. Change the settings in the Media Condition Factors dialog (optional).
6. Click **Finish** to create your media pool and exit the wizard.

Tip: You can modify an already configured media pool. However, its media type cannot be modified.

Modifying a Media Pool

You can modify media pool properties to better suit your needs: you can change the name of a media pool, its description, the media usage and allocation policy, or the media condition factors. You cannot

change the media type.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Pools**. A list of configured media pools is displayed in the Results Area.
3. Right-click the name of the media pool that you want to modify and click **Properties**.
4. In the General property page, you can change the name of the media pool in the Pool Name text box or change the description in the Description text box.
5. Click the **Allocation** tab to change the settings for Media Usage Policy and Media Allocation Policy, to (de)select the usage of a free pool, to enable or disable the **Move free media to free pool** option, or to select the **Magazine Support** option.
6. Click the **Condition** tab to change settings in the Media Condition Factors dialog or to set the media condition factors to default.
7. Click **Apply** to confirm.

Deleting a Media Pool

By deleting a media pool from the Data Protector configuration, you stop using this media pool for backups. You cannot delete a media pool that is used as a default pool for backup devices. In this case, change the media pool for all devices or delete the devices.

If you try to delete a media pool that is not empty, you will be prompted to export or move all media in the pool first.

If you delete a media pool that is used in a backup specification, the media pool is removed from the specification.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Pools**. A list of configured media pools appears in the Results Area.
3. Right-click the media pool you want to delete and then click **Delete**. Confirm the action.

The media pool is no longer displayed in the list of the configured media pools.

Media Life Cycle

The media life cycle begins with the usage of the media and ends when the maximum usage criteria is reached. It typically consists of the following:

Preparing media for backups

This includes formatting/initializing media and assigning them to a media pool by formatting (unused media and used non-Data Protector media) or importing (used Data Protector media). When dealing with already-used media, consider using the recycle/unprotect and export functionality.

Using media for backups

This includes how media are selected for a backup, which media condition factors are checked (for example, the number of overwrites), how new backups are appended to the media, and when the data protection expires.

Vaulting media to a safe place

Vaulting media includes preparing media for safe storage and the actual storage. To prepare for vaulting, you need to set up the appropriate data protection and catalog protection policies, to create a list of vault locations, to specify and modify media locations, to eject media, and in some cases, scan devices.

Data Protector supports vaulting on various levels:

- Data protection and catalog protection policies.
- Easy selecting and ejecting of media from the library.
- Media location tells you the physical location where the media are stored.
- A report shows media used for backup within a specified time-frame.
- A report shows which backup specification have used specified media during the backup.
- A report shows media stored at a specific location with data protection expiring at a specific time.
- Displays a list of media needed for a restore and the physical locations where the media are stored.
- Filtering of media from the media view based on specific criteria, such as time written to the media or media with expired protection.

It is recommended to make a copy of the backed up data for vaulting purposes, and keep the original on site to enable a restore. Data Protector enables interactive or automated creation of additional copies of the data on the media.

Retiring media

Once a medium has expired (its maximum usage criteria exceeded), it is marked as Poor and is no longer used by Data Protector.

Media Types

A media type is the physical kind of media, such as DDS or DLT. With Data Protector you select the appropriate media type when configuring devices and Data Protector estimates the available space on the media for the particular media pool.

Supported media types

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

Media Quality

The quality of media influences how media are selected for a backup, as it affects the ability to write to the medium and read the data contained on it. Media in good condition are selected before media in fair condition. Media in poor condition are not selected for a backup.

Media status is based on one of the following media condition factors:

- Good
- Fair
- Poor

You can view the Info property page of a medium for information about the medium quality (condition).

You can change the media condition factors that are used to calculate the condition of a medium in the Options property page of the media pool properties. The new media condition factors are used to calculate the condition of all media in the media pool.

Media quality helps you determine when the medium has to be replaced.

Device error and media quality

If a device fails during backup, the media used for backup in this device are marked as poor. This prevents future errors if the problem was caused by the bad media.

If this error was due to a dirty drive, clean the drive and verify the medium to reset its condition.

It is recommended that you investigate if media are marked poor. You can use Verify to get more information on each medium's condition. It is not recommended to simply recycle the medium.

How Media Are Selected for Backup

Data Protector media management automatically selects the most appropriate media for backup. Basic media selection criteria:

- Media in poor condition are not selected for backup.
- Media in fair condition are used only if no media in good condition are available.
- If available, media in good condition are used first.
- Media are always selected from the specified pool. In case the pool does not contain unprotected media, Data Protector accesses a free pool (if configured).

In addition, the media selection is based on the following factors:

Media allocation policy

You can influence how media are selected for backup using the media allocation policy. You can specify a loose policy, where any suitable media can be used for backup, or a strict policy, where specific media have to be available in a predefined order.

Preallocating media

You can specify the order in which media from a media pool will be used for backup. This order is called a pre-allocation list.

Media condition

Media condition also influences which media are selected for backup. For example, media in good condition are used for backup before media in fair condition. Media in poor condition are not used for backup.

Media that are marked as fair are only used if there are no protected objects on those media. Otherwise, a mount request is issued for free media.

Media usage

The media usage policy controls how new backups are added to the already used media, and influences which media are selected for backup.

Limitation

Backups cannot be appended on media used in Travan devices.

Appendable media must be in good condition, contain some currently protected objects and must not be full. If several devices are used with load balancing, the appendable concept applies on a per-device basis, that is, each device uses an appendable medium as first media in a session. The backup sessions appending data on the same medium do not have to relate to the same backup specification.

Note: If you use the append functionality and the backup requires more than one medium, only the first medium used can contain backed up data from a previous session. Subsequently, Data Protector will use empty or unprotected media only.

The policy can be: Appendable, Non-Appendable, or Appendable on incrementals only.

You can create restore chains for one client on media. These media will contain only one full backup and the incremental backups related to the same client:

- Configure one pool per client with the **Appendable on incrementals only** media usage policy.
- Link a different pool to each client in the backup specification, or create a separate backup specification per client.

Be aware that occasionally media will be created which contain incremental backups only.

Media selection factors

Allocation policy	Allocate unformatted media first	Data Protector Selection order
Loose	OFF	<ol style="list-style-type: none">1. Preallocation list (if specified)2. Appendable (as set in usage policy)3. Unprotected Data Protector Media4. Unformatted Media5. Fair Media
Loose	ON	<ol style="list-style-type: none">1. Preallocation list (if specified)2. Appendable (as set in usage policy)3. Unformatted Media4. Unprotected Data Protector Media5. Fair Media
Strict	Not applicable	<ol style="list-style-type: none">1. Preallocation list (if specified)2. Appendable (as set in usage policy)3. Unprotected Data Protector Media4. Fair Media

Use of Different Media Format Types

Data Protector recognizes and uses two different format types to write data to media:

- Data Protector (for backup devices that are under direct Data Protector control)
- NDMP (for backup devices that are connected to NDMP servers)

The two format types use two different Data Protector Media Agent components (the General Media Agent or the NDMP Media Agent) to communicate with backup devices.

Limitations

- Media that are written by one format type will be recognized as blank or as foreign in a backup device that uses a different format type.
- You cannot back up objects using different format types on the same medium.
- You cannot have two different Data Protector Media Agent components installed on the same system.
- It is strongly recommended that you use different media pools for different media format types.

WORM Media

WORM (write once, read many) is a data storage technology that allows information to be written to a medium a single time and prevents the drive from erasing the data. WORM media are by design not rewritable because they are intended to store data that you do not want to erase accidentally.

How to use WORM media with Data Protector

Detection of the WORM tapes is supported on Windows platforms only. On other platforms Data Protector does not recognize the tape as not rewritable and treats it as any other tape. When attempting to overwrite data on a WORM medium, the following error messages are displayed:

Cannot write to device ([19] The media is write protected.)

Tape Alert [9]: You are trying to write to a write-protected cartridge.

To prevent this, do the following:

- Set the backup protection for WORM media to Permanent.
- Keep WORM media and rewritable media in separate media pools.

Supported WORM media

All Data Protector media operations are supported with supported WORM media. For an up-to-date list of supported WORM tape drives and media, see the latest support matrices at <https://softwaresupport.hpe.com/>.

About Formatting Media

Formatting (initializing) media prepares it for use with Data Protector by saving the information about the media (media IDs, description, and location) in the IDB and by writing this information on the medium (media header), also. When you format media, you also specify to which media pool it belongs.

Formatting with padding blocks

You can extend the size of the media header and fill it up with incompressible data, padding blocks. This becomes useful when creating media copies. The padding blocks are not copied to the target medium. This way you make sure that the target medium does not reach the end of the tape before the source medium.

Tape padding is not required if you copy backed up data using the object copy functionality.

Tape padding is disabled by default. To enable it, set the OB2BLKPadding_n option in the omnirc file on the system with the backup device connected.

When to format media

You need to format media before you use them for backup. However when using the Loose media allocation policy for the media pool, formatting media as a separate step is not required. If global option

InitOnLoosePolicy is set to 1 (default is 0), Data Protector automatically formats new media when they are selected for backup.

Non-Data Protector media must be formatted before backup.

Data Protector media with protected data are not formatted until you remove the protection, after which the old data can be overwritten.

Media label

Upon formatting, Data Protector labels each medium with a unique media label and medium ID. Both are stored in the IDB and allow Data Protector to manage the medium. The media label is a combination of the user-defined description and the barcode of the medium (if the **Use barcode as medium label on initialization** option is selected for your library). The barcode is displayed as a prefix of the medium description. For example, [CW8279]Default DLT_1 is a media label with the Default DLT_1 description and the CW8279 barcode. You can optionally write the barcode as medium label to the medium header on the tape during the initialization of the medium.

Once you have formatted a medium, you cannot change the medium label and location written on the medium itself unless you format it again (which results in overwriting the data). Modifying medium properties only changes this information in the IDB.

Although you can change the label and exclude the barcode number, this is not recommended. In this case you should manually keep track of the actual barcode and the medium label you assigned to the medium.

Recognized Media Formats

Data Protector recognizes common formats of data on a medium, if the medium was already used by some other application. However, it is not recommended that you rely on Data Protector to recognize other media types, as recognition depends on the platforms you use.

To be sure that no Data Protector media are being overwritten, you must select the strict allocation policy.

Data Protector behaves differently according to the recognized format, as shown in the table below.

Data Protector media format categories

Media format	Backup behavior	Possible operations
Unknown or new (blank)	Loose policy: used for backup Strict policy: not used for backup	Format media
Media written with compression, now used without compression	Loose policy: used for backup Strict policy: not used for backup	Format media
Media written without compression, now used with compression	Loose policy: used for backup Strict policy: not used for backup	Format media
Foreign Data Protector (from another	Not used for backup	Import or force format

cell)		media
tar, cpio, OmniBack I, ANSI label	Not used for backup (cannot be guaranteed)	Force format media
Data Protector unprotected media	Used for backup	Export media
Data Protector protected media	Append backups	Recycle (unprotect) media

Note: If you try to read from a medium that was written using hardware compression with a device that does not support hardware compression, Data Protector cannot recognize the medium and read the data. Therefore, the medium will be treated as unknown or new.

Formatting a Medium

You have to format media before you use them for backup. Data Protector media with protected data are not formatted until you remove the protection, after which the old data is overwritten.

Note: You cannot format a file library device until after the first backup has been made to it. This is because before this point the device does not contain any file depots, and you cannot create them manually. File depots created during backup are the equivalent of a medium. According to the file library device's media pool media allocation policy, newly formatted media are automatically deleted.

Use the **Force operation** option to format media in other formats recognized by Data Protector (tar, OmniBack I, and so forth), or to re-format Data Protector media.

Data Protector media with protected data will not be formatted until the protection is removed.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Pools**.
3. In the Results Area, right-click the media pool to which you want to add a medium, and then click **Format** to open the wizard.
4. Select the device where the target medium is located and then click **Next**.
5. Specify the **Medium description** and location for the new medium (optional) and then click **Next**.
6. Specify additional options for the session: you can select the **Eject medium after operation** option or use the **Force operation** option. You can also **Specify medium size** or leave the **Default** option selected.
7. Click **Finish** to start the formatting and exit the wizard.

When the formatting is complete, the media format is set to Data Protector.

Formatting All Media in a Magazine

You have to format media before you use them for backup. Data Protector media with protected data are not formatted until you remove the protection, after which the old data is overwritten.

Prerequisite

To format all media in a magazine in a single step, use a device with the **magazine support** option selected.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Pools**.
3. In the Results Area double-click the desired media pool.
4. Right-click the **Magazines** item and then click **Format Magazine** to open the wizard.
5. Select the library's drive to perform the operation with and then click **Next**.
6. Specify the description and location for the new media (optional) and then click **Next**.
7. Specify additional options for the session: you can use the **Force operation** option and select the **Specify medium size** option or leave the **Default** option selected.
8. Click **Finish** to start the formatting and exit the wizard.

When the formatting is complete, the media format is set to Data Protector.

Formatting a Single Medium in a Magazine

You have to format media before you use them for backup. Data Protector media with protected data are not formatted until you remove the protection, after which the old data is overwritten.

Prerequisite

To format a medium in a magazine, use a device with the **magazine support** option selected.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Pools**.
3. In the Results Area, right-click the media pool to which you want to add a medium, and then click **Format** to open the wizard.
4. Select the device where the target medium is located and the slot with the medium to perform operation on and then click **Next**.
5. Specify the description and location for the new medium (optional) and then click **Next**.
6. Specify additional options for the session: you can use the **Force operation** option and select the **Specify medium size** option or leave the **Default** option selected.
7. Click **Finish** to start the formatting and exit the wizard.

When the formatting is complete, the media format is set to Data Protector.

Formatting Media in a Library Device

You have to format media before you use them for backup. Data Protector media with protected data are not formatted until you remove the protection, after which the old data is overwritten.

If you use a library device, you can select multiple slots using the Ctrl key and format several media in a single step.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand Devices, expand the library device, and then click **Slots**.
3. In the Results Area, right-click the slots that have the media you want to format, and then click **Format** to open the wizard.
4. Select the library's drive to perform the operation with and then click **Next**.
5. Select the media pool to which you want to add the formatted media and then click **Next**.
6. Specify the **Medium description** and location for the new media (optional) and then click **Next**.
7. Specify additional options for the session: you can use the **Force operation** option and select the **Specify medium size** option or leave the **Default** option selected.
8. Click **Finish** to start the formatting and exit the wizard.

When the formatting is complete, the media format is set to Data Protector.

About Importing Media

Media import is the act of adding Data Protector media that is foreign to the cell to a media pool without losing the data on the media. The media must have been exported previously — that is it comes from another Data Protector cell.

When you Import media, the information about backed up data on the media is read into the IDB so that you can later browse it for a restore.

Considerations

- During media import, attribute information such as object or media size is not reconstructed, so the size of the imported objects is shown as 0 kB.
- Depending on the backup device and media you use, importing can take a considerable amount of time.
- Media cannot be imported into free pools.
- If you try to import a removed copy and the original media is not in the IDB, you either need to import the original media first using the Force operation option, or to import the copy using the Import copy as original option.
- When importing WORM media on which data protection has already expired to a Data Protector cell, make sure to specify a new data protection value by using the option **Protection** (by default, the value is set to Permanent). This allows Data Protector to append to the WORM media.

When to import media?

You typically use the import functionality when you move media between Data Protector cells. In this case, information about space on the medium is not updated.

You should import all media used in one backup session at once. If you only add some media from the backup session, you are not able to restore data spanning to other media.

With file library devices, it is only possible to import file depots which previously belonged to the file library device and which have previously been exported. If you want to import media from a file library residing on a host other than the target host, it can only be done to a jukebox device.

Importing a Medium

Import media when you want to add media already used by Data Protector to a media pool so you can later browse the data for restore.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**.
3. In the Results Area, right-click the device to which you want to import a medium, and then click **Import** to open the wizard.
4. Select the media pool to add imported media to and then click **Next**.
5. Select the **Import copy as original option** and decide on the **Logging** option that suits your needs (optional).
6. Click **Finish** to start importing and exit the wizard.

The Session Information message displays the status of your import operation. When the import is complete, the media type is set to Data Protector.

Importing All Media in a Magazine

Import media when you want to add media already used by Data Protector to a media pool so you can later browse the data for restore.

Prerequisite

To import all media in the magazine in one step, use a device with the **magazine support** option selected.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Pools**.
3. In the Results Area, double-click the media pool that has the media in the magazine. The Media and Magazine items appear.

4. Right-click the **Magazines** item and then click **Import Magazine** to open the wizard.
5. Select the library's drive to perform the operation with and then click **Next**.
6. Specify the description for the new media (optional) or leave the **Automatically generate** option set and then click **Next**.
7. Select the **Import copy as original** option and decide on the **Logging** option that suits your needs (optional).
8. Click **Finish** to start the import and exit the wizard.

The Session Information message displays the status of your import operation. When the import is complete, the media type is set to Data Protector.

Importing a Single Medium in a Magazine

Import a medium that has been used by Data Protector when you want to add the medium to a media pool so you can later browse the data for restore.

Prerequisite

To import a medium in the magazine, use a device with the **magazine support** option selected.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Pools**.
3. In the Results Area, double-click the media pool that has the medium in the magazine. The Media and Magazine items appear.
4. Right-click the **Media** item and then click **Import** to open the wizard.
5. Select the library's drive and slot where the target medium is located and then click **Next**.
6. Select the **Import copy as original** option and decide on the **Logging** option that suits your needs (optional).
7. Click **Finish** to start importing and exit the wizard.

The Session Information message displays the status of your import operation. When the import is complete, the media type is set to Data Protector.

Importing Media in a Library Device

Import media when you want to add media already used by Data Protector to a media pool so you can later browse the data for restore.

If you use a library device, you can select multiple slots using the Ctrl key and format several media in a single step.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Devices**, expand the library device, and then click **Slots**.

3. In the Results Area, select the slots that have the media that you want to import.
4. Right-click the selected slots and then click **Import** to open the wizard.
5. Select the library's drive where the exchanger will load media to import and then click **Next**.
6. Select the media pool to which you want to add the imported media and then click **Next**.
7. Select the **Import Copy as Original** option and decide on the **Logging** option that suits your needs (optional).
8. Click **Finish** to start the import and exit the wizard.

The Session Information message displays the status of your import operation. When the import is complete, the media type is set to Data Protector.

Exporting and Importing Media with Encrypted Backups

To restore data from encrypted backup to a client in a different Data Protector cell, you need to import the media and the encryption keys to the destination Cell Manager, as described in the following sections.

Note: Data Protector also provides advanced manual management of encryption keys (such as expiring, reactivating, exporting, importing, and deleting keys) via the command-line interface (CLI). For details, see the `omnikeytool` man page or the *HPE Data Protector Command Line Interface Reference*.

Cell Manager environment or MoM environment without CMMDB

In a Cell Manager environment or in a MoM environment where local MMDBs are used, perform the following steps to export and import a medium with encrypted backup:

Steps

1. On the original Cell Manager, export the medium from the IDB. This operation also exports the relevant encryption keys from the keystore into the file `mediumID.csv`, in the default exported encryption keys directory.
2. Transfer the `mediumID.csv` file to the destination Cell Manager and place it into the directory default imported encryption keys directory.
3. Insert the exported medium into the drive that will be used by the destination Cell Manager.
4. On the destination Cell Manager, import the medium. This operation also imports the keys from the `mediumID.csv` file.

Note: If the key file is not present, you can still import the medium, but the catalog import will abort because of missing decryption keys.

MoM environment with CMMDB

In a MoM environment where the CMMDB is used, all media information is stored on the MoM Manager, but encryption keys IDs used by these media as well as the CDB are stored in a local keystore on each respective Cell Manager. Note that all media management operations need to be done on the MoM Cell Manager.

To export and import a medium with encrypted backup if the CMMDB resides on the MoM Manager, perform the following steps:

Steps

1. Export the medium from the CMMDB. The key IDs are exported into the file *mediumID.csv*, in the default exported encryption keys directory.
2. Transfer the *mediumID.csv* file to the destination Cell Manager and place it into the default imported encryption keys directory.
3. From the MoM Manager, eject a medium from a library.
4. Move a medium from the original media pool to the destination media pool, which is associated with a drive in the destination cell. This operation also imports the catalogue.
5. Insert the exported medium into the drive that will be used by the destination Cell Manager.
6. On the destination Cell Manager, import the medium. This operation also imports the keys from the *mediumID.csv* file.

About Media Copying

The Data Protector media copy functionality enables you to copy media after a backup has been performed. Media copying is a process that creates an exact copy of a medium containing a backup. You can move either the copies or the original media to a safe place for archiving or vaulting purposes, and keep the other set of media on site for restore purposes.

Prerequisites

You need two devices, one for a source medium and one for a target medium. You can also copy media in library devices with multiple drives. In this case, use one drive for the source medium and another for the target medium.

- The source medium and the target medium must be of the same media type.
- If your target media are Data Protector media with data protection, you must first recycle the media and then format them.

Limitations

- You can make multiple copies (target media) of a medium (source medium), but you cannot make copies of media copies.
- You can copy only Data Protector resident media (media in devices).
- As media copying is designed to make exact copies of media that are usually moved to a different location, it is not supported with file libraries. To make copies of data in a file library, use the object copy functionality.
- The media copy operation is not available for media in free pools.
- Device concurrency for NAS devices controlled by an NDMP Server is limited to 1.
- Media copying is not supported for NDMP-Celerra backup sessions.

When to copy media

You can copy a medium as soon as the backup session finishes. However, you need to consider the availability of devices that will be used for copying the media. It is recommended to wait for all backups using specific devices to finish before using the devices for media copying.

Results of Copying Media

The result of copying media is two identical sets of media, the original media set and the copy. Either of them can be used for restore.

After the source medium has been copied, Data Protector marks it as non-appendable to prevent appending new backups. (This would result in the original being different from its copy.) The copy is also marked as non-appendable.

Restoring from a copy

By default, Data Protector restores data from the original media set. However, if the original media set is not available, but a copy is available, the copy is used for the restore.

If neither the original nor a copy is available in the device during restore, Data Protector issues a mount request, displaying both the original and the copy as the media required for restore. You can use any one of these.

If you perform a restore using a standalone device, you can choose to restore from the copy rather than from the original. To do this, insert the copy in the device that will be used for the restore, or select the device containing the copy. However, if you perform a restore using a library device and the original is in the library, Data Protector will use it for the restore.

Note: When copying media, it is possible that the target medium reaches the end of the tape before the source medium. This happens if the source medium was written in streaming mode and you make a copy on a busy system or through a loaded network, which can create blank space where the tape has stopped and started again. You can prevent this by enabling tape padding when you format media.

Copying a Medium

You can copy media for archiving or vaulting purposes. You need to start the copying of each medium separately, as only one medium can be copied in a media copy session.

Copying a medium in a standalone device

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Devices**, right-click the device with the medium you want to copy and click **Copy**.
3. Select the device (library's drive and slot) where the target medium is located and then click **Next**.

4. Select the media pool to which you want to add the medium copy and then click **Next**.
5. Specify the description and location for the medium copy (optional), and then click **Next**.
6. Specify additional options for the session: you can select the **Force operation** option, specify the medium size and medium protection.

Tip: Use the **Force operation** option if the target media have other formats recognized by Data Protector (tar, OmniBack I, and so on) or if they are Data Protector media without protection.

7. Click **Finish** to start copying and exit the wizard.

The Session Information message displays the status of the media copy operation.

Copying a medium in a library device

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, under **Media**, expand **Pools**, and then expand the media pool that has the medium you want to copy. Right-click the medium and click **Copy** to open the wizard.
3. Select a drive for the medium you want to copy and click **Next**. This step is skipped if the library has only one drive.
4. Select the device (library's drive and slot) where the target medium is located and then click **Next**.
5. Select the media pool to which you want to add the medium copy and then click **Next**.
6. Specify the description and location for the medium copy (optional), and then click **Next**.
7. Specify additional options for the session: you can select the **Force operation** option, specify the medium size and medium protection.

Tip: Use the **Force operation** option if the target media have other formats recognized by Data Protector (tar, OmniBack I, and so on) or if they are Data Protector media without protection.

8. Click **Finish** to start copying and exit the wizard.

The Session Information message displays the status of the media copy operation.

Automated Media Copying

Automated media copying is an automated process that creates copies of the media containing backups. Compared to manually started media copying, note the additional limitation:

Limitations

- You cannot use standalone devices for automated media copying; only library devices can be used. You cannot use Backup to Disk (B2D) devices for automated media copying.
- Automated media copying is not supported for NDMP-Celerra backup sessions.

Automated media copying

First you create an automated media copy specification. When the automated media copy session begins, Data Protector generates a list of media, referred to as source media, based on the parameters specified in the automated media copy specification. For each source medium, a target medium is selected to which the data will be copied. The target media are selected from the same media pool as the source media, from a free pool, or from the blank media in a library.

For each source medium, Data Protector selects a pair of devices from the devices that you specified in the automated media copy specification. The automated media copy functionality provides its own load balancing. Data Protector tries to make optimum use of the available devices by utilizing as many devices as possible and selecting local devices if they are available.

Devices are locked at the beginning of the session. The devices that are not available at that time cannot be used in the session, as device locking after the beginning of the session is not possible. Note that at least a pair of devices must be available for each media type for the entire session to complete successfully. If the minimum number of devices necessary for the session cannot be locked, the session fails.

The source medium defines the destination pool of the target medium. This means that the copied media will belong to the same pool as the original media.

The default protection period for the copy is the same as the protection for the original. You can set a different protection period when creating or modifying the automated media copy specification.

The automated media copy functionality does not handle mount or cleanme requests. If a mount request is received, the media pair concerned is aborted, but the session continues. You can manually copy the media that were not copied after the automated media copy session finishes.

If a media error occurs, the device with errors will be avoided within that automated media copy session. However, if there are no other devices available, it will be reused.

Types of automated media copying

There are two types of automated media copying: post-backup media copying and scheduled media copying.

Post-backup media copying

Post-backup media copying takes place after the completion of a backup session. It copies the media used in that particular session.

Scheduled media copying

Scheduled media copying takes place at a user-defined time. Media used in different backup specifications can be copied in a single session. You create an automated media copy specification to define which media will be copied.

Configuring Post-Backup Media Copying

Post-backup media copying is a process that creates a copy of a medium used in a particular backup session after the backup session has finished.

Note: If a backup session is aborted, a post-backup media copy session is started anyway, in case some objects completed successfully.

Limitations

- You can only use library devices.
- The source medium and the target medium must be of the same type.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Automated Operations** and click **Add Post-Backup Media Operation** to open the wizard.
3. In the Backup specification drop-down list, select the backup specification the media of which you want to copy. In the Media operation type drop-down list, select **Media Copy**, and click **Next**.
4. Select the source devices and the destination devices that will be used. For each media type, you must have at least one pair of devices (one source and one destination device). Click **Next**.
5. Specify the number of copies, whether the media will be ejected automatically after the operation, as well as the location and protection for the target media. Click **Finish** to exit the wizard.

Configuring Scheduled Media Copying

Scheduled media copying is a process that creates a copy of a medium used in a particular backup session at a scheduled time. You can schedule several copy operations in a single session. The media will be copied simultaneously if enough devices are available. Otherwise, they will be copied sequentially.

Limitations

- You can only use library devices.
- The source medium and the target medium must be of the same type.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, right-click **Automated Operations** and click **Add Scheduled Media Operation** to open the wizard.
3. In the Media operation name text box, type a name for the operation. In the Media operation type drop-down list, select **Media Copy**, and click **Next**.
4. Select the source devices and the destination devices that will be used. For each media type, you must have at least one pair of devices (one source and one destination device). Click **Next**.
5. Specify the time frame within which you want to search for backup sessions. Click **Next**.
6. Specify the backup specifications of the backups you want to copy. Click **Next**.
7. Specify the required condition and protection of the source media. Click **Next**.

8. Specify the number of copies, whether the media will be ejected automatically after the operation, as well as the location and protection for the target media. Click **Next**.
9. Right-click a date and click **Schedule** to display the Schedule Media Operation dialog box. Specify the options as desired and click **OK**.
10. Click **Finish** to exit the wizard.

Scheduling Media Copying on Specific Dates

You can schedule a media copy operation on a specific date at a specific time.

You can schedule the media copying while you are adding a new scheduled media operation. To modify the scheduled time of an existing scheduled media operation, follow these steps:

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Automated Operations**. All configured automated operations are displayed.
3. Click the scheduled media copy operation for which you want to change the schedule options, and click the **Schedule** tab.
4. In the Schedule page, scroll through the calendar (clicking the single arrows) for the month in which you want to make the changes.
5. Right-click the unwanted dates that are selected, and click **Delete**. Right-click new dates and click **Schedule** to display the Schedule Media Operation dialog box.
6. Specify the options as desired and click **OK**.
7. Click **Apply**.

Tip: You can click **Reset** to remove all previous schedules.

Scheduling Periodic Media Copying

You can schedule a media copy operation so that it is performed periodically.

You can schedule the media copying while you are adding a new scheduled media operation. To modify the schedule of an existing scheduled media operation, follow the steps below.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Automated Operations**. All configured automated operations are displayed.
3. Click the scheduled media copy operation for which you want to change the schedule options, and click the **Schedule** tab.
4. In the Schedule page, right-click a date and click **Schedule** to display the Schedule Media Operation dialog box.
5. Under Recurring, select **Daily**, **Weekly**, or **Monthly**. Specify the **Recurring options** accordingly.

6. Under Time options, select the time when the operation will be performed. Select **Use starting** and specify the starting date.

Note: If you set the recurring to 2 or more (for example, every 2 weeks on Saturday) without setting the starting date, the first copy session may not be scheduled on the first possible date that matches your selection (for example, it will be scheduled on the second Saturday) due to the Data Protector scheduling algorithm. Check the schedule in the Schedule property page.

7. Click **OK** and then **Apply**.

If the chosen time slot is already occupied, Data Protector prompts you that there are scheduling conflicts, and asks if you wish to continue. If you click Yes, the new schedule will be applied where possible (on the days when the time slot is still free). If you click No, the new schedule will be discarded.

Tip: You can click **Reset** to remove all previous schedules.

Disabling and Enabling an AMC Schedule

When you schedule an automated media copy operation, the schedule is enabled by default. You can disable the schedule if you do not want the operation to be performed, and enable it again when desired.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Automated Operations**. All configured automated operations are displayed.
3. Click the scheduled media copy operation for which you want to disable or enable the schedule, and click the **Schedule** tab.
4. In the Schedule page, select or deselect the **Disable schedule** option. Click **Apply**.

Disabling and Enabling AMC on Holidays

By default, scheduled media copy operations are performed on holidays as well. If you do not want these operations on holidays, you can specify this while adding a new scheduled media operation, or modify it for an existing operation by following these steps:

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Automated Operations**. All configured automated operations are displayed.
3. Click the scheduled media copy operation for which you want to change the **Holidays** option, and click the **Schedule** tab.
4. In the Schedule page, select the **Holidays** option to prevent the operation from being performed on holidays. Deselect the option if you want the operation to be performed on holidays.
5. Click **Apply**.

Resetting an AMC Schedule

You can reset the schedule of a scheduled media copy operation. If you do that, the operation will not be performed until you set a new schedule.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Automated Operations**. All configured automated operations are displayed.
3. Click the scheduled media copy operation for which you want to reset the schedule, and click the **Schedule** tab.
4. In the Schedule page, click **Reset**. Click **Apply**.

All schedules are removed.

Tip: You can undo the action by clicking **Undo**.

Scanning a Device

Scan a device to update Data Protector information about the media in the device or after changing the location of the media manually.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**.
3. In the Results Area, right-click the device that you want to scan and click **Scan**.

The Session Information message displays the status of your scanning operation.

Scanning Media in a Library Device

Scan media in selected slots of a library to update Data Protector information about the media in the device.

Depending on the number of selected slots, scanning may take some time. Data Protector must load a medium from each slot into a drive, and then read the media header.

You can select multiple slots using the Ctrl key and scan several media in a single step. However, you can only use one drive.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**.
3. In the Results Area, double-click the library device and then double-click **Slots**.

4. In the Results Area, select the slots that have the media you want to scan.
5. Right-click the selected slots and then click **Scan** to open the wizard.
6. Select the library's drive where the exchanger will load media to scan.
7. Click **Finish** to start the scan and exit the wizard.

The Session Information message displays the status of your scanning operation.

Tip: If you have the Barcode reader support option enabled, you can quickly scan a SCSI library using the **Barcode Scan** option.

Scanning a Drive in a Library Device

Scan a drive of a library device to update the Data Protector information about the media in the drive.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**.
3. In the Results Area, double-click the library device whose drive you want to scan and double-click the target **Drives** icon.
4. Right-click the drive you want to scan and click **Scan**.

The Session Information message displays the status of your scanning operation.

Activating Barcode Reader Support

If a SCSI library device uses media with barcodes, Data Protector can use the barcodes to provide the following barcode support:

- Recognition of cleaning tapes with a CLN prefix.
- Reference to media by their barcodes. Data Protector displays the media barcode as a prefix of the medium description.
- Quick scan of the media in the slots of the library repository using the media barcodes.

Tip: If you select the **Use barcode as medium label on initialization** option in the library properties, the **Use barcode** option is enabled by default in the **Medium description** options during the initialization of the medium. If this option is not selected, the default option is **Automatically generate**. The default option will be used when Data Protector automatically formats a medium.

Note: All barcodes in a cell must be unique, regardless of the type of media or the fact that there are multiple libraries.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Devices**, right-click the target library device, and then click

Properties. The library device Property page opens.

3. Click the **Control** tab and then select the **Barcode reader support** option.
4. To write the barcode to the medium header on tape each time you initialize a medium with this library, select the **Use barcode as medium label on initialization** option.
5. Click **Apply** to confirm.

Barcode Scanning of a Library Device

Use the **Barcode Scan** option to quickly scan a SCSI library. This is considerably faster than a scan of a repository without the barcode functionality.

Prerequisite

You must have the **Barcode reader support** option enabled.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand Devices, right-click the target library device, and then click **Barcode Scan**.

The Session Information message displays the status of your barcode scan operation.

Searching and Selecting Media

You can search and select media in a media pool or in a library device. You can also list media using the List of Media report. Use this function to locate and select specific media without browsing the entire list of media.

Media selection is especially useful for vaulting purposes, such as moving all media written to last week to a vault.

Searching and selecting media in a media pool

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Pools**.
3. In the Results Area, right-click a media pool and then click **Select Media**. The Select Media dialog box appears.
4. Search and select media according to the medium description, media location, session, timeframe, protection, or use Combine Selections options.

Searching and selecting media in a library device

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**.
3. In the Results Area, double-click a library device, right-click **Slots**, and then click **Select Media**. The Select Media dialog box appears.
4. Search and select media according to the medium description, media location, session, timeframe, protection, or use Combine Selections options.

Searching for media using the List of Media report

Steps

1. In the Context List, click **Reporting**, and click the **Tasks** tab.
2. In the Scoping Pane, expand **Pools and Media**, and click **List of Media** to open the wizard.
3. Follow the wizard, specifying the criteria for your search. Click **Finish** to display the results of the search.

Pre-allocation List of Media for Backup

You can specify the order in which media from a media pool will be used for backup. This order is called a pre-allocation list. You specify the pre-allocation list when configuring a backup. The purpose of a pre-allocation list is to control which media will be used for a backup session. You have to match the pre-allocation list with the available media before each backup.

You can also preallocate media when using the object copy or object consolidation functionality.

Depending on the allocation policy of the media pool, Data Protector behaves in two different ways:

- If the pre-allocation list is used in combination with the **Strict** media allocation policy, Data Protector expects the media in a backup device to be available in that order. If the media are not available, Data Protector issues a mount request. If the media mentioned in the pre-allocation list are loaded in a SCSI exchanger, Data Protector handles the media sequence automatically.
- If the pre-allocation list is used in combination with the **Loose** media allocation policy, media in the pre-allocation list are used first. If the media are not available, any suitable media in the library are used.

Preallocating Media for Backup

The following may provide additional information:

- You can also preallocate media when using the object copy or object consolidation functionality.
- A file library media pool has a media usage policy of non-appendable by default. As this policy gives you the benefits of the file library, it is not recommended to change it and to use the pre-allocation list for file library device media.

To preallocate media in a saved backup specification, follow these steps:

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of backup specification (for example, Filesystem). All saved backup specifications are displayed.
3. Double-click the appropriate backup specification and click the **Destination** tab.
4. In the Destination page, right-click the device that is selected for the backup and click **Properties**.
5. In the Device Properties dialog box, select the desired media pool from the Media pool drop-down list.
6. Under Prealloc list, click **Add**.
A list of media of the selected media pool is displayed.
7. Select a medium and click **Add**.
8. Repeat steps 6 and 7 for all desired media. When finished, click **OK** to return to the Destination property page.
9. Repeat steps 4 to 8 if several devices are used for the backup.
10. Click **Apply** to save the changes.

Recycling a Medium

You recycle (unprotect) media when you want to remove the data protection for all backed up data on the media, thus allowing Data Protector to overwrite the media during one of the next backups. Recycling does not actually change data on the medium; it only tells Data Protector that this data is not protected anymore.

Consider the following:

- Recycling removes protection of all objects on a medium. This also includes data from the same object and session that resides on other media.
- The Recycle operation is not available for media in free pools.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Pools**. A list of configured media pools is displayed in the Results Area.
3. Double-click the media pool that has the medium that you want to recycle.
4. Right-click the target medium name and then click **Recycle**. You can also select multiple media at the same time using Ctrl or Shift keys.

When the operation is completed, the protection of the medium is set to None.

Importing the Catalog from Media

Importing the catalog from a medium writes the detail information like file names and file versions into the IDB, enabling you to browse files and directories for restore.

You can also use Import Catalog if catalog protection has expired for a particular object and you can no longer browse its files and directories. If the detailed information on the specified media already exists in the IDB, the data will not be duplicated.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Pools**.
3. In the Results Area, double-click the media pool with the medium from which you want to import the catalog.
4. Right-click the medium and click **Import Catalog**.
5. If there are more drives, select the library drive to import media to and then click **Next**.
6. Select the **Logging** option that suits your needs.
7. Click **Finish** to start the import and exit the wizard.

The Session Information message displays the status of your import operation. When the import is complete, you can browse files and directories for restore.

Verifying a Medium

Verifying a medium checks whether the data format on the medium is valid and updates information about the medium in the IDB. You can verify only resident Data Protector media. Depending on the backup device and media you use, verifying can take a considerable amount of time.

You can verify a medium copy before vaulting it. You can also verify the medium to check whether the backup is usable, if errors were reported during backup.

When verifying media, Data Protector performs the following:

- Checks the Data Protector headers that have information about the medium (media identification, description, and location).
- Reads all blocks on the medium and verifies block format.
- If the **CRC Check** option was used during backup, recalculates the CRC and compares it to the one stored on the medium. In this case, the backup data itself is consistent within each block. This level of check has a high level of reliability.

If the CRC Check option was not used, and the verify operation passed, this means that all the data on the medium has been read. The medium did not cause a read error, so the hardware status of the tape is at the very least acceptable. This level of check can be viewed as partial.

Verifying a medium in a standalone device

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Devices**, right-click the device that has the medium you want to verify and click **Verify**.
3. In the Results Area, you can select the **Eject medium after operation** option. Click **Finish** to verify the medium.

This step is skipped in case of a standalone file device.

The Session Information message displays the status of the verification.

Verifying a medium in a library device

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, under **Devices**, expand the library device and then expand **Slots**. Right-click the slot with the medium you want to verify and click **Verify**.
3. In the Results Area, select a library drive for performing the verification and click **Finish**.

The Session Information message displays the status of the verification.

Moving a Medium

You can move a medium from one media pool to another media pool of the same type, if you want to reorganize the backups and rearrange the purpose of each pool. It is also useful when you want to use the medium in a device which is the default device of another media pool.

Note: You cannot move a medium to the free media pool. When using a free pool, media are moved in two instances (behavior depends on the free pool options selected):

- When media are selected (allocated) for backup, they are moved from a free pool to a regular pool.
- When the media protection has expired, media are moved from a regular pool to a free pool.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Pools**.
3. In the Results Area, double-click the media pool from which you want to move a medium. A list of media in the respective pool appears.
4. Right-click the medium that you want to move and then click **Move to Pool** to open the wizard. You can also select multiple media at the same time using Ctrl or Shift keys.

5. Select the media pool to move the media to.
6. Click **Finish** to move the medium and exit the wizard.

Tip: To move media to another cell, export the media from one cell and then import them to the other cell.

Exporting a Medium

Export a medium when you want to move it to another Data Protector cell. Exporting removes information about a medium and its contents from the IDB. Data Protector no longer knows that this medium exists. The data on the medium remains unchanged.

Note: It is recommended to not manually export media on backup to disk devices (B2D) that rely on daily maintenance to clean the storage, because of the non-trivial nature of manually exporting all the media. Allow the daily maintenance to clean the storage.

If you export the original media and still have copies, then one of the copies becomes the original.

Before exporting the medium you must remove its protection by recycling the medium.

You should export all the media of the same backup session. If the data from the session spans several media and you only export one medium, you may not be able to restore the data. Data Protector still knows that the data exists on the media, but some media are no longer available.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Pools**.
3. In the Results Area, double-click the media pool that has the medium that you want to export, right-click the medium name, and then click **Export**.
4. Confirm the action.

The exported medium is no longer displayed in the list of media in the pool.

Copying the Catalog Media Data to the MCF File

Copying media-related catalog data to a file writes the detail information like file names and file versions into media container format (MCF) files that reside on the Cell Manager in the directory *Data_Protector_program_data\Config\Server\export\mcf* (Windows systems) or */var/opt/omni/server/export/mcf* (UNIX systems). These files can then be imported into another Data Protector Cell Manager where the media-related catalog data becomes available for browsing.

Limitations

- You can select only Data Protector media.
- Due to the nature of the Data Protector file library, where media cannot be exported from one and imported into another library, **Copy Catalog to File** and **Import Catalog from File** of such media should be avoided.

Recommendations

- Due to a possibly large amount of catalog data per medium, it is recommended to store the files on a separate partition or mount point.
- You can reduce the size of the files by setting the `EnableMCFCompression` global option to 1. By default, the compression is disabled.

The following may provide additional information:

- The media-related catalog data is not removed from the original Cell Manager.
- This operation creates one MCF file per each medium.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and then expand **Pools**.
3. Expand the media pool with the medium whose catalog you want to copy.
4. Right-click the medium and click **Copy Catalog to File**.
5. Specify the output directory for the MCF file, which will contain media-related catalog data.
6. Click **Finish** to start copying and exit the wizard.

The exported MCF file can be transferred to the destination Cell Manager.

Tip: You can achieve the same result by expanding **Devices**, right-clicking on the slot of the selected device, and then performing steps 5 and 6.

Importing the Catalog Media Data from the MCF Files

Importing media-related catalog data copies from media container format (MCF) files from the original Cell Manager enables you to browse through the files on the destination Cell Manager.

Prerequisites

- Ensure that the MCF files that you want to import are transferred from the original Cell Manager and accessible on the current Cell Manager.

Limitations

- After a medium is imported from a file, it cannot be used by operations that require physical presence of media (for example, restore, medium copy). In order for a medium to be fully usable for Data Protector operations, it has to be physically accessible and scanned by using Data Protector medium scan, otherwise a mount request will be issued.

The following may provide additional information:

- When you import numerous media catalogs from the MCF files, make sure to import all media that are part of a restore chain.
- You can import different types of media from various media pools within one session.

- Data Protector GUI only shows and allows selection of the files with mcf extension. Other files are hidden from the directory tree. However, you can select them via the command-line interface (CLI). For details, see the `omnim` man page or the *HPE Data Protector Command Line Interface Reference*.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media**, right-click **Pools**, and click **Import Catalog from MCF File** to open the wizard.
3. Specify MCF files you want to import.
4. Specify additional options for the session: by default, the **Import to original pool if possible** option is selected. You can select the Prefix for new pools or the **Import Copy as Original** option.
5. Click **Finish** to start importing and exit the wizard.

Modifying Media Description

A media description helps you identify media. The description is written on the media and stored in the IDB. You add a media description when you format new media. If the media was auto-formatted during a backup, you may want to change the automatically created description to something that better suits your needs.

When you modify a media description, Data Protector modifies the description in the IDB and not on the medium itself. If you export and then import media, the description in the IDB is replaced with the description from the media.

The descriptive part of the media label is changed too, but the barcode part remains the same.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Pools**.
3. In the Results Area, double-click the media pool with the medium description that you want to change. The list of media in the media pool is displayed.
4. Right-click the medium with the description you want to change and then click **Properties** to open the General property page for the medium.
5. In the Description text box, type in a new description for the medium.
6. Click **Apply** to confirm.

Modifying Media Location

Specifying a media location helps you find media when they are out of a device. The location information is stored in the IDB. You should enter the location when you initialize media and modify the location whenever you move media to a different place (vault), such as off-site storage ("Shelf 4-Box 3").

The location is never written to the media header.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Pools**.
3. In the Results Area, double-click the media pool with the media location you want to change. A list of media in the media pool displays.
4. Right-click the media with the specified location that you want to change and then click **Change Location** to open the wizard.
5. Specify the new location for the media.
6. Click **Finish** to exit the wizard.

Creating a List of Locations

You can create a list of pre-defined vault locations that you often use. This pre-defined vault location list is available when choosing a location for specific media in different media management tasks (for example, when formatting your media).

