Hewlett Packard
Enterprise

HPE Data Protector

Software Version: 9.08

Command Line Interface Reference

Legal Notices

Warranty

Restricted Rights Legend

Copyright Notice

Trademark Notices

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

Section 9: Introduction

omniintro(9)

omniintro — introduction to the HPE Data Protector commands and command-line utilities

DESCRIPTION

HPE Data Protector is an enterprise backup solution that provides reliable data protection and high accessibility for business data. Data Protector provides extensive media management, unattended backups, post-backup data management, integrations with various databases and supports various backup and other backup-dedicated devices. For information on the Data Protector concepts and functionality, see the Data Protector guides and the Data Protector Help.

COMMANDS

USER COMMANDS (1):

omniabort

> Aborts an active session.
>
> This command is available on systems with the Data Protector `User Interface` component installed.

omniamo

> Starts an automated media operation session.
>
> This command is available on the Data Protector Cell Manager.

omnib

> Backs up filesystems, disk images, the Data Protector Internal Database (IDB), Microsoft Exchange Server single mailboxes and Public Folders, Microsoft Exchange Server 2007, Microsoft Exchange Server 2010/2013, Microsoft SQL Server, Microsoft SharePoint Server 2007/2010/2013, SAP R/3, SAP MaxDB, Oracle, MySQL, PostgreSQL, Informix Server, VMware vSphere, Microsoft Hyper-V, Sybase, Lotus, IBM DB2 UDB, and NDMP objects.
>
> This command is available on systems with the Data Protector `User Interface` component installed.

omnicc

> Handles the Data Protector licensing, reports the number of configured and available Data Protector licenses, installs the licenses, imports and exports Data Protector clients, manages access to secured clients, enables encrypted control communication, and creates a template for the user_restrictions file.
>
> This command is available on systems with any Data Protector component installed.

omnicellinfo

> Displays configuration information about the Data Protector cell.

This command is available on systems with the Data Protector `User Interface` component installed.

omniclus

Manages load balancing in a cluster environment in the event of an application (Data Protector or other) failover.

This command is available on systems with the Data Protector `MS Cluster Support` component installed (Windows systems) and on the Data Protector Cell Manager (UNIX systems).

omnicreatedl

Creates a filesystem backup specification file (datalist); or an HPE P9000 XP Disk Array Family or HPE P6000 EVA Disk Array Family Microsoft Exchange Server 2003 ZDB backup specification file (datalist).

This command is available on systems with the Data Protector `User Interface` component installed.

omnidb

Queries the Data Protector Internal Database (IDB).

This command is available on systems with the Data Protector `User Interface` component installed.

omnidbp4000

Manages the configuration data which the Data Protector HPE P4000 VSS Agent uses to connect to the CIMOM providers.

This command is available on Windows systems with the Data Protector `User Interface` component installed.

omnidbsmis

Executes administrative tasks on the ZDB database (SMISDB) and on a disk array of the HPE P6000 EVA Disk Array Family.

This command is available on systems with the Data Protector `User Interface` component installed.

omnidbvss

Queries the VSS database; manages, browses, and lists the items of the VSS database.

This command is available on systems with the Data Protector `User Interface` component installed.

omnidbxp

Queries the ZDB database (XPDB), manipulates the P9000 XP LDEV exclude file, configures the HPE P9000 XP Disk Array Family command devices usage, and manages the user authentication data which the Data Protector HPE P9000 XP Agent uses to connect to specific disk arrays.

This command is available on systems with the Data Protector `User Interface` component installed.

omnidbzdb

Executes administrative tasks on HPE 3PAR StoreServ Storage, NetApp Storage, EMC VNX, and

EMC VMAX as well as manages the configuration data, which the integration agents use to connect to the CIMOM providers and storage systems.

This command is available on systems with the Data Protector `User Interface` component installed.

`omnidownload`

Downloads information about a backup device and a library from the Data Protector Internal Database (IDB).

This command is available on systems with the Data Protector `User Interface` component installed.

`omniiso`

Primarily serves as a pre-exec script to prepare the ISO image file for One Button Disaster Recovery (OBDR); can also be used as a standalone command to automate your backup and disaster recovery process.

This command is available on systems with the Data Protector `Automatic Disaster Recovery` component installed.

`omnimcopy`

Makes a copy of a Data Protector medium using Data Protector backup devices as the source and destination.

This command is available on systems with the Data Protector `User Interface` component installed.

`omniminit`

Initializes a Data Protector medium.

This command is available on systems with the Data Protector `User Interface` component installed.

`omnimlist`

Lists the contents of a Data Protector medium.

This command is available on systems with the Data Protector `User Interface` component installed.

`omnimm`

Provides media management for Data Protector.

This command is available on systems with the Data Protector `User Interface` component installed.

`omnimnt`

Responds to a Data Protector mount request for a medium.

This command is available on systems with the Data Protector `User Interface` component installed.

`omnimver`

Verifies data on a medium.

This command is available on systems with the Data Protector `User Interface` component installed.

omniobjconsolidate

Consolidates Data Protector backup objects into synthetic full backups.

This command is available on systems with the Data Protector `User Interface` component installed.

omniobjcopy

Creates additional copies of objects backed up with Data Protector on a different media set.

This command is available on systems with the Data Protector `User Interface` component installed.

omniobjverify

Verifies Data Protector backup objects, either interactively or using pre-configured post-backup, or scheduled verification specifications.

This command is available on systems with the Data Protector `User Interface` component installed.

omnir

Restores filesystems, disk images, the Data Protector Internal Database (IDB), Microsoft Exchange Server single mailboxes and Public Folders, Microsoft Exchange Server 2007, Microsoft Exchange Server 2010/2013, Microsoft SQL Server, Microsoft SharePoint Server 2007/2010/2013, MySQL, PostgreSQL, SAP R/3, SAP MaxDB, Informix Server, VMware vSphere, VMware vCloud Director, Microsoft Hyper-V, Lotus, IBM DB2 UDB, and NDMP objects backed up with Data Protector. The command is also used to start the instant recovery process. To restore a Sybase database, see the `syb_tool` man page.

This command is available on systems with the Data Protector `User Interface` component installed.

omnirpt

Generates various reports about the Data Protector environment, for example, about backup, object copy, object consolidation, and object verification sessions in a specific time frame, session specifications, media, Data Protector configuration, and single sessions.

This command is available on systems with the Data Protector `User Interface` component installed.

omnistat

Displays the status of active Data Protector backup and restore sessions.

This command is available on systems with the Data Protector `User Interface` component installed.

omniupload

Uploads information about a backup device from an ASCII file to the Data Protector Internal Database (IDB).

This command is available on systems with the Data Protector `User Interface` component installed.

omniusb

Writes the DR ISO image to a USB drive, and makes the drive bootable

This command is available on systems with the Data Protector `Automatic Disaster Recovery` component installed.

omniusers

Adds or removes Data Protector users to or from an existing Data Protector user group, or lists the configured Data Protector users.

This command is available on non-Windows systems with the Data Protector `User Interface` component installed.

SharePoint_VSS_backup.ps1

Creates backup specifications and starts backup sessions for Microsoft SharePoint Server.

This command is available on Windows systems with the Data Protector `MS Volume Shadow Copy Integration` component installed.

syb_tool

A utility used to get ISQL command needed to restore a Sybase database that was backed up by Data Protector.

This command is available on systems with the Data Protector `Sybase Integration` component installed.

ADMINISTRATIVE COMMANDS (1M):

ob2install

Runs installation, removal, upgrade, or installation check of the specified Data Protector components to/from/on a remote UNIX system using the specified an appropriate Installation Server.

This command is available on the Data ProtectorInstallation Server.

omnib2dinfo

Displays information about ObjectStore and StoreOnceSoftware stores.

This command is available on systems with the Data Protector `User Interface` component installed.

omnicheck

Performs a DNS connections check within a Data Protector cell and lists Data Protector patches installed on Data Protector clients.

This command is available on systems with any Data Protector component installed.

omnidbcheck

Checks the consistency of the Data Protector Internal Database (IDB).

This command is available on the Data Protector Cell Manager.

omnidbinit

Initializes the Data Protector Internal Database (IDB).

This command is available on the Data Protector Cell Manager.

omnidbutil

Handles various Data Protector Internal Database (IDB) maintenance tasks.

This command is available on the Data Protector Cell Manager.

omnidlc

Gathers or deletes Data Protector debug, log, and getinfo files from the Data Protector cell or from a MoM environment.

This command is available on the Data Protector Cell Manager

omnidr

A general purpose Data Protector disaster recovery command. Based on its input, it decides on what type of restore to perform (online restore using `omnir` or offline restore using `omniofflr`), as well as how to perform the restore (whether or not to use live operating system features).

This command is available on systems with the Data Protector `User Interface` component installed.

omnihealthcheck

Checks the status of Data Protector services, the consistency of the Data Protector Internal Database (IDB), and if at least one backup of the IDB exists.

This command is available on the Data Protector Cell Manager.

omniinetpasswd

Manages the local Data Protector Inet configuration on Windows systems where the Inet process must be run under a specific user account, and sets a user account to be used by the Installation Server during remote installation.

This command is available on systems with any Data Protector component installed.

omniintconfig.pl

Configures, updates configuration parameters, and checks the configuration of one or multiple Oracle databases.

This command is available on systems with the Data Protector `User Interface` component installed.

omnikeytool

Manages keys used to encrypt backup data.

This command is available on the Data Protector Cell Manager.

omnimigrate.pl

Migrates the Data Protector Internal Database (IDB) from the format used in earlier versions to the format used in the latest Data Protector version.

This command is available on the Data Protector Cell Manager.

omniofflr

Enables restore of any type of Data Protector backup objects in the absence of operable Data Protector Internal Database (IDB), including the IDB itself.

This command is available on systems with any Data Protector component installed.

omniresolve

Resolves a filesystem object or a list of filesystem objects and writes the results to the standard output or to a Unicode file.

This command is available on systems with any Data Protector integration component installed.

omnirsh

Returns the hostnames of the physical and virtual nodes for the specified cluster hostname, or returns the cell information stored in the cell_info file on the specified cluster.

This command is available on the Data Protector Cell Manager.

omnisetup.sh

Installs or upgrades a Data Protector UNIX Cell Managers, UNIX Installation Servers, UNIX and Mac OS X cli systems locally; installs and removes patch bundles.

This command is available on the Data Protector installation DVD-ROMs for UNIX systems or is provided together with a patch bundle.

omnisrdupdate

Updates the System Recovery Data (SRD) file.

This command is available on systems with the Data Protector User Interface component installed.

omnisv

Starts or stops the Data Protector services or daemons, displays their status, or turns the maintenance mode on or off.

This command is available on the Data Protector Cell Manager

omnitrig

Triggers Data Protector scheduled backups.

This command is available on the Data Protector Cell Manager.

sanconf

Auto-configures a library, modifies an existing library or drive configuration, or removes drives from a library configuration, within a SAN environment.

This command is available on systems with the Data Protector User Interface component installed.

upgrade_cm_from_evaa

Upgrades the EVADB entries created by the HPE EVA Agent (legacy) to the SMISDB entries created by the HPE P6000 / HPE 3PAR SMI-S Agent.

This command is available on the Data Protector Cell Manager.

util_cmd

Sets, retrieves or lists the parameters stored in the Data Protector Oracle, MySQL, SAP R/3, Microsoft Exchange Server 2010/2013, Informix, and Sybase configuration files.

This command is available on systems with any Data Protector component installed.

`util_oracle8.pl`

Configures an Oracle database and prepares the environment for backup, and checks the configuration of an Oracle database.

This command is available on systems with the Data Protector `Oracle Integration` component installed.

`vepa_util.exe`

Configures a VMware ESX(i) Server system, VMware vCenter Server system, VMware vCloud Director, Microsoft Hyper-V system, checks the configuration, configures virtual machines, browses and lists VMware datacenters and VMware vCloud Director organizations.

This command is available on systems with the Data Protector `Virtual Environment Integration` component installed.

COMMAND-LINE UTILITIES (1M):

`cjutil`

Starts, stops, and queries the Windows Change Journal.

This command is available on systems with the Data Protector `Disk Agent` component installed.

`omnicjutil`

Remotely controls and administers the Windows Change Journal on Windows clients.

This command is available on the Data Protector Cell Manager.

`omnistoreapputil`

Acts as a user interface to Storage Appliances, such as IAP and VLS.

This command is available on the Data Protector Cell Manager.

`uma`

Controls the robotics of SCSI compliant autochangers.

This command is available on systems with the Data Protector `General Media Agent` or `NDMP Media Agent` component installed.

RETURN VALUES:

Possible return values of commands are:

`0` - Program completed successfully.

`1` - Program failed, command syntax error.

`2` - Program failed, invalid argument.

`3` - Program failed, internal error.

`4` - Program failed, reason unknown.

Some commands may return additional error messages. These are described in individual reference pages.

COMMANDS FOR LAUNCHING THE Data Protector GUI

`manager`

> Launches the Data Protector GUI with all Data Protector contexts activated or, when additional options are specified, with the specified contexts activated.

> This command is available on systems with the Data Protector `User Interface` component installed.

`mom`

> Launches the Data Protector Manager-of-Managers GUI with all Data Protector contexts activated (with the exception of the Internal Database and Devices & Media contexts) or, when additional context options are specified, with the specified contexts activated.

> This command is available on systems with the Data Protector `Manager-of-Managers User Interface` component installed.

COMMAND LOCATIONS

WINDOWS SYSTEMS:

- user commands (1), administrative commands (1M), command-line utilities (1M):

  Data_Protector_home`\bin`
- commands that launch the Data Protector GUI (5):

  Data_Protector_home`\bin`

HP-UX, SOLARIS, AND LINUX SYSTEMS:

- user commands (1):

  `/opt/omni/bin`
- administrative commands (1M), command-line utilities (1M):

  `/opt/omni/lbin`

  `/opt/omni/sbin`

OTHER UNIX SYSTEMS:

- user commands (1), administrative commands (1M), command-line utilities (1M):

  `/usr/omni/bin`

HPE recommends that you enable invocations of the Data Protector commands from any directory by extending the value of the appropriate environment variable in your operating system configuration with the above paths. Procedures in the Data Protector documentation assume the value has been extended.

DIRECTORY STRUCTURE ON WINDOWS CELL MANAGERS

Data_Protector_home

- Data Protector home directory

Data_Protector_home\bin

- Directory containing Data Protector commands, Disk Agent, Media Agent files, message catalogs, and commands for Cell Manager maintenance

Data_Protector_home\docs

- The Data Protector guides, including the *HPE Data Protector Command Line Interface Reference*, the Data Protector support matrices

Data_Protector_home\help

- The Data Protector Help

Data_Protector_program_data

- Data Protector program data directory

Data_Protector_program_data\Config\client

- Directory containing the client configuration directories and files

Data_Protector_program_data\Config\Server

- Directory containing the following configuration directories:

  barlists - database backup specifications

  barschedules - database backup specification schedules

  cell - the cell configuration

  datalists - backup specifications

  devices - templates for devices

  options - default options

  schedules - backup schedules

  sessions - data about sessions

  snmp - the OpenView/SNMP trap sending configuration

  users - the user configuration

Data_Protector_program_data\Config\Server\dr

- Directory containing the following disaster recovery directories:

  asr - ASR archive files

  p1s - P1S files for Enhanced Automated Disaster Recovery

  srd - SRD files

Data_Protector_program_data\Config\Server\export\keys and *Data_Protector_program_data*\Config\Server\import\keys