Steps

1. In the Context List, click **Devices & Media**.
2. In the Edit menu, click **Locations**.
3. Enter the location you want and click the **Add** button. Repeat the step to enter several locations.
4. Click **Finish**.

Setting the Media Location Priority

If an object version that you want to restore, copy, consolidate, or verify exists on more than one media set, any of the media sets can be used for the operation. By default, Data Protector automatically selects the most appropriate media set. You can influence the media set selection by specifying the media location priority.

If the media location priority is set, Data Protector will use the media set with the highest priority (priority 1 is the highest, priority None is the lowest) if more than one media set equally matches the conditions of the media set selection algorithm.

The media location priority can be overridden on the restore, object copy, object consolidation, or object verification session level.

The following may provide additional information:

- By default, media location priority is considered only if two or more media sets have the same rating. For media location priority to take precedence over other selection factors, set the global option `UserSpecifiedMediaPriorityHasHigherImportance` to 1.
- For media location priority to take effect, the location of each medium must be specified. This can be done for individual or multiple media.
- Media location priority does not consider copies obtained using the media copy functionality. Such a

copy is used only if the original medium (the medium that was used as a source for copying) is unavailable or unusable.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Media** and click **Locations**.
3. In the Results Area, double-click a location to display its properties.
4. In the **Location priority** drop-down list, select one of the available numbers, where 1 means the highest priority.
5. Click **Apply** to confirm your selection.

Vaulting a Medium

It is recommended to make a copy of the backed up data for vaulting purposes, and keep the original on site to enable a restore. Data Protector enables interactive or automated creation of additional copies of the data on the media.

Prerequisites

- You need to have the desired data protection and catalog protection policies set when configuring a backup specification.
- You need to configure a vault in Data Protector. Use the name indicating the physical location where media will be kept.

Steps

1. In the Data Protector Manager, change the location of the media that you want to store.
2. Eject the media from the device and then store the media in the vault.

Erasing a Medium

This functionality is only available for magneto-optical platters. You use it to erase a magneto-optical platter before a backup session, and as a result the backup speed is increased.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**.
3. In the Results Area, double-click the magneto-optical device that has the medium you want to erase.
4. Right-click the medium and then click **Erase** to open the wizard.
5. (Optional) Select the **Eject medium after operation** option.
6. Click **Finish** to erase the medium and exit the wizard.

The Session Information message displays the status of your erase operation.

Detection of Write-Protected Media

Data Protector can detect and handle media that has been mechanically protected by setting the write protection switch on.

The following operations can detect and handle write-protected media:

- Read-only operations, such as list, scan and verify. Read-only operations detect the write-protected media and proceed without any warnings.
- Write operations, such as initialize, erase and backup. Write operations detect the write-protected media and either abort the session or skip the write-protected media. Backup sessions treat write-protected media as unusable media and behave according to the media allocation policy. If the allocation policy is strict, a mount request is issued. If the allocation policy is loose, the medium is skipped.

The detection of a write-protected medium and all changes to the write-protection state of the medium are logged to the `media.log` file.

Note: It is recommended not to use write-protected media with Data Protector.

About Mount Requests

A mount request is a screen prompt that tells you to insert media into a device. Once you respond to the mount request by providing the required media, the session continues.

Data Protector issues a mount request in the following cases:

- The specified media is not available. This can be the case if a pre-allocation list is used for backup or the media necessary for restore is missing media.
- No suitable media is available. This can be the case if the media from a pool that are currently in the library are not suitable, if the medium in a standalone device is not suitable, or if the device is empty.
- The mail slot is open. In this case, you have to shut the mail slot.

The most appropriate media for backup are selected automatically by Data Protector. You have to be aware of the way media are selected for backup.

About Library-Specific Media Management

Data Protector provides some specific media management tasks for complex devices, such as libraries, to simplify management of a large number of media.

Some tasks follow the standard procedure, for example, selecting, copying, recycling or moving media, and modifying media location. Other tasks, such as adding or deleting a slot, and entering, ejecting, verifying, formatting, importing, scanning, or erasing media, may depend on the device type used.

In libraries with the barcode support, Data Protector can generate media descriptions based on barcodes and writes them to the medium header on the tape during initialization.

The use of library media by other applications

Media in a library (especially in very large libraries such as ADIC/GRAU and StorageTek) can be used by many applications, not just by Data Protector, so you have to know which applications use which media to prevent them from being overwritten.

Ideally, you will use the library with Data Protector exclusively and let Data Protector manage the complete library. However, if you have other applications using the library, you should take care to assign non-intersecting subsets of media to Data Protector and other applications. Data Protector maintains its own independent media allocation policy. This implies that if a specific medium has been allocated to Data Protector (added to a Data Protector media pool), it remains under Data Protector control during its lifetime or until it is removed from Data Protector media pool.

For each type of media you have to configure a library in Data Protector. While an ADIC/GRAU or StorageTek system can store many physically different types of media, Data Protector can only recognize a library with a single type of media in it. Therefore you have to create a Data Protector library for every media type in the system.

The following may be helpful:

- For ADIC/GRAU DAS and StorageTek libraries, use Data Protector commands to handle media. If you handle media manually using ADIC/GRAU DAS or StorageTek ACS commands, Data Protector will not be able to track the changes in location or information on the media.
- Manage the whole library with Data Protector. This provides single-point administration where you can track Data Protector and non-Data Protector media in the library.
- Create at least one media pool for each media type, for example, one for 4mm and one for 3480 media type. Depending on your environment, you may want to create more media pools, for example, one for each department.
- Make sure that Data Protector and other applications do not use the same set of media.

About the Data Protector Query Operation Used with ADIC/GRAU DAS or STK ACS Libraries

When the Data Protector query operation is started, all the media configured on the DAS or ACS Library Server is queried, even in cases when these media are configured in Data Protector as belonging to several logical ADIC/GRAU DAS or STK ACS libraries (for the same physical library). Additionally, the Data Protector query operation queries also the media configured on the DAS or ACS Library Server that are configured to be used with applications other than Data Protector. The consequence is that after the query operation is started from Data Protector, the media belonging to other logical ADIC/GRAU DAS or STK ACS libraries than the one for which the query operation was started, are moved to the logical ADIC/GRAU DAS or STK ACS library for which the query operation was started.

Therefore, with ADIC/GRAU DAS or STK ACS libraries, it is not recommended to use the Data Protector query operation. It is recommended to add volsers manually using the Data Protector add volsers operation instead of synchronizing the IDB using the Data Protector query operation.

Note: The information in this section does not apply in case of ADIC/GRAU DAS libraries, when

logical libraries are not configured using Data Protector, but using the ADIC/GRAU DAS utilities. If several logical libraries are configured using the ADIC/GRAU DAS utilities, the Data Protector query operation can safely be used on such libraries.

Adding a Slot

Data Protector provides full support of handling slots and media in media pools used by libraries. Adding a slot configures a location for media in a storage device.

On some libraries, slots are detected and added automatically when the library is configured.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**.
3. In the Results Area, right-click the name of the library, and then click **Properties**.
4. Click the **Repository** tab, specify the slot that you want to use with Data Protector, and then click **Add** to add the slot to the list. Use a dash to enter a multiple slots at a time, for example, 5-12.
Make sure you use a format supported by your library. For example, when adding slots to a SCSI library, do not use letters or leading zeros.
5. Click **Apply** to confirm.

Deleting a Slot

Data Protector provides full support of handling slots and media in media pools used by libraries. Deleting a slot prevents Data Protector from using and accessing the slot in the repository. Information about the slot is removed from the IDB.

Deleting of media slots is enabled only for empty slots on any device.

This action does not affect volsers in the GRAU DAS library but only removes specific media from the IDB. Therefore, Data Protector does not know that these media exist and does not use them.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**.
3. In the Results Area, right-click the name of the library, and then click **Properties**.
4. Click the **Repository** tab, select the slot that you want to remove, and then click **Delete**.
5. Click **Apply** to confirm.

The slot is no longer displayed in the slot list.

Entering a Medium

Entering a medium means physically entering it into a library repository and automatically registering the added media as members of the library.

You can select the slot that you want to use. Entering media does not affect the media pool to which they belong.

It is recommended that you use the Data Protector GUI to enter a medium. If you manually enter a medium using the device's controls, the information in the IDB is no longer consistent, and you have to scan the device to update this information.

Tip: You can enter multiple media into a device in a single action.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**. A list of configured devices is displayed in the Results Area.
3. In the Results Area, double-click the name of the library.
4. Double-click **Slots** to display a list of slots in the Results Area.
5. Right-click the slot (or multiple slots) where you want to enter the media, and then click **Enter** to start the session.

You will be prompted to insert additional media into the device as needed.

Ejecting a Medium

Ejecting a medium means physically transferring it from the repository slot to the Insert/Eject area (also called a mail slot) in a library device.

It is recommended that you use the Data Protector Manager to eject media. If you manually eject a medium using the device's controls, the information in the IDB is no longer consistent. To update this information, scan the device.

When media cannot be ejected because the mail slot is full, Data Protector retries the operation until the mail slot becomes free or until the predefined time limit expires. During this retry, the robotics are accessible to other sessions.

During the eject execution, none of the specified media can be used by other sessions.

Bulk eject of media

You can eject multiple media from a library in a single action. Data Protector instructs you to remove media from a mail slot when it becomes full, to free up space for other media selected for ejection.

Predefined eject of media

With some operations, such as automated media copying, you can specify whether the media will be ejected automatically after the session finishes.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**. The list of configured devices displays in the Results Area.
3. In the Results Area, double-click the name of the library.
4. Double-click the **Slots** item to display the list of slots in the Results Area.
5. Right-click the slot (or multiple slots) to be ejected, and then click **Eject** to open the wizard.
6. Specify a new location for the medium (optional).
7. Click **Finish** to eject the medium and exit the wizard.

The Session Information message displays the status of your eject operation.

Erasing Media in a Library Device

Erasing a medium is only available for magneto-optical platters. You can only erase a magneto-optical platter media before a backup session. This results in an increase in backup speed.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, click **Devices**.
3. In the Results Area, double-click the magneto-optical device that has the media that you want to erase. The Slots and Drives items appear.
4. Double-click **Slots**.
5. Right-click the slots with the media that you want to erase and then click **Erase** to open the wizard.
6. Select the drive in the library where the exchanger will load media to erase.
7. Click **Finish** to erase the media and exit the wizard.

The Session Information message displays the status of your erase operation.

Adding Volsers Manually

With ADIC/GRAU DAS or STK ACS libraries, you can manually add volsers to a library configured in Data Protector instead of querying the library. With ADIC/GRAU DAS or STK ACS libraries, when several logical libraries are configured for the same physical library, this is the recommended way of adding volsers to a library configured in Data Protector. With ADIC/GRAU DAS libraries, however, when logical libraries are not configured using Data Protector, but using the ADIC/GRAU DAS utilities,

the Data Protector query operation can safely be used on such libraries instead of adding volsers manually.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, browse for the library to which you want to add volsers and expand it.
3. Right-click **Slots** and select **Add Volser(s)** from the pop-up menu.
4. In the Prefix text box, enter the volser's prefix. It usually consists of three letters.
In the From text box, enter the start number of the range of volsers that you want to add to the library.
In the To text box, enter the end number of the range of volsers that you want to add to the library.
5. Click Finish to add the volsers to IDB.

Querying the ADIC/GRAU DAS and StorageTek ACSLM Hosts

To get information about a repository in the ADIC/GRAU or StorageTek libraries from the server, you can query the DAS or ACSLM host (server). A query responds with the contents of the media database of the server, and then synchronizes the information in the IDB with what is actually in the repository.

This is especially useful if you have used GRAU DAS or StorageTek ACS commands to manage media, as this results in inconsistencies with the IDB - Data Protector does not know the latest status of media in the library repository.

Limitation

Volsers scan may not complete successfully if the ADIC/GRAU library is configured with more than 3970 volsers in a repository. A workaround for this problem is to configure multiple logical ADIC/GRAU libraries in order to separate the slots from the large repository into several smaller repositories.

With ADIC/GRAU DAS and STK ACS libraries, when several logical libraries are configured for the same physical library, it is not recommended to query the DAS or STK ACSLM Server. Add volsers manually. With ADIC/GRAU DAS libraries, however, when logical libraries are not configured using Data Protector, but using the ADIC/GRAU DAS utilities, the Data Protector query operation can safely be used on such libraries.

Steps

1. In the Context List, click **Devices & Media**.
2. In the list of configured devices, right-click the library you want to query, and then click **Query**.

This action queries the DAS or ACSLM host for information.

Chapter 9: Backup

About Backup

A backup is a process that creates a copy of system data on backup media. This copy is stored and kept for future use in case the original is destroyed or corrupted.

A backup session is based on the backup specification and can be started interactively. During a backup session, Data Protector reads the backup objects, transfers their data through the network, and writes them to the media residing in the devices.

Make sure that the data you back up is consistent. For example, you might shut down an application before backup or put it into "backup" mode to avoid data changes during a backup. If you back up data that is inconsistent, you may encounter unexpected results when you restore and attempt to use the data.

Advanced features of Data Protector backup include:

- Automatically balancing the usage of devices (load balancing)
- Backing up shared disks
- Scheduling unattended backups
- Combining full and incremental backups to save time and media
- Allowing backups to be organized in many different ways
- Backing up to multiple locations simultaneously using the object mirror functionality

Procedures in Data Protector Help assume that you use the default Backup View (By Type) that is set according to the type of data available for the backup or template.

For information on how to back up database applications such as Oracle, SAP R/3, Microsoft Exchange Server, Microsoft SQL Server, Informix Server, IBM DB2 UDB or Sybase, see the *HPE Data Protector Integration Guides*.

Setting the Backup View

You can set your Backup View according to your needs. The default Backup View is **By Type**.

Steps

1. In the Context List, click **Backup**.
2. In the View menu, select one of the available views.

The Backup context displays according to the view you have chosen.

Full and Incremental Backups

A basic approach to improving backup performance is to reduce the amount of backed up data. You should take full advantage of time and resources when planning your full and incremental backups. There is often no need to perform full backups of all the systems on the same day.

Consider the following about the backup types:

	Full backup	Incremental backup
Resources	Takes more time to complete than incremental backup and requires more media space.	Backs up only changes made since a previous backup, which requires less time and media space.
Device handling	If you use a standalone device with a single drive, you need to change the media manually if a backup does not fit on a single medium.	It is less likely that the backup will require additional media.
Restore	Enables simple and quick restore.	A restore takes more time because of the number of media needed.
IDB impact	Occupies more space in the IDB.	Occupies less space in the IDB.

Note: You must set appropriate data protection to ensure all the needed full and incremental backups are available for restore. If the data protection is not properly set, some media might get overwritten, which results in a broken restore chain.

Conventional Incremental Backup

How conventional incremental backup works

Before running an incremental backup of a backup object, Data Protector compares the trees in the backup object with the trees in the valid restore chain of this object. If the trees do not match (for example, an additional directory in the backup object was selected for backup since the last backup or multiple backup specifications with the same backup object and different trees exist), a full backup is automatically performed. This ensures that all the selected files are included in the backup.

Detection of changes

With conventional incremental backup, the main criterion for determining whether a file has changed or not since a previous backup is the file's modification time. However, there are cases where this criterion is not effective. For example, if a file has been renamed, moved to a new location, or if some of its attributes have changed, its modification time does not change. Consequently, the file is not always backed up in an incremental backup. Such files are backed up in the next full backup.

Whether a file with a changed name, location, or attributes is backed up in an incremental backup or not also depends on the setting of the following options in the backup specification. The preferred setting improves detection of changes.

Windows systems: Do not use archive attribute

By default this option is not selected (archive attribute is used). This is the preferred setting.

UNIX systems: Do not preserve access time attributes

By default this option is not selected (access time attributes are preserved). Preferably this option is selected.

You can perform a conventional incremental backup using the Windows NTFS Change Log Provider. In such a case, a Windows Change Journal is used to generate a list of files that have been modified since the last full backup and a file tree walk is not performed. Using the Change Log Provider improves the overall incremental backup performance in the same way as it improves the performance of the enhanced incremental backup. In case the Change Log Provider cannot be used for some reason, a regular conventional incremental backup is performed.

To reliably detect and back up renamed and moved files, as well as files with changes in attributes, use enhanced incremental backup.

Enhanced Incremental Backup

With conventional incremental backup, the main criterion for determining whether a file has changed or not since a previous backup is the file's modification time. However, there are cases where this criterion is not effective. For example, if a file has been renamed, moved to a new location, or if some of its attributes have changed, its modification time does not change. Consequently, the file is not always backed up in an incremental backup. Such files are backed up in the next full backup.

Enhanced incremental backup reliably detects and backs up also renamed and moved files, as well as files with changes in attributes.

Detection of certain changes (such as changes of permissions or ACLs) also depends on the setting of the following options in the backup specification. The preferred setting enables maximal detection of changes with enhanced incremental backup.

- **Windows systems:** Do not use archive attribute

By default this option is not selected (archive attribute is used). This is the preferred setting.

- **UNIX systems:** Do not preserve access time attributes

By default this option is not selected (access time attributes are preserved). The preferred setting is when this option is selected.

Enhanced incremental backup also eliminates unnecessary full backups of an entire backup object when some of the trees selected for backup change. For example, if an additional directory is selected for backup since the last backup, a full backup of this directory (tree) is performed, whereas the backup of the rest is incremental.

In addition, you can perform enhanced incremental backup using the Windows NTFS Change Log Provider. In such a case, a Windows Change Journal is used to generate a list of files that have been modified since the last full backup and a file tree walk is not performed. Using the Change Log Provider improves the overall incremental backup performance, especially in environments that contain millions of files only a few of which have changed.

Why use enhanced incremental backup

Use enhanced incremental backup:

- To ensure incremental backup of files with changes in name, location, or attributes.
- To eliminate unnecessary full backups if some of the selected trees change.
- To enable subsequent object consolidation (synthetic backup).

Impact on disk space consumption

Enhanced incremental backup uses a small database on each client that is backed up. The database is created per file system mount point. The enhanced incremental backup repository is located in the following directory:

- **Windows systems:** *Data_Protector_home\enhincrd\MountPointDir*

The mount point directory (MountPointDir) is obtained from the mount point by replacing any ":" (colon) and "\" (backslash) characters with the "_" (underscore) characters, and omitting the trailing ":" or "\".

- **HP-UX and Linux systems:** */var/opt/omni/enhincrd*

The impact on disk space on the client is typically less than 1% of the size of the files selected for backup. Ensure that the enhanced incremental backup database is regularly purged. You can do this by setting the OB2_ENHINC_DELETE_INTERVAL and OB2_ENHINC_DELETE_THRESHOLD omnirc options.

Disk Agent concurrency

Multiple Disk Agents may access the enhanced incremental backup database simultaneously. To avoid possible problems with the backup, configure the Disk Agents behavior by setting the following omnirc options:

- OB2_ENHINC_LOCK_TIMEOUT
- OB2_ENHINC_SQLITE_MAX_ROWS
- OB2_ENHINC_MAX_MEMORY_LIMIT

Limitations

- Enhanced incremental backup is only supported on a directory level. If you select individual files for backup, the enhanced incremental mode will not be used.
- Hardlink detection is disabled when enhanced incremental mode is used.

Incremental Backup Using Change Log Provider

With a conventional and enhanced incremental backup, a list of files to be backed up is generated by performing a file tree walk. This process can take a considerable amount of time, especially when the directory structure is large and contains millions of files. The Windows NTFS Change Log Provider, based on the Windows Change Journal, addresses this issue by querying the Change Journal for a list of changed files rather than performing a file tree walk. The Change Journal reliably detects and records all changes made to the files and directories on an NTFS volume, which enables Data Protector to use

the Change Journal as a tracking mechanism to generate a list of files that have been modified since the last full backup. This is very beneficial for the environments with large filesystems, where only a small percentage of files change between backups. In this case, the process of determining changed files completes in a much shorter period of time.

Each NTFS volume has its own Change Journal database. Whenever a change to a file or directory is made, a record is appended to the journal. The record identifies the file name, the time and the type of the change. Note that the actual changed data is not kept in the journal. If the journal file gets too big, the system purges the oldest records at the start of the journal. If the data required for backup has been purged from the Change Journal, Data Protector performs a full backup and issues a warning that the Change Journal could not be used.

Whether a file is backed up in an incremental backup that uses the Change Log Provider depends on setting the Use native Filesystem Change Log Provider if available option in a backup specification. If this option is specified, Data Protector attempts to use the Change Journal. If the Change Journal is not active, Data Protector issues a warning. In case this occurs during the enhanced incremental backup, a full backup is performed instead. In case it occurs during the conventional incremental backup, a regular incremental backup is performed instead. The options Do not preserve access time attributes and Do not use archive attribute are automatically set and cannot be disabled.

Prerequisites

- Make sure the Change Journal is activated on a needed volume by using the `omnicjutil -query` command. If the Change Journal is not active, start it by running `omnicjutil - start`. For more information on the `omnicjutil` command, see the *HPE Data Protector Command Line Interface Reference* located in the `Mount_point/DOCS/C/MAN` directory on the DVD-ROM.
- Make sure at least one full backup (the option **Use native Filesystem Change Log Provider if available** selected in the backup specification) exists before starting an enhanced incremental backup using the Change Log Provider.

Performance and Disk Space Consumption

To achieve the best Change Log Provider performance, use incremental backups when starting the backup (the backup type is Incr). Incr1-9 is supported as well, but some performance degradation is possible.

When turned on, the Change Journal consumes some CPU time and disk space. The disk space consumption is limited to 4 GB. You can set the maximum size of the Change Journal, as well as the size to be truncated for the journal when it reaches its maximum size. For more information, see the *HPE Data Protector Command Line Interface Reference*.

To optimize the Change Log Provider performance, you can specify the number of entries the Change Log Provider can hold in memory using the `OB2_CLP_MAX_ENTRIES` omnirc option. For detailed information, see the *HPE Data Protector Troubleshooting Guide*.

In the following cases, Data Protector performs a full backup and ignores setting the Change Log Provider option in a backup specification:

- If the Change Journal is not active on the client system.
- If the needed data has been purged from the Change Journal.

- If the Change Journal ID is different from what it used to be (this means that another application has deleted and then recreated the Change Journal).

By default, the Change Log Provider does not create the Enhanced Incremental Repository when it is first executed. This means that the first time a Change Log Provider error occurs, a full backup is performed, which creates the Enhanced Incremental Repository. This behavior can be changed through the `OB2_CLP_CREATE_EI_REPOSITORY` omnirc option. See the *HPE Data Protector Troubleshooting Guide* for more information.

Considerations

- Data Protector does not have exclusive access to the Change Journal. This means that, by activating or deactivating the Change Journal, other applications can affect Data Protector. If a Change Journal is disabled on a given volume, no file and directory changes are logged into the journal. By default, an NTFS volume has its Change Journal disabled, so you must explicitly activate it using the `cjutil` or the `omnicjutil` command. At the same time, any other application can activate or disable the volume's journal at any time. For more information on the Change Journal, see the Windows documentation.

Note that on Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012 the Change Journal is active by default.

- Using the Change Log Provider is beneficial in the environments with a small percentage of changes on a filesystem. A backup of a filesystem with many changes (for example, with many temporary files created and deleted soon after creation) is faster with a normal tree walk.
- The Windows Change Journal API does not provide detailed information about attributes. All attribute changes are grouped together. Using the API, you cannot determine if an entry in the Change Journal is caused by an archive attribute being unset or by a change in the last accessed time.

The Change Log Provider does not unset the archive attribute. The normal Data Protector behavior is to unset the archive attribute after the file is backed up. For this reason, when the Change Log Provider is used, the option **Do not use archive attribute** is automatically selected.

The normal Data Protector behavior is to reset the last access time after the file is backed up (because the backup process always changes the last access time). The Change Log Provider does not reset it, therefore the option **Do not preserve access time attributes** is automatically selected.

The reason for automatic selection of these two options is to avoid situations when the same files are backed up several times. If the archive attribute is unset or the last access time is reset, an entry appears in the Change Journal and the files are backed up in the next session even if they have not been changed.

- You need to occasionally monitor the `NextUsn` number using the `cjutil - query` command and restart the Change Journal when `NextUsn` approaches the `MaxUsn` number.
- If a backup specification has been changed, all new trees are backed up completely. This means that a normal tree walk is performed for all new trees and the Change Log Provider is used for the old ones.
- If a directory underneath the backup space is renamed, a normal tree walk is performed on that directory.

Limitations

- Only backup of Windows NTFS is supported.

Synthetic Backup

Synthetic backup is an advanced backup solution that eliminates the need to run regular full backups. After an initial full backup, only incremental backups are run, and subsequently merged with the full backup into a new, synthetic full backup. This can be repeated indefinitely, with no need to run a full backup again. In terms of restore speed, such a backup is equivalent to a conventional full backup.

Data Protector performs synthetic backup with an operation called object consolidation.

How to perform synthetic backup

The synthetic backup procedure consists of the following steps:

1. In the backup specification that is used for the full backup and incremental backups, enable the **Enhanced incremental backup** option.
2. Perform a full backup.
3. Configure subsequent incremental backups to be written to one file library or B2D devices (except Smart Cache).
4. When at least one incremental backup exists, perform object consolidation. How frequently you perform object consolidation depends on your backup strategy.

Virtual full backup

Virtual full backup is an even more efficient type of synthetic backup. This solution uses pointers to consolidate data rather than copy the data. As a result, the consolidation takes less time and avoids unnecessary duplication of data.

The procedure is basically the same as for regular synthetic backup, with the following additional requirements:

- All backups must be written to one file library: the full backup, incremental backups, and the resulting virtual full backup.
- The file library must use distributed file media format.

Note: Virtual full backup enables you to reduce space consumption, as objects share the same data blocks. However, in case of a corruption of a data block, multiple objects might be affected. For better reliability, keep the file library on a RAID disk.

Standard Backup Procedure

A standard backup procedure consists of several parts:

- Selecting the data to be backed up.
- Selecting where to back it up to.
- Selecting how many additional backup copies (mirrors) to create.
- Starting or scheduling a backup session.

This is done while creating a backup specification. Details of how to back up are defined by setting various options, either using defaults or setting them to meet your specific needs.

To change these predefined settings, specify:

- the backup options for all objects in the target backup specification, such as pre-exec and data protection
- the dates and times that you want backups to be performed

Prerequisites

- You need to have a Disk Agent installed on every system that is going to be backed up, unless you use NFS (on UNIX systems) or you perform network share backup (on Windows systems) for backing up these systems.
- You need to have at least one backup device configured in the Data Protector cell.
- You need to have prepared media for your backup.
- You need to have appropriate user rights for performing a backup.

Filesystem backup

For each filesystem, you can restrict the backup to specific directory trees. For each directory tree you can:

- Exclude any sub-tree or file
- Back up files that match a specific wildcard pattern
- Skip files that match a specific wildcard pattern

Some files are permanently in use, for example, by software applications. These files should be excluded from the filesystem backup and should be backed up in a special way.

Creating a Backup Specification

A backup specification defines the clients, disks, directories, and files to be backed up; the tape devices or drives to be used; the number of additional backup copies (mirrors); backup options; and the timing information (when you want backups to be performed). A backup specification can be as simple as backing up one disk to a standalone DDS drive or as complex as specifying a backup for 40 large servers to a tape library with 8 drives.

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 load balanced or not). This size should not exceed 80 kB.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**.
3. Right-click the type of item that you want to back up (for example, **Filesystem**), and click **Add Backup**.
4. In the **Create New Backup** dialog box, select one of the available templates, the backup type, and specify other options as desired. Click **OK** to open the wizard.
5. In the case of zero downtime backup, the Configuration page is displayed. Configure the integration and then click **Next**.
6. In the case of integration backup, select the client and the application database. Click **Next**.
7. In the Source property page, expand the system that contains the objects that you want to back up and then select what you want to back up.

On UNIX systems, if you intend to perform instant recovery, select all filesystems inside the volume group to be backed up. Otherwise, instant recovery will not be possible using the Data Protector GUI or (if you perform instant recovery using the Data Protector CLI) data can be corrupted.

Click **Next**.

8. In the Destination property page, select the device(s) you will use for your backup.

You can also specify whether you want to create additional copies (mirrors) of the backup during the backup session. Specify the desired number of mirrors by clicking the **Add mirror** and **Remove mirror** buttons. Select separate devices for the backup and for each mirror. It is not possible to mirror objects backed up using the ZDB to disk or NDMPbackup to IAP functionality.

Tip: If the backup is load balanced, you can set the order in which Data Protector will use the devices by right-clicking a selected device and clicking **Order devices**.

Click **Next**.

9. In the Options property page, you can set the backup options. Backup options are available according to the type of data being backed up. For example, all backup options available for a filesystem backup are not available for a disk image backup. Click **Next**.
10. In the Schedule property page, specify the dates and times that you want your backups performed (optional). Click **Next**.
11. In the Backup summary page, review the summary of the backup specification. It is recommended that you first save the backup specification and then start a preview. Preview is not available for the Data Protector Internal Database backup, the backup sessions of specific Data Protector application integrations, and zero downtime backup (ZDB). Click **Next**.
12. At the end of the Backup wizard, you can save, start, or preview the configured backup. The following happens:
 - If you save the configured backup, it appears in the Backup context of the Scoping Pane as a new backup specification. You can later preview or start the saved backup without any modifications, or you can modify it and then preview or start it.

- If you start or preview the configured backup, the Session Information message displays the status of your backup.

Tip: You can create multiple backup specifications by copying an existing specification and then modifying one of the copies.

Modifying a Backup Specification

You can modify an already configured and saved backup specification.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Click the backup specification that you want to modify.
4. In the Source property page, as well as in the other property pages (Destination, Options, and Scheduling), modify your backup specification, and then click **Apply**.

Once you have modified your backup, you can preview or start it in the **Actions** menu.

Note: Preview is not available for the Data Protector Internal Database backup, the backup sessions of specific Data Protector application integrations, and zero downtime backup (ZDB).

Tip: When you modify a backup specification, perform backup and then select the object for restore, only files and directories backed up in the last version are selected for restore. To change the backup version, right-click the object and then click **Select Version**.

Previewing and Starting a Backup

You can preview a backup to verify your choices. Previewing does not read data from disk(s) selected for backup, nor does it write data to the media in the device configured for the backup. However, it checks the communication through the used infrastructure and determines the size of data and the availability of media at the destination.

You can start an existing (configured and saved) backup after you have given Data Protector all the information for the backup.

Limitations

- Preview is not available for the Data Protector Internal Database backup and the backup sessions of specific Data Protector application integrations.
- Preview is not available for zero downtime backup (ZDB).

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Select the backup specification that you want to start or preview.
4. In the **Actions** menu, click **Preview Backup** if you want to preview it or **Start Backup** to start it.
5. In the Preview or Start Backup dialog box, select the backup type (Full or Incremental; some other backup types are available for specific integrations) and the Network load.
In the case of ZDB to disk+tape or ZDB to disk (instant recovery enabled), specify the Split mirror/snapshot backup option.
6. Click **OK** to preview or to start the backup.

The Session Information message displays the status of your backup.

Tip: When configuring a new backup, you can start an interactive backup or an interactive preview at the end of the Backup wizard.

Aborting a Backup

Aborting a backup session terminates a backup session. A backup copy will only exist for data that was backed up before you aborted the session.

Steps

1. In the Actions menu, click **Abort** to abort a backup session.
If you abort a backup session while it is still determining the sizes of the disks that you have selected for the backup, it will not abort immediately. The backup will be aborted once the size determination is completed.

Tip: You can abort one or more currently running sessions in the Data Protector Monitor context.

Restarting Failed Backups

During a backup session, some systems may not be available because they were shut down, there are temporary network connectivity problems, and so on. These circumstances result in some systems not being backed up or being backed up only partially — in other words, some objects fail. You can restart a problematic session after you have resolved the impeding issues. This action restarts only the failed objects.

Prerequisite

- You either have to be in the Data Protector Admin user group or have the Data Protector Monitor user right.

Considerations

- For failed filesystem and Oracle Server integration backup sessions, you can also use resume session functionality to continue backup right from the point where the session failed.

Limitations

- You cannot restart failed sessions that were run interactively, meaning they were based on unsaved backup specifications.
- It is not possible to restart several sessions at the same time.

Do not change a backup specification before restarting a failed backup session. Otherwise, it is not possible to restart all objects.

Steps

1. If you are using an ordinary Cell Manager, in the Context List, click **Internal Database**.
If you are using a Manager-of-Managers, in the Context List, select **Clients** and expand **Enterprise Clients**. Select a Cell Manager with the problematic session. From the Tools menu, select **Database Administration** to open a new Data Protector GUI window with the Internal Database context displayed.
2. In the Scoping Pane, expand **Internal Database** and click **Sessions**.
A list of sessions is displayed in the Results Area. Status of each session is denoted in the Status column.
3. Right-click a failed, an aborted, or a session that completed with failures or errors and select **Restart Failed Objects** to back up the objects that failed.
4. Click **Yes** to confirm.

Copying a Backup Specification

You can copy an already configured and saved backup specification.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. In the Results Area, right-click the backup specification that you want to copy and then click **Copy As**. The Copy Backup As dialog box opens.
4. In the Name text box, enter the name for the copied backup specification. Optionally, from the Group drop-down list, select the backup specification group for your copied backup specification to belong to.
5. Click **OK**.

The copied backup specification is displayed in the Backup context of the Scoping Pane and in the Results Area under the new name.

Deleting a Backup Specification

You can delete an already configured and saved backup specification.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Right-click the backup specification that you want to delete and then click **Delete**. Confirm your choice.

The backup specification is removed from the Backup context of the Scoping Pane.

Advanced Backup Tasks

You can control a backup in many ways. Data Protector offers a set of advanced backup tasks for Windows and UNIX systems.

Prerequisites

- You need to have a Disk Agent installed on every system that is going to be backed up, unless you use NFS (on UNIX systems) or you perform network share backup (on Windows systems) for backing up these systems.
- You need to have at least one backup device configured in the cell.
- You need to have prepared media for your backup.
- You need to have appropriate user rights for performing a backup.
- You have to consider the standard backup procedure before proceeding.

What are advanced backup tasks?

Advanced backup tasks include specifying certain options that are not used by default or taking some actions that do not follow the standard backup procedure.

- [Selecting Network Shared Disk for Backup](#)
- [Selecting Only Specific Files \(Matching\) for Backup](#)
- [Skipping Files for a Backup](#)
- [Selecting the Location for the Shortcut for Starting a Backup](#)
- [Backing Up Using Multiple Disk Agents](#)
- [Client Backup With Disk Discovery](#)
- [Disk Image Backup](#)
- [Web Server Backup](#)

Selecting Network Shared Disk for Backup

You can back up data on Windows shared disks. You have to use a regular Data Protector Disk Agent client to back up other remote systems via shared disks.

Backup using the shared disk method is a workaround to back up systems that cannot be backed up otherwise. This method is not recommended as the main backup approach.

Back up a filesystem located on a Windows system shared in network:

- if the system is not a part of the Data Protector cell and does not have the Data Protector Disk Agent installed.
- if you want to back up platforms not directly supported by Data Protector, such as Windows for Workgroups, Windows 3.1 systems or Windows NT.

Tip: To reduce the network load, a Disk Agent client should be a Media Agent client as well. Otherwise, data is transferred over the network twice.

Prerequisite

You must change the Data Protector Inet account on the Disk Agent client in order to have the right permissions to access the shared disk that you want to back up. This account has to have permission to access both the local client system and the remote shared disks. For Windows versions earlier than Windows Vista and Windows Server 2008 systems, the account must be a specific user account, not the local system account.

Once you have set the user account for the Inet service, you can back up the shared disks as though they were residing on the local system.

Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012

You must add a user account with permissions to access the shared disk that you want to back up. This account must be a local system account.

This prerequisite must be fulfilled before changing the Data Protector Inet account on the Disk Agent client. Run the following command on the Data Protector client where Disk Agent will be running:

```
omniinetpasswd -add User@Domain [Password]
```

Requirements

- You have to map the shared drives using the Backup wizard.
- Use the Windows GUI, because browsing of Windows systems is not supported in the UNIX GUI.

Limitations

- Backing up shared disks does not back up all file attributes. Only what is visible on the sharing host can be backed up. The data can be restored but some file/directory attributes may be missing.

- Backing up writers that store their data on network shared volumes using the VSS functionality is not supported. Additionally backing up network shares or remote network folders with Disk Agent and Use Shadow Copy option enabled is also not supported on Windows Server 2012.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**.
3. Right-click the type of item that you want to back up (for example, **Filesystem**) and then click **Add Backup**.
4. In the **Create New Backup** dialog box, select one of the available templates and then click **OK** to open the wizard.
5. In the Source property page, select **Network share backup** in the drop-down list (available if the GUI is running on Windows systems).
6. Click **Map Network Share** to open the **Browse Network Shares** window.
7. In the Client system drop-down list, select the client system with the Disk Agent that you will use for your backup.
8. In the Shared directories box, select or specify the shared disk and then click **OK**. If you want to select more disks, use **Apply**.
9. In the Source property page, select or specify the shared filesystems that you want to back up. Click **Next**.
10. In the Destination property page, select the device(s) you will use for your backup.

You can also specify whether you want to create additional copies (mirrors) of the backup during the backup session. Specify the desired number of mirrors by clicking the **Add mirror** and **Remove mirror** buttons. Select separate devices for the backup and for each mirror. It is not possible to mirror objects backed up using the ZDB to disk or NDMP backup functionality.

Tip: If the backup is load balanced, you can set the order in which Data Protector will use the devices by right-clicking a selected device and clicking **Order devices**.

Click **Next**.

11. In the Options property page, you can set the backup options. Backup options are available according to the type of data being backed up. For example, all backup options available for a filesystem backup are not available for a disk image backup.
On Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012, perform additional steps:
 - a. Under the Backup Specification Options, click the **Advanced** button.
 - b. In the Backup Options dialog box, under Ownership, enter the information about the user account with permissions to access the shared disk that will be backed up.
 - c. Click **OK**.
12. Click **Next**.
13. In the Schedule property page, specify the dates and times that you want your backups performed (optional). Click **Next**.
14. In the Backup summary page, review the summary of the backup specification. It is

recommended that you first save the backup specification and then start a preview. Click **Next**.

15. At the end of the Backup wizard, you can save, start, or preview the configured backup. The following happens:
 - If you save the configured backup, it appears in the Backup context of the Scoping Pane as a new backup specification. You can later preview or start the saved backup without any modifications, or you can modify it and then preview or start it.
 - If you start or preview the configured backup, the Session Information message displays the status of your backup.

One Disk Agent is started for each disk you back up. This may reduce your backup performance if you start too many backups at the same time.

Selecting Only Specific Files (Matching) for Backup

By using wildcard characters, you can back up files matching specific criteria.

Note: This functionality is not supported with Data Protector NDMP server integration.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Select the backup specification with the target object.
4. Click the **Backup Object Summary** tab.
5. In the Backup Object Summary page, right-click a backup object and then click **Properties**.
6. Click the **Trees/Filters** tab and then click the **Filter** button.
7. In the Onlys textbox, enter the criteria you want to use to back up only specific files and then click the **Add** button.
Repeat this step if you want to use more criteria.
8. Click **OK**.

Skipping Files for Backup

By using wildcard characters, you can skip files matching specific criteria from being backed up.

Note: Skipping files for backup is not supported with Data Protector NDMP server integration.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of

backup specification (for example, **Filesystem**). All saved backup specifications are displayed.

3. Select the backup specification with the target object.
4. Click the **Backup Object Summary** tab.
5. In the Backup Summary page, right-click a backup object and then click **Properties**.
6. Click the **Trees/Filters** tab and then click the **Filter** button.
7. In the Skips textbox, enter the criteria you want to use to skip some files (like *.tmp) and then click the **Add** button.

Repeat this step if you want to use more criteria.

8. Click **OK**.

Selecting the Location for the Shortcut for Starting a Backup

You can create a shortcut of the selected backup specification on the disk that you can later use to run the backup without using the Data Protector GUI. Double-clicking it opens the command prompt and runs the `omnib` command for the selected backup specification.

Limitations

- Shortcut for starting a backup is supported only on Windows systems.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of backup specification (for example, **Filesystem**).
3. Right-click the selected backup specification and click **Select the Location for the Shortcut**. The Save As dialog box appears.
4. Enter the name and select the location for the shortcut, then click **Save**.

The shortcut for starting a selected backup appears at the selected location on the disk.

Backing Up Using Multiple Disk Agents

When you back up large objects, you can speed up your backup by using multiple Disk Agents.

The following may provide additional information:

- In the backup specification, you have to manually define which directories/files will be backed up using a new Disk Agent. You should take care to avoid overlapping the same data.
- If more than one Disk Agent is concurrently accessing the same disk, the performance of retrieving data from the disk will drop. This can be different when using disk arrays.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**.
3. Right-click the type of item that you want to back up (for example, **Filesystem**) and then click **Add Backup**.
4. In the Create New Backup dialog box, select one of the available templates and then click **OK** to open the wizard.
5. In the Source property page, do not make a selection of directories/files that are located on the same logical disk or mountpoint if you want them to be backed up using multiple Disk Agents. However, you can select directories/files to be backed up using one Disk Agent. Click **Next**.
6. In the Destination property page, select the device(s) you will use for your backup. Click **Next**.
You can also specify whether you want to create additional copies (mirrors) of your backup during the backup session. Specify how many mirrors you want to create and which device(s) will be used for this purpose by clicking **Add mirror** and **Remove mirror**. The devices used for creating object mirrors must not be the same as the devices you use for backup. Object mirroring is not supported for ZDB to disk and for NDMP backup.
7. In the Options property page, specify further options as desired and click **Next**.
8. In the Schedule property page, specify the dates and times when you want your backups performed (optional). Click **Next**.
9. In the Backup summary page, click **Manual add**.
10. In the Select Backup Object dialog box, select the type of object to be backed up (for example, **Windows filesystem**). Click **Next**.
11. In the General Selection dialog box, select the client system and the mountpoint to be backed up. It is also necessary to enter a description. Click **Next**.
12. In the Trees/Filters Selection dialog box, specify the directories/files to be backed up or excluded from backup. What you select here will be backed up using one Disk Agent. Click **Next**.
13. In the General, Advanced, and Windows Specific Object Options dialog boxes, specify further options as desired and click **Next** and in the last one **Finish**.
14. Repeat steps 9-13 for directories/files on the mountpoint to be backed up using another Disk Agent.
15. In the Backup summary page, review the summary of the backup specification and then click **Next**.
16. At the end of the Backup wizard, you can save, start, or preview the configured backup.

Handling of Small Reoccurring Backups

When you need to perform reoccurring backups of numerous small objects, you need to run numerous backup sessions. During each backup session, media are loaded and unloaded in the drive. Not only is such backup slow, but it also causes media to deteriorate. To use media more economically and save time, it is recommended to create a file library device and use it to perform small reoccurring backups to disk instead of tape. You can then use the object copy functionality to move the data from the disk to a tape medium.

Using this method, a backup will be performed faster and media will be used more economically because they will be loaded and unloaded only once, during the object copy session.

To perform frequent backups of numerous small objects, perform the following tasks:

1. Configure a file library device. Set the block size of each writer to the block size of the device that will be used in the second stage.
2. Create one backup specification for all small objects. Use the file device created in the first step for the backup.
3. Perform or schedule the backup.
4. Use the object copy functionality to move the backed up data to tape.

Disk Image Backup

You can perform disk image backup on UNIX and Windows platforms.

A disk image backup of a disk is a high-speed backup where Data Protector backs up the disks, disk partitions, or logical volumes without tracking the file and directory structure stored on these data sources. Data Protector stores the disk image structure at a character level.

You can perform either a disk image backup of specific disk sections or a complete disk.

Note: On Windows systems, disk image backup is performed by using VSS writers. This ensures that the volume remains unlocked during the backup and can be accessed by other applications. This is especially important when backing up System volume. The VSS backup of disk images is enabled by default. To customize the VSS disk image backup, use the following omnirc options: OB2_VSS_RAW_BACKUP, OB2_VSS_RAW_BACKUP_ALLOW_FALLBACK, and OB2_VSS_SNAPSHOT_TIMEOUT.

When to use a disk image backup?

- When you have many small files and a high backup speed is required.
- When a full disk backup is needed, for example, for disaster recovery or before a major software update. On Windows systems, disk image backup can be used when preparing to the EADR and OBDR.
- When a direct disk-to-disk connection is not possible and you want to duplicate a filesystem to another disk. The latter must be identical to the original disk.

How to specify a disk image section?

On UNIX systems

- To specify a disk image section, use the following format: `/dev/rdisk/Filename`, for example: `/dev/rdisk/c2t0do`
- To specify a raw logical volume section, use the following format: `/dev/vgNumber/rlvolNumber`, for example: `/dev/vg01/rlvol1`

On Windows systems

You can specify a disk image section in two ways: the first way selects a particular volume, and the second way selects an entire disk. In the case of zero downtime backup, use the second way:

- `\\.\DriveLetter`, for example: `\\.\E:`

Note: When a drive letter is specified for the volume name, the volume is not being locked during the backup. A volume that is not mounted or mounted as an NTFS folder cannot be used for disk image backup.

- `\\.\PHYSICALDRIVE#`, where # is the current number of the disk you want to back up. For example: `\\.\PHYSICALDRIVE3`

Where to find a disk image section?