- Directories containing encryption keys

Data_Protector_program_data\server\db80

- The Data Protector Internal Database (IDB)

Data_Protector_program_data\server\db80\idb

- The IDB tablespaces

Data_Protector_program_data\server\db80\dcbf

- The Detail Catalog Binary Files (DCBF) part of the IDB

Data_Protector_program_data\server\db80\keystore

- The encryption keystore database

Data_Protector_program_data\server\db80\keystore\catalog

- The keyid catalog

Data_Protector_program_data\server\db80\logfiles

- The IDB archived log files and the IDB recovery file (obdrindex.dat)

Data_Protector_program_data\server\db80\msg

- The Data Protector session messages

Data_Protector_program_data\server\db80\smisdb

- The ZDB database (SMISDB)

Data_Protector_program_data\server\db80\smisdb\p4000\login

- The data which the Data Protector HPE P4000 VSS Agent uses to connect to the configured CIMOM providers

Data_Protector_program_data\server\db80\smisdb\p10000\login

- The data which the Data Protector HPE 3PAR VSS Agent and the Data Protector HPE P6000 / HPE 3PAR SMI-S Agent use to connect to the configured CIMOM providers for the HPE 3PAR StoreServ Storage disk arrays

Data_Protector_program_data\server\db80\smisdb\netapp\login

- the data which the Data Protector NetApp Storage Provider uses to connect to the NetApp Storage system

Data_Protector_program_data\server\db80\smisdb\emcvmax\login

- the data which the Data Protector EMC VMAX Storage Provider uses to connect to the EMC VMAX Storage system

Data_Protector_program_data\server\db80\smisdb\emcvnx\login

- the data which the Data Protector EMC VNX Storage Provider uses to connect to the EMC VNX Storage system

Data_Protector_program_data\server\db80\vssdb

- The VSS database (VSSDB)

Data_Protector_program_data\server\db80\xpdb

- The ZDB database (XPDB)

Data_Protector_program_data\log and *Data_Protector_program_data*\log\server

- Log files

Data_Protector_program_data\log\server\auditing

- Audit logs

Data_Protector_program_data\tmp

- Temporary and debug log files

DIRECTORY STRUCTURE ON UNIX CELL MANAGERS

/etc/opt/omni/client

- Directory containing the client configuration directories and files

/etc/opt/omni/IS

- Directory, containing the Installation Server configuration directories and files.

/etc/opt/omni/server

- Directory containing the following configuration directories:

barlists

    database backup specifications

barschedules

    database backup specification schedules

cell

    the cell configuration

datalists

    backup specifications

devices

    templates for devices

options

    default options

schedules

    backup schedules

sessions

    data about sessions

sg

    scripts for Service Guard support

snmp

    the OpenView/SNMP trap sending configuration

users

    the user configuration

/etc/opt/omni/server/dr

- Directory containing the following disaster recovery directories:

  `asr`

    ASR archive file

  `p1s`

    P1S files for Enhanced Automated Disaster Recovery

  `srd`

    SRD files

`/opt/omni`

- Data Protector home directory. It contains the following Data Protector executable directories:

  `bin`

    Data Protector user commands

  `lbin`

    Disk Agent and Media Agent files and some administrative commands

  `sbin`

    Cell Manager and Data Protector Internal Database (IDB) administrative commands

`/opt/omni/doc`

- The Data Protector guides, including the *HPE Data Protector Command Line Interface Reference*, the Data Protector support matrices

`/opt/omni/help`

- The Data Protector Help

`/opt/omni/lib`

- Directory containing the following directories:

  `/opt/omni/lib/man`

    Data Protector man pages

  `/opt/omni/lib/nls`

    message catalogs

`/var/opt/omni`

- Directory containing the following directories:

  `/var/opt/omni/log` and `/var/opt/omni/server/log`

    log files

  `/var/opt/omni/server/export/keys` and `/var/opt/omni/server/import/keys`

    encryption keys

  `/var/opt/omni/server/log/auditing`

    audit logs

  `/var/opt/omni/server/sessions`

    data about sessions

`/var/opt/omni/tmp`

> temporary files

`/var/opt/omni/server/db80`

- Directory containing the following Data Protector Internal Database (IDB) directories:

`/var/opt/omni/server/db80/idb`

> the IDB tablespaces

`/var/opt/omni/server/db80/dcbf`

> the Detail Catalog Binary Files (DCBF) part of the IDB

`/var/opt/omni/server/db80/keystore`

> the encryption keystore database

`/var/opt/omni/server/db80/keystore/catalog`

> the key ID catalog

`/var/opt/omni/server/db80/logfiles`

> the IDB archived log files and the IDB recovery file (`obdrindex.dat`)

`/var/opt/omni/server/db80/msg`

> the Data Protector session messages

`/var/opt/omni/server/db80/smisdb`

> the ZDB database (SMISDB)

`/var/opt/omni/server/db80/smisdb/p4000/login`

> the data which the Data Protector HPE P4000 VSS Agent uses to connect to the configured CIMOM providers for the HPE 3PAR StoreServ Storage disk arrays

`/var/opt/omni/server/db80/smisdb/p10000/login`

> the data which the Data Protector HPE 3PAR VSS Agent and the Data Protector HPE P6000 / HPE 3PAR SMI-S Agent use to connect to the configured CIMOM providers for the HPE 3PAR StoreServ Storage disk arrays

`/var/opt/omni/server/db80/smisdb/netapp/login`

> the data which the Data Protector NetApp Storage Provider uses to connect to the NetApp Storage system

`/var/opt/omni/server/db80/smisdb/emcvmax/login`

> the data which the Data Protector EMC VMAX Storage Provider uses to connect to the EMC VMAX Storage system

`/var/opt/omni/server/db80/smisdb/emcvnx/login`

> the data which the Data Protector EMC VNX Storage Provider uses to connect to the EMC VNX Storage system

`/var/opt/omni/server/db80/xpdb`

> the ZDB database (XPDB)

SEE ALSO

cjutil(1M), ob2install(1M), omniabort(1), omniamo(1), omnib(1), omnib2dinfo(1), omnicc(1), omnicellinfo
(1), omnicheck(1M), omnicjutil(1M), omniclus(1), omnidb(1), omnidbcheck(1M), omnidbinit(1M),
omnidbp4000(1), omnidbrestore(1M), omnidbsmis(1), omnidbutil(1M), omnidbvss(1), omnidbxp(1),
omnidbzdb(1), omnidlc(1M), omnidownload(1), omnidr(1M), omnigui(5), omnihealthcheck(1M),
omniinetpasswd(1M), omniiso(1), omniintconfig.pl(1M), omnikeytool(1M), omnimcopy(1), omniminit(1),
 omnimigrate.pl(1M), omnimlist(1), omnimm(1), omnimnt(1), omnimver(1), omniobjconsolidate(1),
omniobjcopy(1), omniobjverify(1), omniofflr(1M), omnir(1), omniresolve(1M), omnirpt(1), omnirsh(1M),
omnisetup.sh(1M), omnisrdupdate(1M), omnistat(1), omnistoreapputil(1M), omnisv(1M), omnitrig(1M),
omniupload(1), omniusb(1), omniusers(1), sanconf(1M), SharePoint_VSS_backup.ps1(1), syb_tool(1),
uma(1M), upgrade_cm_from_evaa(1M), util_cmd(1M), util_oracle8.pl(1M), vepa_util.exe (1M)

Section 1: User commands

omniabort(1)

omniabort — aborts an active session
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omniabort -version | -help

omniabort -session *SessionID*

DESCRIPTION

This command aborts an active session, identifying it by the *SessionID*. A list of all active sessions and their session IDs is available using the `omnistat` command.

OPTIONS

`-version`

Displays the version of the `omniabort` command.

`-help`

Displays the usage synopsis for the `omniabort` command.

`-session` *SessionID*

Specifies the *SessionID* of the session to be aborted. Use the `omnistat` command to get the *SessionID* of the session.

NOTES

When using this command to abort the check for unrequired incrementals, manually terminate the `omniabort` utility afterwards.

EXAMPLES

1. To abort a session with the SessionID "R-2013/05/13-12", execute:

   omniabort -session R-2013/05/13-12
   omniabort -sess 12

SEE ALSO

omnistat(1)

omniamo(1)

omniamo — starts an automated media operation session
(this command is available on the Data Protector Cell Manager)

SYNOPSIS

omniamo -version|-help

omniamo -amc *ConfigurationName* {-post_backup|-scheduled}

DESCRIPTION

This command starts an automated media operation session for the specified post-backup or scheduled configuration. Before starting a post-backup operation, you must export the session ID of the backup session that used the media you want to copy.

Windows systems: set SESSIONID= *SessionID*

UNIX systems: export SESSIONID= *SessionID*

Use this command if you want to immediately start an automated media operation. Also, if an automated media operation has failed, you can use this command to start the operation again.

OPTIONS

-version

   Displays the version of the omniamo command.

-help

   Displays the usage synopsis for the omniamo command.

-amc *ConfigurationName* {-post_backup|-scheduled}

   Starts the post-backup or scheduled automated media copy operation with the specified name.

EXAMPLES

1. To start the scheduled automated media copy operation with the configuration name "MediaCopy1", execute:

   omniamo -amc MediaCopy1 -scheduled

2. To start the post-backup automated media copy operation with the configuration name "MyFiles" and session ID 2011/09/13-0001 on Windows, execute:

   set SESSIONID=2011/09/13-0001

   omniamo -amc MyFiles -post_backup

3. To start the post-backup automated media copy operation with the configuration name "MyDocs" and session ID 2011/09/13-0002 on UNIX, if you are using an sh-like shell, execute:

```
SESSIONID=2011/09/13-0002
```

```
export SESSIONID
```

```
omniamo -amc MyDocs -post_backup
```

4. To start the post-backup automated media copy operation with the configuration name "MyBackup" and session ID 2011/09/13-0003 on UNIX, if you are using a csh-like shell, execute:

```
export SESSIONID=2011/09/13-0003
```

```
omniamo -amc MyBackup -post_backup
```

SEE ALSO

omnib2dinfo(1M), omnidownload(1), omnimcopy(1), omniminit(1), omnimlist(1), omnimm(1), omnimnt(1), omnimver(1), omniupload(1), sanconf(1M), uma(1M)

omnib(1)

omnib — backs up filesystems, disk images, the Data Protector Internal Database (IDB), Microsoft Exchange Server single mailboxes and Public Folders, Microsoft Exchange Server 2007, Microsoft Exchange Server 2010/2013, Microsoft SQL Server, Microsoft SharePoint Server 2007/2010/2013, SAP R/3, SAP MaxDB, Oracle, MySQL, PostgreSQL, Informix Server, VMware vSphere, Microsoft Hyper-V, Sybase, Lotus, IBM DB2 UDB, and NDMP objects
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnib -version | -help

omnib -filesystem *Client:MountPoint Label* -device *BackupDevice* [*MIRROR_OPTIONS...*] [*GENERAL_OPTIONS*] [*FILESYSTEM_OPTIONS*] [-public]

omnib -filesystem *Client:MountPoint Label* -device *BackupDevice* -ndmp *NDMP_Server_Type* [*NDMP_OPTIONS*] [-public]

omnib -winfs *Client:MountPoint Label* -device *BackupDevice* [*MIRROR_OPTIONS...*] [*GENERAL_OPTIONS*] [*FILESYSTEM_OPTIONS*] [*WINFS_OPTIONS*] [-public]

omnib -host *Client:/ Label* -device *BackupDevice* [*MIRROR_OPTIONS...*] [*GENERAL_OPTIONS*] [*FILESYSTEM_OPTIONS*] [-public [-storedrim]]

omnib -rawdisk *Client Label SectionList* -device *BackupDevice* [*MIRROR_OPTIONS...*] [*GENERAL_OPTIONS*] [-public]

omnib -restart *SessionID*

omnib -datalist *Name* [*BACKUP_SPECIFICATION_OPTIONS*]

omnib -resume *SessionID* [-no_monitor]

omnib -sap_list *ListName* [-barmode *SapMode*] [*LIST_OPTIONS*]

omnib -sapdb_list *ListName* [-barmode *SapdbMode*] [*LIST_OPTIONS*]

omnib -oracle8_list *ListName* [-barmode *Oracle8Mode*] [*LIST_OPTIONS*]

omnib -sybase_list *ListName* [-barmode *SybaseMode*] [*LIST_OPTIONS*]

omnib -informix_list *ListName* [-barmode *InformixMode*] [*LIST_OPTIONS*]

omnib -mssql_list *ListName* [-barmode *MSSQLMode*] [*LIST_OPTIONS*]

omnib -msese_list *ListName* [-barmode *MSExchangeMode*] [*LIST_OPTIONS*]

omnib -e2010_list *ListName* [-barmode *E2010Mode*] [*LIST_OPTIONS*]

omnib -lotus_list *ListName* [-barmode *LotusMode*] [*LIST_OPTIONS*]

omnib -msvssw_list *ListName* [-barmode *VSSMode*] [*LIST_OPTIONS*]

omnib -mbx_list *ListName* [-barmode *MSMailboxMode*] [*LIST_OPTIONS*]

omnib -db2_list *ListName* [-barmode *DB2Mode*] [*LIST_OPTIONS*]

omnib -mssps_list *ListName* [-barmode *MSSPSMode*] [*LIST_OPTIONS*]

omnib -mssharepoint_list *ListName* [-barmode *MSSharePointMode*] [*LIST_OPTIONS*]

omnib -idb_list *ListName* [-barmode *IDBMode*] [*LIST_OPTIONS*]

omnib -veagent_list *ListName* [-barmode *VirtualEnvironmentMode*] [*LIST_OPTIONS*]

omnib -integ MySQL *ListName* [-barmode *MySQLMode*]

omnib -integ PostgreSQL *ListName* [-barmode *PostgreSQL*]

MIRROR_OPTIONS

 -mirror *BackupDevice* [-pool *MediaPool* -prealloc *MediaList*]

GENERAL_OPTIONS

 -preview

 -pool *MediaPool*

 -prealloc *MediaList*

 -protect {none | weeks *n* | days *n* | until *Date* | permanent}

 -report {warning | minor | major | critical}

 -pre_exec *Pathname*

 -post_exec *Pathname*

 -compress

 -encode [aes256]

 -load {low | medium | high}

 -crc

 -no_monitor

 -keepcatalog {weeks *n* | days *n* | until *Date*}

 -variable *VariableName VariableValue*

 -priority *NumValue*

FILESYSTEM_OPTIONS

 -trees *TreeList*

 -only *MatchPattern*

 -exclude *TreeList*

 -skip *MatchPattern*

 -lock

 -touch

 -[no_]log | -log_dirs | -log_file

 -mode {Full | Incremental[1-9]}

 -enh_incr

   -clp

   -[no_]hlink

   -size *FromRange ToRange*

WINFS_OPTIONS

  -no_share[_info]

  -[no_]nthlinks

  -[no_]archatt

  -[no_]vss [fallback]