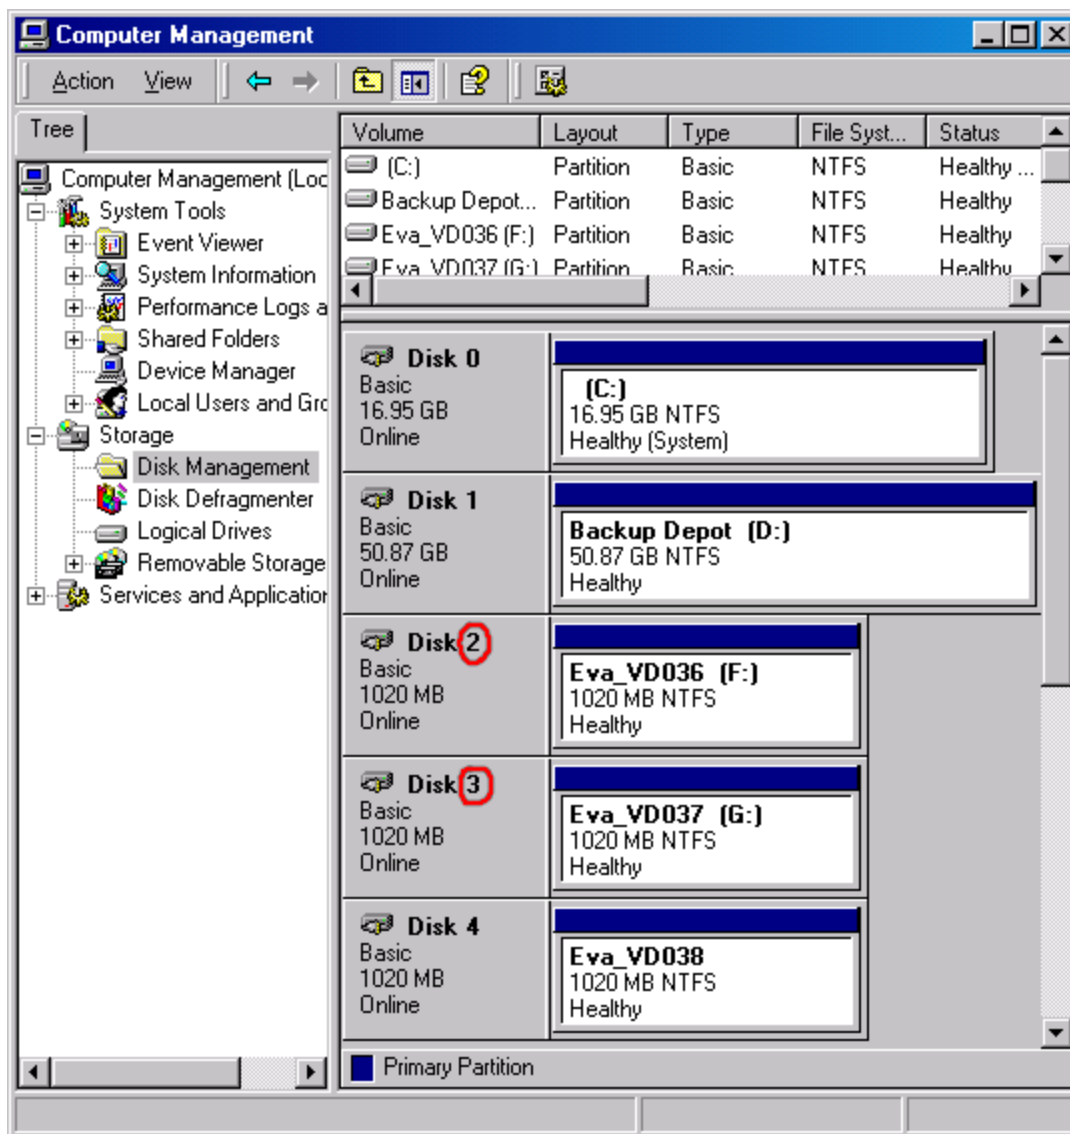
On UNIX systems

The disk image sections are usually listed in the `/dev/rdisk` directory. Raw logical volumes can be found in `/dev/vgNumber`. On HP-UX systems, raw logical volumes can be found in `/dev/vgNumber`. The first letter of a raw logical volume is `r`, such as `/dev/vg01/r1vol2`.

On Windows systems

You can find the current numbers of your disks (as well as the drive letters) by clicking **Administrative Tools** from the Control Panel and then **Computer Management, Storage, Disk Management**.

The numbers representing disks (physical drive number) on Windows system



Note: On Windows systems, the numbers representing disks can change if the system is restarted.

Client Backup with Disk Discovery

For a client backup with disk discovery, you specify a client as a data source. If another disk is mounted later, it will be included in the backup. In contrast to a filesystem backup, where you have to specify any newly added disk or mounted filesystem that is not yet specified in the backup specification, this is unnecessary if you use disk discovery.

Data Protector contacts the client at backup time and finds all filesystems on the disks that are attached to that system. Each detected filesystem (also CONFIGURATION on Windows systems) is then backed up as a regular filesystem. The description for each filesystem object is generated and the filesystem mountpoint is appended to the description of the client backup.

When backing up using disk discovery, Data Protector only backs up real disks. Therefore, on UNIX systems, Data Protector does not discover NFS, CD-mounted filesystems, and removable mountpoints. Also on Windows systems, Data Protector does not discover CDs and drives with removable media.

When to use disk discovery

This backup type is particularly useful in dynamic environments, where configurations change rapidly. It is recommended under the following conditions:

- If you back up workstations with relatively small disks that are frequently mounted or dismounted.
- If you would like to back up the data following a mountpoint into one directory, regardless of how many filesystems are mounted. For example, /home/data, where /home/data/disk1 and /home/data/newdisk/disk2 can be mounted or dismounted frequently and independently of each other.
- If you back up a whole system to prepare for disaster recovery.

Backup specification

When creating a backup specification that will define a disk discovery backup, click the check box next to the client system name and not next to the disks (volumes) of the system. Once you have selected the client system, you can check the configured backup type in the Backup Object Summary property page. Under the Type label, you should see *Client System*.

Web Server Backup

To back up a web server, use the standard backup procedure for backing up files, directories, and clients. Additionally, you need to consider the following:

- When performing a client backup, Data Protector backs up the whole web server, but not the data stored on other clients/servers. To back up data on other clients/servers, you need to select them for backup, also.
- When performing a filesystem backup, you need to know where all the files and directories of the web server and its respective clients are located. Always include web configuration files and root directories.
- Data Protector backs up all files in a static state. If files are changed during the backup, the changes are not backed up.

In case a database, such as Oracle or Informix Server, is included on a web server, use the backup procedure specific to the database.

Enabling Wake ONLAN Support

If your Windows systems support Wake ONLAN, you can use the Data Protector Wake ONLAN support.

When a Backup Session Manager fails to connect to a client that is configured to use Wake ONLAN support, it sends a wake-up request according to the Wake ONLAN protocol, and retries to connect to

the client. This enables the full use of power-saving features of desktop systems, which would otherwise interfere with the backup process.

You can enable Wake ONLAN support for computers equipped with a Wake ONLAN-compatible LAN interface, such as the HPE NightDIRECTOR series. The Wake ONLAN (WOL) option is available in the BIOS setup.

When you install a Disk Agent on a Windows client and add it to a cell, the client's MAC address is automatically discovered. You can also manually change the MAC address.

Steps

1. In the Context List, click **Clients**.
2. In the Scoping Pane, browse for the desired client, right-click it, and click **Properties**.
3. Click the **Advanced** tab.
4. Select the Enable Magic Packet option. If needed, change the MAC Address.
5. Click **Apply**.

About Backup Templates

Data Protector backup templates can help you simplify the handling of (many) backup specifications and related options. A template has a set of clearly specified options for a backup specification, which you can use as a base for creating and modifying backup specifications.

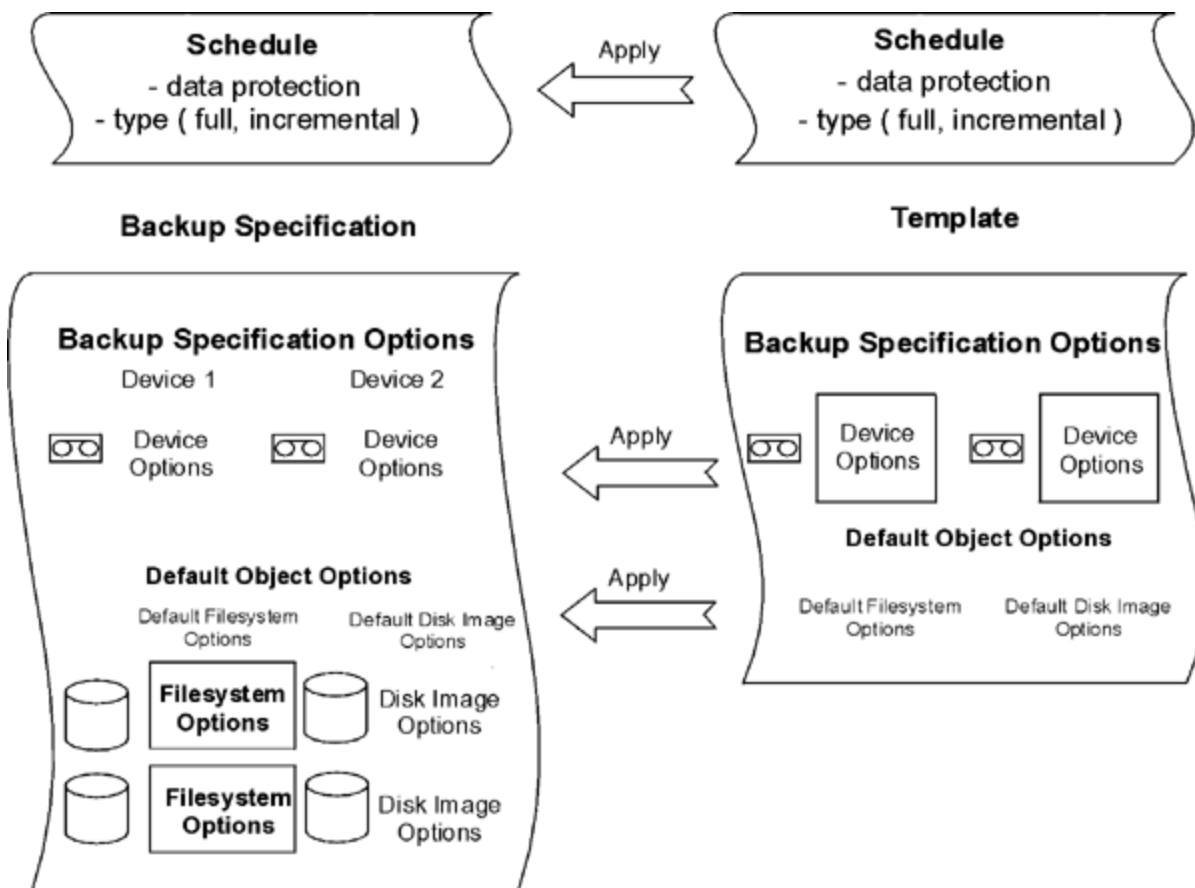
The purpose of a template is to configure multiple backup specifications with different objects that are used in the same way (common option setting for particular areas like device options or/and filesystem options).

Data Protector offers you default templates for different types of data (Filesystem, Exchange, and so on) without specifying objects, devices, options, and a schedule. In blank backup templates, such as Blank Filesystem Backup, Blank Informix Backup, and so on, there are no objects or devices selected. Backup specification options and object options have Data Protector default values, and there is no backup schedule.

Templates are created and modified in a way similar to backups, except that elements, such as objects, are not selected within the backup template. A template can be applied later to existing backup specifications or it can be used when creating a new backup. If you later change the template, you have to apply it again if you want the changes to take effect.

Tip: Moving the cursor above a template displays a pop-up window with a description of the template.

Backup options scheme



Creating a New Backup Template

You can create a new backup template with special settings for the environment with special needs.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Templates**.
3. Right-click the type of template that you want to create (for example, **Filesystem**) and then click **Add Template** to open the wizard.
4. Follow the wizard and decide on the backup device that you want to use, backup options you want to set, as well as on scheduling.

The new template is available when creating a new backup specification or when applying a template to one or several backup specifications.

Modifying a Backup Template

You can modify a backup template. If you want your backup specification to change according to the template, you have to reapply it because the backup specifications are not automatically updated.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Templates** and then the type of the template that you want to modify (for example, **Filesystem**). All saved templates of that type are displayed.
3. Right-click the template that you want to modify and then click **Properties**.
4. In the template's property pages, modify the template that you have selected and then click **Apply**.

After you have modified your backup template, you can apply it to a backup specification or use it for creating a new backup specification.

Copying a Backup Template

You can copy a backup template.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Templates** and then expand the appropriate type of backup template (for example, **Filesystem**). All saved backup templates are displayed.
3. In the Results Area, right-click the template that you want to copy and then click **Copy As**. The Copy Backup As dialog box opens.
4. In the Name text box, enter a name for the copied template. Optionally, from the Group drop-down list, select a different group for your copied template.
5. Click **OK**.

The copied backup template is displayed in the Scoping Pane and in the Results Area.

Deleting a Backup Template

You can delete a backup template.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Templates** and then expand the appropriate type of backup template (for example, **Filesystem**). All saved backup templates are displayed.
3. Right-click the template that you want to delete and then click **Delete**. Confirm your choice.

The backup template is removed.

Applying a Backup Template to a Backup Specification

You can apply a template to one or several backup specifications. In this case, you can select which option groups should be applied.

Note: If you apply a backup template to an existing backup specification and select the Filesystem and/or Schedule options, the protection settings from the template will replace the previous data protection settings in the respective parts of the backup specification.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**.
3. Right-click a saved backup specification and then click **Apply Template**.
4. In the **Apply Template** dialog box, select the template that you want to apply to the backup specification.

Tip: You can deselect some of the template's options (**Trees**, **Backup options**, **Device**, and so on), so that they will not be applied to the selected backup specification.

Note: To apply a template to an integration backup specification, the backup specification you would like to apply should not be opened in the Results Area. If you first click on the backup specification to open it, and then try to apply the template to this backup specification, the **Apply Template** option will not be available.

5. Click **OK** to apply the template to the backup specification.


Once you have applied the template options, you can still modify your backup specification and change any setting.

About Backup Options

Data Protector offers a comprehensive set of backup options that let you fine tune a backup. All these options have default values (selected or not selected) that are appropriate in most cases.

The availability of backup options depends on the type of data being backed up. For example, not all backup options available for a filesystem backup are available for a disk image backup. Common and specific application options in the options property page for Exchange, SQL, and so on, are described in the context-sensitive Help for a specific backup type.

Backup options scheme



Available backup options

The following set of options is available when backing up the data:

Backup specification options

These options apply to the entire backup specification, regardless of the type of the backup objects.

Filesystem options

These options apply to each object of a filesystem backup. You can also change options for specific objects. Specific object settings override default settings.

Disk image options

These options apply to each object of a disk image backup. You can also change options for a specific object. Specific object settings override default settings.

Device options

These options define the behavior of backup devices. If you do not set the device options, the values are read from the device definition.

Schedule options

For each individual or periodic scheduled backup, you can specify the backup type (full or incremental; some other backup types are available for specific integrations), network load, and data protection. With ZDB, you can select ZDB to disk+tape or ZDB to disk (if instant recovery is enabled).

Data protection that is specified in the Schedule Backup dialog overrides protection settings anywhere else in the backup specification.

Most Frequently Used Options

The following is the list of options that are usually modified according to specific backup policies.

- [Data protection](#)
- [Catalog protection](#)
- [Logging](#)
- [Load balancing](#)
- [Ownership](#)

Data protection: How long data is kept on the media

Configuring protection policies is extremely important for the safety of the data and for successful management of your environment. You have to specify how long your backed up data is kept on the medium based on your company data protection policies. For example, you may decide that data is out of date after three weeks and can be overwritten with a subsequent backup.

You can specify data protection in different places. Different combinations are available, depending on whether you are running an interactive backup, starting a saved backup specification, or scheduling a backup. The default value is Permanent.

Interactive backups

When configuring an interactive backup, you can change the default data protection for the entire backup. Additionally, you can specify different data protection periods for individual backup objects. The protection that is specified on the backup object level overrides the default protection setting.

Backups using a saved backup specification

When starting saved backups using the GUI, the data protection is applied as described for interactive backups.

When starting saved backups using the CLI, you can also specify data protection. This will override all data protection settings in the backup specification.

Scheduled backups

You can specify a different period of protection for each individual or periodic scheduled backup. The data protection specified in the Schedule Backup dialog overrides all other data protection settings in the backup specification. If you leave the default protection, data protection is applied as described for interactive backups.

Catalog protection: How long data is kept in the IDB

You can set catalog protection and data protection independently. When the data protection ends and a medium is overwritten, the catalogs for the objects are removed regardless of the catalog protection setting.

Catalog protection, together with logging level, has a big impact on growth of the IDB, convenience of browsing data for restore, and backup performance. It is important that you define a catalog protection policy that is appropriate to your environment. Catalog protection has no effect if the logging level is set to No Log.

If catalog protection is permanent, the information in the IDB is removed only when media are exported or deleted. In this case, the size of the IDB grows linearly until the data protection period is reached even if the number of files in the cell does not change.

The default value is **Same as data protection**. This means that you can browse and select files or directories as long as the media are available for restore.

Due to operating system limitations, the latest protection date that can be set is Jan 18th, 2038.

Expired catalog protection

Once the catalog protection expires, the information is not immediately removed from the IDB. Data Protector removes it automatically once per day. Since the information in the IDB is organized on a per-medium basis, it is removed only when the catalog protection expires for all objects on the medium.

When catalog protection expires, you are still able to restore, but you must specify filenames manually.

Catalog protection and backup

Catalog protection settings do not have any impact on backup performance.

Catalog protection and restore

When catalog protection expires, data is restored as if it were backed up using the No Log option.

Logging: Changing details about data stored in the IDB

Data Protector logging level defines the amount of detail on files and directories that is written to the IDB during backup. Four logging levels are available:

- Log All
- Log Files
- Log Directories
- No Log

HPE recommends to use different logging levels in the same cell. A cell often consists of some mail (or similar) server that generates a large number of files on a daily basis, database servers that store all information in a handful of files, and some user workstations. Since the dynamics of these systems are rather different, it is difficult to prescribe one setting that suits them all. HPE recommends to create several backup specifications with the following logging level settings:

- For e-mail servers, use the Log Directories option.
- For database servers, use the No Log option since browsing of individual files makes no sense in this case.
- For workstations, use the Log Files option so that you can search for and restore different versions of the files.
- The Log All option allows to view the file attributes such as modification time and ACLs.

Logging level and backup speed

The backup speed is approximately the same regardless of the logging level chosen.

Logging level and browsing for restore

Changing the level of stored information affects your ability to browse files using the Data Protector GUI during a restore. If the **No Log** option is set, browsing is not possible; if the **Log Directories** option is set, browsing of directories is possible; if the **Log Files** option is set, full browsing is possible but file attributes (size, creation and modification dates, and so on) are not displayed.

If you know the names of the files you want to restore, you can always manually specify them instead of browsing for them, regardless of the effective logging level.

Logging level and restore speed

The restore speed is approximately the same when the corresponding backup session was run with either **Log All**, **Log Directories**, or **Log Files** logging level.

If backup session was run with **No Log** logging level, the restore speed may reduce when restoring single files. In this case, Data Protector has to read all data from the beginning of the object before finding a file to be restored.

In case of a full system restore, the entire backup object is read anyway, so the logging level does not play an important role.

Load balancing: Balancing the usage of backup devices

Use the Load balancing option when you want to back up a large number of objects to a number of available devices and you want Data Protector to keep all devices busy all the time. You should use Load balancing to minimize the impact of unavailable devices on the backup.

Clear the Load balancing option when you want to back up a small number of objects, when the objects are backed up on simple devices (such as DDS), when you want to manually select the devices to which objects will be backed up, or when you want to know which media objects will be backed up on.

Objects are assigned to an available device from the list of devices specified in the load balanced backup specification. The first device is started and the number of selected objects for is defined with

its concurrency. The next device is started and objects are selected until there are no more objects in the list or the maximum number of devices are running.

If a device becomes unavailable, only the objects that are being backed up to it at failure time are aborted. All objects backed up to the device before the failure time are actually backed up. If there are any other devices specified in the backup specification and the maximum number of devices has not been used, a new device will start. A device may become unavailable because it:

- failed during a backup
- stopped during a backup
- is in use by another session
- cannot be started at all

Objects to be backed up are reached according to the following criteria:

- Objects that reside on the client connected to the backup device have a higher priority.
- Objects are selected so that the number of Disk Agents per client is kept as low as possible.
- The size of objects does not play a role in assigning an object to a device.

The following rules should be considered when applying device options from a template:

- If the load balancing option is not selected in the template, then the devices are not used with the backup specification.
- If the load balancing option is selected in both the template and the backup specification, then the device options are applied.
- If the load balancing option is selected only in the template, then the device options will be applied only if the backup specification has no devices.

Ownership: Who is able to restore

Who is a backup session owner?

Each backup session and all the data backed up within it is assigned an owner. The owner can be the user who starts an interactive backup, the account under which the CRS process is running, or the user specified as the owner in the backup specification options.

If a user starts an existing backup specification without modifying it, the backup session is not considered as interactive.

If a modified backup specification is started by a user, the user is the owner unless the following is true:

- The user has the Switch Session Ownership user right.
- The backup session owner is explicitly defined in the backup specification, where the username, group or domain name, and the system name are specified.

If a backup is scheduled on a UNIX Cell Manager, the session owner is `root:sys` unless the above conditions are true.

If a backup is scheduled on a Windows Cell Manager, the session owner is the user specified during the installation, unless the above conditions are true.

Why change the backup owner?

You may want to change the backup owner in case the administrator configures and schedules a backup specification, and operators are allowed to run it, but they cannot modify or save it. If the Private backup option is set for all objects, the operators will not be able to restore anything, but can still manage backups and restart failed sessions.

If the backup configuration is changed and not saved, the backup is treated as an interactive backup and the owner is not changed. If you interactively start an incremental backup and you are not the owner of the full backup, you will get another full backup instead of the incremental.

Who can restore a private object?

Unless an object is marked as Public, only the following users can restore the object:

- Members of the Admin and Operator user group.
- The backup session owner who has the Start Restore user right. Other user rights may be required, such as Restore to Another Client.
- Users that have the See Private Objects user right.

The right to see and restore private objects can be granted to groups other than admin or operator as well.

Backup Specification Options

These options apply to the entire backup specification, regardless of the type of backup objects.

Basic option is **Load balancing**. By default, this option is enabled in the Create New Backup dialog. If you disabled it there, you can select it later in the Destination property page of the backup specification, in the Backup tab.

For more information on backup specification options, see the HPE Data Protector Help.

General backup specification options

- Description
- On client
- Post-exec
- Pre-exec
- Reconnect broken connections
- Ownership

Clustering backup specification options

Automatic session restart

If a failover of the cluster-aware Data Protector happens during backup, all running and pending backup sessions fail. The following options define the Data Protector behavior after the failover:

- Do not restart backups at failover
- Restart backup of all objects
- Restart backup of failed objects

Abort session and abort ID parameters

When some cluster-aware application other than Data Protector is running on other node than Data Protector and fails over to the node, where Data Protector is running, it is possible to control the load on this system. The following options used together with the `omniclus` command define the Data Protector behavior after the failover.

- Do not check elapsed session time
- Abort if less than
- Abort if more than
- Do not check abort ID
- Check against abort ID

EMC Symmetrix backup specification options

Client systems

- Application system
- Backup system

Mirror type

- TimeFinder
- Symmetrix Remote Data Facility
- Combined [SRDF + TimeFinder]

EMC Symmetrix split pre-exec and post-exec

- Split pre-exec
- Split post-exec

EMC Symmetrix options

- Run discovery of Symmetrix environment
- Re-establish links before backup
- Re-establish links after backup

HPE P9000 XP Disk Array Family backup specification options

Client systems

This set of options can only be modified after the backup specification has been saved.

- Application system
- Backup system

Mirror type

- HPE Business Copy P9000 XP
- HPE Continuous Access P9000 XP
- Combined (HPE Continuous Access P9000 XP + HPE Business Copy P9000 XP)
- MU number(s)

Replica management options

- Keep the replica after the backup
- Track the replica for instant recovery

At the start of the session

- Synchronize the disks if not already synchronized
- Abort the session if the mirror disks are not already synchronized

At the end of the session

- Prepare the next mirror disk for backup (resynchronize)

Application system options

- Dismount the filesystems on the application system
- Stop/quiesce the application command line
- Restart the application command line

Backup system options

- Use the same mountpoints as on the application system
- Root of the mount path on the backup system
- Add directories to the mount path
- Automatically dismount filesystems at destination mountpoints
- Leave the backup system enabled
- Enable the backup system in read/write mode

HPE P6000 EVA Disk Array Family backup specification options

Client systems

- Application system
- Backup system

Replication mode

- HPE Business Copy P6000 EVA
- HPE Continuous Access P6000 EVA + HPE Business Copy P6000 EVA

Replica handling during failover scenarios

- Follow direction of replication
- Maintain replica location

Snapshot management options

- Snapshot source
- Snapshot type
- Redundancy level
- Delay the tape backup by a maximum of n minutes if the snapclones are not fully created

Mirrorclone preparation / synchronization

- At the start of the session
- At the end of the session

Replica management options

- Keep the replica after the backup
- Number of replicas rotated

- Track the replica for instant recovery

Application system options

- Dismount the filesystems on the application system before replica generation
- Stop/quiesce the application command line
- Restart the application command line

Backup system options

- Use the same mountpoints as on the application system
- Root of the mount path on the backup system
- Add directories to the mount path
- Automatically dismount filesystems at destination mountpoints
- Leave the backup system enabled
- Enable the backup system in read/write mode

Filesystem Options

These options apply to each object of a filesystem backup.

The basic option is Protection.

There are several sets of **Advanced** filesystem options:

- Filesystem options
- Other filesystem options
- WinFS filesystem options

For more information on filesystem options, see the HPE Data Protector Help.

Filesystem options

- Catalog protection
- Post-exec
- Pre-exec
- Public
- Report level

Other filesystem options

- Backup files of size
- Backup POSIX hard links as files
- Backup POSIX hard links as files

- Copy full DR image to disk
- **Data security**
 - None
 - AES 256-bit
 - Encode
- Display statistical info
- Do not preserve access time attributes
- Enhanced incremental backup
- Use native Filesystem Change Log Provider if available
- Lock files during backup
- Logging

Data Protector logging level defines the amount of detail on the backed up files and directories that is written to the Internal Database during backup. Four logging levels are available:

- Log All
- Log Files
- Log Directories
- No Log
- Software compression

WinFS filesystem options

- Asynchronous reading
- Back up share information for directories
- Detect NTFS hardlinks
- Do not use archive attribute
- **Open files**
 - Number of retries
 - Time out
- Report open locked files as
- **MS Volume Shadow Copy Options**
 - Use Shadow Copy
 - Allow fallback

Disk Image Options

These options apply to all disk image objects that you select for backup.

The basic option is Protection.

For more information on disk image options, see the HPE Data Protector Help.

You can set the following **Advanced** disk image options:

- Catalog protection
- **Data security**
 - None
 - AES 256-bit
 - Encode
- Display statistical info
- Post-exec
- Pre-exec
- Public
- Report level
- Software compression

Device Options

You can set these options for the currently selected backup device in a specific backup specification. These options are a subset of the options you set while configuring a backup device or changing its properties. The options listed are valid for a particular backup specification. These options overwrite options set in the Devices & Media context, which apply for the respective device in general.

For more information on device options, see the HPE Data Protector Help.

Device properties - General

- CRC check
- Concurrency
- Drive-based encryption
- Media pool
- Prealloc list
- Rescan

Schedule Options

When scheduling a backup, you can set additional options. For each scheduled backup, you can specify the backup type (full or incremental; some other backup types are available for specific integrations), network load, and data protection. With ZDB, you can select ZDB to disk+tape or ZDB to disk (if instant recovery is enabled).

In advanced scheduler, you can also set priority, estimated duration, and end of recurrence.

Data protection specified in scheduler overrides protection settings anywhere else in the backup specification.

For more information on schedule options, see the HPE Data Protector Help.

Session options

- **Backup type**
 - Full
 - Incremental
- Network load
- Backup protection
- Priority
(advanced scheduler only)
- Estimated duration
(advanced scheduler only)

Split mirror/snapshot backup

(available with ZDB, but only in the case of ZDB to disk+tape or ZDB to disk (instant recovery enabled))

Setting Backup Options

You can set backup options while you are creating a new backup specification. In this case you get to the Options property page by following the wizard.

You can also set backup options for the backup specification that you have already configured and saved.

Note: Object options (filesystem and disk image options) can be set on two levels. First, you can set the *default object options* for all filesystems and for all disk image objects in the backup specification separately. Then you can set them differently for a *specific object*. These settings will override the defaults. For example, to compress data from all clients, except for one with a slow CPU, enable the **Compression** option when setting filesystem options. Then select the slow client and clear the **Compression** option for this client.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the backup specification for which you want to set backup options and click the **Options** tab.
4. In the Options page, set the options as desired. Click one of the **Advanced** buttons to set advanced options (according to the type of options that you want to set).
Besides backup specification options, you can set for example filesystem options, disk image options, and so on, depending on which type of data the backup specification is configured for.
5. Find the option that you need and then select or deselect it or enter the needed information.
6. Click **OK** and then click **Apply** to save the changes.

Specifying Data Protection

You can specify data protection when running interactive backups, starting saved backup specifications, or scheduling backups. The default value is Permanent.

Note: Due to operating system limitations, the latest protection date that can be set is Jan 18th, 2038.

Specifying data protection on the backup specification level

You can specify data protection when you are creating a new backup specification, or modifying an existing one.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the backup specification for which you want to set backup options and click the **Options** tab.
4. If you are backing up filesystems, specify the Protection option under Filesystem Options. For integrations, click **Advanced** under Common Application Options, and specify the **Protection** option in the Options tab.
5. Click **OK** and then click **Apply** to save the changes.

Specifying data protection for individual backup objects

You can specify a different protection period for filesystem and disk image objects.

You can specify data protection for individual objects when you are creating a new backup specification, or modifying an existing one.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the backup specification for which you want to set backup options and click the **Backup Object Summary** tab.
4. Right-click an object and click **Properties**.
5. Click the **Options** tab and specify the Protection option.
6. Click **OK** and then click **Apply** to save the changes.

Specifying data protection for scheduled backups

You can specify a different period of protection for each individual or periodic scheduled backup. The data protection specified in the Schedule Backup dialog overrides all other data protection settings in the backup specification.

You can specify data protection for a scheduled backup while scheduling a backup.

Specifying data protection using the CLI

When you run a backup using the CLI, you can also specify data protection. This will override all data protection settings in the backup specification.

Steps

1. Enter the following command:

```
omnib -datalist Name -protect ProtectionPeriod
```

where *Name* is the name of the backup specification.
For example, to run a backup with protection of two weeks, enter:

```
omnib -datalist MyBackup -protect weeks 2
```

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

Changing Options for a Specific Object

You can apply options to specific objects or manually change default options.

You can apply these options while you are creating a new backup specification. In this case you get to the Backup Object Summary page by following the wizard.

You can also apply options for backup specifications that you have already configured and saved.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the backup specification for which you want to apply options for a specific object and click the **Backup Object Summary** tab.
4. In the Backup Object Summary page, you can change object properties, the order of the objects, or mirror options.

To change object properties:

- a. Right-click the object and then click **Properties**.
- b. In the Object Properties dialog box, change the options for the specific object. Depending on the object selected, some of the following tabs are displayed: **General**, **Options**, **Other**, **Trees/Filters**, **WinFS Options**, **Options**, and **Database**. Click the appropriate tab to modify the options.
- c. Click **OK** to apply the changes.

To change the order of objects:

- a. Right-click an object and click **Move up** or **Move down**. Repeat the procedure until you have the desired order.
- b. Click **Apply**.

To change mirror options:

- a. Select an object and click **Change Mirror**.
- b. To change the device for a mirror, make sure the mirror is selected, highlight the mirror, and select a device from the **Device** drop-down list. You can also deselect a mirror for the selected backup object.

Changing Backup Device Options

You can set backup device options and the order of devices while you are creating a new backup specification. In this case you get to the Destination property page by following the wizard.

You can also set backup device options for the backup specification that you have already configured and saved.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the backup specification for which you want to change the device options, and click the **Destination** tab.
4. In the Destination property page, you can change device options.

- To change devices for a backup that is load balanced, deselect devices and select others.
 - To change devices for a backup that is not load balanced, select all devices that you want to use. Then click the **Backup Object Summary** tab, select the desired object and click **Change device**.
 - To change devices for a mirrored object, select all devices that you want to use for a specific mirror. Then click the **Backup Object Summary** tab, select the desired object and click **Change mirror**.
 - To change the order of devices (if the backup is load balanced), right-click any selected device and click **Order devices**.
 - To set other device properties, right-click any selected device and click **Properties**.
5. Specify the desired option(s) and click **OK**.
 6. Click **Apply**.

Setting Schedule Backup Options

When scheduling a backup, you can set further options. These options are only valid for scheduled backups and not for those started interactively. Data protection that is specified in the Schedule backup dialog overrides protection settings anywhere else in the backup specification.

You can set schedule backup options while you are creating a new backup specification for a scheduled backup. In this case you get to the Schedule property page by following the wizard.

You can also set schedule backup options when scheduling a backup in a backup specification that you have already configured and saved.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the appropriate backup specification and click the **Schedule** tab.
4. In the Schedule page, scroll through the calendar (clicking the single arrows) for the month in which you want to schedule your backup.
5. Right-click the date on which you want to run the backup, and click **Schedule** to display the Schedule Backup dialog box.
6. After specifying the Recurring and Time options, set the Session options as desired. Select a backup type (Full or Incremental; some other backup types are available for specific integrations), network load, and backup protection for the scheduled backups.

In the case of ZDB to disk+tape or ZDB to disk (instant recovery enabled), specify the Split mirror/snapshot backup option.

Click **OK**.

7. Repeat steps 4 to 6 for all backups that you want to schedule.
8. Click **Apply** to save the changes.

Tip: To modify schedule options for individual backups, perform steps 5 and 6. If you wish to change the time of backup as well, you need to remove the unwanted scheduled backup(s). To modify other options, this is not necessary, because the new backup will override the previous one.

About Pre- and Post-Exec Commands

What are pre- and post-exec commands?

Pre- and post-exec commands are used to perform additional actions before and/or after a backup or restore. Such actions include checking the number of files to back up, stopping some transaction processing, or shutting down an application before backup and restarting it afterwards. Pre- and post-exec commands are not supplied by Data Protector. You must write your own scripts to perform the required actions. They can be written as executables or batch files on Windows systems, or as shell scripts on UNIX systems. All the commands that run within the batch file must return an exit code 0 to signify success or greater than 0 to signify a failure.

There is a special behavior for backup objects of the `Client System` type (host backup). Even if pre- and post-exec commands are specified once, each is started once per each filesystem (or logical drive).

Configuring pre- and post-exec commands for backup

Pre- and post-exec commands can be configured on two levels:

Backup specification

The pre-exec command is executed before the backup session starts. The post-exec command is executed when the backup session stops. You specify these commands as backup options for the entire backup specification. By default, pre- and post-exec commands for the backup session are executed on the Cell Manager, but you can choose another system.

Backup object

The pre-exec command for a backup object starts before the object is backed up. The post-exec command for the backup object is executed after the object is backed up. You specify these commands as backup options for objects. Pre- and post-exec commands for an object are executed on the system where the Disk Agent that backs up the object is running.

How are pre- and post-exec commands run?

1. The pre-exec command for the entire backup specification starts and completes.
2. For each object in the backup specification:

- a. The pre-exec starts and completes.
 - b. The object is backed up.
 - c. The post-exec (for each object in the backup specification) starts and completes.
3. The post-exec command for the entire backup specification starts and completes.

Pre- and Post-Exec Commands for a Backup Specification

Pre- and post-exec commands can be written as executables or batch files on Windows systems, or as shell scripts on UNIX systems. All the commands that run within the batch file must return an exit code 0 to signify success or greater than 0 to signify a failure.

Pre- and Post-exec characteristics

- [Start-up and location of the commands](#)
- [Environment variables](#)
- [SMEXIT values](#)
- [Considerations for pre- and post-exec commands](#)

Start-up and location of the commands

Pre- and post-exec commands for a backup session are started before and after the backup session, respectively. They are executed on the Cell Manager by default, but you can choose another system.

Windows systems

Pre- and post-exec scripts are started by the Data Protector CRS when executed on the Cell Manager; and under the Data Protector Inet Service account (by default, the local system account) when executed remotely.

The scripts on the Cell Manager and on other systems must be located in the `Data_Protector_home\bin` directory and the user must specify only the filename or relative pathname.

Only `.bat`, `.exe`, and `.cmd` are supported extensions for pre- and post-exec commands. To run a script with unsupported extension (for example, `.vbs`), create a batch file that starts the script. Then configure Data Protector to run the batch file as a pre- or post-exec command, which then starts the script with the unsupported extension.

If you use quotes (") to specify a pathname, do not use the combination of backslash and quotes (\"). If you need to use a trailing backslash at the end of the pathname, use the double backslash (\\).

Note: The direct usage of `perl.exe` is prohibited.

UNIX systems

Pre- and post-exec scripts are started by the backup session owner, unless the backup session owner has Backup as root permission; the commands are then started under root.

On the Cell Manager or a remote UNIX client, 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`

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.

Environment variables

The following environment variables are set by Data Protector and can be used only in pre- and post-exec scripts for a backup specification on the Cell Manager and not if the command is executed on any other system.

For more information on environment variables, see the HPE Data Protector Help.

- DATALIST
- MODE
- OWNER
- PREVIEW
- RESTARTED
- SESSIONID
- SESSIONKEY
- SMEXIT

SMEXIT values

Value	Description
0	All files were successfully backed up.
10	All agents completed successfully, but not all files were backed up.
11	One or more agents failed or there was a database error.
12	None of the agents completed the operation; session was aborted by Data Protector.
13	Session was aborted by a user.

Considerations for pre- and post-exec commands

- On Windows systems, you have to specify the full filename, including the extension (for example, `.exe` or `.bat`).
- By specifying the script name, if you need to use single (on UNIX systems) or double (on Windows systems) quotes because of spaces in a path, never use the combination of both. Either use single quotes or double quotes. For example, "S'ilvousplat.bat" is wrong, S'ilvousplat.bat is allowed.
- Upon successful completion the exit value of a pre- or post-exec command must be zero.

- If a pre-exec command fails (returns a value less than 0), the status of the backup session is set to Failed and the session is aborted. A post-exec command is not executed.
- If a post-exec command fails (returns a value less than 0), the backup session status is set to Completed with errors.
- If a post-exec command returns a value less than 0 and the omnib command 11, the backup status is set to Completed with failures.
- Post-exec command is always executed, unless the session is aborted and the pre-exec command is not executed or not set. If the OB2FORCEPOSTEXEC omnirc option is set, the post-exec command is always executed.
- By default, pre- and post-exec commands are NOT executed during a preview of the backup. This behavior is defined by the ExecScriptOnPreview option in the global options file.
- Pre- and post-exec commands are handled in the same way as commands entered at the command prompt. However, special characters ?, *, ", |, <, and > are not allowed.
- The execution of pre- and post-exec commands is implemented using the pipe mechanism. All processes started in the pre- or post-exec functions have to finish before processing continues.
- While pre- or post-exec commands are running, the backup session cannot be aborted.
- Pre- and post-exec commands run in the background mode. Therefore, do not use any commands that require user interaction.
- Time-out is provided. Pre- and post-exec scripts have to send some output at least every 15 minutes by default or the scripts are aborted. You can change this time interval by modifying the ScriptOutputTimeout global option.
- Any output of the pre- and post-exec commands is written to the IDB and shown in the Data Protector GUI.
- You can disable the session pre- and post-exec command execution on the Cell Manager by setting the SmDisableScript global option to 1.
- You can disable the remote session pre- and post-exec command execution on any client by adding the line OB2REXECOFF=1 in the omnirc file.
- You can secure the client by specifying which Cell Managers are allowed to access the client. Only permitted Cell Managers will be able to execute pre- and post-exec commands on the client.
- On UNIX systems, the text written by a command to stdout is sent to the Session Manager and written to the database. A stderr is redirected to /dev/null. You can redirect it to stdout to get error messages logged to the database.

Specifying Pre- and Post-Exec Commands for a Backup Specification

To specify pre- and post-exec commands for a saved backup specification, perform the following steps:

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the backup specification for which you want to specify pre- and post-exec

- commands and click the **Options** tab.
- Under Backup Specification Options, click **Advanced**.
 - In the Backup Options dialog box, General tab, write the filename or pathname in the **Pre-exec** and/or **Post-exec** text box.
 - Click **OK** and then click **Apply** to save the changes.

Pre- and Post-exec Commands for a Specific Backup Object

Pre- and post-exec commands can be written as executables or batch files on Windows systems and shell scripts on UNIX systems. All the commands that run within the batch file must return an exit code 0 to signify success or greater than 0 to signify a failure.

Start-up and location of the commands

Pre- and post-exec commands for an object are executed before and after the backup of the object, respectively. You can specify these commands for all objects in a backup specification, or for each individual object. When backing up integrations, for example Oracle, the database is considered as an object, so the commands are executed before and after the database backup. These commands are executed on the system where the Disk Agent is running.

<p>Windows systems:</p>	<p>Pre- and post-exec scripts for a backup object are started under the Data Protector Inet Service account (by default, the local system account).</p> <p>The exec scripts for backup objects can reside in any directory on the system where the Disk Agent is running. However, for client backups, they must reside in <i>Data_Protector_home\bin</i>. If the scripts are located in the <i>Data_Protector_home\bin</i>, specify only the filename, otherwise the full pathname must be specified.</p> <p>Only .bat, .exe, and .cmd are supported extensions for pre- and post-exec commands. To run a script with unsupported extension (for example, .vbs), create a batch file that starts the script. Then configure Data Protector to run the batch file as a pre- or post-exec command, which then starts the script with the unsupported extension.</p> <p>If you use quotes (") to specify a pathname, do not use the combination of backslash and quotes (\"). If you need to use a trailing backslash at the end of the pathname, use the double backslash (\\).</p>
<p>UNIX systems:</p>	<p>Pre- and post-exec scripts are started by the backup session owner, unless the backup session owner has Backup as root permission; the commands are then started under root.</p> <p>The exec commands for backup objects can reside in any directory on the system where the Disk Agent is running. However, for client backups, they must reside in the default Data Protector administrative commands directory. If the commands are located in the the default administrative commands directory, specify only the filename, otherwise the full pathname must be specified.</p>

Environment variable

For the post-exec command Data Protector sets the BDACC environment variable.

Considerations for pre- and post-exec commands

- If you perform a client system (host) backup, the pre-exec script is started before the first filesystem backup of the particular system, while the post-exec script is started after the backup. In this case, BDACC cannot be exported because the variable is related to a single filesystem object, not to a whole client system (host).
- On Windows systems, you have to specify the full filename, including the extension (for example, .exe or .bat).
- By specifying the script name, if you need to use single (on UNIX systems) or double (on Windows systems) quotes because of spaces in a path, never use the combination of both. Either use single quotes or double quotes. For example, "S'ilvousplat.bat" is wrong, S'ilvousplat.bat is allowed.
- Upon successful completion the exit value of a pre- or post-exec command must be zero.
- If a pre-exec command fails (returns a non-zero value), the backup of this object is aborted. The status of the object is set to *Aborted* and Disk Agent stops processing but the post-exec command is executed (unless the post-exec command is dependent on the BDACC environment variable). No backup of the object exists.
- If a post-exec command fails (returns a non-zero value), the status of the object is set to *Aborted*. The backup of the object exists and data can be restored.
- If there is no executable script on the client or if the path of the script is wrong, Data Protector displays an error message that the script failed and the session is aborted.
- By default, pre- and post-exec commands are NOT executed during a preview of the backup. This behavior is defined by the `ExecScriptOnPreview` global option.
- Pre- and post-exec commands are handled in the same way as commands entered at the command prompt. However, special characters `?`, `*`, `"`, `|`, `<`, and `>` are not allowed.
- While the pre- or post-exec commands are running, the backup session cannot be aborted.
- The pre- and post-exec processes operate in the background mode. Therefore, do not use commands that require user interaction in the pre- and post-exec commands.
- Time-out is provided. Pre- and post-exec scripts have to send some output at least every 15 minutes by default or the scripts are aborted. You can change this time interval by modifying the `ScriptOutputTimeout` global option.
- Any output of the pre- and post-exec commands is written to the IDB and shown in the Data Protector graphical user interface.
- Pre- and post-exec commands should send some output to the Disk Agent at least every 120 minutes by default, or the backup of the object is aborted. You can change this time interval by modifying the `SmDaIdleTimeout` global option.
- On UNIX systems, the text written by a command to `stdout` is sent to the Session Manager and written to the database. A `stderr` is redirected to `/dev/null`. You can redirect it to `stdout` to get error messages logged to the database.

Security considerations

Pre- and post-exec commands are potentially dangerous because they enable numerous possible exploits if they are used by unauthorized personnel. If you are not using them, it is advisable to disable them. Also, if you are using pre- and post-exec scripts, keep them in a secured location to prevent unauthorized personnel from modifying them.

By setting the `StrictSecurityFlag` global option to `0x0100`, only users having the **Backup as root** or **Restore as root** permissions are allowed to run pre-/post-exec commands.