  -async

BACKUP_SPECIFICATION_OPTIONS

  -select *SelectList*

  -mode {Full | Incremental[1-9]}

  -protect {none | weeks *n* | days *n* | until *Date* | permanent}

  -preview

  -disk_only

  -load {low | medium | high}

  -crc

  -no_monitor

LIST_OPTIONS

  -barcmnd *Command*

  -protect {none | weeks *n* | days *n* | until *Date* | permanent}

  -load {low | medium | high}

  -crc

  -no_monitor

  -test_bar

  -disk_only

NDMP_OPTIONS

  -ndmp_user *UserName*

  -ndmp_passwd *Password*

  -ndmp_env *FileName*

  -ndmp_bkptype {dump | nvb | SMTape}

  -[no_]log -log_dirs -log_file

  -mode {full | incremental1}

  -pool *MediaPool*

   -prealloc *MediaList*

   -protect {none | weeks *n* | days *n* | until *Date* | permanent}

   -report {warning | minor | major | critical}

   -variable *VariableName VariableValue*

OTHER OPTIONS

   NDMP_Server_Type= Generic | NetApp | Celerra | BlueArc | Hitachi | HPX9000

   SapMode= full | incr

   SapdbMode= full | diff | trans

   Oracle8Mode= full | incr1 | ... | incr4

   SybaseMode= full | trans

   InformixMode= full | inf_incr1 | inf_incr2

   MSSQLMode= full | copy | diff | trans

   MSSPSMode= full | diff | trans

   MSExchangeMode= full | incr

   E2010Mode= full | copy | incr | diff

   LotusMode= full | incremental

   VSSMode= full | copy | incr | diff

   MSMailboxMode= full | incr | incr1

   DB2Mode= full | incr | delta

   MSSharePointMode= full | diff | incr

   IDBMode= full | incr

   VirtualEnvironmentMode= full | diff | incr

   MySQLMode= full | incr | trans

   PostgreSQLMode= full | incr

   Date= [*YY*]*YY*/*MM*/*DD* (1969 < *YYYY* < 2038)

DESCRIPTION

The omnib command uses a backup specification (list of file or database objects) to back up data objects. The following Data Protector functionality is supported:

Session management

Controls the backup sessions. The Session Manager reads the backup specification or uses the command options to determine what to back up and how many copies of the backup objects to create (object mirroring), then initiates the Disk and Media Agents for disks and backup devices which will be used in the session. Once the session has completed, the Session Manager updates the MMDB with the session information.

Media management

Provides easy and efficient management of large sets of media by grouping media, tracking their status, implementing a media rotation policy, supporting the barcode recognition, vaulting the media, automating the library device operations, storing the media related information in a central place and sharing this information among several Data Protector cells.

Data compression

Writes data to media in a compressed format.

Data encryption

Writes data to media in an encrypted format using the Advanced Encryption Standard (AES) algorithm.

Backup monitoring

When the backup command is executed, it sends a request (specifying the backup objects) to the Session Manager. When the Session Manager (SM) accepts the request, it assigns a unique SessionID to the session. You can use this SessionID to monitor the progress of the session using the Monitor context of the Data Protector GUI or the `omnistat` command. You can also use the `omniabort` command to terminate a session.

> **Note:** During the Internal Database Backup (PostgreSQL) in `Incr` mode, the configuration files are backed up as full.

OPTIONS

`-version`

Displays the version of the `omnib` command

`-help`

Displays the usage synopsis for the `omnib` command

`-filesystem` *Client:MountPoint Label*

Specifies the client, mount point and label of the filesystem to be backed up.

`-winfs` *Client:MountPoint Label*

Specifies the client, mount point and label of the Windows filesystem to be backed up.

`-host` *Client:/ Label*

Specifies the client to be backed up as a set of filesystems defined at backup time. The label is used as a prefix for each of these filesystem labels. Client backup is useful for systems with filesystem configuration that often changes.

`-rawdisk` *Client Label SectionList*

Specifies the client, sections (pathnames of disk image sections) and label of the node to be backed up.

`-datalist` *Name*

Specifies the name of the backup specification file for the backup. The backup specification contains the data objects (filesystems and disk image sections) to be backed up.

`-restart` *SessionID*

Tries to restart a failed session, specified by its sessionID.

-resume *SessionID*

Resumes a failed or aborted backup session. This option is applicable to a filesystem backup and Oracle Server integration backup. While resume of a filesystem backup creates an incremental backup of the failed session, the Oracle Server integration resumes the backup by creating a new session using the same backup specification as the failed session. In both cases, only the data that has not been backed up in the failed session is backed up.

-sap_list *ListName*

Specifies the name of the SAP R/3 backup specification file for the backup. The SAP R/3 backup specification contains the SAP R/3 objects to be backed up.

-barmode *SapMode*

For SAP R/3 objects, the possible modes are `full` and `incr`. The default value for this option is `full`.

-sapdb_list *ListName*

Specifies the name of the SAP MaxDB backup specification file for the backup. The SAP MaxDB backup specification contains the SAP MaxDB objects to be backed up.

-barmode *SapdbMode*

For SAP MaxDB objects, the possible modes are `full`, `diff`, and `trans`. The `full` option triggers a full backup of the SAP MaxDB instance, the `diff` option triggers a differential backup, and the `trans` option triggers an archive logs backup. The default value for this option is `full`.

-oracle8_list *ListName*

Specifies the name of the Oracle backup specification file for the backup. The Oracle backup specification contains the Oracle objects to be backed up.

-barmode *Oracle8Mode*

For Oracle objects you can specify `full` for full backup or `incr1` through `incr4` for incremental backups.

-sybase_list *ListName*

Specifies the name of the Sybase backup specification file for the backup. The Sybase backup specification contains the Sybase objects to be backed up.

-barmode *SybaseMode*

For Sybase objects you can specify `full` for full database backup or `trans` for transaction backup. The default value for this option is `full`.

-informix_list *ListName*

Specifies the name of the Informix Server backup specification file for the backup. The Informix Server backup specification contains the Informix Server objects to be backed up.

-barmode *InformixMode*

For Informix Server objects you can specify the following modes:

`full`: full backup of dbspaces specified during the backup specification creation time,

`inf_incr1`: first incremental backup,

`inf_incr2`: second incremental backup.

The default value for this option is `full`.

`-mssql_list` *ListName*

Specifies the name of the Microsoft SQL Server backup specification file for the backup. The Microsoft SQL Server backup specification contains the Microsoft SQL Server objects to be backed up.

`-barmode` *MSSQLMode*

For Microsoft SQL Server objects you can specify `full` for a full database backup, `copy` for a copy-only full backup, `diff` for a differential database backup or `trans` for a transaction log backup. The default value for this option is `full`.

In Microsoft SQL Server log shipping configurations, transaction log backup cannot be performed. A differential database backup is started when a transaction log backup is requested.

In Microsoft SQL Server availability group configurations, when you trigger a full or a differential backup of a database belonging to an availability group secondary replica, the backup type is automatically changed to a copy-only full backup.

`-integ PostgreSQL` *ListName*

Specifies the name of the PostgreSQL backup specification file for the backup. The PostgreSQL backup specification contains a list with the PostgreSQL objects to be backed up.

`-barmode` *PostgreSQL*

For PostgreSQL objects, you can specify full for a full backup or incr for an incremental backup. Note, that an incremental backup cannot be run without a previously successful full backup.

If this option is not specified, Data Protector performs a full backup.

`-msese_list` *ListName*

Specifies the name of the Microsoft Exchange Server 2007 backup specification file for the backup. The Microsoft Exchange Server 2007 backup specification contains the Microsoft Exchange Server 2007 objects to be backed up.

`-barmode` *MSExchangeMode*

For Microsoft Exchange Server 2007 objects you can specify `full` for full database and log files backup or `incr` for incremental backup of log files. The default value for this option is `full`.

`-e2010_list` *ListName*

Specifies the name of the Microsoft Exchange Server backup specification file for the backup. The Microsoft Exchange Server backup specification contains the Microsoft Exchange Server 2010/2013 objects to be backed up.

`-barmode` *E2010Mode*

For Microsoft Exchange Server 2010/2013 objects, you can specify `full` for a full backup, `copy` for a copy backup, `incr` for an incremental backup, or `diff` for a differential backup.