You can disable pre- and post-exec scripts for any backup object by adding the line `OB2OEXECOFF=1` in the `omnirc` file on the specific client. To disable the remote session pre- and post-exec command execution on any client, add the `OB2REXECOFF=1` into the `omnirc` file on the specific client.

You can secure the client by specifying which Cell Managers are allowed to access the client. Only permitted Cell Managers will be able to execute pre- and post-exec commands on the client.

Specifying Pre- and Post-Exec Commands for Backup Objects

Specifying pre- and post-exec commands for all objects

To specify pre- and post-exec commands for all objects in a saved backup specification, perform the following steps:

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the backup specification for which you want to specify pre- and post-exec commands and click the **Options** tab.
4. Under Filesystem Options (Disk Image Options in a saved backup specification for disk image backup), click **Advanced**.
5. In the Filesystem Options (Disk Image Options for disk image backup) dialog box, Options tab, write the filename or pathname in the **Pre-exec** and/or **Post-exec** text box.
6. Click **OK** and then click **Apply** to save the changes.

Specifying pre- and post-exec commands for individual objects

To specify pre- and post-exec commands only for individual objects in a saved backup specification, perform the following steps:

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the backup specification for which you want to specify pre- and post-exec commands and click the **Backup Object Summary** tab.

4. Right-click an object and click **Properties**.
5. In the Object Properties dialog box, click the **Options** tab.
6. Write the filename or pathname in the **Pre-exec** and/or **Post-exec** text box.
7. Click **OK** and then click **Apply** to save the changes.

Specifying pre- and post-exec commands for integrations

When backing up integrations, for example Oracle, the database is considered as an object, so the commands are executed before and after the database backup. The commands are executed on the application client.

To specify pre- and post-exec commands for an integration in a saved backup specification, perform the following steps:

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of backup specification (for example, **Oracle Server**). All saved backup specifications are displayed.
3. Double-click the backup specification for which you want to specify pre- and post-exec commands and click the **Options** tab.
4. Under Application Specific Options, click **Advanced**.
5. In the Application Specific Options dialog box, write the filename or pathname in the **Pre-exec** and/or **Post-exec** text box.
6. Click **OK** and then click **Apply** to save the changes.

About Backup Schedule

You can configure unattended backups by scheduling backup sessions to execute at specific times.

- When the scheduled backup is started, Data Protector tries to allocate all needed resources, such as licenses, devices, and access to the IDB. If one of the needed resources is not available, the session is queued while Data Protector is trying to get the needed resources for the queued session every minute until the time-out period is reached. The time out can be modified by changing the `SmWaitForDevice` global option.

When Data Protector gets the resources, the queued sessions are started. The queued sessions may not be started in the order they are displayed.

- To prevent Cell Manager overload, the number of concurrent backup sessions in a cell is conservatively limited by default. If more sessions than the effective limit are scheduled at the same time, and the effective limit is lower than the maximum configurable limit, the overflow sessions are queued. The limit can be modified using the `MaxBSessions` global option.

On the other hand, the concurrently invoked sessions that fall above the maximum configurable limit are not started, and relevant errors are logged into the Data Protector Event Log.

Scheduling and priority (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.

Note: The Schedule priority and Pause lower priority jobs options are not supported in the CMMDB environment.

- 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.
- The Advanced Scheduler maintains an internal job queue to manage the priorities. If multiple jobs are sharing the same file library as the target device, the Advanced Scheduler will only dispatch one job at a time, and will only dispatch the next one when the prior one completes and frees up the device. The highest priority job will be dispatched first. If multiple jobs have the same priority, then the one with the earliest schedule time will be dispatched first. If they have the same priority and schedule time, then one of them will be picked up randomly and dispatched.
- The Advanced Scheduler does not track jobs started by the basic scheduler, including ones that might use the same file library device. In this case, the Monitor context will display all of the sessions from the basic scheduler, and only one session from Advanced Scheduler. For VMware backup sessions, if the basic scheduler already has a backup session running, and the Advanced Scheduler also schedules the same VMware backup session, the Monitor context will only display the basic scheduler session.

Scheduling and priority example

The following is an example of how Advanced Scheduler handles the backup sessions based on priority and pausing.

There are three sessions to be scheduled, where:

- Job1 has a priority of 2000 with the **Pause lower priority jobs** option enabled.
 - Job2 has a priority of 4000.
 - Job3 has a priority of 3000 with the **Pause lower priority jobs** option enabled.
1. Job2 is currently running.
 2. Schedule Job1 and Job3 for the same time. The Job1 session has the option to pause other sessions enabled. Thus, the Job2 session will be paused in favor of Job1.
 3. Once the Job1 session completes, the Job3 session runs.

The paused Job2 session will remain paused until it is able to run per the schedule and priority. This session may never get the opportunity to run, if there are other higher priority sessions.

Care should be taken when scheduling high priority jobs with the **Pause lower priority jobs** option enabled.

Scheduling options

For each scheduled backup, you can specify the backup type (full or incremental; some other backup types are available for specific integrations), network load, and data protection. In the Advanced Scheduler you can also specify the priority, estimated duration and recurrence pattern.

In the case of ZDB to disk+tape or ZDB to disk (if instant recovery is enabled), you can specify the **Split mirror/snapshot backup** option.

Each backup specification can be scheduled multiple times with different option values. Within one backup specification, you can schedule both ZDB-to-disk and ZDB-to-disk+tape sessions, and specify a different data protection period for each individual or periodic scheduled backup.

For each scheduled object operation, such as object consolidation, object copy and object verification, you can specify the priority, start times, time zones, estimated duration and recurrence patterns in the Advanced Scheduler.

Scheduling and different time zones

All schedules are displayed in the calendar in the time zone of the Cell Manager system. If you specified a backup or object operation session for a different time zone than that of the Cell Manager, the session will run at the specified time in the specified time zone.

Scheduling tips

- To simplify scheduling, Data Protector provides backup specifications for group clients. All clients configured in one backup specification are backed up at the same time in a single backup session.
- Make sure you have sufficient media and devices to run unattended backups smoothly.
- When applying a backup template, the schedule settings of the template override the schedule settings of the backup specification. After applying the template, you can still modify the backup specification and set a different schedule.
- You can modify the basic scheduler granularity by changing the `SchedulerGranularity` global option. By default, the granularity is 1 minute. Finer granularity enables you to execute backup specifications more frequently and helps you to avoid scheduling conflicts. This global option only works for the basic scheduler. Advanced scheduler granularity is set by recurrence pattern and is at least 1 minute or more.
- When Backup and Copy sessions are started, they require memory to be allocated as they are resource intensive, especially on the Media Agent servers. So, you need to ensure that multiple backup and copy sessions do not start at the same time. For example, if you need to start nine backup specifications at approximately 6 PM, you need to start the first three backups at 5.45 PM, the next three at 6 PM, and the last three backups at 6.15 PM. Instead of scheduling all the nine backup specifications to start at 6 PM.

Backing up during holidays (basic scheduler only)

You can set different holidays by editing the Holidays file that resides in the default Data Protector server configuration directory.

By default, Data Protector runs backups on holidays. If you want to change the default behavior, consider the following example. If the date January 1 is registered as a holiday, Data Protector will not back up on that date. If you have scheduled a full backup for January 1st and an incremental for January 2nd, Data Protector will skip running the full backup on January 1st but will run the incremental backup scheduled for January 2nd. The incremental backup will be based on the last full backup.

It is generally not recommended to skip backups on holidays.

Consider the following when editing or adding new entries in the `Holidays` file:

- The first number in each line indicates the consecutive day of the year. The value is ignored by Data Protector, but it must be set between 0 and 366. You can set it to 0 to indicate that the number does not correspond to the date that follows it.
- The date is specified as `Mmm d`, where `Mmm` is the three-letter abbreviation of the month and `d` is day of month as a number (for example, `Jan 1`). Note that the month must be specified in English, regardless of your locale.
- The description of the holiday is optional and is currently not used by Data Protector.

Regardless of the year specified at the top of the file, the holidays specified in the file are always used as-is and must be edited manually if the holidays do not occur on the same dates each year. If you are not using the **Holidays** option for the scheduler, you can remove or comment out the entries in the `Holidays` file to prevent confusion in case of accidental use of a `Holidays` file that is out of date or has not been customized for your country or company specific requirements.

Handling scheduling conflicts (basic scheduler only)

When scheduling periodic backups, it can happen that the chosen backup start time is already occupied by another scheduled backup in the same backup specification. In that case, Data Protector prompts you that there are scheduling conflicts, and asks if you wish to continue. If you click **Yes**, the new schedule will be applied where possible (on the days when the time slot is still free). If you click **No**, the new schedule will be discarded.

Scheduling a Backup on a Specific Date and Time

You can schedule your backup sessions to start automatically on a specific date at a specific time. You usually want to back up on specific dates when configuring exceptions to your regular periodic backups, for example, if you want to back up some data before a specific event.

You can schedule your backup while you are creating a new backup specification by following the wizard. Modifying the schedule of an existing backup specification can also be done as follows:

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the appropriate type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the appropriate backup specification and click the **Schedule** tab.
4. In the Schedule page, scroll through the calendar (clicking the single arrows) for the month in which you want to schedule your backup.

5. Right-click the date on which you want to run the backup, and click **Schedule** to display the Schedule Backup dialog box.
6. Under Time options, specify the starting time for your backup. Under Session options, select a backup type (Full or Incremental; some other backup types are available for specific integrations), network load, and backup protection for the scheduled backups.

In the case of ZDB to disk+tape or ZDB to disk (instant recovery enabled), specify the Split mirror/snapshot backup option.

Click **OK**.

7. Repeat steps 4 to 6 for all backups that you want to schedule.
8. Click **Apply** to save the changes.

If you schedule a backup in a time slot that is already occupied by a scheduled backup, the new scheduled backup overrides the previous one.

Tip: You can use **Reset** to remove all previous schedules.

Scheduling a Periodic Backup

Periodic backups are based on a time period after a specific date. For example, you may configure periodic backups so that a full backup is done on Sunday at 3:00 and repeated every two days. The next full backup would be at 3:00 the following Tuesday. Periodic backups simplify backup configuration for regularly scheduled backups.

You can schedule a periodic backup while you are creating a new backup specification by following the wizard, or you can modify the schedule of an existing backup specification, as is described in the following procedures:

Tip: You can use the **Reset** button to remove all previous schedules.

Using a predefined backup schedule

The predefined backup schedules can be used to simplify your configuration of filesystem backup specification. You can modify the schedules later.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then expand **Filesystem**. All saved backup specifications are displayed.
3. Double-click the appropriate backup specification and click the **Schedule** tab.
4. In the Schedule page, click **Predefined** to display the Choose Predefined Schedule dialog box.
5. Select a suitable backup schedule and click **OK**.
6. Click **Apply** to save the changes.

Configuring a recurring backup

You can schedule a backup so that it starts at a specific time and date on a set schedule. For example, you could schedule a full backup to take place every Friday at 21:00 for the next six months.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications**, and then expand the appropriate type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the appropriate backup specification and click the **Schedule** tab.
4. In the Schedule page, click **Add** to display the Schedule Backup dialog box.
5. Under Recurring, select **Daily**, **Weekly**, or **Monthly**.
6. Under Time options, select the time for the backup to start. To set the starting date, select the **Use starting** option and select a date.

Note: If you set the recurring to 2 or more (for example, every 2 weeks on Saturday) without setting the starting date, the first backup may not be scheduled on the first possible date that matches your selection (for example, it will be scheduled on the second Saturday) due to the Data Protector scheduling algorithm. Check the schedule in the Schedule property page.

7. Under Recurring options, select more precisely when the backups will start.
8. Set the Session options as desired. Select a backup type (Full or Incremental; some other backup types are available for specific integrations. For further details, click **Help** on the Schedule Backup dialog box.), network load, and backup protection for the scheduled backups.

In the case of ZDB to disk+tape or ZDB to disk (instant recovery enabled), specify the Split mirror/snapshot backup option.

Click **OK**.

9. Repeat steps 4 to 8 if you want to schedule another recurring backup.
10. Click **Apply** to save the changes.

If there are scheduling conflicts, Data Protector notifies you so that you can modify the schedule.

Running Consecutive Backups

You can start a backup after another one finished. For example, you can start a backup of the Oracle database after the filesystem backup finished.

Use the post-exec command in the first backup specification to start a consecutive backup.

Steps

1. Schedule the first backup.
2. Click the **Options** tab and click **Advanced** under **Backup Specification Options**.
3. In the Post-exec text box, enter the `omnib` command with the name of the backup specification

that you want to start after the first one is finished (for example, `omnib -datalist name_of_the_backup_specification`) and then click **OK**.

Tip: You can also specify your own script that checks the status of the first backup.

Resetting a Backup Schedule

When you reset your schedule, you clear all the schedule settings for the current year in the backup specification

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the backup specification for which you want to change the device options and then click the **Schedule** tab.
4. In the Schedule property page, click **Reset**.

All the previous schedules are removed.

Tip: You can undo your reset by clicking **Undo**.

Disabling and Enabling a Backup Schedule

By default, the schedule is enabled when added, but you can disable it, leaving the schedule settings intact for later use.

Disabling backup schedules does not influence currently running backup sessions.

Basic and advanced schedulers function independently of each other, including settings such as disabling and enabling schedules.

Steps (Basic Scheduler)

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the backup specification for which you want to disable or enable backup schedule and then click the **Schedule** tab.
4. In the Schedule property page, select the **Disable schedule** option to disable a backup schedule or deselect this option to enable it.
5. Click **Apply**.

Steps (Advanced Scheduler)

1. In the Context List select **Backup**.
2. In the **Actions** menu, click **Advanced Scheduler**.
3. In the Sessions pane, click on desired backup specification or object operation.
4. Click the schedule entry in the Calendar pane, or select the entry in the Schedules pane and click **Edit**.
5. In **Schedule** dialog box, clear the **Schedule enabled** option to disable the schedule, or select it to enable the schedule.
6. Click **Save**.

Disabling and Enabling Backups on Holidays

By default, Data Protector runs backups on holidays. You can change this behavior by selecting the **Holidays** option. The backup on holidays is not performed until you deselect this option.

Steps

1. In the Context List, click **Backup**.
2. In the Scoping Pane, expand **Backup Specifications** and then expand the type of backup specification (for example, **Filesystem**). All saved backup specifications are displayed.
3. Double-click the backup specification for which you want to disable or enable backup schedule during holidays and then click the **Schedule** tab.
4. In the Schedule property page, select the Holidays option to disable backing up during holidays or deselect this option to enable backing up during holidays. You can identify holidays from the Holidays file or as red dates on the Schedule Calendar.
5. Click **Apply**.

Customizing the Schedule Calendar

You can customize the appearance of the calendar that is used for scheduling various tasks, such as a backups, automated media copying, and report generation.

You can customize the calendar when scheduling one of the scheduled operations, or when reviewing the schedule. After you have opened the Schedule property page of the scheduled operation, do the following:

Steps

1. In the Schedule property page, right-click a month name and select the desired option from the pop-up menu.
2. Customize the calendar as desired and click **OK**.

About Backup Specification Groups

Data Protector lets you organize backup specifications into different groups. This is useful if, for example, you administer a large number of backup specifications and want to group them by common characteristics.

Grouping the backup specifications into meaningful groups can facilitate finding and maintaining single backup specifications. This also allows you to apply common options settings from a template to the entire group. For example, if you want to change the list of devices to all backup specifications in the group, you can selectively apply the device settings of a template.

Tip: You can apply common options settings (for example, for devices) from a template to a group of backup specifications. Select all the backup specifications within the group (click on the name of the group and then CTRL+A), right-click a target group, and then click **Apply Template**.

Note: Data Protector GUI can display the 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 load balanced or not). This size should not exceed 80 kB.

Example of backup specification groups

Backup specifications for a big corporation could be organized as follows:

User_files	This group contains backup specifications that perform weekly full backups for all users in each of the 10 departments.
SERVERS_DR	This group contains backup specifications for the company's servers to prepare for disaster recovery. Each time a new server is installed, a new backup specification is created and added to this group.
END_USER_ARCHIVE	This group is used to save backup specifications that are made per end user request. For example, end users who want to free up disk space have to archive their hard disks first.

Viewing Backup Specification Groups

Procedures in Data Protector Help assume that you use the default Backup View (By Type). You can change the view in order to see the backup specifications arranged by groups.

Steps

1. In the Context List, click **Backup**.
2. In the View menu, select **By Group**.

Creating a Backup Specification Group

You can create different backup specification groups using various criteria.

Steps

1. In the Context List, click **Backup**.
2. In the View menu, click **By Group**. The list of available backup groups appears under Backup Specification in the Scoping Pane.
3. Right-click the **Backup Specification** item and then click **Add Group**. The Add New Group dialog box displays.
4. In the Name text box, enter a name for your new group and then click **OK**.

The new backup specification group appears under the Backup Specification item. You can now add backup specifications into the appropriate groups.

Saving a Backup Specification into a Group

You can save a new backup specification into a specific group.

Steps

1. In the Context List, click **Backup**.
2. In the View menu, click **By Group**. The list of available backup groups appears under Backup Specifications in the Scoping Pane.
3. Expand **Backup Specifications**, right-click the group to which you want to add a backup specification, and then click **Add Backup** to open the Backup wizard.
4. Follow the wizard to create a new backup specification. In the final page (the Save, Start, or Preview page) of the wizard, click **Save As**. The Save Backup As dialog box appears.
5. In the Name text box, enter the name of the backup specification.
6. In the Group drop-down list, select the group to which you want to save the backup specification and then click **OK** to save the specification and exit the wizard. By default, the displayed backup group is the one that you right-clicked to start the wizard.

The saved backup specification appears under the selected group.

Moving a Backup Specification or Template Among Groups

You can move a backup specification or template from one backup group to another.

Steps

1. In the Context List, click **Backup**.
2. In the View menu, click **By Group**. The list of available backup groups appears under Backup Specifications and Templates in the Scoping Pane.
3. Expand **Backup Specifications** or **Templates** and the group that has the backup specification or template you want to move.
4. Right-click the backup specification or template that you want to move and then click **Change Group**. The Change Group dialog box appears.
The **Change Group** option is disabled if the backup specification properties are displayed.
5. In the Name drop-down list, select the group to which you want to move the backup specification or template and then click **OK**.

The backup specification or template is displayed under its new group.

Deleting a Backup Specification Group

You can delete a backup specification group that you do not need any more.

Steps

1. In the Context List, click **Backup**.
2. In the View menu, click **By Group**.
3. Expand the **Backup Specification** item and the **Templates** item. The list of available backup groups appears.
4. Expand the group that you want to delete.
A group that contains backup specifications and templates cannot be deleted. You must first delete or move the backup specifications and templates from the group.
5. Right-click the target group and then click **Delete Group**.

The target backup specification group has been removed.

About Windows Systems Backup

The backup procedure is the same as the standard backup procedure, however there are some Windows specific aspects.

Limitation

To run a VSS filesystem backup, your system must have at least one NTFS filesystem.

What is backed up?

A filesystem backup of a disk drive involves reading the directory structure, the contents of the files on the selected disk drive, as well as Windows-specific information about the files and directories.

Windows Server 2012

- Compressed files are backed up and restored compressed
- Encrypted files are backed up and restored encrypted

Windows-specific information

- Full Unicode file names
- FAT16, FAT32, VFAT, and NTFS attributes

Once a file has been backed up, its archive attribute is cleared. You can change this behavior by setting the Do not use archive attribute option among the Advanced filesystem backup options in the backup specification.

- NTFS alternate data streams
- NTFS security data
- Directory share information

If a directory is shared over a network, share information will be backed up by default. During the restore, share information will be restored by default and directory will be shared on the network after the restore. You can change this behavior by clearing the **Backup share information for directories** option.

What is not backed up?

In the backup specification, you can specify the list of files to be excluded or skipped by the backup (private exclude list). In addition to the private exclude list, Data Protector by default excludes the following:

Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012:

- The default Data Protector log files directory from a Windows client backup from a Windows client or Cell Manager (Windows Server 2008 only) backup.
- The default Data Protector temporary files directory from a Windows client backup from a Windows client or Cell Manager (Windows Server 2008 only) backup.
- The Internal Database directory from a Windows Cell Manager (Windows Server 2008 only) backup.
- The files specified in the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup.

Windows Server 2012

- Volumes formatted with the Resilient File System (ReFS)

Other Windows systems

- The default Data Protector log files directory from a Windows client backup.
- The default Data Protector temporary files directory from a Windows client backup.
- The files specified in the registry key `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup`.

For example, the Internal Database directory is excluded from the Cell Manager backup even though it was selected in the backup specification. This is because the IDB must be backed up in a special way to ensure data consistency.

NTFS 3.1 filesystem features

- The NTFS 3.1 filesystem supports reparse points
The volume mount points, Single Instance Storage (SIS), and directory junctions are based on the reparse point concept. These reparse points are selected as any other filesystem object.
- The NTFS 3.1 filesystem supports symbolic links, which were introduced with Windows Vista and Windows Server 2008 operating systems.
Data Protector handles symbolic links in the same way as NTFS reparse points.
- The NTFS 3.1 filesystem supports sparse files as an efficient way of reducing the amount of allocated disk space.
These files are backed up sparse to save tape space. Sparse files are backed up and restored as sparse to the NTFS 3.1 filesystem only.
- Some of the NTFS 3.1-specific features are controlled by system services that maintain their own data records. These data structures are backed up as a part of CONFIGURATION.
- The NTFS 3.1 filesystem supports the Object IDs that are backed up by Data Protector along with other alternate data streams.
- Encrypted files
The Microsoft-encrypted NTFS 3.1 files are backed up and restored encrypted, but their contents can only be properly viewed when they are decrypted.

Reparse points

Reparse points are plain filesystem objects with a unique tag attached, known as a reparse point ID. The NTFS 3.1 directories or files can contain a reparse point, which typically imitates the contents by directing to data from another location.

By default, when Data Protector encounters reparse points, the reparse point IDs are not followed. This is also known as backing up raw reparse points. It impacts the way you configure your backups:

- If you configure a backup using disk delivery, all data will be backed up once.
- If you back up filesystems or drives containing reparse points, you must ensure that the data pointed to by reparse points gets backed up. For example, the Windows directory junction reparse points are not followed, so the junctions have to be backed up separately. Exceptions are SIS reparse points.

The Single Instance Storage (SIS) service regularly checks the files on a disk. If the service detects several identical files, it replaces them with reparse points and stores the data in a common repository, reducing disk space usage.

Reparse points let you mount logical volumes as disk drives. Data Protector treats the mounted volumes as if they were ordinary drives, so they are visible as selectable objects for backup.

Sparse files

Sparse files contain many zero data sets — many more, for example, than compressed files. At backup time, Data Protector automatically skips zero-parts so that the media space on the backup device is allocated for non-zero parts only.

UNIX and Windows sparse files are not compatible.

Warnings when backing up system disks

There are certain files on the system disk that are always busy and cannot be opened by any application, including the Disk Agent. The contents of these files can only be backed up as a part of CONFIGURATION.

When these files are accessed by a filesystem backup, such as when the whole system disk is backed up, Data Protector fails to open them and reports warnings or errors.

While this behavior is correct from the filesystem backup point of view, it can create a manageability problem. Due to the large number of warnings that are always reported, it is likely that a failure of another file may be overlooked.

Exclude the files that are backed up through a CONFIGURATION backup from a filesystem backup to avoid warnings.

Note: When backing up an inactive system disk (for example in a dual-boot situation) the previously listed files are not a part of the currently active CONFIGURATION. These files can be backed up in a filesystem backup, and should not be excluded.

Configuration Backup (Windows)

The special data structures maintained by the Windows operating system are not treated as a part of the filesystem backup. Data Protector lets you back up a special data structure known as CONFIGURATION.

To perform a configuration backup, select the object CONFIGURATION or just parts of it when creating a filesystem backup specification. Event Logs, Profiles, and User Disk Quotas are always backed up if CONFIGURATION is selected in the Backup wizard.

On Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012, CONFIGURATION backup is performed using Microsoft Volume Shadow Copy Service.

Limitations

- Only one CONFIGURATION backup can run on a system at a time.
- Active Directory Service and SysVol should be backed up in pair.

Windows configuration objects

- Active Directory Service
- Certificate Server
- COM+ Class Registration Database (ComPlusDatabase)
- DFS
- DHCP
- DNS Server
- EISA Utility Partition
- Event Logs
- File Replication Service
- Internet Information Server (IIS)
- User Profiles (Documents and Settings)
- Windows Registry
- Removable Storage Management Database
- SystemRecoveryData
- SysVol
- Terminal Services Database
- User Disk Quotas (QuotaInformation)
- WINS server

CONFIGURATION differs between Windows systems.

For some objects, special points have to be considered. These are listed in the sections below.

Active Directory

When backing up the Active Directory service, the File Replication Service (FRS) and Distributed File System (DFS) also get backed up. All configuration information about replicated files and distributed files is stored in the Active Directory.

DFS

Data Protector backs up Windows Distributed File System (DFS) as part of one of the following:

- Windows Registry, if the DFS is configured in a standalone mode
- Windows Active Directory, if the DFS is configured in a domain mode

DHCP and WINS

When Data Protector backs up DHCP and/or WINS databases, the respective service is stopped and then restarted after the database is backed up. It is recommended to schedule the backup of the CONFIGURATION of a server that is running DHCP and/or WINS service during off hours.

DHCP and WINS services also provide their own internal backup copies of their databases. If your environment cannot tolerate occasional shutdowns of these services, you can exclude them from Data Protector CONFIGURATION backups and back up the internal backup copy of the databases via filesystem backup. For details about location of the internal backup copies and how to ensure that these copies are made frequently enough, see the Microsoft MSDN documentation.

Profiles

If the entire system is selected for backup, "Profiles" is backed up twice, (once as a part of filesystem backup and once as a part of CONFIGURATION). To avoid this, exclude the profile data from the filesystem backup. The user profile data resides in the `c:\Documents and Settings` directory:

These directories contain all user profiles configured on the system and are backed up by Data Protector. If a system is configured for multiple users, each defined user has a separate user profile. For example, the `All Users` and the `Default User` profile contain the profile components common to all defined users and the profile components assigned to a newly created user.

Data Protector reads the location of the profiles from the following Registry keys:

```
HKEY_USERS\.\DEFAULT\Software\Microsoft\Windows\
```

```
CurrentVersion\Explorer\Shell Folders
```

(where information about common profile components resides)

```
HKEY_USERS\.\DEFAULT\Software\Microsoft\Windows\
```

```
CurrentVersion\Explorer\User Shell Folders
```

Removable Storage Management Database

On Windows Vista and Windows Server 2008 operating systems, to enable backup of the Removable Storage Management Database configuration object, ensure that Removable Storage Manager is installed on the system which will be backed up.

Terminal Service Database

On Windows Vista and Windows Server 2008 operating systems, to enable backup of the Terminal Service Database configuration object, ensure that the Terminal Server Licensing service is installed on the system which will be backed up.

Windows services

Backing up the Windows services means backing up the data structures used by the respective services. A particular database is exported (dumped) into a file that is then backed up. The Windows

services are always backed up if CONFIGURATION is selected in the Backup wizard.

A Windows service has to be up and running so that Data Protector can detect it and present it as a selectable item in the Backup wizard. If a service is not running at the backup time, the corresponding backup object will fail.

To back up one of the services, select the corresponding folder under CONFIGURATION. If you use Active Directory to publish Certificate revocation lists (CRLs), for example, back up the Active Directory services along with the Certificate Server.

System State Data Backup

The Windows System State consists of several elements related to various aspects of Windows system. These elements are structured under their respective Windows backup object.

The Windows System State is not a selectable backup item. Data Protector enables you to back up individual objects, such as Registry or COM+ Class Registration Database. Backing up the whole CONFIGURATION tree is recommended. On Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012 systems backing up specific volumes or the entire client system using the filesystem backup functionality with the option **Use Shadow Copy** selected is required.

System State includes the following:

- Boot files: Ntldr.exe, Ntdetect.com, and boot.ini
- Registry and COM+ Class Registration Database (ComPlusDatabase)
- System File Protection service kept in the System Volume Information directory

If the services are installed and configured, the System State data of a Windows Server system also includes:

- ActiveDirectoryService
- CertificateServer
- Cluster Service information
- IIS Metadirectory
- RemoteStorageService
- RemovableStorageManagementDatabase
- SystemFileProtection
- SYSVOL directory
- TerminalServiceDatabase

On Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012 systems, System State data also includes data belonging to additional server roles or services that may be installed.

Remote Storage Service

Remote Storage Service is used to automatically move infrequently accessed files from local to remote storage. Remote files are recalled automatically when the file is opened. Although RSS databases are part of System State data, you back them up manually.

Remote Storage Services:

- Remote Storage Engine: *%SystemRoot%\system32\RsEng.exe*
Coordinates the services and administrative tools used for storing infrequently used data
- Remote Storage File: *%SystemRoot%\system32\RsFsa.exe*
Manages operations on remotely stored files
- Remote Storage Notification: *%SystemRoot%\system32\RsNotify.exe*
Notifies the client about recalled data

Remote Storage databases:

Remote Storage databases are located in the following directory:

%SystemRoot%\system32\RemoteStorage

- RSS Engine Database: *%SystemRoot%\system32\RemoteStorage\EngDb*
- RSS Engine Backup Database: *%SystemRoot%\system32\RemoteStorage\EngDb.bak*
- RSS File Database: *%SystemRoot%\system32\RemoteStorage\FsaDb*
- RSS Trace Database: *%SystemRoot%\system32\RemoteStorage\Trace*

Removable Storage Management Database

You can back up the Removable Storage database, but this service is not used for Data Protector media management. The native robotics driver used with robotics media changers has to be disabled before a device is configured by Data Protector.

System File Protection

System File Protection service scans and verifies the versions of all protected system files after you restart your computer. If the System File Protection service discovers that a protected file has been overwritten, it retrieves the correct version of the file and then replaces the incorrect file. Data Protector enables you to back up and then restore protected files without being overwritten. The protected files can be backed up using the Move Busy Files option in a standard filesystem backup procedure.

About UNIX Systems Backup

To perform a backup on a UNIX system, use the standard backup procedure. You need to perform some additional steps when backing up disks using NFS, for VxFS snapshot backup, or for a UNIX disk image backup.

Limitations

- When backing up NFS mounted filesystems, not all file attributes are preserved.
- The maximum size of files you can back up depends on operating system and filesystem limitations.

For a complete list of supported platforms and known limitations, see the *HPE Data Protector Product Announcements, Software Notes, and References*.

What is backed up?

- Data Protector backs up the directory structure, regular files, and special files. Special files are character device files, block device files, UNIX domain sockets, FIFO files, HP-UX network special files, and XENIX special named files.
- Symbolic links are not followed and are backed up as symbolic links.
- Mountpoints are not followed and are backed up as ordinary empty directories.
- If there are multiple hardlinks referencing the same file, the file is backed up only once. You can change this by setting the **Backup POSIX hard links as files** option.
- Basic ACLs (file permission attributes) and time attributes are backed up together with the files on all supported UNIX platforms. However, the support for extended ACLs is limited on some platforms. For details, see the *HPE Data Protector Platform and Integration Support Matrix* at <https://softwaresupport.hpe.com/>. The time of the last access to each file is saved before reading the file and then returned to the original value after the file is backed up. This behavior can be changed by setting the **Do not preserve access time attributes** option.

What should be excluded from a UNIX filesystem backup?

- Internal Database directories, that need to be backed up (online) in a special way.
- Temporary directories

NFS Backup

NFS (Network Filesystem) is a distributed filesystem protocol that allows a computer to access files over a network as if they were on its local disks. NFS lets you back up a filesystem from a remote UNIX system that is locally accessible.

When to use NFS backup?

- When a system is not a part of the Data Protector cell or does not have a Disk Agent installed.
- When you want to back up the system platforms that are not supported by Data Protector.

When configuring a regular filesystem backup, it is recommended that you exclude the NFS mounted filesystems from the backup. This avoids warning messages and a duplicated backup of the same disks if the system where the disks are actually located is also backed up.

Limitations

- You can back up NFS mounted volumes on HP-UX, Solaris, and Linux clients. You cannot back up soft links, character, and device files. For details on supported platforms, see the latest support matrices at <https://softwaresupport.hpe.com/>.
- ACL (Access Control List) attributes are not preserved. NFS does not support ACLs on remote files. Individual manual entries specify the behavior of various system calls, library calls, and commands.

When transferring a file with optional entries over the network or when manipulating a remote file, the optional entries may be silently deleted.

The backup procedure for an HP OpenVMS filesystem is the same as the standard filesystem backup procedure, with some OpenVMS-specific aspects.

Prerequisites

- To back up data on an OpenVMS system, install the OpenVMS Disk Agent on the OpenVMS system.
- To use backup devices connected to an OpenVMS system with Data Protector, install the General Media Agent on the OpenVMS system.

Limitations

- Any file specifications that are entered into the GUI or passed to the CLI must be in UNIX style syntax

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

The string should begin with a slash, followed by the disk, directories, and filename, 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 residing on ODS-5 disks. For example, an OpenVMS file specification of:

`1DGA100:[bUSERS.DOE]LOGIN.COM';1`

must be specified in the form:

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

- 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 wildcard characters for the version number, as follows:

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

- 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 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 image backups are not available on OpenVMS. There is no equivalent to a "BACKUP/PHYSICAL".
- The **Backup POSIX hard links as files** (`-hlink`), **Software compression** (`-compress`) and **Encode** (`-encode`) options are 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.

There is no support for an equivalent to BACKUP/IMAGE. To make a restored copy of an OpenVMS system disk bootable, the OpenVMS WRITEBOOT utility has to be used to write a boot block on to the restored disk.

- Files being backed up 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.
- 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. For example:

```
/SYS$MANAGER/DP_SAVE1.COM
```

- When specifying wildcard characters for **Skip** (-skip) or **Only** (-only) filters, use '*' for multiple characters and '?' for single characters.
- Data Protector file library is not supported on OpenVMS ODS-2 disks.
- On OpenVMS systems, Data Protector does not support disk quotas on volumes and volume sets.

To perform backup of data located on a volume with disk quota enabled, configure the pre-exec script so that it disables disk quota on the involved volume before backup starts, and configure the post-exec script so that it enables the disk quota after backup completes.

What is backed up?

The directory structure and the files are backed up together with the following filesystem information:

- File and directory attributes
- ACL (Access Control List)

Files can be backed up from mounted FILES-11 ODS-2 or ODS-5 volumes only.

About Novell Open Enterprise Server (OES) Backup

The backup procedure for Novell OES is the same as the standard backup procedure, with some Novell OES-specific aspects.

Prerequisites

- The Data Protector Disk Agent must be installed on the Novell OES system.
- The Target Service Agent for File Systems (tsafs) must be loaded in dual mode.
- For NDS/eDirectory backup, the Target Service Agent for Novell Directory Services (tsands) must be loaded.
- For GroupWise backup, the GroupWise Target Service Agent for File Systems (tsafsgw) must be loaded.
- The user account used to log in to Novell OES backup services must be selected and saved to the

file HPLOGIN.NLM. Any user account can be used, but the files and directories for backup will be limited to those of the user account.

- Storage Management Services (SMS) must be installed on the Novell OES system.

Limitations

- Software data compression is not supported. Even when the backup option Software compression is selected, it has no effect on the backed up data.

Backup and restore of compressed files

Novell OES provides file compression. By default, Data Protector backs up and consequently restores compressed files in their compressed format. Such files can only be restored to Novell OES with compressed volumes.

What is backed up?

- Native Linux volumes
- Novell GroupWise data

After backing up each file, the file archive flag is cleared and the archive date/time is set.

Configuring Novell OES

Saving the username and password using the HPLOGIN utility

The HPLOGIN utility is located in the directory `/opt/omni/sbin`. Run this utility to save proper user credentials (username and password) to the file `/root/OMNI$CFG.DAT`.

Steps

1. Change the current working directory to `/opt/omni/sbin`.
2. Run the HPLOGIN utility:

```
./hplogin
```

Loading Target Service Agent for File Systems (tsafs) in dual mode

Steps

1. Configure TSA on the target system. TSA is loaded in Linux mode by default. Change it to dual mode:
 - a. Change the current working directory to `/opt/novell/sms/bin`.
 - b. Check if tsafs is already loaded:

```
./smsconfig -t
```

- c. If it is loaded then unload it:

```
./smsconfig -u tsafs
```

- d. Load TSA in dual mode:

```
./smsconfig -l tsafs --tsaMode=Dual
```

2. The full path name of the tsafs configuration file on Open Enterprise Server Linux is `/etc/opt/novell/sms/tsafs.conf`. When TSA is loaded, it reads the configuration file for its default configuration. Configure this file to automatically load tsafs in dual mode every time TSA is loaded.
3. Edit the file `/etc/opt/novell/sms/tsafs.conf`, change `tsamode` from `Linux` to `dual`, and save the file:

tsamode=Dual

Loading the Target Service Agent for Novell Directory Services (tsands)

You can load the tsands agent manually or configure its automatic loading during the Novell OES startup.

Steps

- To load the agent manually:

- Open a terminal window.
- Change current directory to `/opt/novell/sms/bin`.
- Run the following command to check if the agent is already loaded:

```
./smsconfig -t
```

- d. If the agent is not loaded, load it:

```
./smsconfig -l tsands
```

- To configure automatic loading of the agent:

- a. Add the following line to the configuration file `/etc/opt/novell/sms/smdrd.conf`:

```
autoload: tsands
```

Loading the GroupWise Target Service Agent for File Systems (tsafsgw)

You can load the tsafsgw agent manually or configure its automatic loading during the Novell OES startup.

Steps

- To load the agent manually:
 - a. Open a terminal window.
 - b. Change current directory to `/opt/novell/sms/bin`.
 - c. Run the following command to check if the agent is already loaded:

```
./smsconfig -t
```
 - d. If the agent is not loaded, load it by providing appropriate parameters:

```
./smsconfig -l tsafsgw --home DomainDirectory --home PostOfficeDirectory
```
- To configure automatic loading of the agent:
 - a. Add the following line to the configuration file `/etc/opt/novell/sms/smdrd.conf` (replace argument placeholders with actual values):

```
autoload: tsafsgw --home DomainDirectory --home PostOfficeDirectory
```

About Backup Performance

You should consider backup performance factors when configuring backups. Due to the high number of variables and permutations, it is not possible to give distinct recommendations that fit all user requirements and affordable investment levels. However, the following should be considered when trying to improve a backup or restore performance:

Infrastructure

Infrastructure has a high impact on backup and restore performance. The most important factors are the parallelism of data paths and the use of high speed equipment.

- Network versus local backups and restores
Sending data over the network introduces additional overhead, as the network becomes a component of performance considerations. Data Protector handles the datastream differently for the following cases:
 - Network datastream: disk to memory to network to memory to device
 - Local datastream: disk to memory to device

To maximize performance, local backup configurations are recommended for high-volume datastreams.

- The devices used, the computer systems themselves, and the parallel usage of hardware can also have a noticeable impact on performance.

To work toward maximizing your backup or restore performance, you can:

- set appropriate concurrency to achieve device streaming
- optimize segment and block size
- adjust the number of Disk Agent buffers

- use software or hardware compression
- use disk-based backup devices — file libraries
- plan full and incremental backups
- use advanced backup strategies such as synthetic backup and disk staging
- optimize the distribution of backup objects to media
- disable filesystem scan

Object mirroring and backup performance

Object mirroring has an impact on backup performance. On the Cell Manager and Media Agent clients, the impact of writing mirrors is the same as if additional objects were backed up. On these systems, the backup performance will decrease depending on the number of mirrors. On the Disk Agent clients, there is no impact caused by mirroring, as backup objects are read only once.

Backup performance also depends on factors such as device block sizes and the connection of devices. If the devices used for backup and object mirroring have different block sizes, the mirrored data will be repackaged during the session, which takes additional time and resources. If the data is transferred over the network, there will be additional network load and time consumption.

High Performance Hardware Other than Devices

The speed of the computer systems themselves in reading the disk and writing to the device has a direct impact on performance. The systems are loaded during backup by reading the disk, handling software (de-)compression, and so on.

The diskread data rate and available CPU are important performance criteria for the systems themselves, in addition to the I/O performance and network types.

Hardware Parallelism

You can use several datapaths in parallel as a very efficient method to improve performance. This includes the network infrastructure. Parallelism helps in the following situations:

- If there are several systems backed up locally, that is, with the disk(s) and the related devices connected on the same system.
- If there are several systems backed up over the network. In this case, the network traffic routing needs to be such that the datapaths do not overlap, otherwise the performance is reduced.
- If there are several objects (disks) backed up to one or several (tape) devices.
- If several dedicated network links between certain systems can be used. For example, system_A has 6 objects (disks) to be backed up and system_B has 3 fast tape devices. The solution is to put 3 network links dedicated to backup between system_A and system_B.
- If there are several devices used and the **Load balancing** option is enabled.

Concurrency

The number of Disk Agents started for each Media Agent is called Disk Agent (backup) concurrency and can be modified using the Advanced options for the device or when configuring a backup. The concurrency set in the backup specification takes precedence over the concurrency set in the device definition.

Data Protector provides a default number of Disk Agents that are sufficient for most cases. For example, on a standard DDS device, two Disk Agents send enough data for the device to stream. For library devices with multiple drives where each drive is controlled by one Media Agent, you can set the concurrency for each drive independently.

Performance impact

If properly set, backup concurrency increases backup performance. For example, if you have a library device with four drives, each controlled by a Media Agent and each Media Agent receives data from two Disk Agents concurrently, data from eight disks is backed up simultaneously.

Multiple data streams

You can concurrently back up parts of a disk to multiple devices. This method speeds up the backup and is useful for backing up very large and fast disks to relatively slow devices. Multiple Disk Agents read data from the disk in parallel and send the data to multiple Media Agents.

Note that if one mount point was backed up through many Disk Agents, data is contained in multiple objects. To restore the whole mount point you have to define all parts of the mount point in a single backup specification and then restore the entire session.

Device Streaming

To maximize a device's performance, it has to be kept streaming. A device is streaming if it can feed enough data to the medium to keep it moving forward continuously. Otherwise, the tape has to be stopped, the device waits for more data, reverses the tape a little, resumes writing to the tape, and so on. In other words, if the data rate written to the tape is less than or equal to the data rate that can be delivered to the device by the computer system, then the device is streaming. Device streaming is also dependent on other factors such as network load and the block size of the data written to the backup device in one operation. In network-focused backup infrastructures, this deserves attention. For local backups, where disks and devices are connected to the same system, a concurrency of 1 may suffice, if your disks are fast enough.

How to configure device streaming

To allow the device to stream, a sufficient amount of data must be sent to the device. Data Protector accomplishes this by starting multiple Disk Agents for each Media Agent that writes data to the device.

Block Size

Segments are not written as a whole unit, but rather in smaller subunits called blocks. The device hardware processes data it receives using a block size specific to the device type.

Data Protector uses a default device block size regarding different device types. The block size applies to all devices created by Data Protector and to Media Agent running on the different platforms.

Increasing the block size can improve performance. You can adjust the blocks sent to the device while configuring a new device or when changing the device properties using the Advanced options for the device. A restore adjusts to block size.

Caution: Before increasing the block size for a device controlled by the Data Protector Media Agent running on a particular operating system, make sure the desired block size does not exceed the default maximum block size supported by the operating system. If the limitation is exceeded, Data Protector cannot restore data from such a device. For information if and how the maximum supported block size can be adjusted, see the operating system documentation.

You should change the block size before formatting tapes. The device block size is written on a medium header so that Data Protector knows the size to be used. If the device block size differs from the medium block size, an error occurs.

However, before changing the block size for the device, you need to check the supported block size of the used host adapter. The minimum block size for old SCSI cards, such as Adaptec 2940, used to be 56 kB. The minimum block size that is mainly used with newer SCSI cards is 64 kB.

You can increase the maximum block size on a Windows Media Agent client by modifying its Registry. The procedure depends on the host bust adapter type: SCSI, Fibre Channel, or iSCSI. For details, see the linked example topic.

Before changing the block size for a particular host bus adapter, see the vendor documentation or contact the vendor support.

Segment Size

A medium is divided into data segments, catalog segments, and a header segment. Header information is stored in the header segment, which is the same size as the block size. Data is stored in data blocks of data segments. Information about each data segment is stored in the corresponding catalog segment. This information is first stored in the Media Agent memory and then written to a catalog segment on the medium as well as to the IDB.

Segment size, measured in megabytes, is the maximum size of data segments. If you back up a large number of small files, the actual segment size can be limited by the maximum size of catalog segments. Segment size is user configurable for each device and influences performance during restore and during import of media. You can adjust the segment size while configuring a new device or when changing the device properties using the Advanced options for the device.

Optimal segment size depends on the media type used in the device and the kind of data to be backed up. The average number of segments per tape is 50. The default segment size can be calculated by dividing the native capacity of a tape by 50. The maximum catalog size is limited to a fixed number (12 MB) for all media types.

Data Protector finishes a segment when the first limit is reached. When backing up a large number of small files, the media catalog limit is reached faster, which can result in smaller segment sizes.

Number of Disk Agent Buffers

Data Protector Media Agents and Disk Agents use memory buffers to hold data waiting to be transferred. This memory is divided into a number of buffer areas (one for each Disk Agent, depending on device concurrency). Each buffer area consists of 8 Disk Agent buffers (of the same size as the block size configured for the device).

You can change this value while configuring a new device or when changing the device properties using the Advanced options for the device, although this is rarely necessary. There are two basic reasons to change this setting:

- Shortage of memory: the shared memory required for a Media Agent can be calculated as follows:

$\text{DAConcurrency} * \text{NumberOfBuffers} * \text{BlockSize}$

Reducing the number of buffers from 8 to 4, for instance, results in a 50% reduction in memory consumption, but also results in the creation of performance implications.