Note that an incremental backup session cannot be followed by a differential backup session, nor the other way around. You must first run a full backup session.

If this option is not specified, a full backup is performed.

-lotus_list *ListName*

Specifies the name of the Lotus Notes/Domino Server backup specification file for the backup. The Lotus Notes/Domino Server backup specification contains the Lotus database objects to be backed up.

-barmode *LotusMode*

For Lotus Notes/Domino Server objects you can specify `full` for full database backup or `incr` for a full backup of selected Lotus Notes/Domino objects, if the amount of data changed from the last backup is bigger than the value specified for the backup specification option Amount of log changes (KB) in the Data Protector GUI. In case that transaction logging is enabled, the full backup of all archived transaction logs is also performed. The default value for this option is `full`.

-msvssw_list *ListName*

Specifies the name of the Microsoft VSS backup specification file for the backup. The Microsoft VSS backup specification contains the Microsoft VSS objects to be backed up.

-barmode *VSSMode*

Available backup types primarily depend on the VSS writer that is chosen to be backed up. While some VSS writers support several backup types (for example `full`, `copy`, `incr`, `diff` with Microsoft Exchange Server 2003 writer), others support only `full`. For more information, see the *HPE Data Protector Integration Guide for Microsoft Volume Shadow Copy Service*. Even when supported with the selected VSS writer by Data Protector, not all types might be available at all times.

Data Protector aborts the backup session if an unsupported or unavailable backup type is specified.

-mbx_list *ListName*

Specifies the name of the Microsoft Exchange Server single mailbox backup specification file for the backup. The Microsoft Exchange Server single mailbox backup specification contains single mailboxes to be backed up.

-barmode *MSMailboxMode*

For Microsoft Exchange Server single mailboxes, you can specify `full` for a full mailbox backup, `incr` for an incremental mailbox backup, or `incr1` for an incremental1 mailbox backup. The default value for this option is `-full`.

-db2_list *ListName*

Specifies the name of the IBM DB2 UDB backup specification file for the backup. The IBM DB2 UDB backup specification contains the IBM DB2 UDB objects to be backed up.

-barmode *DB2Mode*

For IBM DB2 UDB objects you can specify `full` for full database backup, `incr` for incremental database backup, or `delta` for delta database backup. The default value for this option is `full`.

-mssps_list *ListName*

Specifies the name of the Microsoft SharePoint Portal Server backup specification file for the backup. The Microsoft SharePoint Portal Server backup specification contains the Microsoft SharePoint Portal Server objects to be backed up.

-barmode *MSSPSMode*

For Microsoft SharePoint Portal Server objects you can specify the following modes:

`full`: full backup,

`diff`: differential database backup of Microsoft SQL Server databases and full backup of other Microsoft SharePoint Portal Server objects,

`trans`: transaction log backup of Microsoft SQL Server databases and full backup of other Microsoft SharePoint Portal Server objects.

The default value for this option is `full`.

`-mssharepoint_list` *ListName*

Specifies the name of the Microsoft SharePoint Server 2007/2010/2013 backup specification file for the backup. The Microsoft SharePoint Server 2007/2010/2013 backup specification contains the Microsoft SharePoint Server 2007/2010/2013 objects to be backed up.

`-barmode` *MSSharePointMode*

For Microsoft SharePoint Server 2007/2010/2013 objects you can specify the following modes:

`full`: full backup,

`diff`: a Microsoft SQL Server differential backup of the database, and backup of the index files that have been changed since the last full backup,

`incr`: a backup of transaction logs (.log) that have been created since the last transaction log backup of the Microsoft SQL Server database, and backup of the index files that have been changed or created since the last backup of any type.

If this option is not specified, a full backup is performed.

`-idb_list` *IDBList*

Specifies the name of the Internal Database backup specification file for the backup. The Internal Database backup specification contains a list with the Data Protector Internal Database and its related objects to be backed up.

`-barmode` *IDBMode*

For Internal Database objects, you can specify `full` for a full backup or `incr` for an incremental backup. Note that an incremental backup cannot be run without a previously successful full backup.

If this option is not specified, a full backup is performed.

`-veagent_list` *ListName*

Specifies the name of the virtual environment backup specification file for the backup. The backup specification contains the virtual environment objects to be backed up.

`-barmode` *VirtualEnvironmentMode*

For VMware vSphere objects, the available modes are `full`, `diff`, and `incr`. The `full` option triggers a full backup, the `diff` option triggers a differential backup, and the `incr` option triggers an incremental backup.

For Microsoft Hyper-V objects, the available modes are `full` and `incr`. The `full` option triggers a full backup and the `incr` option triggers an incremental backup. Under specific circumstances, the incremental backup session falls back and Data Protector performs a full backup instead. For more information, see the *HPE Data Protector Integration Guide*.

If this option is not specified, Data Protector attempts to start a full backup.

`-integ MySQL` *ListName*

Specifies the name of the MySQL backup specification file for the backup. The MySQL backup specification contains a list with the MySQL objects to be backed up.

`-barmode` *MySQLMode*

For MySQL objects, you can specify `full` for a full backup, `incr` for an incremental backup, or `trans` for transaction log backup. Note, that an incremental backup cannot be run without a previously successful full backup.

If this option is not specified, Data Protector performs a full backup.

`-device` *BackupDevice*

Specifies the backup device to be used for the backup.

`-iap`

Use to perform backups to the HPE Integrated Archive Platform (IAP).

`-public`

If you use this option, you allow other users to see and restore your data. By default for filesystem backups, only the Data Protector administrator and the user who created a backup can see and restore the data.

`-storedrim`

If this option is specified, a disaster recovery OS image is created and saved to the Cell Manager's disk at the end of the backup session.

The image is stored in P1S files directory with the filename `ClientName.img`.

Note that you can obtain the image from a disk much faster than from a backup medium.

MIRROR_OPTIONS

`-mirror` *BackupDevice*

Specifies one or several backup devices to be used for object mirroring. Different backup devices should be specified for the backup and for each mirror.

`-pool` *MediaPool*

Instructs the Session Manager to use an alternate media pool for object mirroring. By default, the default media pool for the backup device is used.

`-prealloc` *MediaList*

Specifies a list of media to be used for object mirroring. If the Media Allocation policy for the pool is set to "strict", the media in the Prealloc list are used in the sequence shown in the list. If one of these media is unavailable, a mount prompt is issued. NOTE: If the Media Allocation Policy is "strict", you must specify a Prealloc list.

GENERAL_OPTIONS

`-preview`

Checks the backup objects, backup devices and options you selected, without performing the

backup. The check includes: backup objects, status of the backup device, available media, and the approximate amount of data which will be backed up.

-pool *MediaPool*

Instructs the Session Manager to use an alternate media pool for the backup. By default, the default media pool for the backup device is used.

-prealloc *MediaList*

Specifies a list of media to be used for the backup. If the Media Allocation policy for the pool is set to "strict", the media in the Prealloc list are used in the sequence shown in the list. If one of these media is unavailable, a mount prompt is issued. NOTE: If the Media Allocation Policy is "strict", you must specify a Prealloc list.

-protect {none|weeks *n*|days *n*|until *Date*|permanent}

Sets the level of protection for the backup session. The media containing this backup session cannot be overwritten until the protection expires. By default, the protection is permanent.

-report {warning |minor|major|critical}

Sets the level of error notification for the session. Errors are classified (in ascending order) as: warning, minor, major and critical. When you select a level, errors of this level and higher are displayed in the Monitor window. For example, if major is selected, only major and critical errors are reported. By default, all errors are reported.

-pre_exec *Pathname*

Instructs the Session Manager to execute this command before starting the backup session. The complete *Pathname* of the command should be specified. The command is executed on the Session Manager system.