- Streaming

If the available network bandwidth varies significantly during backup, it is important that a Media Agent has enough data ready for writing to keep the device in the streaming mode. In this case, you should increase the number of buffers.

Software Compression

Software compression is done by the client CPU when reading the data from the disk. This reduces the data that gets sent over the network, but it requires significant CPU resources from the client.

By default, software compression is disabled. In general, only hardware compression should be used in order to improve performance. Software compression should only be used for backup of many systems over a slow network where the data can be compressed before sending it over the network.

If software compression is used, hardware compression should be disabled since trying to compress data twice actually expands the data.

Hardware Compression

Most modern backup devices provide built-in hardware compression that can be enabled when you create a device file or SCSI address in the device configuration procedure.

Hardware compression is done by a device that receives the original data from the Media Agent client and writes it to the tape in compressed mode. Hardware compression increases the speed at which a tape drive can receive data, because less data is written to the tape.

Consider the following regarding hardware compression:

- Use hardware compression with caution, because media written in compressed mode cannot be read using the device in an uncompressed mode and the other way round.
- Do not use software and hardware compression at the same time because double compression

decreases performance without giving better compression results.

- HPE Ultrium LTO drives use automatic hardware compression that cannot be disabled. It is recommended to keep the Software compression option disabled when you configure an HPE Ultrium LTO drive with Data Protector.
- When reading from a medium that was written using hardware compression with a device that does not support hardware compression, Data Protector cannot recognize the medium and read the data. Such a medium is treated as unknown or new.

When configuring the device, if you select the SCSI address from the drop-down list, Data Protector automatically determines whether the device can use hardware compression.

On UNIX systems you can enable the hardware compression by selecting a hardware compression device file.

On Windows systems, if the detection is not successful and you manually enter the SCSI address, C to the end of the device/drive SCSI address, for example: `scsi:0:3:0C` (or `tape2:0:1:0C` if tape driver is loaded). If the device supports hardware compression, it will be used, otherwise the C option will be ignored.

To disable hardware compression on Windows systems, add N to the end of the device/drive SCSI address, for example: `scsi:0:3:0N`.

For multipath devices, this option is set for each path separately.

Disk Image Versus Filesystem Backup

When you choose between a disk image backup and a filesystem backup, you should take into account their advantages and disadvantages. In most cases, filesystem backup is recommended.

	Filesystem backup	Disk image backup
Backup consistency	Files can be locked during backup and are backed up in a consistent state. The structure of files and directories is preserved.	Files are not locked during backup and are backed up in a point-in-time state. The structure of files and directories cannot be browsed.
Backup size	The space occupied by the backed up data is the same as the cumulative size of file and folder data at backup time.	The space occupied by the backed up data on backup media is the same as the size of the original backed up volume.
Backup and restore speed	Backup and restore speed is higher when a backed up disk is not full and the number of files is small.	Backup and restore speed is higher when a backed up disk is full and there is a big number of small files.
Restore usability	It is easier to navigate through the restored files, because the structure of files and directories is preserved.	The whole disk or a disk section is restored, the structure of files and directories cannot be browsed.

Note: On Windows systems, you can perform disk image and filesystem backup by using VSS writers. This ensures that the volume remains unlocked during the backup and can be accessed by

other applications. This is important when backing up System volume.

Object Distribution to Media

You can configure a backup such that the backup data ends up on the media in several different configurations. For example, you can configure a backup where one object goes to one medium or where several objects go to several media and each medium contains data from each object.

Under certain conditions, one distribution may be advantageous considering the backup performance, however this may not be the optimal restore configuration. You should define your backup policy such that you optimize the setup for a backup (since it is done frequently) and at the same time have an acceptable restore media situation.

Filesystem Scan

Before Data Protector backs up the files, it performs a scan of the tree selected for backup. This can impact the performance. Because the impact is negligible with quick filesystem scans on Windows systems and filesystem scan functionality on UNIX systems, changing the default settings only for performance purposes is not recommended.

Filesystem scan differs according to the system you want to back up:

System	Filesystem scan functionality	How to disable it?
Windows	Quick filesystem scan (always selected)	You can disable the filesystem scan by setting the OB2NOTREEWALK omnirc option to 1.
	Detect NTFS hardlinks (default: not selected)	Selecting the Detect NTFS hardlinks option results in significant reduction in performance. You should only select it if you have NTFS hardlinks present.
UNIX	Detect hardlinks and calculate size (default: selected)	Selecting the Backup POSIX hard links as files option renders the filesystem scan inactive.

Miscellaneous Performance Hints

You may be able to improve your backup or restore performance by following the hints listed in the table.

What improves performance?	How to improve performance?
Patches	Ensure that you have installed all patches pertaining to performance on the network.

Locality of devices	Use local devices where possible.
LAN cards	<p>You can move an FDDI card up on the bus so that it receives a higher priority. Use ftp to transfer large files between the Media Agent and Disk Agent systems to see how the speed compares to Data Protector performance.</p> <p>Note that network cards configured in half-duplex decrease performance.</p>
High-speed device	You can simulate a high-speed device on the Media Agent client if you suspect that the sustained data flow to the tape device is too low or that the device does not handle it correctly.
Device configuration	You can adjust the blocks sent to the device to increase performance.
CRC check option	You can disable the CRC Check option. If enabled, this option impacts performance due to the CRC calculation, which is performed by the Media Agent client.
Logging and Report Level	<p>You can disable logging by setting it to No Log if an update of the IDB takes too long.</p> <p>You can filter messages by setting the Report level to Critical.</p>
Data Protector Application Clients	You can decrease the <code>SmWaitforNewClient</code> value, if a restore session of the Application clients (Oracle, SAP R/3) takes too long. Set it to a value lower than the default (5 minutes).

Chapter 10: Object Consolidation

About Object Consolidation

The Data Protector object consolidation functionality enables you to merge a restore chain of a backup object into a new, consolidated version of this object. Using this functionality, you no longer need to run full backups. Instead, you can run incremental backups indefinitely and consolidate the restore chain as needed.

During the object consolidation session, Data Protector reads the backed up data from the source media, merges the data, and writes the consolidated version to the target media. The result of an object consolidation session is a synthetic full backup of the object version you specified.

Types of object consolidation

You can start an object consolidation session interactively or specify an automated start of the session. Data Protector offers two types of automated object consolidation: post-backup object consolidation and scheduled object consolidation.

Post-backup object consolidation

Post-backup object consolidation takes place after the completion of a backup session that is specified in the automated object consolidation specification. It consolidates objects selected according to the automated object consolidation specification that were backed up in that particular backup session.

Scheduled object consolidation

Scheduled object consolidation takes place at a user-defined time. Objects backed up during different backup sessions can be consolidated in a single scheduled object consolidation session.

How to Consolidate Objects

First, create an object consolidation specification. In the specification, select the object versions you want to consolidate, the media and devices you want to use, and session options.

Selection of devices

You need separate devices for reading full backups, reading incremental backups, and writing the synthetic full backup. The destination devices can have a larger block size than the source devices. However, to avoid impact on performance, it is recommended that the devices have the same block size and are connected to the same system.

Devices that are not available at the beginning of a session cannot be used in that session. If a media error occurs, the device with errors will be avoided within that session.

Object consolidation options

You can enable source object filtering and specify data protection, catalog protection, and logging level in the object consolidation specification. Equivalents to most of these options are used for backup as well.

Selection of the media set

If an object version that will participate in consolidation has copies residing on different media sets, any of the media sets can be used as a source. By default, Data Protector automatically selects the most appropriate media set. You can influence the media set selection by specifying the media location priority.

The overall process of media selection is the same as for restore. When consolidating objects interactively, you can manually select the media set to be used. You cannot select media when configuring automated object consolidation, as the backup of the objects is often performed at a later time.

Ownership of consolidated objects

Owner of the consolidated backup objects is the owner of the original backup objects, not the Data Protector user who invokes the object consolidation session.

Standard Object Consolidation Tasks

Below are the prerequisites and limitations of the object consolidation functionality:

Prerequisites

- All the backups that will be consolidated were performed with the **Enhanced incremental backup** option enabled.
- All incremental backups that will be consolidated reside in one file library or a B2D device (except Smart Cache).
- The restore chain is complete, meaning that all the object versions that comprise it have the status `Completed` or `Completed/Errors` and all the media holding these object versions are available.
- The necessary backup devices are configured and the media prepared.
- You need a Media Agent installed on every system that will participate in an object consolidation session.
- You need appropriate user rights for starting an object consolidation session. The same user rights apply as for backup.
- To perform a virtual full backup, all the backups (full, incremental, and virtual full) must reside in one file library that uses distributed file media format.

Limitations

- The destination devices must have the same or a larger block size than the source devices.
- The same medium cannot be used as a source medium and as a target medium in the same object consolidation session.
- While the source media are being read, they are unavailable for restore.
- Object consolidation is not available for objects backed up using AES 256-bit encryption.
Object consolidation is supported for all B2D devices, except Smart Cache.

Note: Whenever you change the setting of the **Software compression** or **Encode** option in the backup specification, a full backup must be performed as a basis for subsequent object consolidation.

Consolidating Objects Interactively

You can select objects for interactive consolidation from the Objects or Sessions starting point, depending on your needs. You cannot save an interactive object consolidation specification, you can only start an object consolidation session.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Consolidation**, and then expand **Interactive**.
3. Click **Objects** or **Sessions** to open the wizard.
 - Clicking **Objects** lists objects.
 - Clicking **Sessions** lists sessions in which objects were written to media.
4. Select the points in time of the desired objects to consolidate. You cannot select full backups, as they as such cannot be consolidated.

Selecting a point in time selects the entire restore chain. If several restore chains for the same point in time exist, all of them are selected, but only one will actually be used. Your selection is marked in blue, other incrementals that comprise the restore chain are marked in black, and the corresponding full backup in gray (shaded). The blue check mark indicates the point in time that will be consolidated.

You can select several points in time for consolidation, and the restore chains may overlap. If you select a point in time that already has a black check mark, the check mark will become blue.

To clear a selected restore chain, click the blue check mark. The entire restore chain is cleared, unless some object versions are part of another restore chain, in which case they remain selected with a black check mark.

Click **Next**.

5. Specify the devices that will read the incremental backups and the full backups.

Limit the object consolidation to specific file libraries or B2D devices (except Smart Cache) by selecting them as read devices for incremental backups. Only objects residing in the specified devices will be consolidated.

By default, the read devices for full backups are those used for backup in the selected backup specifications. You can change them here if desired. Click **Next**.

6. Select the destination devices for the object consolidation operation. Data Protector will select the most suitable devices from those you specify here. Click **Next**.
7. Specify options as desired. Click **Next**.
8. A list of media containing the selected objects is displayed.

You can change the media location priority to influence the selection of media in case the same object resides on more than one media set.

Click **Next**.

9. Review the object versions that will participate in the operation. In case of alternative restore chains, it may happen that not all the listed object versions will actually be used. Click **Next**.
10. Review the summary of the selected points in time. To change options for a particular point in time, select it in the list and click **Properties**.
11. Click **Finish** to exit the wizard.

Configuring Post-Backup Object Consolidation

Post-backup object consolidation takes place after the completion of a backup session that is specified by the name of the backup specification in the automated object consolidation specification. It consolidates objects backed up in that particular backup session that match the specified criteria.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Consolidation**, and then expand **Automated**.
3. Right-click **Post Backup** and click **Add** to open the wizard.
4. Select the backup specifications that contain the objects you want to consolidate. Click **Next**.
5. Specify the object filter for the object consolidation operation. Click **Next**.
6. Specify the devices that will read the incremental backups and the full backups.

Limit the object consolidation to specific file libraries or B2D devices (except Smart Cache) by selecting them as read devices for incremental backups. Only objects residing in the specified devices will be consolidated.

By default, the read devices for full backups are those used for backup in the selected backup specifications. You can change them here if desired. Click **Next**.

7. Select the destination devices for the object consolidation operation. Data Protector will select the most suitable devices from those you specify here. Click **Next**.
8. Specify options as desired. Click **Next**.
9. Click **Save as...**, enter a specification name and click **OK** to save the post-backup object consolidation specification.

Scheduling of Object Consolidation

Scheduled object consolidation takes place at a user-defined time. It consolidates objects that match the specified criteria. Objects backed up during different backup sessions can be consolidated in a single scheduled object consolidation session.

When there are many possible restore chains to select from, Data Protector consolidates the one containing the object version with the latest point in time. For example, backup sessions: **Full**, **Incr1**, **Incr2**, **Incr2**, **Incr2** result in three restore chains but Data Protector consolidates only the one consisting of **Full**, **Incr1**, and the latest **Incr2**.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Consolidation**, and then expand **Automated**.
3. Right-click **Scheduled** and click **Add** to open the wizard.
4. Select the backup specifications that contain the objects you want to consolidate. Click **Next**.
5. Specify the time filter for the object consolidation operation. Only objects that were backed up in the specified time frame will be consolidated. Click **Next**.
6. Specify the object filter for the object consolidation operation. Click **Next**.

7. Specify the devices that will read the incremental backups and the full backups.

Limit the object consolidation to specific file libraries or B2D devices (except Smart Cache) by selecting them as read devices for incremental backups. Only objects residing in the specified devices will be consolidated.

By default, the read devices for full backups are those used for backup in the selected backup specifications. You can change them here if desired. Click **Next**.

8. Select the destination devices for the object consolidation operation. Data Protector will select the most suitable devices from those you specify here. Click **Next**.
9. Specify options as desired. Click **Next**.
10. Right-click a date and click **Schedule** to display the Schedule Consolidation dialog box. Specify the options as desired and click **OK**. Click **Next**.
11. Click **Save as...**, enter a specification name and click **OK** to save the scheduled object consolidation specification.

Copying an Object Consolidation Specification

You can copy an already configured and saved object consolidation specification.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Consolidation, Automated**, and then **Post Backup**. All saved object consolidation specifications are displayed.
3. In the Results Area, right-click the object consolidation specification that you want to copy and then click **Copy As**. The Copy As dialog box opens.
4. In the Name text box, type the name for the copied object consolidation specification.
5. Click **OK**.

The copied object consolidation specification is displayed in the Object Operations context of the Scoping Pane and in the Results Area under the new name.

Chapter 11: Copy

About Duplicating Backed Up Data

Duplicating backed up data brings several benefits. You can copy data to improve its security and availability, or for operational reasons.

Data Protector provides the following methods of duplicating backed up data: object copy, object mirror, media copy, and replication on Backup to Disk (B2D) devices.

	Object copy	Replication	Object mirror	Media copy
What is duplicated	Any combination of object versions from one or several backup sessions, object copy sessions, or object consolidation sessions	A set of objects from a backup session, object copy session, or object consolidation session	A set of objects from a backup session	An entire medium
Time of duplication	Any time after the completion of a backup	Any time after the completion of a backup	During backup	Any time after the completion of a backup
Media type of source and target media	Can be different	Data can be replicated only to B2D devices of the same type	Can be different	Must be the same
Size of source and target media	Can be different	The target device must have enough space for the deduplicated data	Can be different	Must be the same

Appendability of target media	Yes	No	Yes	No ¹
Result of the operation	Media containing the selected object versions	An identical copy stored on the target B2D device	Media containing the selected object versions	Media identical to the source media

About Object Copying

What is object copy?

The Data Protector object copy functionality enables you to copy selected object versions to a specific media set. You can select object versions from one or several backup sessions, object copy sessions, or object consolidation sessions. During the object copy session, Data Protector reads the backed up data from the source media, transfers the data, and writes it to the target media.

The result of an object copy session is a media set that contains copies of the object versions you specified.

The following characterizes the object copy functionality:

- **Start of session**
An object copy session can be started interactively or automatically.
- **Selection of media**
As source media, you can use original media sets containing backups, media sets containing object copies, or media sets that are media copies.
However, the selection of the media set is not possible after the start of the object copy session. In case of a mount request, you need to provide the specific medium that is requested by Data Protector, or its identical copy (created using the media copy functionality).
- **Media type**
You can copy objects to media of a different type. Furthermore, the block size of the destination device can be the same or larger than the block size of the source device.
- **Media policy**
You can append data to media already containing backups or object copies.
- **Protection policy**
You can set the protection periods for the source objects and the object copies independently.

You can start an object copy session interactively or specify an automated start of the session.

¹ You can use only unformatted media, empty media, or media with expired protection as target media. After the operation, both the source and the target media become non-appendable.

You can also use a combination of the duplication methods. For example, you can create object copies or media copies of data that is the result of object mirroring. Or, you can copy entire media containing object copies.

Automated object copying

In an automated object copy specification, you can specify one or more criteria for the selection of object versions that will be copied:

- Backup specifications - to copy only object versions backed up using specific backup specifications.
- Object copy specifications - to copy only object versions copied using specific object copy specifications.
- Object consolidation specifications - to copy only object versions consolidated using specific object consolidation specifications.
- Data protection - to copy only protected object versions.
- Number of existing copies - to copy only object versions that do not have more than the specified number of successful copies.
- Libraries - to copy only object versions located on the media in the specified libraries.
- Time frame (only in a scheduled object copy specification) - to copy only object versions backed up in the specified period of time.

Data Protector offers two types of automated object copying: post-backup object copying and scheduled object copying.

Post-backup object copying

Post-backup as well as post-copy and post-consolidation object copying, which are subsets of post-backup object copying, take place after the completion of a session that is specified in the automated object copy specification. They copy objects selected according to the automated object copy specification that were written in that particular session.

Scheduled object copying

Scheduled object copying takes place at a user-defined time. Objects from different sessions can be copied in a single scheduled object copy session.

How to Copy Objects

First, create an object copy specification. In the specification, select the objects you want to copy, the media and devices you want to use, session options, and the media location priority that influences how Data Protector selects the media set in case the same object resides on several media sets.

Selection of devices

You need separate devices to be used with the source media and the target media. The destination devices can have a larger block size than the source devices. However, to avoid impact on performance, it is recommended that the devices have the same block size and are connected to the same system or to a SAN environment.

Object copying is load balanced by default. Data Protector makes optimum use of the available devices by utilizing as many devices as possible.

If you do not specify the source devices to be used in the object copy specification, Data Protector uses the default devices. By default, the devices that were used for writing the objects are used as source devices. You can change the source devices if desired. If destination devices are not specified per object, Data Protector selects the most suitable devices automatically from those you selected in the object copy specification.

Devices are locked at the beginning of the session. Devices that are not available at that time cannot be used in the session, as device locking after the beginning of the session is not possible. If a media error occurs, the device with errors will be avoided within that copy session.

Object copy options

You can enable source object filtering and specify data protection, catalog protection, and logging level for object copies in the object copy specification. Equivalents to most of these options are used for backup as well.

Depending on your policy, backed up objects and their copies may have the same or different option values specified. For example, you can specify the **No Log** value for a backup object to increase the backup performance, and then specify the **Log All** value for the same object in a subsequent object copy session.

To create identical copies of backed up objects, specify the same logging level for object copies. Consider that each object copy with a logging level higher than No Log has an impact on the IDB size.

Selecting the media set to copy from

If an object version that you want to copy exists on more than one media set, which has been created using one of the Data Protector data duplication methods, any of the media sets can be used as a source for copying. By default, Data Protector automatically selects the media set that will be used. You can influence the media set selection by specifying the media location priority.

The overall process of media selection is the same as for restore. When copying objects interactively, you can manually select the media set to copy from when your starting point is Objects or Sessions. You cannot select media when configuring automated object copying, as the backup of the objects is often performed at a later time.

Object copy completion status

Copy objects

You can copy objects that have the status `Completed` or `Completed/Errors`, provided that all the media on which they reside are logged in the IDB. If the copy operation is successful, the status of the copied object is the same as the status of the corresponding backed up object.

If you have aborted an object copy session, or if it failed for other reasons, the object copies that are results of such a session have the status `Failed`. An object copy with the status `Failed` cannot be copied again; its data and catalog protection are set to `None`.

Source objects

If an object copy session fails, the source objects that were copied are left unchanged.

If an object copy session completes with errors, the source objects that were copied successfully have their data and catalog protection set to the values specified in the source object options.

If you abort an object copy session, the data and catalog protection of all source objects are left unchanged. In this case, if you want to change the protection for any of the copied objects, you must do this manually within the IDB.

Ownership of object copies

Owner of the copied backup objects is the owner of the original backup objects, not the Data Protector user who invokes the object copy session.

Standard Object Copy Tasks

Below are the prerequisites and limitations of the object copy functionality:

Prerequisites

- You need to have a Media Agent installed on every system that will participate in an object copy session.
- You need to have at least two backup devices configured in the Data Protector cell.
- You need to have media prepared for the object copy session.
- You need to have appropriate user rights for performing an object copy session.

Limitations

- It is not possible to copy objects backed up using the ZDB to disk or NDMP backup functionality.
- It is not possible to create multiple copies of one object version in one object copy session.
- The destination devices must have the same or a larger block size than the source device.
- The same medium cannot be used as a source medium and as a target medium in the same object copy session.
- During object copying, the media used as sources are unavailable for restore.
- It is not possible to demultiplex SAP MaxDB, DB2 UDB, or SQL integration objects.
- It is not possible to copy objects backed up, copied, or consolidated during sessions that were run interactively from the last page of the wizard.
- It is not possible to start two or more object copy sessions from the same object copy specification in parallel.

Consider the following:

- The Data Protector SAP MaxDB, DB2 UDB, and Microsoft SQL Server integrations have interdependent data streams. Hence the object copy operation must preserve the layout of objects on media to enable a restore. To ensure this, select all objects of these integrations with the same

backup ID for copying. Otherwise, a restore from the copy will not be possible.

- The minimum number of devices required for copying SAP MaxDB, DB2 UDB, or Microsoft SQL Server integration objects equals the number of devices used for backup. The concurrency of the devices used for backing up and copying these objects must be the same.
- If you select the **Change data and catalog protection after successful copy** option when copying objects from a ZDB to disk+tape session, be aware that after the period you specify, the source objects can be overwritten. After the media are overwritten, instant recovery from this backup using the GUI will no longer be possible.
- If you abort an object copy session, the data and catalog protection of all source objects are left unchanged. In this case, if you want to change the protection for any of the copied objects, you must do this manually within the IDB.

Copying Objects Interactively

After an object has been backed up, you can copy it to a new media set.

You can select objects for interactive copying from the Media, Objects, or Sessions starting point, depending on your needs. You cannot save an interactive object copy specification, you can only start an object copy session.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Copy**, then **Object copy**, and then **Interactive**.
3. Click **Media**, **Objects**, or **Sessions** to open the wizard.
 - Clicking **Media** lists media pools and media.
 - Clicking **Objects** lists types of backed up data, such as Filesystem, Database, and so on.
 - Clicking **Sessions** lists sessions in which objects were written to media.
4. Select the objects to copy.

If you selected Sessions in the previous step, you can right-click an integration object and click **Select Backup Set** to select all integration objects with the same backup ID.

Note: From Data Protector 9.09 onwards for VMware backups, the virtual machine disks are considered as objects that run in parallel. The disk objects of the virtual machine are listed but disabled in the **Media** list to understand virtual machine disks backed to the media. The copy or verify operation is performed on the virtual machine objects and all its associated disk objects are considered internally.

From Data Protector 9.09 onwards for VMware integration, the **Next** option is enabled only after selecting the virtual machine object in the **Media** list.

Click **Next**.

5. The devices used for writing the selected objects are used as source devices in the object copy operation by default. You can change the source devices here if desired. Select the original device and click **Change**. The name of the new device appears under Device Status. The new device will

be used only for this session.

For more information on a device, right-click the device and click **Info**.

Specify what Data Protector should do if the selected devices are not available during object copy (for example, if they are disabled or already in use). Select either **Automatic device selection** or **Original device selection**.

Click **Next**.

6. Select the destination devices for the object copy operation.

You can specify devices per object in the Summary page from the list of devices specified here. If you do not specify a device for each object, Data Protector will select the most suitable devices from this list.

Click **Next**.

7. Specify the source object options, target object options, and target media options as desired.

Click **Next**.

Optionally, to enable replication between two B2D devices instead of copying, select **Use replication**. Once Use replication is selected, the Replicate to a foreign cell gets enabled.

8. A list of media containing the selected objects is displayed.

If your starting point was Objects or Sessions, media location priority is also listed. You can change the media location priority to influence the selection of media in case the same object resides on more than one media set.

Click **Next**.

9. Review the summary of the selected objects. To change options for a particular object, select the object in the list and click **Properties**.

You can specify source object options, target object options, and the destination device. If the Objects or Sessions starting point was used, you can manually select which copy of the object version will be used if more than one copy exists.

10. Click **Finish** to start the copy session.

Configuring Post-Backup Object Copying

Post-backup object copying takes place after the completion of a backup session, object copy session, or object consolidation session that is specified by the name of the backup, object copy, or object consolidation specification in the automated object copy specification. It copies objects from that particular session that match the specified criteria.

A post-backup object copy session does not start if the backup session failed. If the backup session has been aborted, but contains completed objects, a post-backup object copy session copies the completed objects by default. To disable the copying of aborted sessions, set the global option `CopyStartPostBackupOnAbortedSession` to 0.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Copy**, then **Object copy**, and then **Automated**.
3. Right-click **Post Backup** and click **Add** to open the wizard.

4. Select the backup, object copy, or object consolidation specifications that contain the objects you want to copy. Click **Next**.
5. Specify the object filter for the object copy operation. Only objects that match the specified criteria will be copied. Click **Next**.
6. Specify the library filter for the object copy operation. Only objects residing on media in the specified libraries will be copied. Click **Next**.
7. The devices used for backup in the selected backup specifications are used as source devices in the object copy operation by default. You can change the source devices here if desired. Click **Next**.
8. Select the destination devices for the object copy operation. Data Protector will select the most suitable devices from those you specify here. Click **Next**.
9. Specify the source object options, target object options, and target media options as desired. Click **Next**.
Optionally, to enable replication between two B2D devices instead of copying, select **Use replication**.
10. Click **Save as...**, enter a specification name and click **OK** to save the post-backup object copy specification.

Scheduling of Object Copying

Scheduled object copying takes place at a user-defined time. Objects from different backup sessions, object copy sessions, or object consolidation sessions can be copied in a single scheduled object copy session.

Tip: You can also schedule object copy sessions with advanced settings with the Advanced Scheduler.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Copy**, then **Object copy**, and then **Automated**.
3. Right-click **Scheduled** and click **Add** to open the wizard.
4. Select the backup, object copy, or object consolidation specifications that contain the objects you want to copy.
You can view backup specifications also by backup group. This way, if you add or remove a backup specification to or from a certain backup group, the object copy functionality automatically acknowledges the change and you do not need to modify the object copy specification manually.
Note that if you change from the group view to any other view, a warning message is displayed that changing the view will remove all current selections. If you continue, all previous selections are cleared.
Click **Next**.
5. Specify the object filter for the object copy operation. Only objects that match the specified criteria will be copied. Click **Next**.
6. Specify the library filter for the object copy operation. Only objects residing on media in the

specified libraries will be copied. Click **Next**.

7. The devices used for backup in the selected backup specifications are used as source devices in the object copy operation by default. You can change the source devices here if desired. Click **Next**.
8. Select the destination devices for the object copy operation. Data Protector will select the most suitable devices from those you specify here. Click **Next**.
9. Specify the source object options, target object options, and target media options as desired. Click **Next**.
Optionally, to enable replication between two B2D devices instead of copying, select **Use replication**.
10. Right-click a date and click **Schedule** to display the Schedule Copy dialog box. Specify the options as desired and click **OK**. Click **Next**.
11. Click **Save as...**, enter a specification name and click **OK** to save the scheduled object copy specification.

Restarting Failed Object Copy Sessions

Due to network connectivity problems or system unavailability, it may occur that during an object copy session some objects fail. You can restart a problematic session after you have resolved the impeding issues. This action restarts only the failed objects.

Prerequisites

- You either have to be in the Data Protector Admin user group or have the Data Protector Monitor user right.

Limitations

- You cannot restart failed sessions that were run interactively, meaning they were based on unsaved object copy specifications.
- It is not possible to restart several sessions at the same time.

Do not change an object copy specification before restarting a failed object copy session. Otherwise, it is not possible to restart all objects.

Steps

1. If you are using an ordinary Cell Manager, in the Context List, click **Internal Database**.
If you are using Manager-of-Managers, in the Context List, select **Clients** and expand **Enterprise Clients**. Select a Cell Manager with the problematic session. From the Tools menu, select **Database Administration** to open a new Data Protector GUI window with the Internal Database context displayed.
2. In the Scoping Pane, expand **Internal Database** and click **Sessions**.
A list of sessions is displayed in the Results Area. Status of each session is denoted in the Status column.
3. Right-click a failed, an aborted, or a session that completed with failures or errors, and select

Restart Failed Objects to copy the objects that failed.

4. Click **Yes** to confirm.

Copying an Object Copy Specification

You can copy an already configured and saved object copy specification.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Copy**, **Object copy**, **Automated**, and then **Post Backup**. All saved object copy specifications are displayed.
3. In the Results Area, right-click the object copy specification that you want to copy and then click **Copy As**. The Copy As dialog box opens.
4. In the Name text box, type the name for the copied object copy specification.
5. Click **OK**.

The copied object copy specification is displayed in the Object Operations context of the Scoping Pane and in the Results Area under the new name.

Advanced Object Copy Tasks

Additional copies of backed up data are created for multiple purposes:

- **Vaulting**
You can make copies of backed up, copied, or consolidated objects and keep them in several locations.
- **Freeing media**
To keep only protected object versions on media, you can copy such object versions, and then leave the medium for overwriting.
- **Demultiplexing of media**
You can copy objects to eliminate interleaving of data.
- **Consolidating a restore chain**
You can copy all object versions needed for a restore to one media set.
- **Migration to another media type**
You can copy your backups to media of a different type.
- **Support of advanced backup concepts**
You can use backup concepts such as disk staging.

Freeing a Medium

A medium can contain backed up objects with different periods of protection. It may happen that only a small amount of media space is occupied by a protected object. However, you cannot reuse such a medium until the protection of all objects has expired.

To rationalize the usage of media, you can free media that contain only some protected objects using the object copy functionality. Protected objects are copied to a new media set and the medium can be reused. You can also free media from failed objects. These objects are not copied in the object copy session.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Copy**, then **Object copy**, and then **Interactive**.
3. Click **Media** to open the wizard.
4. In the Objects page, select the option **Enable selection of protected objects only**. Expand the media pool(s) and select the media that you want to free. Click **Next**.
5. The devices used for writing the selected objects are used as source devices in the object copy operation by default. You can change the source devices here if desired. Click **Next**.
6. Select the destination devices for the object copy operation.
You can specify devices per object in the Summary page from the list of devices specified here. If you do not specify a device for each object, Data Protector will select the most suitable devices from this list.
Click **Next**.
7. In the Options page, under Source object options, select **Change data and catalog protection after successful copy** to remove the protection of the source objects after these objects are copied. Select **Recycle data and catalog protection of failed source objects after successful copy** to remove the protection of failed source objects (these objects will not be copied). Specify other options as desired. Click **Next**.
8. A list of media containing the selected objects is displayed. Click **Next**.
9. Review the summary of the selected objects. To change options for a particular object, select the object in the list and click **Properties**. You can specify source object options, target object options, and the destination device.
10. Click **Finish** to start the copy session.

Demultiplexing a Medium

Multiplexed media contain interleaved data of multiple objects. Such media may arise from backup sessions with the device concurrency more than 1. Multiplexed media may compromise the privacy of backups and require more time for restore.

Using the object copy functionality, you can demultiplex media. Objects from a multiplexed medium are copied to several media.

Limitation

Data Protector reads the source medium only once. To enable demultiplexing of all objects on the medium, the minimum number of destination devices needed for the operation is the same as the device concurrency that was used for writing the objects. If fewer devices are available, some objects will still be multiplexed on the target medium.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Copy**, then **Object copy**, and then **Interactive**.
3. Click **Sessions** to open the wizard.
4. Expand the desired session(s) and select the object(s) to copy. Click **Next**.
5. Perform this step if you do not want the demultiplexing operation to occupy the device(s) configured for regular backups and if you want to use only one device for reading the data during the demultiplexing operation.

Map the source device(s) to a single device.

Skip this step if a standalone file device was used as a source device. If a file jukebox device or a file library device were used as a source device, make sure to map the source device(s) to device (s) in the same file jukebox or in the same file library.

Right-click each device and click **Change Device**. Select the new device and click **OK**.

6. Click **Next**.
7. Select the destination devices for the object copy operation. The number of devices needed depends on the device concurrency that was used when writing the objects.

Right-click each selected drive and click **Properties**. Set the **Concurrency** option to 1. Click **OK**.

You can specify devices per object in the Summary page from the list of devices specified here. If you do not specify a device for each object, Data Protector will select the most suitable devices from this list.

Click **Next**.

8. Specify the source object options, target object options, and target media options as desired. Click **Next**.
9. A list of media containing the selected objects is displayed.

You can change the media location priority to influence the selection of media in case the same object resides on more than one media set.

Click **Next**.

10. Review the summary of the selected objects. To change options for a particular object, select the object in the list and click **Properties**.

You can specify source object options, target object options, and the destination device. You can also manually select which copy of the object version will be used if more than one copy exists.

11. Click **Finish** to start the copy session.

Consolidating a Restore Chain

Using the object copy functionality, you can copy a restore chain of an object version to a new media set. A restore from such a media set is faster and more convenient, as there is no need to load several media and seek for the needed object versions.

Note: Data Protector also provides a more powerful feature called object consolidation. While object copy enables you to copy all the backups of a restore chain into a sequence, object

consolidation merges the backups into a new object version, a synthetic full backup.

Limitation

The selection of a restore chain is not available for integration objects.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Copy**, then **Object copy**, and then **Interactive**.
3. Click **Objects** to open the wizard.
4. In the Objects page, expand a type of data, then a client and its logical disks or mount points to display the object versions. Right-click the object(s) to copy and click **Select Restore Chain**. Click **Next**.
5. The devices used for writing the selected objects are used as source devices in the object copy operation by default. You can change the source devices here if desired. Click **Next**.
6. Select the destination devices for the object copy operation.
You can specify devices per object in the Summary page from the list of devices specified here. If you do not specify a device for each object, Data Protector will select the most suitable devices from this list.
Click **Next**.
7. Specify the source object options, target object options, and target media options as desired. Click **Next**.
8. A list of media containing the selected objects is displayed.
You can change the media location priority to influence the selection of media in case the same object resides on more than one media set.
Click **Next**.
9. Review the summary of the selected objects. To change options for a particular object, select the object in the list and click **Properties**.
You can specify source object options, target object options, and the destination device. You can also manually select which copy of the object version will be used if more than one copy exists.
10. Click **Finish** to start the copy session.

Migrating to Another Media Type

You can use the object copy functionality to migrate backed up data to another media type of the same or a larger block size.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Copy**, then **Object copy**, and then **Interactive**.
3. Click **Media** to open the wizard.

4. Select the objects to copy and click **Next**.
5. The devices used for writing the selected objects are used as source devices in the object copy operation by default. You can change the source devices here if desired. Click **Next**.
6. Select the destination devices for the object copy operation.
You can specify devices per object in the Summary page from the list of devices specified here. If you do not specify a device for each object, Data Protector will select the most suitable devices from this list.
Click **Next**.
7. Specify the source object options, target object options, and target media options as desired.
Click **Next**.
8. A list of media containing the selected objects is displayed. Click **Next**.
9. Review the summary of the selected objects. To change options for a particular object, select the object in the list and click **Properties**.
You can specify source object options, target object options, and the destination device.
10. Click **Finish** to start the copy session.

About Disk Staging

What is disk staging?

The concept of disk staging is based on backing up data in several stages. The backup stages consist of backing up data to a medium of one type and then copying the data to a medium of a different type. Typically the functionality might be used as follows

1. The data is backed up to a medium with high performance and accessibility, but limited capacity (for example, system disk). Such backups are usually kept accessible for the period of time when a fast restore is the most likely to be required.
2. After a certain period of time, the data is moved to a medium with lower performance and accessibility, but high capacity for storage, using the object copy functionality.

You could perform disk staging in this way using a scheduled object copy specification configured specifically for the purpose.

An alternative approach might be as follows:

1. Create a backup specification to back up the data to the high performance medium with the protection set to the overall period for which restore capability is required.
2. Create an automated post-backup copy specification to copy the backed-up data to the lower performance medium, and reset the retention period for the original backup to the critical period for which fast restore capability is required. By default the secondary copy will be retained for the protection period specified in the original backup specification.

With this method there is the extra security of having both copies during the critical period.

Why implement disk staging

The use of the disk staging concept brings the following benefits:

- It improves the performance of backups and restores.
- It reduces costs of storing the backed up data.
- It increases the data availability and accessibility for restore.

Disk staging and small reoccurring backups

Disk staging can also be used to eliminate the need for frequent backups of numerous small objects to tape. Such backups are inconvenient due to frequent loading and unloading of media. The use of disk staging can reduce backup time and prevent media deterioration.

Troubleshooting Object Operations Sessions

Object copy problems

Fewer objects are copied than expected

Problem

With post-backup or scheduled object copy, the number of objects that match the selected filters is higher than the number of objects that are actually copied.

The following message is displayed:

Too many objects match specified filters.

Action

- Tighten the criteria for object version selection.
- Increase the maximum number of objects copied in a session by setting the global option `CopyAutomatedMaxObjects`.

Not all objects in the selected library are copied

Problem

With post-backup or scheduled object copy, some objects that reside on media in the selected library are not copied. This happens if an object does not have a complete media set in the selected library.

Action

Insert the missing media into the selected library, or select the library that has a complete media set for these objects.

Mount request for additional media is issued

Problem

In an interactive object copy session from the Media starting point, you selected a specific medium. A mount request for additional media is issued. This happens if an object residing on the medium spans to another medium.

Action

Insert the required medium into the device and confirm the mount request.

When creating an object copy, the protection end time is prolonged

Problem

When creating an object copy, the protection end time is not inherited from the original object. The protection length is copied, but the start time is set at the object copy creation time and not at the object creation time. This results in a longer protection then for the original. The more time passes between the original backup and the object copy session, the bigger the difference between the protection end times.

For example, if the object was created on September 5, with the protection set to 14 days, the protection will expire on September 19. If the object copy session was started on September 10, the object copy protection will expire on September 24.

In some cases, such behavior is not desirable and the protection end time must be preserved.

Action

Set the global option `CopyDataProtectionEndtimeEqualToBackup` to 1 to ensure that the object copy protection end time is equal to backup object protection end time. By default, the option is set to 0. Increase the maximum number of allowed files.

Replicating session with multiple objects stops responding

Problem

When replicating a session onto another device, the session stops responding. The session output provides the following information:

```
[Normal] From: BMA@company.com "d2d1_1_gw1 [GW 26177:1:15198446278003495809]"
Time: 3/21/2013 9:13:06 AM

COMPLETED Media Agent "d2d1_1_gw1 [GW 26177:1:15198446278003495809]"
```

The problem is known to occur in a dual IP stack network configurations with HP-UX Media Agent.

Action

When configuring a dual IP stack network, add a separate entry for IPv6 localhost addresses to the `/etc/hosts` file on the Media Agent client.

For example, you have the following entry in your `hosts` file:

```
::1 localhost loopback
```

To resolve the issue, add the following line for IPv6 addresses:


```
::1 ipv6-localhost ipv6-loopback
```

Replication session on Data Domain Boost devices is unable to respond to Abort operation during retry period

Problem

When replicating a session from one Data Domain Boost backup device to another when the device does not have enough available streams, the replication session is unable to respond to Abort operations during the retry period.

Action

The problem is known to occur when the `omnirc DP_DDBOOST_SLEEP_SECOND_FOR_STREAM_LIMIT` is set to 0, which is not supported.

This variable defines how many seconds the replication session will wait before beginning another retry when the Data Domain Boost device does not have enough available streams. If the interval is too large or is set to 0, the session will be unable to respond to Abort operations.

The default for `DP_DDBOOST_SLEEP_SECOND_FOR_STREAM_LIMIT` is 60 seconds.

See the `omnirc` file for a complete description of `DP_DDBOOST_SLEEP_SECOND_FOR_STREAM_LIMIT`.

Object consolidation problems

Object consolidation of many points in time opens too many files

Problem

If you start an object consolidation operation with many points in time, Data Protector reads all media necessary to complete the operation. This opens all files at the same time. When Data Protector opens more files than the number allowed by your operating system, a message similar to the following one is displayed:

```
|Major| From: RMA@computer.company.com "AFL1_ConsolidateConc2_bs128" Time: time  
/omni/temp/Cons_Media/AFL1/
```

```
0a1109ab54417fab351d15500c6.fd
```

```
Cannot open device ([24] Too many open files)
```

Action

Increase the maximum number of allowed files.

HP-UX systems:

1. Set the maximum number of open files using the System Administration Manager (SAM):
 - a. Select **Kernel Configuration > Configurable parameters** and then, **Actions > Modify Configurable Parameter**.
 - b. Enter the new **maxfiles_lim** and **maxfiles** values in the **formula/value** field.

2. Restart your computer after applying the new values.

Solaris systems:

1. Set the maximum number of open files by editing the `/etc/system` file. Add the following lines:

```
set rlim_fd_cur=value  
set rlim_fd_max=value
```
2. Restart your computer after applying the new values.

Object consolidation to B2D devices fails in the second attempt

Problem
After the first object consolidation, if you perform an incremental backup and then perform the second object consolidation, the operation fails.
Action
To ensure that the second consolidation succeeds, perform a full backup after the first object consolidation. Thereafter, perform an incremental backup, which can be consolidated later.

About Replication

The Data Protector replication functionality enables you to replicate objects between two Backup to Disk (B2D) devices capable of replication, without transferring data through Media Agents. You can select a backup session, object copy session, or object consolidation session. During the replication session, Data Protector reads the object from the session being replicated and initiates the replication from the source B2D device to the target device.

The result of a replication session is a copy of all objects from the session you specified.

The following characterize the replication functionality:

- Start of session
A replication session can be started interactively or automatically.
- Selection of target devices
You can filter the devices capable of replication and select an appropriate one.
- Protection policy
You can set the protection periods for the source objects and the object copies independently.

You can start a replication session interactively or specify an automated start of the session.

Automated replication

In an automated replication specification, you can specify one or more criteria for the selection of object versions that will be copied:

- Backup specifications - to copy only object versions backed up using specific backup specifications.

- Object copy specifications - to copy only object versions copied using specific object copy specifications.
- Object consolidation specifications - to copy only object versions consolidated using specific object consolidation specifications.
- Data protection - to copy only protected object versions.
- Number of existing copies - to copy only object versions that do not have more than the specified number of successful copies.
- Libraries - to copy only object versions located on the media in the specified libraries.
- Time frame (only in a scheduled object copy specification) - to copy only object versions backed up in the specified period of time.

Data Protector offers two types of automated replication: post-backup replication and scheduled replication.

Post-backup replication

Post-backup as well as post-copy and post-consolidation replication, which are subsets of post-backup replication, take place after the completion of a session that is specified in the automated object copy specification. They copy objects selected according to the automated replication specification that were written in that particular session.

Scheduled replication

Scheduled replication takes place at a user-defined time. Objects from different sessions can be replicated in a single scheduled replication session.

Limitations

- You can select only backup, object copy, object consolidation, or object replication sessions for replication. Selecting individual objects is not supported.
- Different block sizes on the source or target device are not supported.
- When configuring interactive sessions, you can select only one session at a time.

Considerations

- Because replication is session based, settings for individual objects may be overridden. For example, if you already have a number of copies of an object included in the session, Data Protector ignores the option **Include only objects with number of copies less than** and replicates all objects in the session including this object even if this results in the object having more copies than allowed with this option.
- By default, Data Protector selects the original object version (when multiple copies of same object are found) as the source device. In some circumstances, the original version may not be capable of replication because it may be a different media type.

Select the correct source device by selecting the library capable of replication or select the specific library.

How to enable replication

You can enable replication from one B2D device to another *when you create an object copy specification*:

1. Make sure that the source and target devices are capable of replication. Use the filter **Capable of replication** to filter the devices, or explicitly select the specific B2D devices.
2. When setting the copy operation options, select **Use replication**.

For the detailed procedure, see standard object copy tasks.

Automated Replication Synchronization

The Data Protector replication feature enables you to replicate objects between two Backup to Disk (B2D) devices that are capable of replication, without transferring data through the Media Agents. The Automated Replication Synchronization feature is an extension of the normal replication, which allows you to replicate backup metadata between two deduplication appliances that are managed by different Cell Managers. This feature allows you to easily exchange backup data and other metadata between two deduplication devices.

Prerequisites

Ensure that the Data Protector user (under whose account the CRS is running) on the source Cell Manager has access to the target Cell Manager.

Considerations

For Integration backups, do not perform an Automated Replication Synchronization procedure from partially failed backup sessions (backup sessions that are completed with errors). The replication will be successful, but the restore from the replicated session may fail.

Limitations

- Consider all the limitations that apply to the normal replication feature.
- The target Cell Manager should have the same or newer version as that of the source Cell Manager.
- The devices in the source Cell Manager and the foreign Cell Manager (target Cell Manager), which are selected for replication, must point to the same physical device and datastore.
- If the source and foreign Cell Managers use Encrypted Control Communication, ensure that **Trust Relationship** is enabled. For more information on establishing trust, see [Enabling encrypted control communication for all cells in a MoM environment, using the CLI](#).
- The maximum number of media that can be replicated at once depends on the available free connections on the target device. For instance, if the target device has 100 free connections, it is recommended that not more than 100 media be replicated at the same time. Also, if you want to use the target device for other operations, the number of media that can be replicated at the same time must be fewer than the available free connections.