-post_exec *Pathname*

Instructs the Session Manager to execute this command after the backup session. The complete *Pathname* of the command should be specified. The command is executed on the Session Manager system.

-compress

Instructs the General Media Agent to write data to media in the compressed format.

-encode [aes256]

Instructs the General Disk Agent to write data to media in encoded format.

If the aes256 option is specified, data is written to media in encrypted format, using the Advanced Encryption Standard (AES) algorithm.

-load {low|medium|high}

Specifies the level of network traffic generated by a session during a time period. High level generates as much traffic as allowed by the network, resulting in a faster backup. Low level has less impact on the network performance, but results in a slower backup. By default, this option is set to high.

-crc

Instructs the General Media Agent to write a CRC checksum at the end of every block on the medium. If this option is used, you can later verify the CRC checksum on the medium by using the

`omniver` command.

-no_monitor

By default, the command monitors the session and displays the status of the session during the session. If this option is used, the SessionKey is displayed and the command is disconnected from the session.

-keepcatalog {weeks *n* | days *n* | until *Date*}

This option specifies file catalog retention time. If you do not want to save the file catalog at all, use the -no_log option. By default, this option is set to the same value as specified by the `protection` option.

-variable *VariableName VariableValue*

This option lets you specify a variable name and its value for proper operation of some platforms and integrations. Setting user definable variables (a variable name and its value) enables flexible operation on some platforms and integrations with Data Protector. The list of variables and their values that are configurable with Data Protector is dynamic and comes with Data Protector patches.

-priority *NumValue*

In case multiple running sessions request access to a specific device at the same time, this option determines the order in which the sessions will be queued. The *NumValue* can be any value from 1 (the highest priority) to 6000 (the lowest priority). In case the option is not specified, the default value of 3000 is set.

If a low priority session is running when a high priority session starts queuing, the currently running session is allowed to finish. When more sessions request access to a device with the same priority, any of these sessions might acquire access first.

FILESYSTEM_OPTIONS

-trees *TreeList*

Specifies the trees to be included in the backup. If this option is not used, the filesystem is backed up from the mount point level downwards. When specifying several trees, separate each *Tree* with a space. *Tree* must start with a /. Note that when specifying trees on UNIX systems, the complete tree must be specified including the mountpoint, whereas on Windows systems, trees must be specified without volumes (drives). For example: -tree /usr/temp (UNIX system) or -tree \temp (Windows system). This option is not supported with Data Protector NDMP server integration.

-only *MatchPattern*

Specifies that only files that match the *MatchPattern* will be backed up. This option is not supported with Data Protector NDMP server integration.

-exclude *TreeList*

Specifies trees not to be backed up. This option is not supported with Data Protector NDMP server integration.

-skip *MatchPattern*

Specifies that files matching the *MatchPattern* will not be backed up. This option is not supported

with the Data Protector NDMP server integration.

-lock

Instructs the Disk Agent to lock each file before backing it up. If the file is in use (and cannot be locked), the session manager displays a warning that this file cannot be locked and backs up the file anyway. This warning is also logged to the catalog database. By default, files are not locked at backup.

-no_log

Disables the logging of backed up files to the catalog database. By default, the filename and backup history of each backed up file is written to the catalog database.

-log

The default option. All detailed information about backed up files and directories (filenames, file versions, and attributes) are logged to the Data Protector Internal Database (IDB). This allows you to browse directories and files before restore and in addition look at the file attributes. Data Protector can fast position on the tape when restoring a specific file.

-log_dirs

If this option is specified, only the directories are logged into the database. By default, the filename and backup history of each backed up file is written to the catalog database.

-log_file

All detailed information about backed up files and directories (filenames and file versions) is logged to the Data Protector Internal Database (IDB). This information allows you to search for backed up files and allows Data Protector to fast position the tape.

It also does not take much space since some information on file details (file attributes) is not logged to the database.

-mode {Full|Incremental[1-9]}

Specifies the type for the backup session. Full type backs up all specified files. Incremental[1-9] backs up only a subset of the specified files, based on whether or not the files were modified since the last Full or lower-level Incremental backup. Default is the Full type. The level of incremental backup is based on the level number which is specified.

For example, an incremental level 3 backs up only those files (of the specified files) which were modified since the last incremental level 2 or lower backup.

-touch

Whenever a file is opened, read, or locked, which happens during backup, the file's access time attribute changes. By default, after backup, Data Protector resets the file's access time attribute to the value it had before backup. However, on UNIX, this resetting of the access time attribute modifies the file's change time.

If the -touch option is specified, Data Protector does not reset access time attributes. Then, on UNIX, Data Protector can also use the file's change time (inode modification time) as an incremental backup criterion. As a result, files with a changed name, location, or attributes are backed up in an incremental backup.

-no_hlink

If this option is specified, then hard link detection is disabled and hard links are backed up as normal

files. This speeds up the first traversal of the filesystem.

-enh_incr

This option enables enhanced incremental backup. Unlike conventional incremental backup, enhanced incremental backup reliably detects and backs up files with changes in name, location, and attributes. It is also a prerequisite for subsequent object consolidation (synthetic backup).

NOTE: After you select this option, incremental backup will run in the enhanced mode only after a full backup is performed.

-clp

This option enables using the Windows NTFS Change Log Provider with enhanced incremental backups and conventional incremental backups. A list of files to be backed up will be generated by querying the Windows Change Journal rather than performing a file tree walk.

-size *FromRange ToRange*

Limits backup to those files only, of which sizes are in the specified range. The sizes are set in kB. If you set *ToRange* to 0, all files larger then *FromRange* will be backed up.

WINFS_OPTIONS

-no_share[_info]

If this option is specified, share information for directories on Windows systems is *not* backed up. By default, if a directory was shared on the network when a backup was run, the share information for directory is backed up, unless the -no_share[_info] option is specified.

Backing up share information for shared directories enables you to automatically share such directories after restore.

-[no_]nthlinks

If this option is specified then NTFS hard link detection is disabled and NTFS hard links are backed up as normal files. This speeds up the first traversal of the filesystem.

-[no_]archatt

By default, Data Protector uses the archive attribute as an incremental backup criterion and also clears the file's archive attribute after the file is backed up. The archive attribute is automatically set by the system when the file's content, properties, name, or location changes.

If archive attributes cannot be cleared, an error is reported. This affects future incremental backups, so that the files are backed up, although they have not changed. This may happen when backing up removable media with write protection. In the case of ZDB, archive attributes are cleared on the replica and this is not reflected on the source volume. As a result, in the next incremental ZDB session, when a new replica is created, the archive attributes appear again and the corresponding files are backed up although they may not have changed. To enhance the incremental ZDB behavior, specify the -[no_]archatt option.

If the -[no_]archatt option is specified, Data Protector ignores archive attributes and detects changed files using other criteria, such as the file's modification time.

-[no_]vss [fallback]

If the -vss option is specified, the VSS filesystem backup is performed. If the shadow copy creation on the system where the VSS filesystem backup is running, fails, the backup also fails by

default. However, you can avoid backup failure by specifying the `fallback` option. In this case, the backup will continue as the normal filesystem backup.

NOTE: On Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012, VSS file system backup is used even if the `-vss` is not specified. To ensure that VSS is not used, specify `-no_vss`.

`-async`

If this option is specified, Disk Agent performs asynchronous reading from the disk without using Windows cache manager. Concurrent reads of the same file are started simultaneously. If this option is not specified, synchronous reading from the disk is performed.

BACKUP_SPECIFICATION_OPTIONS

`-select` *SelectList*

Specifies which objects (of those in the backup specification) to back up. The *SelectList* is the list of objects to be backed up.

`-mode {Full | Incremental[1-9]}`

Specifies the type for the backup session. `Full` type backs up all specified files. `Incremental[1-9]` backs up only a subset of the specified files, based on whether or not the files were modified since the last full or lower-level `Incremental` backup. Default is the `Full` type. The level of incremental backup is based on the level number which is specified. For example, an incremental level 3 backs up only those files (of the specified files) which were modified since the last incremental level 2 (or lower) backup.