For StoreOnce and Data Domain Boost devices, check the available data connections and replication streams respectively. For more information on the supported streams, see the respective device manuals.

- Automated Replication Synchronization list using an older GUI is not supported. You may see the following error message: "Error parsing Copy Specification file. The file may be corrupted or invalid." This message indicates the Data Protector GUI from older versions does not support the new list.
- The **Include only objects with number of copies less than** option is not supported for the Automated Replication Synchronization procedure.

Automated Replication Synchronization involves two procedures:

1. [Importing the foreign Cell Manager](#)
2. [Performing an Object Copy session](#)

Importing the foreign Cell Manager

The first step in triggering the Automated Replication Synchronization is to import the foreign Cell Manager into the source Cell Manager. To import the foreign Cell Manager:

1. In the Context List, click **Clients**.
2. In the Scoping Pane, right-click **Clients** and click **Import Client**.
3. Type the name of the client or browse the network to select the client (on Windows GUI only) that you want to import. If you are importing a Cell Manager that manages a deduplication appliance, select **Data Protector Foreign Cell Server**.

Note: The above step is relevant if you are performing a Automated Replication Synchronization procedure.

4. Click **Finish** to import the client.

The name of the imported client is displayed in the results area.

Note: You can perform only the Automated Replication Synchronization action on the imported cell manager. You will not be able to perform any other operation using the cell manager.

Performing an Object Copy session

After you import the foreign Cell Manager to the source Cell Manager, you can perform an Object Copy session to copy the backup data and other metadata into the foreign Cell Manager. You can perform a Scheduled, Post-backup or Interactive Object Copy based on your requirements.

To perform an Object Copy session:

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, navigate to **Copy > Object copy > Automated**.
3. Right-click **Scheduled** and click **Add** to open the wizard. You can also execute an interactive or post-backup Object Copy session.
4. Select the backup, object copy, or object consolidation specifications that contain the objects you want to copy. Click **Next**.

5. Specify the object filter for the object copy operation. Only objects that match the specified criteria will be copied. Click **Next**.
6. Specify the library filter for the object copy operation. Only objects residing on media in the specified libraries will be copied. Click **Next**.
7. The devices that are used for backup in the selected backup specifications are used as source devices in the object copy operation by default. You can change the source devices here if desired. Click **Next**.
8. Select the destination devices for the object copy operation. Data Protector will select the most suitable devices from those you specify here. Click **Next**.

Select the **Show Capable of replication** check box to select only those devices that have backup to disk (deduplication) devices. Replication is possible only on Backup to Disk devices.

9. Specify the source object options, target object options, and target media options as desired.

Select **Use replication** to enable replication between two B2D devices instead of copying.

Select **Replicate to foreign cell** to enable replication of objects to the foreign cell server that you imported earlier (This Cell Manager contains the second deduplication device).

Click **Next**.

10. Select the foreign cell server that you imported earlier from the drop down menu. This lists the devices that are linked to the backup to disk store.

All the devices that are created from the target cell manager, and having the same store name are displayed here. Therefore, ensure that you select the device having the correct store name for replication.

Select the required device or gateway and click **Next**.

11. Right-click a date and click **Schedule** to display the Schedule Copy dialog box. Specify the options as desired and click **OK**. Click **Next**.
12. Click **Save as...**, enter a specification name and click **OK** to save the scheduled object copy specification.

Run the scheduled object copy session to complete the Automated Replication Synchronization procedure.

About Object Mirroring

The Data Protector object mirror functionality enables writing the same data to several media sets simultaneously during a backup session. You can mirror all or some backup objects to one or more additional media sets.

The result of a successful backup session with object mirroring is one media set containing the backed up objects and additional media sets containing the mirrored objects. The mirrored objects on these media sets are treated as object copies.

Benefits of object mirroring

The use of the object mirror functionality serves the following purposes:

- It increases the availability of backed up data due to the existence of multiple copies.
- It enables easy multi-site vaulting, as the backed up data can be mirrored to remote sites.
- It improves the fault tolerance of backups, as the same data is written to several media. A media failure on one medium does not affect the creation of the other mirrors.

Limitations

- It is not possible to mirror objects backed up using the ZDB to disk or NDMP backup functionality.
- It is not possible to mirror an object to the same device more than once in a single session.
- Block size of the devices must not decrease within a mirror chain. This means the following:
 - The devices used for writing mirror 1 must have the same or a larger block size than the devices used for backup.
 - The devices used for writing mirror 2 must have the same or a larger block size than the devices used for writing mirror 1, and so on.

How to use object mirroring

You specify object mirroring when configuring a backup specification. In the backup specification, select the objects you want to mirror, and then specify the number of mirrors. To be able to specify more than 5 mirrors, increase the value of the `MaxNumberOfMirrors` global option.

Specify separate devices for the backup and for each mirror. When a backup session with object mirroring starts, Data Protector selects the devices from those you specified in the backup specification. To avoid impact on performance, it is recommended that the devices have the same block size and are connected to the same system or to a SAN environment. The minimum number of devices required for mirroring SAP MaxDB, DB2 UDB, or Microsoft SQL Server integration objects equals the number of devices used for backup.

Object mirroring is load balanced by default. Data Protector makes optimum use of the available devices by utilizing as many devices as possible. When you perform an object mirror operation from the command line, load balancing is not available.

Copying a Medium

You can copy media for archiving or vaulting purposes. You need to start the copying of each medium separately, as only one medium can be copied in a media copy session.

Copying a medium in a standalone device

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Devices**, right-click the device with the medium you want to copy and click **Copy**.
3. Select the device (library's drive and slot) where the target medium is located and then click **Next**.

4. Select the media pool to which you want to add the medium copy and then click **Next**.
5. Specify the description and location for the medium copy (optional), and then click **Next**.
6. Specify additional options for the session: you can select the **Force operation** option, specify the medium size and medium protection.

Tip: Use the **Force operation** option if the target media have other formats recognized by Data Protector (tar, OmniBack I, and so on) or if they are Data Protector media without protection.

7. Click **Finish** to start copying and exit the wizard.

The Session Information message displays the status of the media copy operation.

Copying a medium in a library device

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, under **Media**, expand **Pools**, and then expand the media pool that has the medium you want to copy. Right-click the medium and click **Copy** to open the wizard.
3. Select a drive for the medium you want to copy and click **Next**. This step is skipped if the library has only one drive.
4. Select the device (library's drive and slot) where the target medium is located and then click **Next**.
5. Select the media pool to which you want to add the medium copy and then click **Next**.
6. Specify the description and location for the medium copy (optional), and then click **Next**.
7. Specify additional options for the session: you can select the **Force operation** option, specify the medium size and medium protection.

Tip: Use the **Force operation** option if the target media have other formats recognized by Data Protector (tar, OmniBack I, and so on) or if they are Data Protector media without protection.

8. Click **Finish** to start copying and exit the wizard.

The Session Information message displays the status of the media copy operation.

Scheduling Media Copying on Specific Dates

You can schedule a media copy operation on a specific date at a specific time.

You can schedule the media copying while you are adding a new scheduled media operation. To modify the scheduled time of an existing scheduled media operation, follow these steps:

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Automated Operations**. All configured automated operations are displayed.

3. Click the scheduled media copy operation for which you want to change the schedule options, and click the **Schedule** tab.
4. In the Schedule page, scroll through the calendar (clicking the single arrows) for the month in which you want to make the changes.
5. Right-click the unwanted dates that are selected, and click **Delete**. Right-click new dates and click **Schedule** to display the Schedule Media Operation dialog box.
6. Specify the options as desired and click **OK**.
7. Click **Apply**.

Tip: You can click **Reset** to remove all previous schedules.

Scheduling Periodic Media Copying

You can schedule a media copy operation so that it is performed periodically.

You can schedule the media copying while you are adding a new scheduled media operation. To modify the schedule of an existing scheduled media operation, follow the steps below.

Steps

1. In the Context List, click **Devices & Media**.
2. In the Scoping Pane, expand **Automated Operations**. All configured automated operations are displayed.
3. Click the scheduled media copy operation for which you want to change the schedule options, and click the **Schedule** tab.
4. In the Schedule page, right-click a date and click **Schedule** to display the Schedule Media Operation dialog box.
5. Under Recurring, select **Daily**, **Weekly**, or **Monthly**. Specify the **Recurring options** accordingly.
6. Under Time options, select the time when the operation will be performed. Select **Use starting** and specify the starting date.

Note: If you set the recurring to 2 or more (for example, every 2 weeks on Saturday) without setting the starting date, the first copy session may not be scheduled on the first possible date that matches your selection (for example, it will be scheduled on the second Saturday) due to the Data Protector scheduling algorithm. Check the schedule in the Schedule property page.

7. Click **OK** and then **Apply**.

If the chosen time slot is already occupied, Data Protector prompts you that there are scheduling conflicts, and asks if you wish to continue. If you click Yes, the new schedule will be applied where possible (on the days when the time slot is still free). If you click No, the new schedule will be discarded.

Tip: You can click **Reset** to remove all previous schedules.

Customizing the Schedule Calendar

You can customize the appearance of the calendar that is used for scheduling various tasks, such as a backups, automated media copying, and report generation.

You can customize the calendar when scheduling one of the scheduled operations, or when reviewing the schedule. After you have opened the Schedule property page of the scheduled operation, do the following:

Steps

1. In the Schedule property page, right-click a month name and select the desired option from the pop-up menu.
2. Customize the calendar as desired and click **OK**.

Chapter 12: Object Verification

About Object Verification

The Data Protector object verification functionality allows you to verify backup objects. Using this functionality, you no longer have to interactively verify only single complete backup media. Now you can verify single or multiple objects, on single or multiple media, interactively, in scheduled sessions or in post-operation sessions.

The objects verified can be original backup objects, object copies and consolidated objects.

Data verification

During an object verification session, Data Protector verifies the data of individual backup objects in a similar way to that used when verifying a medium.

Delivery to host

By default, the target host, on which the data verification process is performed, is the original backup source host. This verifies the ability of Data Protector to deliver the backup data from the media agent host to that host. Alternatively, a different target host can be specified, or verification can be performed on the Media Agent host, avoiding any network involvement.

Types of object verification session

You can start an object verification session interactively or specify an automated start to the session. Data Protector offers two types of automated object validation: post-backup object verification and scheduled object validation.

Post-backup object verification

Post-backup object verification is performed immediately after the completion of backup, object copy, or consolidation sessions and verifies the objects created during those sessions. Objects to be verified are specified in a post-backup object verification specification. This specifies the backup, object copy and/or consolidation specifications defining the objects created and provides criteria for filtering the objects. Multiple backup, object copy and/or consolidation specifications can be included in a single post-backup object verification specification.

Scheduled object verification

Scheduled object verification is performed at times specified in the Data Protector scheduler and verifies backup, copy or consolidation object versions created during a specified time period. The objects to be verified, and the valid time period for object version creation, are specified in a scheduled object verification specification. This specifies the backup, object copy and/or consolidation specifications defining the objects

created and provides criteria for filtering the objects. Multiple backup, object copy and/or consolidation specifications can be included in a single scheduled object verification specification.

How to Verify Objects

First, start an interactive session, or create an object verification specification. Select the backup objects that you want to verify, source devices, media, and verification target host.

Selection of backup objects

Automated operation

For automated object verification specifications, you can select objects to verify by selecting backup, object copy or consolidation specifications and then filtering according to protection, number of copies, available libraries or time frame (scheduled only). In this case, it is not possible to select individual object versions for verification: Data Protector verifies all object versions matching the filter criteria.

Interactive operation

For interactive sessions, you can select individual objects from media, sessions or the Objects selection wizard listings in the IDB. In this case, it is possible to select individual copies of the required object versions for verification.

Selection of a source device

By default, Data Protector performs automatic device selection. Alternatively, you can force original device selection or select a new device.

Selection of target host

By default, Data Protector performs the verification process on the source host, that is, the host on which the source objects of the original backup were located, verifying the object data and its delivery. You can also specify an alternative remote host, or the Media Agent host, verifying the object data only. Note that the selected target host must have a Data Protector Disk Agent installed.

Scheduling

Scheduling for scheduled verification operations is performed in the same way as for backups, using the standard Data Protector scheduler.

Standard Object Verification Tasks

Below are the prerequisites and limitations of the object verification functionality:

Prerequisites

- You need a Media Agent installed on every system that will act as a source host in object verification sessions.
- You need a Disk Agent installed on every system that will act as a target host in object verification sessions.
- All Disk Agents involved in object verification processing must be at A.06.11 or later.
- The necessary devices should be configured and the media prepared.
- You need appropriate user rights on both the source and destination hosts to be able to run an object verification session: These are Start restore and Restore from other users user rights.
- If the destination host is a UNIX host, you must have Restore as root permissions.

Limitations

- While the source media are being read, they are unavailable for restore.
- Object verification for application integration objects consists of verifying that the object data is delivered to the target host and that it is consistent from the Data Protector format point of view. No application integration specific checks are performed.
- Object verification is not available for objects backed up using ZDB to disk or the disk part of ZDB to disk+tape.
- The use of Web Reporting with object verification is not supported.

Verifying Objects Interactively

You can select objects for interactive verification from the Media, Objects or Sessions starting point, depending on your needs. You cannot save an interactive object verification specification, you can only start an object verification session.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Verification**, and then expand **Object Verification**.
3. Expand **Interactive**.
4. Click **Media**, **Objects** or **Sessions** to open the wizard.
 - Clicking **Media** lists available media to which objects have been written.
 - Clicking **Objects** lists the objects that have been written to the available media.
 - Clicking **Sessions** lists the sessions in which objects have been written to the available media.
5. Select the objects that you want to verify.

Note: From Data Protector 9.09 onwards for VMware backups, the virtual machine disks are considered as objects that run in parallel. The disk objects of the virtual machine are listed but disabled in the **Media** list to understand virtual machine disks backed to the media. The copy or verify operation is performed on the virtual machine objects and all its associated disk objects are considered internally.

From Data Protector 9.09 onwards for VMware integration, the **Next** option is enabled only after selecting the virtual machine object in the **Media** list.

Click **Next**.

6. Select the source device from which the objects will be read. By default, automatic device selection is selected.

You can also force original device selection, or you can substitute another drive by right-clicking **Original Device** and selecting **Change Device**.

Click **Next**.

7. Select the target host for the object verification operation. This host must have a Data Protector Disk Agent at the required version level installed.

By default, the original backup source host is selected. You can also select the Media Agent host (on which the selected source device is installed) or an arbitrary host from the cell that has a Disk Agent at the required version level installed. Click **Next**.

8. A list of media containing the selected objects is displayed. You can change the media location priority to influence the selection of media in cases where the same object resides on more than one media set.

Click **Next**.

9. A summary of the object versions selected for verification is displayed.

- To display details for a particular object version, select it in the list and click **Properties**.
If more than one copy of an object version exists, by default Data Protector selects the one most suitable for verification. You can manually select which copy to verify in Properties.
Click **OK**.
- To remove an object version from the list, select it in the list and click **Delete**.

10. Click **Finish** to close the wizard and start the verification.

Configuring Post-Backup Object Verification

Post-backup object verification is configured to take place after the completion of a backup session, object copy session, or object consolidation session.

The names of the backup, object copy, and/or consolidation specifications concerned are selected in an automated object verification specification. When a session using any of these selected specifications is run, after completion of that session, Data Protector verifies the objects produced during the session, using the criteria specified in the object verification specification.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Verification**, and then expand **Object Verification**.
3. Expand **Automated**, right-click **Post Backup** and select **Add** to open the wizard.
4. Select the backup specifications that you want to be immediately followed by the object verification specification. Click **Next**.
5. Select the object copy specifications that you want to be immediately followed by the object verification specification. Click **Next**.
6. Select the consolidation specifications that you want to be immediately followed by the object verification specification. Click **Next**.
7. Specify an object filter for the object verification operation, if required. Only objects meeting the specified criteria will be verified. Click **Next**.
8. Specify a library filter for the object verification operation, if required. Only objects contained on media in the specified libraries will be verified. Click **Next**.
9. Select the source device from which the objects will be read. By default, Data Protector uses automatic device selection.

Alternatively, you can force selection of the original device. This means that, if the device is not available, Data Protector will wait until it is available. You can also substitute another drive for the original by right clicking **Original Device** and selecting **Change Device**, for instance, following replacement of the original device by a new one.

Click **Next**.

10. Select the target host for the object verification operation. This host must have a Data Protector Disk Agent installed.

You can select:

- the host on which the original backup object was produced (default selection). This also verifies the Data Protector components in the network path.
- the Media Agent host, that is, the host with the source device, without any network involvement.
- an alternative remote host, verifying the Data Protector components in the network path to that host.

Click **Next**.

11. Click **Save as...**, enter a specification name and click **OK** to save the verification specification.

Configuring Scheduled Object Verification

Scheduled object verification takes place at a user-defined time. Objects generated by different backup sessions, object copy sessions, or object consolidation sessions can be verified in a single scheduled object verification session.

Tip: You can also schedule object verifications sessions with advanced settings with the Advanced Scheduler.

Steps

1. In the Context List, click **Object Operations**.
2. In the Scoping Pane, expand **Verification**, and then expand **Object Verification**.
3. Expand **Automated**, right-click **Scheduled** and select **Add** to open the wizard.
4. Select the backup specifications defining the output objects for which you want to schedule verification. Click **Next**.
5. Select the object copy specifications defining the output objects for which you want to schedule verification. Click **Next**.
6. Select the consolidation specifications defining the output objects for which you want to schedule verification. Click **Next**.
7. Specify an object filter for the object verification operation, if required.
This lets you filter the available objects according to protection, numbers of copies or time of creation. All object versions that meet the filter criteria will be verified.
Click **Next**.
8. Specify a library filter for the object verification operation if required. Only objects contained on media in the specified libraries will be verified. Click **Next**.
9. Select the source device from which the objects will be read. By default, Data Protector uses automatic device selection.
Alternatively, you can force selection of the original device. This means that, if the device is not available, Data Protector will wait until it is available. You can also substitute another drive for the original by right clicking **Original Device** and selecting **Change Device**, for instance, following replacement of the original device by a new one.
Click **Next**.
10. Select the target host for the object verification operation. This host must have a Data Protector Disk Agent installed.
You can select:
 - the host on which the original backup object was produced (default selection). This also verifies the Data Protector components in the network path.
 - the Media Agent host, that is, the host with the source device, without any network involvement.
 - an alternative remote host, verifying the Data Protector components in the network path to that host.Click **Next**.
11. Click on the date from which you want the operation to be performed and click **Add** to display the Schedule Verification dialog box.
12. Specify the required time and frequency for verification sessions. For example:

You can schedule recurring sessions over periods of days, weeks or months, if required.

Click **OK**.

13. Click **Save as...**, enter a specification name and click **OK** to save the verification specification.

Customizing the Object Verification Environment

You can customize the object verification environment by modifying the message level and session status generated when there are no objects to verify for a verification session. To achieve this, modify the `SessionStatusWhenNoObjectToVerify` global option.

Chapter 13: Restore

About Restore

A restore is a process that recreates the original data from a backup copy to a disk. This process consists of the preparation and actual restore of data and, optionally, some post-restore actions that make that data ready for use.

For more information on the concept of restore, see the *HPE Data Protector Concepts Guide* and the *HPE Data Protector Integration Guides*.

Depending on the platform, the way you specify these features and available options can vary.

For information on how to restore with application integrations to applications such as Oracle, SAP R/3, Microsoft Exchange Server, Microsoft SQL Server, Informix Server, IBM DB2 UDB or Sybase, see the *HPE Data Protector Integration Guides*.

Standard Restore Procedure

A standard restore procedure consists of several phases.

1. [Selecting the data to restore.](#)
2. [Finding the necessary media.](#)
3. [Starting the restore session.](#)

Other settings are predefined according to the backup process, but can be modified.

Prerequisite

To perform a restore you must have the appropriate user rights. These rights are defined according to the user group.

Selecting the Data to Restore

You can browse for data to restore in two possible ways: either from the list of the backed up objects or from the list of sessions. The difference is in the scope of directories and files presented for restore:

- **Restore Objects** with a list of backed up objects classified by client systems in the cell and by different data types (such as Filesystem, Disk Image, Internal Database, and so on). You can browse all the directories, files and versions, which were backed up and are still available for restore.
- **Restore Sessions** with a list of filesystem sessions with all objects backed up in these sessions. You can choose to view only sessions from the last year, last month, or last week. You can browse all objects that were backed up in this session (like any drives from all clients named in the backup specification), and all versions of this restore chain. By default, the entire restore chain of the selected directories or files is restored, but you can also restore data from a single session only.

Prerequisite

In order to browse objects and select directories or specific files, the corresponding backups must have been done using a logging level of directory, filenames, or log all.

Selecting the data from the list of the backed up objects

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, under Restore Objects, expand the appropriate data type (for example, **Filesystem**).
3. Expand the client system with the data you want to restore and then click the object (mountpoint on UNIX systems, drive on Windows systems) that has the data.
4. In the Source property page, expand the object and then select directories or files that you want to restore.

By default, when you select a whole directory, only directories and/or files from the last backup session are selected for restore. Directories and files in the same tree structure that have not been backed up in the same backup session are shaded. If you want to restore the data from any other backup session, right-click the selected directory and click **Restore Version**. In the Backup version drop-down list, select the backup version that you want to restore from.

Tip: If you repeat the steps above and select data under more than one object (mountpoint or drive), you can perform a parallel restore.

Selecting the data from the list of the backup sessions

Limitations

- You cannot perform the restore of an online database integration from a specific backup session.
- You cannot use "Restore Sessions" mode to perform a restore from a copy session.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand the **Restore Sessions** to display clients and then objects, backed up on a particular client. Click an object to open the object's property pages.
3. In the **Source** page, select directories and files to be restored.

By default, the entire restore chain is restored (**Show full chain** is selected). To restore only data from this session, select **Show this session only**.

4. Specify the restore destination and set the restore options.
5. Click **Restore** to start the restore session.

Tip: To perform a parallel restore, repeat steps 2 to 4 for additional objects before starting the restore.

Selecting a Specific Backup Version

After selecting the data that you want to restore, you can select its backup version.

Selecting the backup version for each file or directory separately

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, under Restore Objects, expand the appropriate data type (for example, Filesystem).
3. Expand the client system with the data you want to restore and then click the object that has the data.
4. In the Source property page, select the object to restore. By default, the latest backup version is selected for restore.
5. Right-click the object and click **Restore Version**.
6. In the Backup version drop-down list, select the backup version that you want to restore. Click "..." if you need more information on the backup versions. The "..." button is available if the backup was performed using a logging level that logs attributes.
7. Click **OK**.

After you have selected a version for restore, only the files and directories from this version are shown as available for restore in the Source property page. Other files and directories are grayed and will not be restored.

Selecting the backup version for several files or directories simultaneously

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, under Restore Objects, expand the appropriate data type (for example, **Filesystem**).
3. Expand the client system with the data you want to restore and then click the object that has the data.
4. In the Source property page, select multiple objects to restore. By default, the latest backup version is selected for restore.
5. Click the **Restore Summary** tab, select all objects, right-click the selection and then click **Select Version By Time**.
6. Click the **Select version by date and time** option, and select the day from the pop-up menu.

7. You can enter the time by clicking on the displayed time in the **Select version by date and time** drop-down list.
8. Under **Differences in backup time**, make any necessary adjustments in case there is no backup version corresponding to your date and time selection for any of the selected objects.
9. Under **If selected date and time doesn't match with selected criteria**, make any necessary adjustments in case there is no backup version corresponding to your date and time selection and to **Differences in backup time** correction for any of the selected objects.
10. Click **OK**.

After you have specified the criteria for restore, the backup versions corresponding to your selection are shown in the Source property page next to every object to be restored.

Handling File Conflicts

You can choose how to resolve conflicts between the file version currently on the disk and the version from the backup.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, under Restore Objects, expand the appropriate data type (for example, **Filesystem**).
3. Expand the client system with the data you want to restore and then click the object that has the data.
4. In the Source property page, select the disk, directories, or files to be restored.
5. Click the **Destination** tab and then, under File Conflict Handling, select one of the available options:
 - **Keep most recent**
 - **No overwrite**
 - **Overwrite**

Selecting a Device to Restore From

By default, Data Protector restores selected data with the same devices that were used during backup. However, you can select alternative devices for your restore.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, under Restore Objects, expand the appropriate data type (for example, **Filesystem**).
3. Expand the client system with the data you want to restore and then click the object that has the data.

4. In the Source property page, expand the object and then select what you want to restore.
5. Click the **Devices** tab to open the Devices property page.

The devices that were used during backup are listed here.

To restore your data with an alternative device, select the original device and click **Change**. In the Select New Device dialog box, select the alternative device and click **OK**. The name of the new device appears under Device Status. The new device will be used only for this session.

For more information on a device, right-click the device and click **Info**.

Specify what Data Protector should do if the selected devices are not available during restore (for example, if they are disabled or already in use). Select either **Automatic device selection** or **Original device selection**.

Finding Media Needed to Restore

After selecting the data that you want to restore, you need to get a list of media containing the data. This is essential if you use standalone devices or if you keep media outside the library.

If an object version that you want to restore exists on more than one media set, you can influence the selection of the media set that will be used for the restore by setting the media location priority, or manually select the media set that will be used.

If you use synthetic backup, there is often more than one restore chain of the same point in time of an object. By default, Data Protector selects the most convenient restore chain and the most appropriate media within the selected restore chain.

Note: Copies obtained using the media copy functionality are not listed as needed media. A medium copy is used only if the original medium (the medium that was used as a source for copying) is unavailable or unusable.

Limitations

- With some integrations, it is not possible to set the media location priority in the Restore context. The GUI does not display the Media tab for these integrations.
- You cannot manually select the media set when restoring integration objects.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, under Restore Objects, expand the appropriate data type (for example, **Filesystem**).
3. Expand the client system with the data you want to restore and then click the object that has the data.
4. In the Source property page, expand the object and then select what you want to restore.
5. Click the **Media** tab to open the Media property page. The needed media are listed. For more information on a medium, right-click it and click **Info**.

If an object version that you want to restore exists on more than one media set, all media that contain the object version are listed. The selection of the media set depends on the Data Protector internal media set selection algorithm combined with the media location priority setting.

- To override the media location priority setting, select a location and click **Change priority**. Select a different priority for the location and click **OK**.
- To manually select the media set from which you want to restore, click the **Copies** tab. In the Copies property page, select the desired object version and click **Properties**. Select the **Select source copy manually** option, select the desired copy from the drop-down list, and click **OK**.

6. If necessary, insert the media into the device.

Tip: You can also list the media needed for restore, including media containing object copies of the selected objects, by clicking **Needed media** in the Start Restore Session dialog box. This dialog box appears when you start the restore.

Previewing and Starting a Restore

Prerequisites

- Ensure that the needed media is available or loaded in the device.

Limitations

- Preview is not available for the Data Protector Internal Database restore and the restore sessions of Data Protector application integrations.

Steps

1. Select what you want to restore and specify options in the restore property pages, including the selection of the device to be used.
2. Check which media are required for the restore.
3. In the **Actions** menu, click **Preview Restore** if you want to preview it or **Start Restore** to actually start the restore process. You can also click **Preview** or **Restore** button on a **Property** page.
4. In the Start session wizard, review your selection and specify the **Report level**, **Network load**, and **Enable resumable restore** options.

The Restore Monitor shows the progress of the restore.

Aborting a Restore

Aborting a restore session stops the restore. Data processed before the session was aborted is restored to the specified location.

Steps

1. To abort a restore session, click **Abort** in the **Actions** menu.

Tip: You can abort restore sessions from the Data Protector Monitor context.

Restore Location Options

By default, Data Protector restores the data to the same client and directory from which it was backed up. You can change these default settings in the Destination property page by specifying where to restore the data to:

- with appropriate user rights you can restore to another client system
- you can restore to another directory

The general restore location can be set on a per-object basis.

Additionally, Data Protector offers you the **Restore As/Into** option to specify a different location for individual files and directories from the same backup object.

Selecting Restore Location

After selecting the data that you want to restore, you can select the location to restore the data to. You can restore the data to another client system and change the directory path. This applies to the entire object to be restored.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, under Restore Objects, expand the appropriate data type.
3. Expand the client system with the data you want to restore and then click the object that has the data.
4. In the Source property page, select the object to restore.
5. Click the **Destination** tab and then, in the **Target client** drop-down list, select the client system that you want to restore on the new client. By default, Data Protector uses the original directory structure to restore: if the data was backed up from the C:\temp directory on system A, it restores the data to the C:\temp directory on system B.
6. You can change the directory path for your restore by selecting the **Restore to new location** option and then entering or browsing for a new anchor directory. The directory path at backup time is appended to the new anchor directory: if data was backed up from the C:\sound\songs directory and you enter \users\bing as a new path, the data is restored to the C:\users\bing\sound\songs directory.

Specifying Restore Location for Individual Files and Directories

You can specify an individual restore path for any directory or file within each object. The individual location specified under the **Restore As/Into** option overrides the location specified in the Destination

property page.

This capability is available for the initially selected tree node (directory) and for tree nodes that are not hierarchically dependent on any already selected tree nodes. A selected tree node is indicated by a blue check mark, and a dependent tree node is indicated by a black check mark.

Restore into

Restore into appends the path from the backup to the new location selected here. The new location has to be an existing directory.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, under Restore Objects, expand the appropriate data type.
3. Expand the client system with the data you want to restore and then click the object that has the data.
4. In the Source property page, select the object to restore.
5. Right-click the specific file or directory and then click **Restore As/Into**.
6. Under the Destination tab, in the Restore drop-down list, select **Into**.
7. As an option on Windows systems, you can select another drive in the Drive text box to restore the data to. If you want to restore to another client system, click **Browse**.
8. In the Location text box, enter a new path for the file or directory. The original path is added to the new one: if the `colors.mp3` file was backed up from the `C:\sound\songs` directory and you enter `\users\bing` as a new path, the file is restored to the `C:\users\bing\sound\songs` directory.
9. Click **OK**.

Restore as

Restore as replaces the path from the backup with the new location selected here. The destination path can be a new directory or an existing one. You can rename the files and directories as you restore them.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, under Restore Objects, expand the appropriate data type.
3. Expand the client system with the data you want to restore and then click the object that has the data.
4. In the Source property page, select the object to restore.
5. Right-click the specific file or directory and then click **Restore As/Into**.
6. Under the Destination tab, in the Restore drop-down list, select **As**.
7. As an option on Windows systems, you can select another drive in the Drive text box to restore the data to. If you want to restore to another client system, click **Browse**.
8. In the Location text box, enter a new path for the file or directory. For example, if the `colors.mp3`

file was backed up from the C:\sound\songs directory and you enter \users\bing\colors.mp as a new path, the file is restored to the C:\users\bing directory.

Caution: Consider the risk of deleting data with the **Overwrite** option enabled when:

- specifying to restore under a name that already exists
- entering an existing path without specifying the file or directory name.

For example, when you enter a new path \users\bing in the Location text box to restore file colors.mp, but you didn't enter the name of the file, then colors.mp file will be restored as bing. What used to be the bing directory is deleted and substituted with the restored file.

9. Click **OK**.

About Resuming Failed Sessions

Backup and restore sessions that failed for any of the following reasons can be resumed using the Data Protector resume session functionality:

- Network connectivity problem
- Fatal Disk Agent problem
- Fatal Media Agent problem
- Fatal Session Manager problem
- Fatal media problem (for example, torn tape)
- The Abort command invoked from the GUI

However, you have to resolve the impeding problem first.

When you resume a failed session, Data Protector continues with the backup or restore, starting right from where the failed session left off. The resumed session inherits all the options from the original session.

Not all session types can be resumed. Data Protector can resume the following:

- Filesystem backup sessions
- Filesystem restore sessions
- Data Protector Oracle Server integration backup sessions
- Data Protector Oracle Server integration restore sessions

Filesystem backup sessions

The resume session functionality for filesystem backup sessions is based on the checkpoint file information that is written into the Internal Database. When a backup session fails, the last backed up file is marked as a checkpoint in the Internal Database. Thus, the backup session can continue from the point of failure when the session is resumed. The file at the point of failure is backed up from the beginning, while the remaining data is appended to the original backup session as its incremental backup. The resumed session automatically inherits the options of the original session.

In case the file marked as a checkpoint is deleted from the filesystem, the resume functionality can still determine what data has not been backed up yet. A failed backup session can be resumed multiple times until it is completed successfully.

In the graphical user interface, the session can be resumed using the context menu of the failed session. In command-line interface, the session can be resumed using the `omnib -resume` option.

Limitations

- Resume is not supported for disaster recovery.
- Resume is not supported for sessions containing NDMP medium data format objects.
- Objects backed up with the following backup client systems are not resumable: Solaris 9, SCO OpenServer, and OpenVMS.

Filesystem restore sessions

The resume session functionality for filesystem restore sessions is based on checkpoint files that are created during a restore session and contain information about which restore options are used in the session and which files have been successfully restored. As soon as a new file is restored, the corresponding checkpoint file is updated.

By default, the checkpoint files are created on both the Cell Manager and the destination client (the checkpoint file that contains information about restore options is created only on the Cell Manager).

On the Cell Manager, the checkpoint files are created in:

Windows systems: `\config\server\sessions\checkpoint`

UNIX systems: `/var/opt/omni/server/sessions/checkpoint`

On clients, the checkpoint files are created in the default Data Protector temporary files directory, within the Checkpoint subdirectory.

How the functionality works

When you resume a failed restore session, Data Protector reads information from the checkpoint files and continues with the restore from where the failed restore session left off. Actually, when you resume a restore session, its checkpoint files are moved to the checkpoint file directory of the resumed restore session, where they continue to be updated. Consequently, a failed restore session can only be resumed once. If you try to resume the failed session for the second time, the operation fails because its checkpoint files are no longer there.

Considerations

- In cluster environments, ensure that the checkpoint files are created on a shared disk, so that both cluster nodes can access the files. To change the location for the checkpoint files, use the `OB2CHECKPOINTDIR` omnirc option. The option must be set on both cluster nodes and must point to the same directory.
- You can disable the creation of checkpoint files by clearing the option **Enable resumable restore** before you start a restore session (the option can be found in the Start Restore Session dialog box, at the end of the restore wizard). However, if such a restore session fails, you will not be able to resume it because the checkpoint files will be missing. Successfully completed sessions also cannot be resumed since Data Protector deletes the checkpoint files at the end of such sessions.

- A resumed restore session that did not complete successfully is also resumable. This is due to the fact that a resumed restore session inherits the checkpoint files of the original session. Consequently, it inherits all the restore options used in the original session, including the option **Enable resumable restore**.
- When a restore session is removed from the IDB (by default, a session is removed after 30 days), its checkpoint files are purged as well. Checkpoint files are also purged when you initialize the IDB using the `omnidbinit` command.
- If the **No overwrite** option was used to restore one or more objects in a failed session, the `omnirc` option `OB2NOOVERWRITE_TRAVERSEDIOBJ` must be set to 1 before you resume that session.

Limitations

- If a restore session failed because the destination client crashed, the resume session functionality may not work correctly. It all depends on whether or not the checkpoint files were successfully flushed from the memory to the disk when the client crashed.
- If a restore session failed right when hard-linked files were being restored, the resume session functionality may not be able to restore the remaining hard-linked files. This is due to the fact that, during backup, Data Protector backs up a hard-linked file only once. For other files that are hard-linked to it, it backs up only the reference to the file. Consequently, restore of hard-linked files is interconnected so the files must be restored all together. Note that this problem does not occur if the restore session fails before the hard-linked files start to be restored or after they have been successfully restored.
- Suppose you want to restore a tree that has been backed up in the following sessions: Full, Incr, and Incr. If the restore session fails because the tree backup object created in one of the backup sessions is not available (for example, the backup media used in the last Incr backup session are corrupted), you must provide the copy of that backup object. If such an object copy does not exist, you cannot resume the failed restore session, even if a synthetic full backup of the missing backup object exists.

Data Protector Oracle Server integration backup and restore sessions

The resume session functionality for Data Protector Oracle Server integration backup and restore sessions is described in the *HPE Data Protector Integration Guide*.

Resuming Failed Sessions

Backup and restore sessions that failed (for example, due to network connectivity problems) can be resumed using the Data Protector resume session functionality. When you resume a failed session, Data Protector continues with the backup or restore, starting right where the failed session left off.

Prerequisites

- You either have to be in the Data Protector Admin user group or have the Data Protector Monitor user right.

Steps

1. If you are using an ordinary Cell Manager, in the Context List, click **Internal Database**.
If you are using a Manager-of-Managers, in the Context List, select **Clients** and expand **Enterprise Clients**. Select a Cell Manager with the problematic session. From the Tools menu, select **Database Administration** to open a new Data Protector GUI window with the Internal Database context displayed.
2. In the Scoping Pane, expand **Internal Database** and click **Sessions**.
A list of sessions is displayed in the Results Area. Status of each session is denoted in the Status column.
3. Right-click a failed session, and select **Resume Session**.

Advanced Restore Tasks

You can control a restore in many ways. Data Protector offers a set of the advanced restore tasks for the Windows and UNIX system.

Prerequisites

- To perform a restore you need to have the appropriate user rights. These rights are defined according to the user group.
- You have to consider the standard restore procedure before proceeding.

Advanced restore tasks

Advanced restore tasks include specifying rarely used options or taking some actions that do not follow the standard restore procedure. To restore the data you will still have to perform most of the standard restore steps.

The way you follow the standard restore procedure depends on the advanced task you want to perform. For example, you can restore your data without browsing. In this case, you need to specify the desired files in a different way, but can still follow the standard restore procedure in other steps.

- [Skipping Files for Restore](#)
- [Selecting Only Specific Files \(Matching\) for Restore](#)
- [Selecting Open Files for Restore](#)
- [Denying Access to Files During Restore](#)
- [Searching for a File to Restore](#)
- [Selecting a Windows Shared Disk for Restore](#)
- [Restoring Objects in Parallel](#)
- [Disk Image Restore](#)
- [Restore from Media in a Vault](#)
- [Web Server Restore](#)
- [Restore Without Browsing](#)

Skipping Files for Restore

Data Protector allows you to skip files that were backed up, but you do not wish to restore. By using wildcard characters you can skip files matching a specific pattern.

Note: Skipping files for restore is not supported with Data Protector server integration.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand the appropriate data type (for example, **Filesystem**).
3. Expand the client system with the data you want to restore and then click the object (mountpoint on UNIX systems, drive on Windows systems) that has the data.
4. In the Source property page, select the directory that you want to restore.
5. Right-click the directory and then click **Properties**.
6. Click the **Skip** tab.
7. In the text box, enter the file name or the criteria used to match the files to be skipped (for example, *.mp3) and then click **Add**. In this example, no mp3 files would be restored. To use more criteria, repeat this step.
8. Click **OK**.

Selecting Only Specific Files (Matching) for Restore

Data Protector allows you to restore only those files from the backup that match a specific pattern. By using wildcard characters, you can specify the pattern to be used.

Note: This functionality is not supported with Data Protector NDMP server integration.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand the appropriate data type (for example, Filesystem).
3. Expand the client system with the data you want to restore and then click the object (mountpoint on UNIX systems, drive on Windows systems) that has the data.
4. In the Source property page, select the directory that you want to restore.
5. Right-click the directory and then click **Properties**.
6. Click the **Restore Only** tab.
7. In the text box, enter the file names or enter the criteria to match the files to be restored, for example, *.mp3, and then click **Add**. This will restore only mp3 files. For more criteria, repeat this step.
8. Click **OK**.

Selecting Open Files for Restore

By default, Data Protector does not restore the files that are in use by some other application (open files). You can restore open files following the steps below.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand the appropriate data type (for example, **Filesystem**).
3. Expand the client system with the data you want to restore and then click the object (mountpoint on UNIX systems, drive on Windows systems) that has the data.
4. In the Source property page, expand the object and then select what you want to restore.
5. Click the **Options** tab, and then select the **Move busy files** option.

Denying Access to Files During Restore

By default, Data Protector does not lock files during restore. You can change this default behavior.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand the appropriate data type (for example, **Filesystem**).
3. Expand the client system with the data you want to restore and then click the object (mountpoint on UNIX systems, drive on Windows systems) that has the data.
4. In the Source property page, expand the object and then select what you want to restore.
5. Click the **Options** tab, and then select the **Lock files during restore** option.

Searching for a File to Restore

If you do not know the full path of a file that you want to restore, you can search for the file in the IDB, provided that the logging level at backup time was set to Log Files or Log All. You can search for files and directories using the **Restore by Query** task if you know at least a part of the file name.

Steps

1. In the Context List, click **Restore**.
2. Click the **Tasks** navigation tab at the bottom of the Scoping Pane. The predefined restore tasks are listed in the Scoping Pane.
3. Click **Restore by Query** to open the wizard.
4. Specify a part of the file name, using wildcard characters.

For example, type *.exe to search for all backed up files with this extension.

When specifying non-ASCII characters, ensure that the current encoding in the Data Protector GUI and the encoding that was used when the file was created match. Otherwise, Data Protector will not find the files.

In the environment with a UNIX Cell Manager, the wildcard character `?` will not produce the desired results if you want to find a multi-byte character with it. You need to specify multiple wildcard characters `?`. For example, if 3 bytes are used to represent the multi-byte character in the current encoding, add `???` to your string.

If the directories are available, compare only the base name with patterns. If the directories are not available, compare the full path name with patterns.

5. Optionally, specify other parameters. Click **Next**.
6. Optionally, specify the desired time frame and modification time. Click **Next**.

Data Protector will list all files and directories matching the specified criteria.

7. From the list of files matching the selection criteria, select the files that you want to restore. To specify further options, click the appropriate tab. To specify the **Report level**, **Network load**, and **Enable resumable restore** options, click **Next**. To start the restore, click **Finish**.

Selecting a Windows Shared Disk for Restore

Data Protector allows you to restore to a shared disk, even if the data was not originally backed up from the shared disk.

Reasons to restore a UNIX or Windows filesystem to a Windows shared disk:

- If the system is not a part of the Data Protector cell and does not have the Data Protector Disk Agent installed.
- If you want to restore to platforms not directly supported by Data Protector, such as Windows for Workgroups or Windows 3.1 systems.
- If you want to make the data available from several systems.

When you restore your data to a different filesystem type to the one from which it was backed up (UNIX system to Windows system, for example), filesystem-specific attributes may be lost.

Prerequisite

You must change the Data Protector Inet account on the Disk Agent client in order to have the right permissions to access the shared disk that you want to restore to. This account has to have the permission to access both the local client system and the remote shared disks. It must be a specific user account, not the system account

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, expand the appropriate data type.
3. Expand the client system with the data you want to restore, and then click the object that has the data.
4. In the Source property page, expand the object, and then select what you want to restore.

5. Click the **Destination** tab.
6. In the **Target client** drop-down list, select the Windows client system with the Disk Agent that you will use for restore.

Tip: You can skip the remaining steps if you enter the network path manually by specifying the UNC share name of the remote disk (`\\COMPUTER_NAME\SHARE_NAME`, for example, `\\TUZLA\TEMP`) in the **Restore to new location** text box.

You have to do this if you are using the GUI on a UNIX system, since it is not possible for the system to confirm the existence of a Windows shared drive, or to browse it. Therefore, you must confirm yourself that it is available and correctly specified, or the restore may fail.

7. Select the **Restore to new location** option and then click **Browse** to display the **Browse Drives** dialog box.
8. Expand **Microsoft Windows Network** and select the shared disk to which you want to restore the data.
9. Click **OK**.

Restoring Objects in Parallel

A parallel restore allows you to restore data concurrently from multiple objects to multiple disks or filesystems while reading the media only once, thus improving the speed of the restore.

Prerequisite

At backup time, the data from the different objects must have been sent to the same device using a concurrency of 2 or more.

Limitation

You cannot restore the same object in parallel. For example, if you select for the same restore an object under **Restore Objects** and then select the session that includes the same object under **Restore Sessions**, the object will be restored only once and a warning will be displayed.

Steps

1. Select the data as you would for a single restore. You can also specify the restore destination, options, and so forth.
2. Go back to the Restore context in the Scoping Pane and repeat step 1 for data under other objects you want to restore.
3. In the **Actions** menu, click **Start Restore**. You are informed that you selected multiple objects.
4. Select the **All selected objects (parallel restore)** option and click **Next**.
5. In the Start session wizard review your selection. Click **Next**.
6. Specify the **Report level**, **Network load**, and **Enable resumable restore** options and click **Finish** to start the restore of objects in parallel.

Disk Image Restore

A disk image restore is a fast restore of a corresponding disk image backup. Data Protector restores the complete image of a disk, sector-by-sector instead of only restoring selected files or directories.

To restore a UNIX or Windows disk image, expand the **Disk Image** object under the Restore context and then use the standard restore procedure.

Prerequisites

- The backup to be restored has to be of disk image type.
- On UNIX systems, you need to dismount a disk before a disk image restore and mount it back after the restore using the pre- and post-exec commands (for example, pre-exec: `umount /dev/rdisk/disk1`, post-exec: `mount /dev/rdisk/disk1 /mount_dir`).
- If you want to restore a disk image on a disk other than the disk from which you backed it up, the new disk must be of the same size or larger.

Restore from Media in a Vault

Restoring from a medium that comes from a vault is very similar to restoring from any other medium. Depending on how the data and catalog protection policies are defined, however, you may need to do some additional steps:

- If you have a library, enter the medium and scan it.
- If the catalog protection for the medium is still valid, restore the data by selecting what you want to restore using the Data Protector user interface.
- If the catalog protection for the medium has expired, Data Protector does not have detailed information about the data backed up. Restore the data by manually specifying the files or directories that you want to restore.

Tip: To re-read the detailed information about the files and directories from the medium after the catalog protection has expired, export the medium, import it back, and specify that you want to read the Detail Catalog data. After that, you will be able to browse the files and the directories in the Data Protector user interface.

Web Server Restore

To restore a web server, use the standard restore procedure for restoring files, directories, and clients. Additionally, you need to consider the following:

- All data should be restored to the original location.
- Configuration files and root directories should always be included.
- During restore, the web server should be down, however the operating system must be up and running. Restart the web server after the restore.

In case a database, such as Oracle or Informix Server, is included on the web server, use the restore procedure specific for the database.

Restore Without Browsing

When the catalog protection for the data has expired or when the backup was done using the No Log or Log Directories option, you can manually specify a file or a directory for restore.

In case you do not know a file or a directory name, you can restore the entire object and then extract the parts that you need or you can use the **Restore only** feature to restore only files which match a specific pattern and then extract the parts that you need from them.

Restoring the Entire Object and Extracting the Needed Parts

When you are not able to browse for a file or directory you want to restore, you can restore the entire object and then extract only the parts you need.

Prerequisite

To restore the entire object, you need a temporary storage area as large as the entire object.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, under Restore Objects, expand the appropriate data type (for example, **Filesystem**).
3. Expand the client system with the data you want to restore, and then click the object you want to restore.
4. Click the **Destination** tab. Select a temporary directory that is large enough to store the entire object.
5. Specify options in the other restore property pages, including the selection of the device to be used.
6. In the **Actions** menu, click **Preview Restore** if you want to preview it or **Start restore** to actually start the restore process.
7. In the Start session wizard, review your selection and specify the **Report level**, **Network load**, and **Enable resumable restore** options. The Restore Monitor shows the progress of the restore.
8. When the restore is finished, you can extract the needed parts of data from the restored object and copy them to the desired location. Note that you do this outside Data Protector.

Restoring Parts of the Backed Up Object Using Restore-Only Pattern Match

When you are not able to browse for a file or directory you want to restore, the directory (or a file or a higher level directory) can be hit using a pattern match that avoids the restore of most unwanted parts of the object. By using wildcard characters, you can specify the pattern to be used.

Note: This functionality is not supported with Data Protector NDMP server integration.

Prerequisites

- You need to use a fairly specific pattern definition for this feature to be beneficial.
- You need a temporary storage area for the restored parts. Its size depends on the size of the restored object parts, which is connected to the precision of the matching pattern used.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, under Restore Objects, expand the appropriate data type (for example, **Filesystem**).
3. Expand the client system with the data you want to restore and then click the object you want to restore.
4. In the Source property page, right-click the object you want to restore from and then click **Properties**.
5. Click the **Restore Only** tab and in the text box specify the pattern to match the files to be restored (for example, "*order*40*.ppt") and then click **Add**. You should add several such patterns to specify as precisely as possible the type of files to be restored.
6. Click **OK**.
7. Click the **Destination** tab. Select a temporary directory that is large enough to store the parts of the backed up object.
8. Specify options in the other restore property pages, including the selection of the device to be used.
9. In the **Actions** menu, click **Preview Restore** if you want to preview it or **Start restore** to actually start the restore process.
10. In the Start session wizard, review your selection and specify the Report level, Network load, and Enable resumable restore options. The Restore Monitor shows the progress of the restore. If you selected a "Warning" report level, Data Protector issues a Warning message because the list of files and directories is not in the IDB catalog. This does not influence the restore.
11. When the restore is finished you can extract the needed parts of data from the restored object and copy them to the desired location. Note that you do this outside Data Protector.

Restoring the File or Directory Manually

When you are not able to browse for a file or directory you want to restore, you can specify a file or a directory manually. This happens when the catalog protection for your data has expired, or when backup was done using the **No log** option.

Prerequisite

To add a file or a directory manually, you need to know the exact path and the name of the file or the directory. The file and path names are case-sensitive.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, under Restore Objects, expand the appropriate data type (for example, **Filesystem**).
3. Expand the client system with the data you want to restore, right-click the object that has the file or directory that you want to restore manually, and then click **Properties**.
4. Click the **Restore Summary** tab and then enter the missing part of the path and the name of the file or directory you want to restore in the text box.
5. Click **Add** to confirm. The Version window appears.
6. From the Version drop-down list, select the backup version you want to restore and then click **OK**. The object name and version are displayed.
7. Specify options in the other restore property pages, including the selection of the device to be used.
8. In the **Actions** menu, click **Preview Restore** if you want to preview it or **Start restore** to actually start the restore process.
9. In the Start session wizard, review your selection and specify the **Report level**, **Network load**, and **Enable resumable restore** options.

The Restore Monitor shows the progress of the restore. If you selected a "Warning" report level, Data Protector issues the Warning message because the list of files and directories is not in the IDB catalog. This does not influence the restore.

Restore Options

Data Protector offers a set of comprehensive restore options that allow fine-tuning of a restore. All these options have default values which are appropriate in most cases.

The following list of options is set on a per-object basis. The restore options are available according to the type of data being restored.

For detailed information on restore options, see HPE Data Protector Help.

General restore options

- **Show full chain.** Displays all the files and directories in the restore chain. By default, this option is selected and the entire restore chain is restored.
- **Show this session only.** Displays only the files and directories backed up in this session. This enables you to restore files and directories from an incremental backup session without restoring the entire restore chain. By default, this option is disabled.
- **Target client.** By default, you restore to the same client system from which the data was backed up. You can select another system in your cell from the drop-down list. The Disk Agent is started on the selected client system and the data is restored there.

You need to have the **Restore to other clients** user right to be able to restore to another client system.

- **Omit deleted files.** For this option to function properly, the time on the Cell Manager and the time on the system where data is restored must be synchronized.

If this option is selected, Data Protector recreates the state of the backed up directory tree at the time of the last incremental backup session while preserving files that were created or modified afterwards. Files that were removed between the full backup (the initial session defining the restore chain) and the chosen incremental backup are restored and later deleted at folder restore during subsequent incremental restore.

If this option is not selected, Data Protector also restores files that were included in the full backup image and were removed between the full backup (the initial session defining the restore chain) and the chosen incremental backup.

When using the **Restore As** or **Restore Into** functionality with this option enabled, carefully choose the restore location to prevent accidental removal of existing files.

Default: not selected.

- **Move busy files.** This option is relevant if a file on the disk is being used by an application when a restore wants to replace this file. It only applies to the files that are locked by an operating system when they are used by the application or other process. The option is used with the **Keep most recent** or **Overwrite** options.

By default, this option is disabled.

On UNIX systems, Data Protector moves the busy file `filename` to `#filename` (adds a hash in front of the filename). The application will keep using the busy file until it closes the file. Subsequently, the restored file is used.

On Linux systems, this option is not supported.

On Windows systems, the file is restored as `filename.001`. All applications keep using the old file. When the system is rebooted, the old file is replaced with the restored file.

- **List restored data.** Displays the names of the files and directories in the monitor window as the objects are being restored. By default, this option is disabled.
- **Display statistical information.** Reports statistical information (such as size and performance) for each object that is backed up or restored. You can view the information in the monitor window. By default, this option is disabled.
- **Omit unrequired object versions.** This option applies if you select directories for restore and the backup was performed with the logging level **Log All** or **Log Files**.

If this option is selected, Data Protector checks in the IDB for each backup in the restore chain if there are any files to restore. Backups with no object versions to restore are skipped. Note that this check may take some time.

If this option is not selected, each backup in the restore chain is read, even if there was no change since the previous backup.

To restore empty directories, clear this option.

Default: selected.

- **Restore sparse files.** Restores sparse files in their original compressed form. This is important because sparse files can consume additional disk space unless they are restored in their original form. By default, this option is disabled.

This option applies to UNIX sparse files only. Windows sparse files are always restored as sparse.

- **Lock files during restore.** Denies access to files during the restore. By default, this option is

disabled.

- **Restore time attributes.** Preserves the time attribute values of each restored file. When this option is disabled, Data Protector sets the time attributes of the restored objects to the current date and time. By default, this option is enabled.
- **Restore protection attributes.** Preserves the original protection attributes of each restored file. If this option is disabled, Data Protector applies the protection attributes of the current restore session. By default, this option is enabled.

On Windows systems, this option applies to file attributes only. Security information is always restored, even when this option is disabled.

- **Restore share info for directories.** Specifies that share information for directories will be restored. By default, this option is selected.

When restoring a directory that was shared on the network when it was backed up, the directory will also be shared after restore if this option is selected, provided that the backup was made with the **Backup share information for directories** option selected.

Pre- and post-exec commands

- **Pre-exec.** Allows you to enter a command to be executed before the restore of each object is initiated. This command must return success for Data Protector to proceed with the restore.

The pre-exec command is executed on the client system where the Disk Agent is running. The scripts on the Cell Manager and on other systems must be located in the `Data_Protector_home\bin` directory and the user must specify only the filename or relative pathname.

Note that on Windows systems, if your directory names are longer than 8 characters, write the pathname either in quotes or in the short 8.3 MS-DOS compatible form. If you use quotes (") to specify a pathname, do not use a combination of backslash and quotes (\"). If you need to use a trailing backslash at the end of the pathname, use a double backslash (\\).

Note that only `.bat`, `.exe`, and `.cmd` are supported extensions for pre-exec scripts on Windows systems. To run a pre-exec script with an unsupported extension (for example, `.vbs`), create a batch file (`.bat`) that starts the script. Then configure Data Protector to run the batch file as a pre-exec command which then starts the script with the unsupported extension.

- **Post-exec.** Allows you to enter a command to be executed after the restore of each object is completed. The post-exec command is executed on the client system where the Disk Agent is running.

Device selection

- **Automatic device selection.** Applicable when the original devices are not available for a restore or an object copy. Select this option to enable Data Protector to automatically replace unavailable devices with other devices that are selected for the restore or object copy and have the same device tag as the original device. If there are not enough available devices to replace the original devices, the restore or object copy is started with fewer devices than were used during backup.

By default, Data Protector attempts to use the original device first. If the original device is not selected for a restore or an object copy, then a global option is considered. To use alternative devices first or to prevent the use of the original device all together, modify the global option `AutomaticDeviceSelectionOrder`.

For the Data Protector SAP MaxDB, DB2 UDB, Microsoft SQL Server, and Microsoft SharePoint Server 2007/2010/2013 integration, ensure that the number of available devices is equal to or greater than the number of devices that were used during backup.

Default: selected.

- **Original device selection.** Applicable when the original devices are not available for a restore or an object copy at the moment. Select this option to instruct Data Protector to wait for the selected devices to become available.

This is the preferred option for the Data Protector SAP MaxDB, IBM DB2 UDB, Microsoft SQL Server, and Microsoft SharePoint Server 2007/2010/2013 integration.

Default: not selected.

Handling file conflicts

- **Keep most recent.** If this option is selected, the most recent versions of files are kept. If a file on the disk is newer than the backed up version, the file is not restored. If a file on the disk is older than the backed up version, the file is overwritten with the newer version from the backup. By default, this option is enabled.
- **No overwrite.** If this option is selected, files that exist on the disk are preserved. This means that they are not overwritten by other versions of these files from the backup. Only non-existing files are restored from the backup. By default, this option is disabled.
- **Overwrite.** If this option is selected, existing files on the disk are replaced with files from the backup. By default, this option is disabled.

Active Directory specific options

Replication mode

- **Authoritative.** This is a Windows Server specific option dealing with active directory restore. The Active Directory database is not updated after the restore and the restored data overwrites the existing data in the target destination. An authoritative restore can only be performed by running `ntdsutil.exe` from the command prompt after the restore session has finished.
- **Nonauthoritative.** The Active Directory database is updated after the restore using standard replication techniques. The **Nonauthoritative** replication mode is the default option.
- **Primary.** The Primary replication mode allows you to keep the NT directory Service online and is used when you restore `FileReplicationService` along with the Active Directory service. This option must be used when all replication partners for a replicated share have been lost. With regard to the Certificate Server and the Active Directory Server, **Primary** is the same as **Authoritative**.

Setting Restore Options

After selecting the data that you want to restore, you can set the restore options. Restore options have default values that are appropriate in most cases. They are available according to the type of data being restored. For example, all restore options available for a filesystem restore are not available for a disk image restore.

Steps

1. In the Context List, click **Restore**.
2. In the Scoping Pane, under Restore Objects, expand the appropriate data type (for example, **Filesystem**).
3. Expand the client system with the data you want to restore and then click the object (mountpoint on UNIX systems, drive on Windows systems) that has the data.
4. In the Source property page, select the data to restore.
5. Click the **Options** tab to open the Options property page. Select or deselect an option by clicking the box next to it.

About Windows Systems Restore

When restoring a Windows filesystem, Data Protector restores the data within the files and directories, as well as Windows-specific information about the files and directories.

The following Windows-specific information is restored:

- Full Unicode file names
- FAT16, FAT32, VFAT
NTFS attributes
- Sets of alternate data streams.
- Share information

If a directory is shared on a network during backup, the share information is stored on the backup medium. The directory will be shared on the network after the restore by default (unless a shared directory with the same share name already exists). To prevent restoring share information for directories that are being restored, deselect the Restore share information for directories option.

File Conflict Handling options apply also for the restore of the directory share information. For example, if the No overwrite restore option is used for the restore, the directory share information for directories that exist on the disk, is preserved.

- NTFS alternate data streams
- NTFS security data

NTFS 3.1 filesystem features

- The NTFS 3.1 filesystem supports reparse points
The volume mount points, Single Instance Storage (SIS), and directory junctions are based on the reparse point concept. These reparse points are selected as any other filesystem object.
- The NTFS 3.1 filesystem supports symbolic links, which were introduced with Windows Vista and Windows Server 2008 operating systems.
Data Protector handles symbolic links in the same way as NTFS reparse points.
- The NTFS 3.1 filesystem supports sparse files as an efficient way of reducing the amount of allocated disk space.

These files are backed up sparse to save tape space. Sparse files are backed up and restored as sparse to the NTFS 3.1 filesystem only.

- Some of the NTFS 3.1-specific features are controlled by system services that maintain their own data records. These data structures are backed up as a part of CONFIGURATION.
- Encrypted files
The Microsoft-encrypted NTFS 3.1 files are backed up and restored encrypted, but their contents can only be properly viewed when they are decrypted.
- Compressed files are backed up and restored compressed.

Consider the filesystem restore limitations when restoring to a different filesystem type than where the backup was performed.

Restoring objects backed as shared disks

Objects that were backed up as shared disks are associated with the Disk Agent client that was used to back them up. If the environment has not changed, you can restore the shared disk as you would a local Windows filesystem. By default, the same Disk Agent client that was used to back up the shared disk is used to restore the data to the original location.

Windows Filesystem Restore Limitations

You can restore your data to a different filesystem type than the one the backup was performed on.

From	To				
	FAT32	FAT16	CDFS	UDF	NTFS 3.1 ¹
FAT32	FC	FC	N/A	N/A	FC
FAT16	FC	FC	N/A	N/A	FC
CDFS	FC	FC	N/A	N/A	FC
UDF	FC	FC	N/A	N/A	FC
NTFS 3.1 ²	*	*	N/A	N/A	FC

Legend

FC	Full Compatibility. The file attributes are entirely preserved.
*	Reparse points, sparse files and encrypted files are not restored. Files are restored without security information and alternate data streams.

¹ It is used on Windows XP, Windows Vista, Windows 7, Windows 8, Windows Server 2003, Windows Server 2008, and Windows Server 2012.

² It is used on Windows XP, Windows Vista, Windows 7, Windows 8, Windows Server 2003, Windows Server 2008, and Windows Server 2012.

The table shows that NTFS 3.1 filesystem objects can only be adequately restored to the NTFS 3.1 filesystem. The filesystem-specific attributes and alternate data streams are lost when restoring into a different filesystem version.

- A Windows reparse point, such as a directory junction or a volume mountpoint, can be restored to an NTFS 3.1 filesystem only. UNIX reparse points cannot be restored to a NTFS 3.1 filesystem.
- When you restore an NTFS 3.1 filesystem that contains SIS reparse points, a full disk condition may occur. This happens if the original file is restored into multiple target files that can take up more space than available.
- Sparse files are restored as sparse to the NTFS 3.1 filesystem only.
- User Disk Quotas cannot be restored using Data Protector.
- If a user attempts to restore a sparse file to a non-NTFS 3.1 filesystem, Data Protector will issue a warning. A sparse file restored to a filesystem other than NTFS 3.1 will not include zero sections.
- The Microsoft encrypted NTFS 3.1 files can be restored to the NTFS 3.1 filesystem only, because other filesystem drivers cannot decrypt them.

Configuration Restore

To restore the Windows CONFIGURATION, select the CONFIGURATION object or parts of it and follow the standard restore procedure.

The CONFIGURATION consists of data structures that influence system operation. Therefore, the system must be prepared for such a restore. The prerequisites depend on the contents of the CONFIGURATION item and the Windows operating system version.

Limitations

- Active Directory Service and SysVol should be restored in pair.
- User Disk Quotas cannot be restored using Data Protector. The backed up information can be restored manually, using Microsoft utilities.
- Although Data Protector allows you to restore single configuration objects, it is **not recommended** to do so. It is highly recommended that you perform a full configuration restore as part of the **Disaster Recovery** procedure.

Windows configuration objects

For more detailed information on configuration objects, see HPE Data Protector Help.

- Active Directory Service
- Certificate Server
- COM+ Class Registration Database (ComPlusDatabase)
- DFS
- DHCP
- DNS Server
- Event Logs

- File Replication Service
- Internet Information Server (IIS)
- User Profiles (Documents and Settings)
- Windows Registry
- Removable Storage Management Database
- SystemRecoveryData
- SysVol
- Terminal Services Database
- User Disk Quotas (QuotaInformation)
- WINS server

Restart the system after the restore of the whole CONFIGURATION object is finished in order for the restored data to become effective.

Some objects require special considerations and tasks.

Active Directory

To restore the Active Directory service, you have to restart the system using the Directory Services Restore Mode start-up option. When the system is started in the Directory Services Restore Mode, the domain user accounts cannot be used. You have to configure the Data Protector Inet and the `crs` service (for a Cell Manager) to log on using the local system account and then restart the services. When restoring the Active Directory, the File Replication Service (FRS) and Distributed File System (DFS) are also restored.

You can restore the Active Directory in one of three replication modes (Windows specific options):

- nonauthoritative
- authoritative
- primary

Note: To perform an **Authoritative** restore, you also need to run `ntdsutil.exe` after the restore session has finished. For example, to perform a typical authoritative restore, at a command prompt enter `ntdsutil`, then `authoritative restore`, then `restore database`. Restart the server and wait for replication to take place.

Tip: You can also create a post-exec command to perform the additional action needed for the Active Directory authoritative restore. For example, to perform an authoritative restore of an entire directory, use the following line:

```
ntdsutil "popups off" "authoritative restore" "restore database" quit quit
```

DFS

Data Protector restores Windows Distributed File System (DFS) as part of one of the following:

- Windows Registry, if the DFS is configured in a standalone mode
- Windows Active Directory, if the DFS is configured in a domain mode

Profiles

- A user profile cannot be restored successfully if the respective user is logged on, either interactively or as a service. If the user is logged on at the time of the restore, Data Protector will fail to restore the file NTUSER.DAT which contains the user's registry hive.

You have to log off the system and stop all the services that are running under the user account whose profiles you want to restore. The restore session can be started from another system or by logging on the restore target system as a different user.

- To restore all user profiles at once, you must stop any services that do not run under the local system account, and log off from the system. Then start the restore session remotely, using Data Protector GUI on another client.
- A user profile can only be restored when its location is already defined on the system. Individual files of existing user profiles or deleted profiles can be still restored as long as they exist among the system's profiles. If a user profile was deleted from the Control Panel, or the user profile no longer exists on the system for some other reason, the restore fails with the following error:

```
[84:208] Configuration object not recognized by the system => not restored.
```

To restore such user profile, you must first recreate it by logging on as that user. The system assigns a directory for the user's profile and creates a default profile. To keep the restored files unmerged, you can delete the files in the newly created profile before running a restore session. Then log off and start the restore session by logging on as a different user or by using another system. The system may assign a different name to the user. In this case, use the **Restore As** option to restore the files to the newly assigned location.

- When user profiles are restored, files are always overwritten, regardless of the File Conflict Handling options in the restore specification. Also, the **Omit deleted files** option is not available. Files that exist on the disk, but were not present at the time of the backup, will remain in the user profile after the restore.
- User profiles can also be restored using the **Restore As** option. You can specify a temporary location for the files and then manually copy the desired files to the user's profile directory. Or, you can restore directly over the user's profile directory, possibly making use of the **Move busy files** option, which allows you to restore a user profile even if it is in use by a logged on user. However, note that in this case the files that are in use will only be replaced after the system is rebooted.

Registry

If you select the whole Windows Registry for a restore, some of the Registry keys are not restored and some are treated in a special way during a restore. This is because these keys are used by the operating system. You can find them under the following Registry key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BackupRestore\KeysNotToRestore
```

Removable Storage Manager Database

The RSM service must be running on all systems with connected removable storage devices (except for CD-ROMs).

Server configuration objects

The target system must have the respective server installed and running. For all servers, except Certificate Server, the data is restored online.

Certificate Server data is restored offline. Stop the Certificate Server Services before starting a restore. You can restore the Certificate Server only using the authoritative mode.

SysVol

You can perform restore of SysVol directory in one of three modes:

- nonauthoritative
If at least one domain controller in the domain is available and working, files are restored to their original location. The restored data is not propagated to other domain controllers.
- authoritative
Perform authoritative restore if critical SysVol data is deleted from the local domain controller and the deletion is propagated to other domain controllers.
- primary
If all domain controllers in the domain are lost and you want to rebuild domain controller from backup, the FRS is informed that you are restoring primary files and files are restored to their original location.

Windows TCP/IP services

On a Windows system that runs a Microsoft TCP/IP protocol and is configured as a WINS Server, a DHCP Server, or a DNS Server, you can restore the services that manage network communication.

To restore Windows TCP/IP services, expand the CONFIGURATION item and select WNS, DHCP, or DNSServerDatabase.

Each of these services is automatically stopped before the restore.

When the restore has finished, restart the system.

System State Data Restore

If you use Active Directory, which is always a part of the System State, you have to start the system in the Directory Services Restore Mode.

From the Data Protector point of view, the System State consists of some specific filesystem objects and CONFIGURATION objects. On Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012, the System State also includes data belonging to additional server roles or services that may be installed. As opposed to selecting objects in the Backup wizard, different objects for restore are selected in separate Restore wizards.

In the Source property page, select:

- the System State objects that belong to CONFIGURATION:
 - ActiveDirectoryService
 - CertificateServer
 - Cluster Service information
 - IIS Metadirectory
 - RemoteStorageService
 - RemovableStorageManagementDatabase
 - SystemFileProtection
 - SYSVOL directory
 - TerminalServiceDatabase
- SystemVolumeInformation (including System File Protection service)
- boot files (they are located on the system drive)
- volumes on which data belonging to particular server roles or services resides or even the entire client system (in case of Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012)

When the restore is finished, restart the system.

Remote Storage Service

Remote Storage Service (RSS) is used to automatically move infrequently accessed files from local to remote storage. Remote files are recalled automatically when the file is opened.

Although the RSS databases are part of System State data, you restore them manually. The RSS database must be restored offline. You can provide pre- and post-exec scripts to stop and restart the service, or you can stop and restart it manually before and after the restore, respectively.

Select the following directories for restore:

%SystemRoot%\system32\RemoteStorage

%SystemRoot%\system32\NtmsData

System File Protection

System File Protection service scans and verifies the versions of all protected system files after you restart your computer. If the System File Protection service discovers that a protected file has been overwritten, it retrieves the correct version of the file and then replaces the incorrect file. Data Protector enables you to back up and then restore protected files without being overwritten.

About UNIX Systems Restore

When restoring files to the original location from which the backup was performed, Data Protector restores the files, including file attributes.

System specific data, such as ACL (Access Control List) on UNIX systems, is restored only on the same filesystem type and operating system from which the backup was made.

UNIX systems specific information

When restoring VxFS data, use the Restore As option and restore it to the desired location.

About HP OpenVMS System Restore

Use the standard restore procedure to restore HP OpenVMS filesystems.

Limitations

- For files and directories saved on any other operating system platform not all file attributes are restored and no ACL is restored in this case.
- Directories that are created during a restore but have not been included in a save will get the attributes of the first file restored in the directory unless disabled by the `-no_protection` option.
- Any file specifications that are entered into the GUI or passed to the CLI must be in UNIX style syntax

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

The string should begin with a slash, followed by the disk, directories, and filename, 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. For example, an OpenVMS file specification of:

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

must be specified in the form:

`/1DGA100/Users/Doe/Login.Com.1`

- There is no implicit version number. You always have to specify a version number. Only file versions selected for the restore will be restored. 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 wildcard characters for the version number, as follows:

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

- 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.

- If the **Restore Time Attributes** (-notouch) 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.
- A file saved as a soft link will be restored using the equivalent of a DCL SET FILE/ENTER command. No data will be restored in this case. The soft link entered points to the primary path/filename for this file from the time the file was saved. If the primary path/filename does not exist or was not restored, the creation of the soft link will fail.

To make a restored copy of an OpenVMS system disk bootable, the OpenVMS WRITEBOOT utility has to be used to write a boot block after the disk has been restored.

- 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 (that is, upper and lower case letters, Unicode characters, and so on) may not be restored to an ODS-2 disk.
- Files being restored are always locked regardless of whether the **Lock Files during Restore** (-lock) option is enabled or disabled.
- 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. For example:

```
/SYS$MANAGER/DP_SAVE1.COM
```

- If the **Restore Protection Attributes** (-no_protection) option is disabled, the files are created with the default owner, protection, and ACL.
- When specifying wildcard characters for **Skip** (-skip) or **Only** (-only) filters use '*' for multiple characters and '?' for single characters.
- On OpenVMS systems, Data Protector does not support disk quotas on volumes and volume sets.

To perform restore of data located on a volume with disk quota enabled, configure the post-exec script so that it disables disk quota on the involved volume before restore starts, and configure the pre-exec script so that it enables the disk quota after restore completes.

Filesystem information restored

The directory structure and the files are restored, together with the following filesystem information:

- File and directory attributes
- ACL (Access Control List) if available (see Limitations)
- Secondary file entries

During an OpenVMS filesystem backup, files with multiple directory entries are backed up once using the primary path name. Secondary path entries are saved as soft links.

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...].

During a filesystem restore, these extra path entries are restored.

Files can be restored to mounted FILES-11, ODS-2, or ODS-5 volumes only.

Chapter 14: Monitoring, Reporting, Notifications, and Data Protector Event Log

About Monitoring

Data Protector monitoring allows you to manage running sessions and to respond to mount requests. You can view the status of sessions, their type, owner, and session ID; the start time of the sessions as well as the names of the corresponding backup specifications.

When you run an interactive backup, restore, object copy, object consolidation, object verification, or media management session, a monitor window opens, showing the objects, backup devices, and messages generated during the session. Even if the user interface is closed, the session continues.

The **"size"** parameter in monitor context will display the data size units in **"KB, MB, GB, TB, and PB"**.

You can change the level of reported messages during a backup or restore session by changing the **Report level** option during configuration of a backup specification or during startup of a restore session.

You can monitor several cells at the same time using the Manager-of-Managers functionality.

Viewing Currently Running Sessions

You can view currently running sessions in the Monitor context.

Note: A currently running session is displayed in the Monitor context after the pre-exec script has finished.

At refresh intervals (by default 5 seconds), the list of currently running sessions is automatically updated with new sessions. To change the default refresh interval, in the File menu, click **Preferences**, and then click the Monitor tab. You can specify the refresh interval in seconds for the Cell Manager and for the MoM.

Prerequisite

You either have to be added to the Admin user group or granted the Monitor user rights.

Steps

1. In the Context List, click **Monitor**.

In the Results Area, the status of current sessions is displayed.

Tip: You can sort the sessions (by *status*, *type*, *owner*, and so on) by clicking the corresponding column header. For VMware integration, you can sort the sessions by *VM name* and *item name* as well. Here, *VM name* refers to the name of the virtual machine in vCenter and *item name* refers to the name of the disk object or configuration associated with the virtual machine.

2. Double click the running session you want to view.

Tip: To remove all completed or aborted sessions from the Results Area of the Monitor context, in the Scoping Pane, click **Current Sessions** and then select **Clear Sessions** from the Action menu. To remove a particular finished or aborted session from the current sessions list, right-click the session and select **Remove From List**. All completed or aborted sessions are automatically removed from the Results Area of the Monitor context if you restart the Data Protector GUI.

Viewing Finished Sessions

You can view completed or aborted sessions in the Internal Database context.

Prerequisite

You either have to be added to the Admin user group or granted Monitor user rights.

Steps

1. In the Context List, click **Internal Database**.
If you are running Manager-of-Managers, select **Monitor** in the Context List and then select a Cell Manager of your choice. From the Tools menu, select **Database Administration** to open a new Data Protector GUI with the Internal Database context selected.
2. In the Scoping Pane, expand **Sessions** to display all the sessions stored in the IDB. The sessions are sorted by date. Each session is identified by a session ID consisting of a date in a YY/MM/DD format and a unique number.
3. Right-click the session and select **Properties** to view details on a specific session.
4. Click the **General**, **Messages**, or **Media** tab to display general information on the session, session messages, or information on media used for this session, respectively.

Aborting Running Sessions

You abort a session if you want to stop a backup, restore, or media management operation. A backup copy or restored data only exist for data that was backed up or restored before you aborted the session.

Prerequisite

You either have to be added in the `admin` user group or granted the Monitor user rights.

Steps

1. In the Context List, click **Monitor**. The progress and status of current sessions appear in the Results Area.
If you are running a Manager-of-Managers, expand the **Enterprise Monitor** in the Scoping Pane and then select the Cell Manager you want to monitor. The progress and status of current sessions appear in the Results Area.

2. Click the column headings to sort the sessions.
3. Right-click the session and select **Abort**.

If you abort a backup session while it is still determining the sizes of the disks that you have selected for the backup, it does not abort immediately. The backup is aborted once the size determination (treewalk) is completed.

Tip: If you started a backup, restore, or media management session interactively, you can also abort the session in the Data Protector Backup, Restore, or Devices & Media context respectively.

About Reporting

Data Protector reports provide various information on your backup environment. For example, you can check the status of the last backup, object copy, object consolidation, or object verification, check which systems in your network are not configured for backup, check on the consumption of media in media pools, check the status of devices and more.

You can configure reports and report groups using the Data Protector GUI or any Web browser with Java support. Report groups allow you to easily manage reports, to schedule the reports in the report group, and to define the criteria for grouping the reports in report groups.

Parameters allow you to customize reports. Some parameters allow multiple selections. If no optional input parameters (optional selections) are specified when configuring a report, a default value is set, which is *all* in the case of objects and *no time limit* in the case of time frames. To configure a report or report group you need to provide:

- name for the report
- type of the report
- send method
- recipient(s)
- format

All other input parameters (selections) depend on the type of the report.

Note: The VADP Reporting feature is enabled by default. In order to disable it, set `EnableDPAforVM` global variable to 0.

Features

- You can gather various reports in a report group, which can be scheduled, started interactively, or triggered by a notification.
- Reports can be started using the Data Protector GUI, the Data Protector CLI, the Data Protector web reporting interface, the Data Protector scheduler, a notification event, or a post-exec script that includes a Data Protector CLI command that starts the report.
- Reporting is also available for a multiple-cell configuration when you use the Manager-of-Managers (MoM) functionality.
- The output of the reports is provided in various formats and can optionally display input parameters (selections), also.

Reports Formats

You can generate Data Protector reports in various formats.

If you start each report individually, the report is displayed in the Data Protector Manager and you do not have to choose the report format.

If you gather reports into report groups, you have to specify the format and the recipients of each report.

You can choose from the following report formats:

- **ASCII** - A report is generated as plain text.
- **HTML** - A report is generated in HTML format. This format is useful for viewing using a web browser. For example, you can check if your systems have been backed up by clicking a link and viewing the report on the Intranet.
- **Short** - A report is generated as plain text, in summary form showing only the most important information. This is the suggested format for broadcast messages.
- **Tab** - A report is generated with fields separated by tabs. This format is useful if you plan to import the reports into other applications or scripts for further analysis, such as Microsoft Excel.

The actual output of a report varies depending on the selected format. Only the Tab format displays all fields for all reports, other formats may sometimes display only selected fields.

Reports Types

Depending on the information about your backup environment that you want to retrieve, you can generate various types of reports:

Configuration reports

Configuration reports provide information on the configuration of the Data Protector cell, on devices not used for backup, on systems not configured for backup, and so on.

Cell Information

<i>Description:</i>
<p>Lists Data Protector cell-related information (number of clients, backup specifications, media management server, licensing server).</p> <p>The VADP feature introduced in Data Protector 8.14 provides enhanced reports for Virtual Machines. The VMware virtual machines are represented as Data Protector clients called VADP clients. The VADP clients display the information on the Guest OS of the virtual machine. If the VM tools are installed and running, and VM is powered on, the Host information section of the output displays information, such as the operating system, IP address, or hostname.</p> <p>The VM hostname must display the DNS name, if it is configured on a virtual machine.</p>

	<p>The VM hostname must display the IP address, if VM has no DNS name and IPv4 is available.</p> <p>The VM hostname must display the VM name, if DNS name or IP address is not available (or) if the VM has only the IPv6 address.</p>
<i>Required selections:</i>	none
<i>Optional selections:</i>	none
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	cell_info

Client Backup

<i>Description:</i>	<p>Lists information about the specified clients like: filesystems not configured, all objects, all objects with a valid backup and their backup times and average sizes.</p> <p>Note that Client Backup reports do not include information about application integration backup objects and backup specifications.</p>
<i>Required selections:</i>	hostname
<i>Optional selections:</i>	none
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	host

Clients not Configured for Data Protector

Note: Generating this report can take some time depending on the condition of the network. This type of report cannot be aborted.

<i>Description:</i>	Lists clients in the selected domains that are not part of the current cell.
<i>Required selections:</i>	network range(s)
<i>Optional selections:</i>	none
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	hosts_not_conf

Configured Clients not Used by Data Protector

<i>Description:</i>	Lists all configured clients that are not used for backup and do not have any device configured.
---------------------	--------------------------------------------------------------------------------------------------

<i>Required selections:</i>	none
<i>Optional selections:</i>	none
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	hosts_unused

Configured Devices not Used by Data Protector

<i>Description:</i>	Lists configured destination devices that are not used for backup, object copy, or object consolidation at all.
<i>Required selections:</i>	none
<i>Optional selections:</i>	none
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	dev_unused

Licensing

<i>Description:</i>	Lists all licenses and the available number of licenses.
<i>Required selections:</i>	none
<i>Optional selections:</i>	none
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	licensing

Look up Schedule

<i>Description:</i>	Lists all backup, object copy, object consolidation, or verification specifications that are scheduled to start in the next specified number of days, up to one year in advance.
<i>Required selections:</i>	number of days
<i>Optional selections:</i>	none
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	lookup_sch

IDB report

IDB report provides information on the size of the IDB.

IDB Size

<i>Description:</i>	Provides a table that contains information about the Media Management Database, Catalog Database, Archived Log Files, Datafiles, statistics for Detail Catalog Binary Files directories, SMBF (msg directory), and low IDB disk space.
<i>Required selections:</i>	none
<i>Optional selections:</i>	none
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	db_size

The **Used** columns in this report show the percentage of used items for each IDB part. This figure is calculated as the current number of items divided by the number of maximum items for particular IDB part in percents. In case the number of items is unlimited, this figure is always 0%.

To find out whether certain parts of IDB are running out of space, you can additionally configure the IDB Space Low notification.

Pools and media reports

Pools and media pools reports provide information on media pools and used media.

Extended List of Media

<i>Description:</i>	Lists all media matching the specified search criteria. For each medium, it provides information about the medium ID, medium label, media location, media condition, media protection, used and total space (MB), the time when the medium was last accessed, the media pool and media type, session specifications that have used the medium for backup, object copy, or object consolidation, as well as the session type and subtype.
<i>Required selections:</i>	none
<i>Optional selections:</i>	session specification(s), backup specification group, description, location (s,) pool name(s), media type (DDS, DLT, and so forth), condition, expiration, timeframe, library device(s)
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	media_list_extended

List of Media

<i>Description:</i>	Lists all media matching the specified search criteria. For each medium, it provides information about the medium ID, medium label, media location,
---------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------

	media condition, media protection, used and total space (MB), the time when the medium was last accessed, the media pool and media type.
<i>Required selections:</i>	none
<i>Optional selections:</i>	description, location(s), pool name(s), media type (DDS, DLT, and so forth), condition, expiration, timeframe, library device(s)
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	media_list

List of Pools

<i>Description:</i>	Lists all pools matching the specified search criteria. For each pool it provides information about the pool name, description, media type, total number of media, number of full and appendable media containing protected data, number of free media containing no protected data, number of poor, fair, and good media.
<i>Required selections:</i>	none
<i>Optional selections:</i>	pool name(s), location(s), media type (DDS, DLT, and so forth), library device(s), timeframe
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	pool_list

Media Statistics

<i>Description:</i>	Reports statistics on the media matching the search criteria. The following information is provided: number of media; number of scratch media; number of protected, good, fair, and poor media; number of appendable media; total, used, and free space on media.
<i>Required selections:</i>	none
<i>Optional selections:</i>	description, location(s), pool name(s), media type (DDS, DLT, and so forth), condition, status, expiration, timeframe, library device(s)
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	media_statistics

Session specification reports

Session specification reports provide information on backups, object copy, object consolidation or object verification, such as average size of backed up objects, schedule of sessions, filesystems not configured for backup, and so on.

Average Backup Object Sizes

<i>Description:</i>	<p>Displays the average size of an object in the specified backup specification. It displays the size of the full and the incremental backup of the object.</p> <p>The VADP feature introduced in Data Protector 8.14 provides enhanced reports for Virtual Machines. The VMware virtual machines are represented as Data Protector clients called VADP clients. The new object name format for VADP clients is as follows:</p> <p><hostname>:/<vCenter>/<path>/<vmname> [<UUID>]</p> <p>Here, <hostname> is the DNS of the guest virtual machine. If the DNS name is unknown, the IP address or VM name is used.</p>
<i>Required selections:</i>	none
<i>Optional selections:</i>	backup specification(s), backup specification group, number of days (counted from the moment of starting the report backwards)
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	obj_avesize

Filesystems Not Configured for Backup

<i>Description:</i>	Lists all disks (filesystems), that are not configured in any of the selected backup specifications.
<i>Required selections:</i>	none
<i>Optional selections:</i>	backup specification(s), backup specification group
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	fs_not_conf

Object's Latest Backup

<i>Description:</i>	<p>Lists all objects in the IDB. For each object, it displays the last full and the last incremental backup time, the last full and the last incremental object copy time, and the last object consolidation time.</p> <p>The VADP feature introduced in Data Protector 8.14 provides enhanced reports for Virtual Machines. The VMware virtual machines are represented as Data Protector clients called VADP clients. The new object name format for VADP clients is as follows:</p> <p><hostname>:/<vCenter>/<path>/<vmname> [<UUID>]</p>
---------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

	<p>Here, <hostname> is the DNS name of the guest virtual machine. If the DNS name is unknown, the IP address or VM name is used.</p> <p>You can narrow the scope of objects listed using the backup specification filters and/or object creation time filter (see <i>Optional Selections</i>). However, consider the following particularities:</p> <ul style="list-style-type: none"> • Objects of the Filesystem type (filesystem objects) that do not match the condition in the object creation time filter are listed anyway. However, in this case, their object creation time fields remain empty. • If you clear certain filesystem objects from a backup specification, these filesystem objects will not be included in the report even if the objects exist in the IDB. <p>The above considerations are not applicable for objects of the Bar type (integration objects).</p>
<i>Required selections:</i>	none
<i>Optional selections:</i>	backup specification(s), backup specification group, number of days (counted from the moment of starting the report backwards)
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	obj_lastbackup

Objects Without Backup

<i>Description:</i>	Lists all objects that are part of a backup specification and do not have a valid backup (successfully completed backup, the protection has not yet expired). This report is not available for backup specifications for integrations.
<i>Required selections:</i>	none
<i>Optional selections:</i>	backup specification(s), backup specification group, number of days (counted from the moment of starting the report backwards)
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	obj_nobackup

Session Specification Information

<i>Description:</i>	Displays information about all selected backup, object copy, object consolidation, and object verification specifications, such as type (for example, IDB, MSESE, E2010), session type, session specification name, group, owner, and pre- and post-exec commands.
<i>Required Selections</i>	none

<i>Optional selections:</i>	session specification(s), backup specification group
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	dl_info

Session Specification Schedule

<i>Description:</i>	Lists the next start time for each specified backup, object copy, object consolidation, and object verification specification up to one year in advance.
<i>Required selections:</i>	none
<i>Optional selections:</i>	session specification(s), backup specification group
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	dl_sched

Trees in Backup Specifications

<i>Description:</i>	<p>Lists all trees in the specified backup specification. It also shows names of drives and the name of a tree.</p> <p>The VADP feature introduced in Data Protector 8.14 provides enhanced reports for Virtual Machines. The VMware virtual machines are represented as Data Protector clients called VADP clients. The report displays all the VM names for VMware objects.</p>
<i>Required selections:</i>	none
<i>Optional selections:</i>	backup specification(s), backup specification group
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	dl_trees

Sessions in timeframe reports

Sessions in timeframe reports provide information on backup, object copy, object consolidation or object verification sessions that ran in a specified period of time.

Client Statistics

<i>Description:</i>	<p>Lists clients and their backup status statistics. Only the clients that match the search criteria are listed.</p> <p>The VADP feature introduced in Data Protector 8.14 provides enhanced reports for Virtual Machines. The VMware virtual machines are represented</p>
---------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

	as Data Protector clients called VADP clients, wherein the VM name is the client name.
<i>Required selections:</i>	timeframe
<i>Optional selections:</i>	backup specification(s), backup specification group, hostname(s)
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	host_statistics

Device Flow

<i>Description:</i>	Graphically presents the usage of each device. A flow chart of the backup, object copy, and object consolidation sessions matching the search criteria is shown. If you set the <code>RptShowPhysicalDeviceInDeviceFlowReport</code> global option to 1, the same physical devices (presented by their lock names or serial numbers) are grouped together. If there is no lock name or serial number specified, the logical name is displayed.
<i>Required selections:</i>	timeframe
<i>Optional selections:</i>	session specification(s), backup specification group
<i>Supported formats:</i>	HTML
<i>omnirpt option:</i>	device_flow

Extended Report on Used Media

<i>Description:</i>	Provides extended information on destination media that have been used by backup, object copy, and object consolidation sessions in the specific time frame, as well as the session type and subtype.
<i>Required selections:</i>	timeframe
<i>Optional selections:</i>	session specification(s), backup specification group
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	used_media_extended

List of Sessions

<i>Description:</i>	Lists all sessions and their statistics in the specified timeframe.
<i>Required selections:</i>	timeframe
<i>Optional selections:</i>	session specification(s), backup specification group

<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	list_sessions

Object Copies

<i>Description:</i>	<p>Displays the number of valid copies of object version in the specified timeframe. The number of copies includes the original object version.</p> <p>The VMware virtual machines are represented as Data Protector clients called VADP clients. The new object name format for VADP clients is as follows:</p> <p><hostname>:/<vCenter>/<path>/<vmname> [<UUID>]</p> <p>Here, <hostname> is the DNS name of the guest virtual machine. If the DNS name is unknown, the IP address or VM name is used.</p>
<i>Required selections:</i>	timeframe
<i>Optional selections:</i>	session specification(s), backup specification group, number of copies
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	obj_copies

Report on Used Media

<i>Description:</i>	Lists destination media that have been used during the backup, object copy, and object consolidation sessions in the specified timeframe together with their statistics.
<i>Required selections:</i>	timeframe
<i>Optional selections:</i>	session specification(s), backup specification group
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	used_media

Session Errors

<i>Description:</i>	Displays a list of error messages that occurred during a backup, object copy, object consolidation, or object verification session. The messages are grouped by client.
<i>Required selections:</i>	timeframe
<i>Optional selections:</i>	session specification(s), backup specification group, hostname(s), message level

<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	session_errors

Session Flow

<i>Description:</i>	<p>Graphically presents the duration of each session for the specified timeframe. A flow chart of the backup, object copy, object consolidation, and object verification sessions matching the search criteria is shown.</p> <p>Colors in the chart represent the following overall status of the sessions:</p> <ul style="list-style-type: none">• Red: Session failed or was aborted.• Green: Session completed successfully or completed with errors.• Yellow: Session completed with failures.• Blue: Session is queuing or a mount request is issued.
<i>Required selections:</i>	timeframe
<i>Optional selections:</i>	session specification(s), backup specification group
<i>Supported formats:</i>	HTML
<i>omnirpt option:</i>	session_flow

Session Statistics

<i>Description:</i>	Shows backup, object copy, or object consolidation status statistics in the selected timeframe.
<i>Required selections:</i>	timeframe
<i>Optional selections:</i>	session specification(s), backup specification group
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	session_statistics

Single session reports

Single session reports provide detailed information on a specific session.

Session Devices

<i>Description:</i>	Provides information about all destination devices that were used in the selected session.
<i>Required selections:</i>	session ID

<i>Optional selections:</i>	none
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	session_devices

Session Media

<i>Description:</i>	Provides information about all destination media that were used in the selected session.
<i>Required selections:</i>	session ID
<i>Optional selections:</i>	none
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	session_media

Session Object Copies

<i>Description:</i>	<p>Displays the number of valid copies in a selected backup, object copy, or object consolidation session.</p> <p>The VADP feature introduced in Data Protector 8.14 provides enhanced reports for Virtual Machines. The VMware virtual machines are represented as Data Protector clients called VADP clients. The new object name format for VADP clients is as follows:</p> <p><hostname>:/<vCenter>/<path>/<vmname> [<UUID>]</p> <p>Here, <hostname> is the DNS name of the guest virtual machine. If the DNS name is unknown, the IP address or VM name is used.</p>
<i>Required selections:</i>	session ID
<i>Optional selections:</i>	none
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	session_objcopies

Session Objects

<i>Description:</i>	<p>Lists all backup, object copy, or object consolidation objects and their statistics that were part of a selected session.</p> <p>The VADP feature introduced in Data Protector 8.14 provides enhanced reports for Virtual Machines. The VMware virtual machines are represented as Data Protector clients called VADP clients. The Session Objects report displays the VM name and VM path.</p>
---------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

<i>Required selections:</i>	session ID
<i>Optional selections:</i>	none
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	session_objects

Session per Client

<i>Description:</i>	<p>Provides information about each client that was part of the selected backup session. Using the Generate multiple reports option, this report can be split into smaller reports, one for each client.</p> <p>The VMware virtual machines are represented as Data Protector clients called VADP clients. The new object name format for VADP clients is as follows:</p> <p><hostname>:/<vCenter>/<path>/<vmname> [<UUID>]</p> <p>Here, <hostname> is the DNS name of the guest virtual machine. If the DNS name is unknown, the IP address or VM name is used.</p>
<i>Required selections:</i>	session ID
<i>Optional selections:</i>	message level
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	session_hosts

Single Session

<i>Description:</i>	Displays all relevant information about a single Data Protector backup, object copy, or object consolidation session.
<i>Required selections:</i>	session ID
<i>Optional selections:</i>	message level
<i>Supported formats:</i>	all formats
<i>omnirpt option:</i>	single_session

Reports Send Methods

You can choose among various send methods when configuring or starting a report or a report group.

Broadcast message send method

The broadcast message send method lets you send a broadcast message with the output of the report to specified systems.

Broadcast messages can be sent (to Windows systems only) by specifying the system to which the broadcast message should be sent. Broadcast messages are limited to 1000 characters, so the short format is preferred.

E-mail send method

You can send an e-mail with the output of the report to specified recipients. Make sure you provide the full e-mail address of the recipient.

Due to security features of Microsoft Outlook, using the e-mail send method may cause the CRS service to stop responding. For details and solutions, see the *HPE Data Protector Product Announcements, Software Notes, and References*. Alternatively, use e-mail (SMTP) as the e-mail send method.

Note: If Microsoft Exchange Server 2007 is installed on the Data Protector Cell Manager, the e-mail reporting send method does not work. Use the e-mail (SMTP) send method instead.

On Windows systems

To send an e-mail report from a Windows system, you need to have a mail profile. You can either use an existing mail profile or create a new one, named OmniBack.

To use an existing mail profile, add the following line to the Data Protector `omnirc` file:

```
OB2_MAPIPROFILE=existing_MAPI_profile_name
```

The display of HTML e-mail report on Windows depends on the e-mail client settings. Many email clients display the report as plain ASCII text. To ensure the report displays correctly as HTML, open it in a Web browser.

On UNIX systems

The e-mail subsystem has to be configured and running on a UNIX system; no additional configuration is needed.

Due to the operating system limitations, international characters in localized e-mail reports can be displayed incorrectly on UNIX systems, if they are passed between systems using a different locale.

E-mail (SMTP) send method

You can send an e-mail with the output of the report to specified recipients using the SMTP protocol. Make sure you provide the full e-mail address of the recipient.

This is the recommended e-mail send method.

By default, the address of the SMTP server used for sending the reports is set to the Cell Manager IP address. To change the address, edit the `SMTPServer` global option. The SMTP server must be accessible from the Cell Manager system, but does not need to be part of the Data Protector cell.

On Windows systems

On how to configure your existing Microsoft Exchange Server to support SMTP, see Microsoft Exchange Server documentation.

The display of HTML e-mail report on Windows depends on the e-mail client settings. Many e-mail clients display the report as plain ASCII text. To ensure the report displays correctly, open it in a Web browser.

On UNIX systems

Due to the operating system limitations, international characters in localized e-mail reports may display incorrectly on UNIX if they are passed between systems using a different locale.

External send method

The external script send method allows you to process the output of the report in your own script. The script receives the output as standard input (STDIN). The recommended format for script processing is the tab format.

The script, which is located on the Cell Manager system, must reside in the `/opt/omni/sbin` (HP-UX systems) or `Data_Protector_home\bin` (Windows systems) directory. Provide only the name of the script, not the entire path.

Note that only `.bat`, `.exe`, and `.cmd` are supported extensions for external scripts on Windows systems. To run a script with an unsupported extension (for example, `.vbs`), create a batch file that starts the script. Then configure Data Protector to run the batch file as an external script, which then starts the script with the unsupported extension.

You can also use this delivery method to perform a scheduled eject of the specified media.

Log to file send method

The log to file send method lets you post a file with the output of the report.

The file is posted to the Cell Manager system. You have to specify the name of the file to which you want to post the report. The file will be overwritten if it exists.

SNMP send method

The SNMP trap send method allows you to send a report as an SNMP trap. The SNMP trap can be further processed by applications that use SNMP traps.

Note: The SNMP send method is appropriate only for reports that do not exceed the maximum size of the configured SNMP trap. Otherwise, the report gets fragmented.

On Windows systems

SNMP traps are sent to the systems configured in the Windows SNMP traps configuration. You need to configure Windows SNMP traps to use the SNMP send method on the Cell Manager.

On UNIX systems

On a UNIX Cell Manager, SNMP traps are sent to the systems configured in the report.

Configuring Report Groups Using the Data Protector GUI

You can run Data Protector reports individually (interactively) or you can group them into report groups and then start the report group. You can add individual reports to an already configured report group. Mount Request Report and Device Error Report can be used only in a report group and are not available as interactive reports.

Using the Data Protector GUI, a report group allows you to:

- Start all the reports at once (interactively).
- Schedule the group to start the reports at a specified time.
- Start the group when triggered by a notification.

To display the input parameters (selections) in the output of a report, select the **Show selection criteria in report** option in the Report wizard. This option is not available for reports that have no required or optional input parameters (selections). The output of the report displays only required parameters and optional parameters with changed default values.

Prerequisites

- You either have to be added in the `admin` user group or granted the Reporting and notifications user rights.
- The Data Protector user under whose account the CRS service is running should not be removed. This user is configured by default at installation time. On a Windows Cell Manager, this is the user under whose account the installation was performed. On a UNIX Cell Manager, this is the `root` user of the Cell Manager.

Configuration phases

Configuring a report group

Steps

1. In the Context List, select **Reporting**.
2. Right-click **Reports**, and then click **Add Report Group** to open the wizard.
3. Name the report group and then click **Next**.
4. Optionally, schedule the report group using the Data Protector scheduler.

5. Click **Finish** to add the report group and exit this wizard. Follow the Add Report wizard to add reports.

Tip: To trigger a report group by a notification, configure a report group and then configure the notification to use the Use Report Group send method.

Adding a report to a report group

Steps

1. In the Reporting context, expand **Reports**, right-click a report group, and click **Add Report** to open the Add Report wizard. If configuring a report immediately after the report group configuration procedure, skip this step.
2. In the Results Area, select a type of report from the list.
3. In the Name text box, enter the name of the report and select a report in the Type drop-down list. Click **Next**.
4. The wizard options are available according to the selected report. For example, all wizard options available for the IDB Size report are not available for the List of Media report. Click **Next** as many times as needed to reach the last page of the wizard.
5. In the Send method drop-down list, select a sending method for the report, then enter the recipient of the report in the Email address text box. In the Format drop-down list, select the format of the report. Click **Add** to add the recipient to the group of configured recipients.
Repeat this step for any number of recipients.
6. Click **Finish** to add the report to the report group and exit the wizard.

Repeat this procedure for all the reports you want to add to a report group.

Running Report Groups Using the Data Protector GUI

You can run all the reports in a report group together.

Prerequisites

- You have to be either added in the Admin user group or granted the Reporting and notifications user rights.
- The Data Protector user under whose account the CRS service is running should not be removed. This user is configured by default at installation time. On a Windows Cell Manager, this is the user under whose account the installation was performed. On a UNIX Cell Manager, this is the root user of the Cell Manager.

Steps

1. In the Context List, select **Reporting**.
2. In the Scoping Pane, browse for and right-click the report group you want to start and then click **Start**.
3. Click **Yes** to confirm.

Running Individual Reports Using the Data Protector GUI

You can run individual reports interactively or you can group them into report groups and then run all the reports in the report group together.

Mount Request Report and Device Error Report can only be used in a report group and are not available as interactive reports.

Prerequisites

- You have to be in the Admin user group or have the Reporting and notifications user rights.
- The Data Protector user under whose account the CRS service is running should not be removed. This user is configured by default at installation time. On a Windows Cell Manager, this is the user under whose account the installation was performed. On a UNIX Cell Manager, this is the root user of the Cell Manager.

Steps

1. In the Context List, select **Reporting**.
2. Click the **Tasks** tab below the Scoping Pane.
3. In the Scoping Pane, browse for the desired type of report and select a report to open the wizard.
4. The wizard options are available according to the selected report. For example, all wizard options available for the IDB Size report are not available for the List of Media report. Click **Next** as many times as needed to reach the last page of the wizard.
5. At the end of the Report wizard, click **Finish** to display the output of the report.

Running Reports and Report Groups Using the Data Protector CLI

You can generate Data Protector reports using the command-line interface (CLI). The CLI allows you to include Data Protector reports in other scripts you are using. You can generate individual reports, start report groups, define report formats and send methods.

Prerequisites

- You have to be either added in the Admin user group or granted the Reporting and notifications user rights.
- The Data Protector user under whose account the CRS service is running should not be removed. This user is configured by default at installation time. On a Windows Cell Manager, this is the user under whose account the installation was performed. On a UNIX Cell Manager, this is the root user of the Cell Manager.

Steps

1. Use the `omnirpt` command to generate reports. For a detailed description of the command, see the `omnirpt man` page or the *HPE Data Protector Command Line Interface Reference*.

Creating a New Mail Profile

To send an e-mail report or notification from a Windows system, you need to have a mail profile. To create a new mail profile, named *OmniBack*, on a Windows system with Microsoft Outlook 2002 installed, use the procedure below.

Due to security features of Microsoft Outlook, using the e-mail send method may cause the CRS service to stop responding. For details and solutions, see the *HPE Data Protector Product Announcements, Software Notes, and References*. Alternatively, use the e-mail (SMTP) send method.

Steps

1. In the Windows Control Panel, double-click the **Mail** icon.
2. In the Mail Setup - Outlook dialog box, click **Show Profiles**.
3. In the Mail dialog box, click **Add**.
4. In the New Profile dialog box, type *OmniBack* in the Profile Name text box and click **OK** to start the E-Mail Accounts wizard.
5. Select **Add a new e-mail account** and click **Next**.
6. In the Server Type page, select **Microsoft Exchange Server** and click **Next**.
7. In the Exchange Server settings page, type the name of the local Microsoft Exchange Server system, your username, and click **Next**.
8. Click **Finish** to end the wizard.

Configuring Windows SNMP traps

On a Windows Cell Manager, SNMP traps are sent to the systems configured in the Windows SNMP traps configuration. On Windows systems, to send a notification or a report using the SNMP send method, you need to configure Windows SNMP traps.

On a UNIX Cell Manager, SNMP traps are sent to the systems configured in the notification or report; no additional configuration is needed.

Prerequisites

1. On Windows XP and Windows Server 2003 systems, a Windows installation DVD-ROM is required.

Steps

1. From the directory *Data_Protector_home\bin* invoke the *omnisnmp* command.
It will create the appropriate Data Protector entry in the System registry under *CurrentControlSet\Services\SNMP\Parameters\ExtensionAgents*.
2. **Windows XP, Windows Server 2003:**

- a. In the Control Panel, select **Network Connections**.
- b. In the Advanced menu, select **Optional Networking Components** to start the wizard.
- c. Select **Management and Monitoring** tools and click **Next**.
- d. Follow the wizard to install the management and monitoring tools.

Windows 7, Windows 8:

- a. In the Control Panel, select **Programs and Features**.
- b. Select **Turn Windows features on or off**.
- c. Select **Simple Network Management Protocol (SNMP)** and click **OK**.

Windows Server 2008, Windows Server 2012:

- a. In the Start menu, right-click **Computer** and select **Manage**.
 - b. Select **Features** and click **Add Features**.
 - c. In the Features tree, select **SNMP Services** and then **SNMP Service**.
 - d. Click **Next** and then **Install**.
3. Open **Control Panel, Administrative Tools, Services**.
 4. Right-click **SNMP Service** and select **Properties**.
 - a. Select the **Traps** tab. Enter `public` in the Community name text box and the hostname of the application management server in the Trap Destinations text box.
 - b. Select the **Security** tab. Under Accepted community names, select the community `public`, click **Edit** and set Community rights to `READ CREATE`.

If you select the **Under Accept SNMP packets from these hosts** and authentication fails, then use the IP address instead of host name in the Trap Destinations text box.
 - c. Confirm your changes.
 5. Invoke `omnisnmp`.

About Notifications

Data Protector allows you to send notifications from the Cell Manager when specific events occur. For example, when a backup, object copy, object consolidation, or object verification session is completed, you can send an e-mail with the status of the session.

You can set up a notification so that it triggers a report.

You can configure notifications using the Data Protector GUI or any Web browser with Java support.

Input parameters let you customize notifications. Some input parameters allow multiple selections. All other input parameters depend on the type of the notification. Depending on the send method, the recipient can be any of the following:

- a system
- an e-mail address
- an SNMP trap
- a script
- a file

- a configured report group
- the Data Protector Event Log

By default, notifications are configured with default values and are sent to the Data Protector Event Log. To send additional notification using some other sending method and/or other input parameters values, the configuration values must be changed.

To access the Data Protector notification functionality. You either have to be added in the `admin` user group or granted the **Reporting and notifications** user rights.

Notification Types - Events that Trigger Notifications

There are two main types of notifications.

- Notifications that are triggered when an event occurs
- Notifications that are scheduled and started by the Data Protector checking and maintenance mechanism

Alarm

<i>Event/notification name:</i>	Alarm
<i>What triggers the notification:</i>	Data Protector Internal critical conditions, such as Automated Media Copy upgrade, Upgrade Core Part End, Upgrade Detail Part End, Purge End, abort of session, Disk Agents upgrade during UCP, and so on.
<i>Default message level:</i>	Warning
<i>Message displayed:</i>	Alarm: <i>ALarm_message</i>

Expired Certificates

<i>Event/notification name:</i>	ExpiredCertificates
<i>What triggers the notification:</i>	The certificate stored on Cell Manager certificate directory is expired or not yet valid. The Cell Manager certificate directory stores all client certificates for Secure Control Communication.
<i>Default message level:</i>	Warning
<i>Message displayed:</i>	Certificate <i>certificate_name</i> expired or not yet valid.

Csa Start Session Failed

<i>Event/notification name:</i>	CsaStartSessionFailed
---------------------------------	-----------------------

<i>What triggers the notification:</i>	The backup session that ends with the error message: Could not start a new backup session.
<i>Default message level:</i>	Major
<i>Message displayed:</i>	CsaStartSession failed for datalist <i>datalist_name</i> .

Device Error

<i>Event/notification name:</i>	DeviceError
<i>What triggers the notification:</i>	An error on the device Device (default: <Any>).
<i>Default message level:</i>	Critical
<i>Message displayed:</i>	Error on device <i>Device</i> occurred.

End of Session

<i>Event/notification name:</i>	EndofSession
<i>What triggers the notification:</i>	A backup, copy, consolidation, or object verification session specified in the session specification Session Specification (default: <Any>) that ends with the message Session Status (default: Completed with errors).
<i>Default message level:</i>	Warning
<i>Messages displayed:</i>	Backup session <i>session_ID</i> of session specification <i>backup_specification</i> , backup specification group <i>group</i> completed with overall status <i>session_overall_status</i> ; <i>session_type</i> session <i>session_ID</i> of session specification <i>session_spec</i> , completed with overall status <i>session_status</i> .

File Library Disk Usage

<i>Event/notification name:</i>	FileLibraryDiskUsage
<i>What triggers the notification:</i>	A lack of free disk space for the file library Name of the File Library (default: All).
<i>Default message level:</i>	Warning
<i>Message displayed:</i>	The <i>File Library Device</i> is low in disk space in the <i>File Library Path</i> directory.

Health Check Failed

<i>Event/notification name:</i>	HealthCheckFailed
<i>What triggers the notification:</i>	<p>A non-zero value returned by the <code>omnihealthcheck</code> command. The command returns zero if the following is true:</p> <ul style="list-style-type: none">• The Data Protector services (CRS, MMD, <code>hdpd-idb</code>, <code>hdpd-idb-cp</code>, <code>hdpd-as</code>, KMS, <code>omnitrig</code>, and <code>omniinet</code>) are active.• The Data Protector Media Management Database (MMDB) is consistent.• At least one backup of the IDB exists. <p>For more information on this command, see the <code>omnihealthcheck</code> man page or the <i>HPE Data Protector Command Line Interface Reference</i>. By default, Data Protector starts the Health Check (which runs the <code>omnihealthcheck</code> command) once a day.</p>
<i>Default message level:</i>	Critical
<i>Message displayed:</i>	Health check message: <code>healthcheck_command</code> failed.

IDB Backup Needed

<i>Event/notification name:</i>	IDBBackupNeeded
<i>What triggers the notification:</i>	Too many successive incremental IDB backups or insufficiently frequent full IDB backup.
<i>Default message level:</i>	Warning
<i>Message displayed:</i>	There are <i>n</i> successive incremental IDB backups. The last backup of the Data Protector Internal Database was done on <i>MM/DD/YY hh:mm:ss</i> .

IDB Corrupted

<i>Event/notification name:</i>	IDBCorrupted
<i>What triggers the notification:</i>	Corruption of a part of the IDB.
<i>Default message level:</i>	Critical
<i>Message displayed:</i>	Corruption in the <i>IDB_part</i> part of the Data Protector Internal Database has been detected (<i>error_message</i>).

	<p>Values for error messages are:</p> <ul style="list-style-type: none">• Verification of datafile(s) failed.• KeyStore is corrupted.• Media and Media in position tables are not consistent.• Database is not in consistent state.• Database schema is not consistent.
--	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

IDB Limits

<i>Event/notification name:</i>	IDBLimits
<i>What triggers the notification:</i>	Reaching the limit of any of the MMDB or CDB parts.
<i>Default message level:</i>	Major
<i>Message displayed:</i>	The <i>IDB_part</i> part of the Data Protector Internal Database has reached its limit.

IDB Reorganization Needed

<i>Event/notification name:</i>	IDBReorganizationNeeded
<i>What triggers the notification:</i>	One or more IDB entities need to be reorganized due to fragmentation or wasted space.
<i>Default message level:</i>	Warning
<i>Message displayed:</i>	Bloat on table <i>name_of_table</i> detected. Fragmentation of table <i>name_of_table</i> on column <i>uuid</i> detected. Fragmentation of index <i>name_of_index</i> detected.

IDB Space Low

<i>Event/notification name:</i>	IDBSpaceLow
<i>What triggers the notification:</i>	<p>One of the following events:</p> <ul style="list-style-type: none">• The maximum free disk size is below the IDB Disk Free Threshold [MB] (default: 300 MB) value.• The difference between the maximum and current size of all DC directories falls below the DCBF Size Limit Threshold [MB] (default: 500 MB) value.

	<ul style="list-style-type: none"> The maximum free disk size is below the WAL Disk Free Threshold [MB] (default: 300 MB) value. <p>By default, Data Protector checks the IDB Space Low condition once a day.</p>
<i>Default message level:</i>	Major
<i>Message displayed:</i>	Data Protector Internal Database is running out of space.

License Warning

<i>Event/notification name:</i>	LicenseWarning
<i>What triggers the notification:</i>	A need for purchased licenses.
<i>Default message level:</i>	Warning
<i>Message displayed:</i>	<i>n</i> licenses need to be purchased for category <i>name of the license</i> . Run <code>omnicc -check_licenses -detail</code> for more info.

License Will Expire

<i>Event/notification name:</i>	LicenseWillExpire
<i>What triggers the notification:</i>	The forthcoming expiration date of the Data Protector license. The license will expire in number of days specified in License expires in days (default: 10).
<i>Default message level:</i>	Warning
<i>Message displayed:</i>	The first license will expire in <i>License expires in days</i> days.

Mail Slots Full

<i>Event/notification name:</i>	MailSlotsFull
<i>What triggers the notification:</i>	Full mail slots of the device Device (default: <Any>).
<i>Default message level:</i>	Warning
<i>Message displayed:</i>	All mail slots of library <i>Device</i> are full. Remove them immediately.

Mount Request

<i>Event/notification name:</i>	MountRequest
<i>What triggers the notification:</i>	A mount request for the device Device (default: <Any>).
<i>Default message level:</i>	Warning
<i>Message displayed:</i>	Mount request on device <i>Device</i> .

Not Enough Free Media

<i>Event/notification name:</i>	NotEnoughFreeMedia
<i>What triggers the notification:</i>	A lack of free media in the Media Pool . Notice, that if a Media Pool is configured to use a Free Pool , the Number of Free Media from the Free Pool is also considered.
<i>Default message level:</i>	Warning
<i>Message displayed:</i>	Media pool <i>Media Pool</i> contains only <i>number_of_media</i> free media.

Session Error

<i>Event/notification name:</i>	SessionError
<i>What triggers the notification:</i>	A backup, copy, consolidation, or object verification session with a message of the level Single Message Level (default: Major) or higher, displayed in the Monitor window.
<i>Default message level:</i>	Major
<i>Messages displayed:</i>	Backup session <i>session_ID</i> of session specification <i>backup_specification</i> , backup specification group <i>group</i> has errors: <i>number_of_errors</i> . <i>session_type</i> session <i>session_ID</i> of session specification <i>session_spec</i> has errors: <i>number_of_errors</i> .

Start of Session

<i>Event/notification name:</i>	StartofSession
---------------------------------	----------------

<i>What triggers the notification:</i>	A start of a backup, copy, consolidation, or object verification session specified in the session specification Session Specification (default: <Any>).
<i>Default message level:</i>	Normal
<i>Messages displayed:</i>	Backup session <i>session_ID</i> started for session specification <i>backup_specification</i> backup specification group <i>group</i> . <i>session_type</i> session <i>session_ID</i> started for session specification <i>session_spec</i> .

Too Many Sessions

<i>Event/notification name:</i>	TooManySessions
<i>What triggers the notification:</i>	Start of a session when 1000 sessions are already running concurrently.
<i>Default message level:</i>	Warning
<i>Message displayed:</i>	Session cannot start because the maximum number of concurrently running sessions has been reached.

Unexpected Events

<i>Event/notification name:</i>	UnexpectedEvents
<i>What triggers the notification:</i>	An unusually high number of new events in the Data Protector Event Log since the last time the check was made. The number exceeds Number of events (default: 20). By default, Data Protector checks the condition once a day.
<i>Default message level:</i>	Warning
<i>Message displayed:</i>	Data Protector Event Log increased for <i>number_of_events_in_last_day</i> unexpected events in the last day.

Check UNIX Media Agent

<i>Event/notification name:</i>	UnixMediaAgentWarning
<i>What triggers the notification:</i>	The <code>mrgcfg -check_ma</code> command triggers this notification when client devices are using rewind device files instead of no-rewind device files.

<i>Default message level:</i>	Warning
<i>Message displayed:</i>	Media Agents, clients devices may have been configured using rewind device files instead of no-rewind device files. This may lead to problems in SAN environments.

User Check Failed

<i>Event/notification name:</i>	UserCheckFailed
<i>What triggers the notification:</i>	<p>A non-zero value returned by the user-created script/command with the name Command Path located in the default Data Protector administrative commands directory.</p> <p>By default, Data Protector starts the User Check (which runs the script) once a day (default: None).</p>
<i>Default message level:</i>	Major
<i>Message displayed:</i>	User check failed with exit code <i>error_code:error_description</i> .

Notifications Send Methods

You can choose among various send methods when configuring a notification. By default, all notifications are configured to be sent to the Data Protector Event Log. To send a notification using another sending method, also, you must configure an additional notification. The available notification send methods are:

Broadcast Message send method

The broadcast message send method allows you to send a broadcast message with the output of the notification to specified systems after a specified event occurs.

Broadcast messages can be sent to Windows systems only by specifying the target system. Broadcast messages are limited to 1000 characters, so a short format is preferred.

E-mail send method

You can send an e-mail with the output of a notification to specified recipients. Make sure you provide the full e-mail address of the recipient.

Due to security features of Microsoft Outlook, using the e-mail send method may cause the CRS service to stop responding. For details and solutions, see the *HPE Data Protector Product Announcements, Software Notes, and References*. Therefore, the recommended method for sending e-mail notifications is SMTP.

Note: If Microsoft Exchange Server 2007 is installed on the Data Protector Cell Manager, the e-mail notification send method does not work. Use the e-mail (SMTP) send method instead.

On Windows systems

To send an e-mail notification from a Windows system, you need to have a mail profile. You can either use an existing mail profile or create a new one, named `OmniBack`.

To use an existing mail profile, add the following line to the Data Protector `omnirc` file:

```
OB2_MAPIPROFILE=existing_MAPI_profile_name
```

On UNIX systems

The e-mail subsystem has to be configured and running on a UNIX system.

Due to the operating system limitations, international characters in localized e-mail notifications can be displayed incorrectly on UNIX systems, if they are passed between systems using a different locale.

E-mail (SMTP) send method

You can send an e-mail with the output of a notification to specified recipients. Make sure you provide the full e-mail address of the recipient.

This is the recommended e-mail send method.

By default, the address of the SMTP server used for sending the notifications is set to the Cell Manager IP address. To change the address, edit the `SMTPServer` global option. The SMTP server must be accessible from the Cell Manager system, but does not need to be part of the Data Protector cell.

External send method

The external script send method allows you to process the output of the notification in your own script. The script receives the output as standard input (STDIN). The recommended format for script processing is the *tab* format.

The script, which is located on the Cell Manager system, must reside in the default Data Protector administrative commands directory. Provide only the name of the script, no path.

Note that only `.bat`, `.exe`, and `.cmd` are supported extensions for external scripts on Windows systems. To run a script with an unsupported extension (for example, `.vbs`), create a batch file that starts the script. Then configure Data Protector to run the batch file as an external script, which then starts the script with the unsupported extension.

You can also use this delivery method to perform a scheduled eject of the specified media.

Log to File send method

The log to file send method lets you post a file with the output of the notification when a specified event occurs.

The file is posted to the Cell Manager system. You have to specify the name of the file to which you want to post the notification. The file will be overwritten if it exists.

Data Protector Event Log send method

By default, all notifications are sent to the Data Protector Event Log. The Data Protector Event Log is accessible only for Data Protector users in the `admin` user group and to Data Protector users that are granted the Reporting and notifications user rights. You can view or delete all events in the Data Protector Event Log.

SNMP send method

SNMP send method allows you to send an SNMP trap with the output of the notification when a specified event occurs. The SNMP trap can be further processed by applications using SNMP traps.

On Windows systems

On a Windows Cell Manager, SNMP traps are sent to the systems configured in the Windows SNMP traps configuration. You need to configure Windows SNMP traps to use the SNMP send method on Windows systems.

On UNIX systems

On a UNIX Cell Manager, SNMP traps are sent to the systems configured in the notification.

Use report group send method

Use report group send method allows you to run a report group when a specified event occurs.

Configuring Notifications

To configure a notification you need to provide a name for the notification, a type of notification, message level, send method, and recipient. All other input parameters depend on the type of the notification.

Prerequisite

You either have to be added in the `admin` user group or granted the Reporting and notifications user rights.

Steps

1. In the Context List, select **Reporting**.
2. Right-click **Notifications** and click **Add Notification** to open the wizard.
3. The wizard options depend on the notification you selected. For example, all options available for the IDB Space Low notification are not available for the IDB Limits notification. Click **Next** as many times as needed to reach the last page of the wizard.
4. Click **Finish** to exit the wizard.

The notification will be sent using the specified send method when the specified event occurs.

Tip: To trigger a report group by a notification, configure a report group and then configure the notification to use the Use Report Group send method.

About Web Reporting and Notifications

You can use a Web browser to view, configure, and start Data Protector reports and notifications from any system on the network. Using the Web reporting interface, you can configure reports and notifications that are delivered using various reporting methods and formats.

All reporting and notifications functionality accessible using the Data Protector GUI is also accessible by using the Data Protector Web reporting and notifications. It is also available for a multiple-cell configuration when you use the Manager-of-Managers functionality.

When the report is displayed, you can print the report or save it. When you save the report, you can also add this report to an existing or a new report group.

When you install the Data Protector Cell Manager, the Web reporting user (called `java`) is automatically created. By default, no password is needed to use Data Protector Web reporting and notifications. It is strongly recommended to set the password to restrict access to Data Protector Web reporting and notifications functionality.

Requirements

- A supported Web browser must be installed. For a list of Web browsers supported for web reporting and notifications, see the *HPE Data Protector Product Announcements, Software Notes, and References*.
- A supported Java runtime environment must be installed on the system and enabled in the Web browser. For a list of supported Java runtime environments, see the *HPE Data Protector Product Announcements, Software Notes, and References*.
- Use Java 32-bit.
- Deselect `Use SSL2.0 compatible ClientHello` format in the Java control panel.

Limitations

- You cannot edit, view, or delete the saved reports using the Web reporting and notifications interface.
- You cannot start a report group using the Web reporting and notifications interface.
- Whenever multiple input parameters (selections) are to be *typed* in the Web reporting and notifications interface, every parameter has to be wrapped in double quotes if it contains spaces.

Configuring and Launching Web Reporting and Notifications Interface

Prerequisite

A configured and running web server.

Note: The web server does not have to be a Data Protector client.

Steps

1. Set the password for Web reporting as described in [Configuring a Password for Web Reporting](#).
2. Start the Web reporting using the Data Protector GUI as follows:
 - a. In the Context List, click **Reporting**.
 - b. In the Actions menu, click **Web Reporting**.

The URL is opened automatically.

(Or)

Start the Web reporting in a Web browser by manually entering the following URL:

`https://cmhost.domain.com:7116/webreporting/WebReporting.html`

Where,

cmhost.domain.com = actual FQDN name of the cell manager system.

7116 = default Application Server port defined during installation.

The Login to Data Protector Cell Manager page is displayed.

Note: If the browser reports a security certificate while logging in to the Data Protector Cell Manager, click **Continue to the website**. You receive two security warnings before adding the security certificate to the trusted store. When you receive a request for the Java program to run, click **Don't block Java**.

3. In the Login to Data Protector Cell Manager page, proceed as follows:
 - a. In the Cell Manager text box, enter the name of the cell manager system.
 - b. In the Password text box, enter the password set in [Step 1](#).
4. Click **Login**.

Configuring a Password for Web Reporting

When you install the Data Protector Cell Manager, the Web reporting user (called java) is automatically created. By default, the Web user password is random. You need to change the password before accessing the Data Protector Web reporting functionality.

Steps

1. In the Context List, click **Users**.
2. In the **Actions** menu, click **Set Web User Password**.
A Set Web User Password dialog box is displayed.
3. In the Set Web User Password page, proceed as follows:
 - a. In the New Password text box, enter the password.
 - b. In the Confirm New Password text box, re-enter the new password.
 - c. Click **OK** to save the password information.

Configuring Report Groups Using the Web Reporting Interface

You can run Data Protector reports individually (interactively) or group them into report groups. Mount Request Report and Device Error Report can only be used in a report group and are not available as interactive reports.

To display the input parameters (selections) in the output of a report, select the **Display input parameters** option in the report wizard. This option is not available for reports that have no required or optional input parameters (selections). The output of the report displays only required parameters and optional parameters with the changed default values.

Prerequisites

If configured, a web reporting password is required to access the Web reporting and notifications functionality.

Steps

1. Log in to the launched Web Reporting Interface.
2. Expand the Cell Manager item in the Scoping Pane and browse for the desired type of report.
3. Depending on the report selected, provide the needed data. The options that are available depend on which report is selected. For example, all options available for the IDB Size report are not available for the List of Media report.
4. Click **Show** to display the output of the report in the Results Area.
5. Click **Save Report** to add the report to a report group.
6. Type a name for the report. Specify a new or an already configured report group for the report. Click **Define Method & Save Report**.
7. Select the report send method and report format and add the recipients of the report. Click **Save Report** to save the report.

The report group with the report is displayed in the Report Groups folder in the Scoping Pane.

Running Individual Reports Using the Web Reporting Interface

You can run individual reports interactively or you can group them into report groups and then run all the reports in the report group together.

Mount Request Report and Device Error Report can only be used in a report group and are not available as interactive reports.

Prerequisite

If configured, a web reporting password is required to access web reporting and notifications functionality.

Steps

1. Login to the launched Web Reporting Interface.
2. Expand the Cell Manager item in the Scoping Pane. Expand the report type group of your choice and select the report to display the dialog in the Results Area.

Depending on the report selected provide the needed data. The options are available according to the selected report. For example, all options available for the IDB Size report are not available for the List of Media report. Click **Show** to display the output of the report in the results area.

Running Saved Reports Using the Web Reporting Interface

You can run any report saved in a Report Group using the Web reporting interface.

Prerequisite

If configured, a web reporting password is required to access web reporting and notifications functionality.

Steps

1. Login to the launched Web Reporting Interface.
2. Expand the Cell Manager item in the Scoping Pane. Expand **Report Groups** and the report group with the saved report you want to start. Select the saved report you want to start to display the output of the report in the Results Area.

Configuring Notifications Using the Web Reporting Interface

To configure a notification using the Web reporting interface you need to provide a name for the notification, a type of notification, send method, and recipient. All other input parameters depend on the type of the notification.

Prerequisite

If configured, a web reporting password is required to access web reporting and notifications functionality.

Steps


1. Log in to the launched Web Reporting Interface.
2. Expand the Cell Manager item in the Scoping Pane.
3. In the Scoping Pane, select **Notifications** to display the list of already configured notifications in the Results Area.
4. Click **Add Notification** and enter the required data. The options that are available depend on which notification is selected. For example, all options available for the End of Session notification are not available for the Device Error notification.
5. Click **Save Notification** to save the notification.

About Data Protector Event Log

The Data Protector Event Log represents a centralized event management mechanism, dealing with specific events that occurred during the Data Protector operation. The Data Protector event logging mechanism logs two types of events: process-triggered and user-triggered. The events are logged on the Cell Manager in the `Ob2EventLog.txt` file residing in the default Data Protector log files directory.

Viewing the Data Protector Event Log using the Event Log Viewer helps you troubleshoot possible problems.

When the Data Protector graphical user interface is started by a user, if there are new notifications that have not been seen by this user in the Data Protector Event Log, the following message is displayed:



The Data Protector GUI is automatically switched to the Reporting context.

The following may provide additional information:

- You have to be either a member of the `admin` user group or granted the Reporting and notifications user rights.
- The Data Protector Event Log is not refreshed automatically. To view the new messages, refresh it manually by pressing **F5**.

Process-triggered events

An event is logged by the notifications functionality.

User-triggered events

An event is logged when a user performs a certain GUI operation or a set of GUI operations. This set of operations includes modifications of backup, object copy and consolidation specifications, operations on users and user groups, creation and modifications of devices and media related configuration, and remote installation operations.

Logging of user-triggered events is disabled by default. To enable it, you must set the global option `EventLogAudit` to 1.

In a MoM environment, if the global option is set to 1, the events are logged only on the local Cell Manager system.

Accessing Event Log Viewer

You can browse the recorded events by accessing the Data Protector Event Log Viewer.

Prerequisite

You have to be either a member of the `admin` user group or granted the Reporting and notifications user rights.

Steps

1. In the Context List, select **Reporting**.
2. In the Scoping Pane, expand **Reporting**.
3. Select **Event Log** to display it.

Deleting Event Log Viewer Contents

Note: Deleting the Event Log Viewer contents does not delete the contents of the `Ob2EventLog.txt` file.

Prerequisite

You either have to be a member of the Admin user group or granted the Reporting and notifications user rights.

Steps

1. In the Context List, select **Reporting**.
2. In the Scoping Pane, expand **Reporting**.
3. Right-click **Event Log** and select **Empty Event Log** to delete all entries in the Event Log Viewer.

About Auditing

Data Protector provides backup session auditing, which stores non-tamperable and non-overwritable information about all backup tasks that were performed over user-defined periods for the whole Data Protector cell. The auditing information is retrievable on demand in an integral and printable fashion for auditing or administrative purposes.

You can enable auditing information logging and set the retention period for audit log files by modifying the `AuditLogEnable` and the `AuditLogRetention` global options.

Generating an Audit Report

To generate an audit report, follow the steps below.

Note: In a MoM environment, you have to perform audit reports for each Cell Manager separately.

Steps

1. In the Context List, click **Internal Database**.
2. In the Scoping Pane, click the **Auditing** item to open the Auditing page.

3. From the Search interval drop-down list, select one of the values (for example, Last week).
4. Click the **Update** button to display a list of all backup sessions performed during the selected period.
5. Select a specific session from the session list to display detailed information about used media and objects in the middle and bottom part of the Auditing property page.

Chapter 15: Telemetry

Telemetry is a service that is used to capture insights of Data Protector environment for better supportability, product best practices, and account management. The data collected through telemetry service is transmitted to HPE support backend for further analysis to enhance customer experience.

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

- Version information - Data Protector components and its versions. It also provides information on the host OS version.
- Devices or Media Servers - the details associated to a client in the Cell Manager that includes the hostname details, which has the device usage size attached, name of the device, library name, device type, and pool name where the media is placed.
- Device usage size – The usage size of the device.
- Schedules - The schedule provides information grouped by backup and session types and 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.
- Client usage - This information is collected for each client which includes: hostname, application name and total size of data backed up.
- Storage usage - Total data backed up on the device.

Prerequisites

- When configuring telemetry, the customer name and proxy information should be available.
- The system should be configured with java user (java, applet, webreporting) in Data Protector users list.

How Telemetry works?

During telemetry registration, the information is stored in the IDB. Any Windows host where the DP GUI component is installed is a suitable telemetry client. Depending on the configured frequency of upload, the Cell Console (CC) client checks the IDB settings and fetches the telemetry data from the Cell Manager and uploads the data to the backend server.

If the service is offline during the collection process, then telemetry upload does not happen at that time. When service is online, then the client becomes eligible to upload the telemetry data.

In a cell, if there are multiple telemetry clients then only one client performs the upload to the backend according upload frequency. IDB is updated with the upload status and time.

From telemetry clients, if a proxy is required to access the external network then you need to specify the proxy parameters. This can be ignored if a direct connection is possible.

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

Limitations

- The Telemetry Client service is supported only on Windows x64 operating system.
- The Telemetry Client service is not supported on Windows Cluster system.

Telemetry troubleshooting

When system is unable to upload the telemetry data to HPE backend server

1. Problem: Telemetry client not part of any Cell server
Action: Import the client to the Cell Manager.
2. Problem: Data Protector Cell Manager is not accessible
Action: Add web user to Data Protector user list. The system should be configured with java user (java, applet, webreporting) in Data Protector users list.
3. Problem: CAF registration failed
Action: Configure proxy properly from GUI.
4. Problem: Telemetry client service not running
Action: Start the telemetry client service.

Telemetry Registration page

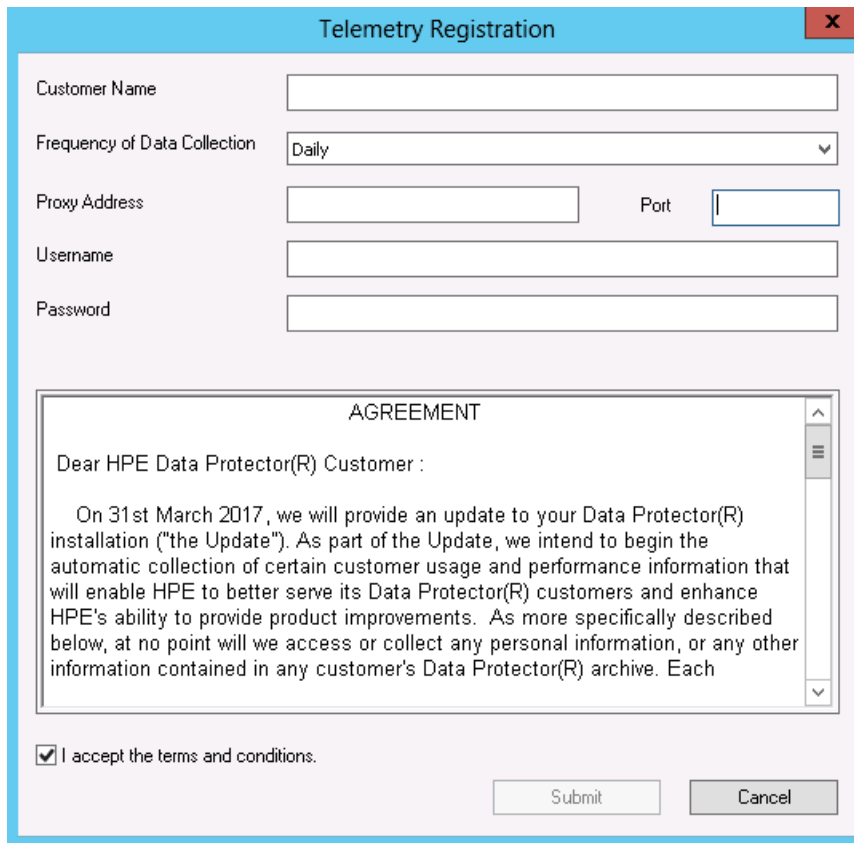
The customer can opt-in to or opt-out from the collection of telemetry data. To browse the telemetry registration page, proceed as follows:

1. In the **Context** List, select **Clients**.
2. Click **Actions** and select **Telemetry Registration**.

To register, enter the following parameters:

- **Customer Name:** Name of the customer.
- **Frequency of Data Collection:** The user can select the frequency in which the data can be collected, that is: Daily, Weekly, Monthly, or Quarterly.
- **Proxy Address:** The address of proxy server configured.
- **Port:** The port of the proxy server.
- **Username:** User name to connect to the proxy server.
- **Password:** Password for the given user name.

Note: The proxy details are configured on the Cell Console (CC) client.



The image shows a 'Telemetry Registration' dialog box with a blue title bar and a red close button. It contains several input fields: 'Customer Name' (text box), 'Frequency of Data Collection' (dropdown menu with 'Daily' selected), 'Proxy Address' (text box), 'Port' (text box), 'Username' (text box), and 'Password' (text box). Below these fields is an 'AGREEMENT' section with a scrollable text area containing a message from HPE. At the bottom, there is a checkbox labeled 'I accept the terms and conditions.' and two buttons: 'Submit' and 'Cancel'.

Telemetry Registration

Customer Name

Frequency of Data Collection: Daily

Proxy Address

Port

Username

Password

AGREEMENT

Dear HPE Data Protector(R) Customer :

On 31st March 2017 , we will provide an update to your Data Protector(R) installation ("the Update"). As part of the Update, we intend to begin the automatic collection of certain customer usage and performance information that will enable HPE to better serve its Data Protector(R) customers and enhance HPE's ability to provide product improvements. As more specifically described below, at no point will we access or collect any personal information, or any other information contained in any customer's Data Protector(R) archive. Each

☒ I accept the terms and conditions.

Submit Cancel

After entering the above fields and selecting the frequency of data collection, accept the terms and conditions and click **Submit**.

Note: The customer related information is gathered, but the host information is masked or replaced with random characters and numbers.

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 Guide				ZDB Guides		GRE Guide		
										Integration VSS	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			
CLI								X										
Concepts, techniques		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			
New features		X							X									
Planning strategy		X		X														
Procedures, tasks	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				
Restore	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

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	<p>This guide describes how to configure and use the Data Protector Granular Recovery Extension for:</p> <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	<p>This guide describes the integrations of Data Protector with the following applications:</p> <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 disk arrays. It is intended for backup administrators or operators. It covers the zero downtime backup, instant recovery, and the restore of filesystems and disk images.

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
Network Data Management Protocol (NDMP) Server	Integration Guide
Oracle Server	Integration Guide, ZDB IG
MySQL Server	Integration Guide
SAP HANA Appliance	Integration Guide
SAP MaxDB	Integration Guide
SAP R/3	Integration Guide, ZDB IG
Sybase Server	Integration Guide

Software application	Guides
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	Concepts, ZDB Admin, ZDB IG
EMC VMAX	Concepts, ZDB Admin, ZDB IG

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 Administrator's Guide (Data Protector 9.09)

Just add your feedback to the email and click send.

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

We appreciate your feedback!