Use incremental level 1 to back up files that were changed since last full backup only. The `Incremental` without level will back up the files that changed since the last backup only (regardless whether it was full or incremental of any level).

`-protect {none | weeks` *n* `| days` *n* `| until` *Date* `| permanent}`

See *GENERAL_OPTIONS*.

`-preview`

Checks the backup objects, backup devices and options you selected, without performing the backup. The check includes: objects due for backup, status of the backup device, available media, and approximate amount of data which will be backed up.

`-disk_only`

A ZDB related option. It instructs Data Protector to perform a ZDB-to-disk session rather than a ZDB-to-tape or ZDB-to-disk+tape session. With ZDB, if the option is not specified, a ZDB-to-tape or ZDB-to-disk+tape session is performed.

`-load {low | medium | high}`

See *GENERAL_OPTIONS*.

`-crc`

Instructs the General Media Agent to write a CRC checksum at the end of every block on the medium. If this option is used, you can later verify the CRC checksum on the medium by using the `omnimver` command.

`-no_monitor`

By default, the command monitors the session and displays the status of the session during the session. If this option is used only the SessionKey is displayed and the command is disconnected from the session.

LIST OPTIONS

`-barcmnd` *Command*

Specifies the command that will be used instead of the command specified with `exec` option in the backup specification. The command should reside in the default Data Protector administrative commands directory.

`-protect` {none|weeks *n*|days *n*|until *Date*|permanent}

See *GENERAL_OPTIONS*.

`-load` {low|medium|high}

See *GENERAL_OPTIONS*.

`-crc`

Instructs the General Media Agent to write a CRC checksum at the end of every block on the medium. If this option is used, you can later verify the CRC checksum on the medium by using the `omnimver` command.

`-no_monitor`

By default, the command monitors the session and displays the status of the session during the session. If this option is used, only the SessionKey is displayed, and the command is disconnected from the session.

`-test_bar`

Enables backup preview mode. Backup preview is only available for backup sessions for Oracle Server, SAP R/3, SAP MaxDB, Microsoft Exchange Server single mailbox, Lotus Notes/Domino Server, IBM DB2 UDB, Informix Server, and Sybase integrations. Zero downtime backup preview is not supported.

This option checks the backup objects, backup devices, and options you selected, without actually performing the backup. The check includes: objects due for backup, status of the backup device, available media, and the approximate amount of data which will be backed up.

`-disk_only`

This ZDB-related option is supported only for specific application integrations, and not for the Internal Database backup. It instructs Data Protector to perform a ZDB-to-disk session rather than a ZDB-to-tape or ZDB-to-disk+tape session.

With ZDB backup specifications, if the option is not specified, a ZDB-to-tape or ZDB-to-disk+tape session is performed.

NDMP_OPTIONS

`-ndmp_user` *UserName*

Sets the username that is used by Data Protector to establish the connection to the NDMP server.

`-ndmp_passwd` *Password*

Sets the password for the username that is used by Data Protector to establish the connection to

the NDMP server.

-ndmp_env *FileName*

Specifies the filename of file with NDMP environment variables for specific NDMP implementations.

-ndmp_bkptype {Dump|NVB|SMTape}

Specifies the backup type for NDMP EMC Celerra backups. Dump is the default backup type, that backs up data at a file level. NDMP volume backup (NVB) is an EMC-specific NDMP backup type. NVB backs up data blocks at a volume level. SMTape backup is an NetApp-specific NDMP backup type. SMTape backs up data blocks at a volume level.

-no_log

Disables the logging of backed up files to the catalog database. By default, the filename and backup history of each backed up file is written to the catalog database.

-log

The default option. All detailed information about backed up files and directories (filenames, file versions, and attributes) are logged to the Data Protector Internal Database (IDB). This allows you to browse directories and files before restore and in addition look at the file attributes. Data Protector can fast position on the tape when restoring a specific file.

-log_dirs

If this option is specified, only the directories are logged into the database. By default, the filename and backup history of each backed up file is written to the catalog database.

-log_file

All detailed information about backed up files and directories (filenames and file versions) is logged to the Data Protector Internal Database (IDB). This information allows you to search for backed up files and allows Data Protector to fast position the tape. It also does not take much space since some information on file details (file attributes) is not logged to the database.

-mode {Full|Incremental[1-9]}

Specifies the mode for the backup session. Full mode backs up all specified files. Incremental[1-9] backs up only a subset of the specified files, based on whether or not the files were modified since the last Full or lower-level Incremental backup. Default is the Full mode. The level of incremental backup is based on the level number which is specified. For example, an incremental level 3 backs up only those files (of the specified files) which were modified since the last incremental level 2 or lower backup.

-pool *MediaPool*

Instructs the Session Manager to use an alternate media pool for the backup. By default, the default media pool for the backup device is used.

-prealloc *MediaList*

Specifies a list of media to be used for the backup. If the Media Allocation policy for the pool is set to "strict", the media in the Prealloc list are used in the sequence shown in the list. If one of these media is unavailable, a mount prompt is issued. NOTE: If the Media Allocation Policy is "strict", you must specify a Prealloc list.

-protect {none|weeks *n*|days *n*|until *Date*|permanent}

Sets the level of error notification for the session. Errors are classified (in ascending order) as: `warning`, `minor`, `major` and `critical`. When you select a level, errors of this level and higher are displayed in the Monitor window. For example, if `major` is selected, only `major` and `critical` errors are reported. By default, all errors are reported.

`-report {warning |minor|major|critical}`

See *GENERAL_OPTIONS*.

`-variable VariableName VariableValue`

This option lets you specify a variable name and its value for proper operation of some platforms and integrations. Setting user definable variables (a variable name and its value) enables flexible operation on some platforms and integrations with Data Protector. The list of variables and their values that are configurable with Data Protector is dynamic and comes with Data Protector patches.

RETURN VALUES

See the man page `omniintro` for return values.

Additional return values of the `omnib` command are:

`10` - There was an error while backing up some files. All agents completed successfully.

`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 user.

EXAMPLES

The following examples illustrate how the `omnib` command works:

1. To do a backup of a tree "/usr" of filesystem "senna" with the label "work", using the compress option, to the backup device "DAT", execute:

   `omnib -device DAT -filesystem senna:/ work -tree /usr -compress`

2. To perform an incremental backup using the backup specification OMNIGROUP, and to make the devices available to this session with the highest priority in case of resource conficts, execute:

   `omnib -datalist OMNIGROUP -mode Incremental -priority 1`

3. To preview a backup of the tree "/Amt3" of the filesystem "Munich", skipping the files with the ".fin" extension, execute:

   `omnib -preview -filesystem Munich:/ -tree /Amt3 -skip "*.fin"`

4. To execute a disk image backup of the section "/dev/rdsk/c201d1s0" on the client "xanadu" to the backup device "Exa" and protecting the session against overwrite for 4 weeks:

   `omnib -rawdisk xanadu section /dev/rdsk/c201d1s0 -dev Exa -protect weeks 4`

5. To execute a full Lotus backup using the "test2" backup specification with the high network load and permanent protection set:

   `omnib -lotus_list test2 -barmode full -protect permanent -load high`

6. To start a full backup using an IBM DB2 UDB backup specification named "TEST", and to set data protection to 10 weeks, execute:

```
omnib -db2_list TEST -barmode full -protect weeks 10
```

7. To start a differential backup using an SAP MaxDB backup specification named "test", and write a CRC checksum at the end of every block on the medium, execute:

```
omnib -sapdb_list test -barmode diff -crc
```

8. To start a differential backup using a Microsoft Exchange Server backup specification named "bSpec1", execute:

```
omnib -e2010_list bSpec1 -barmode diff
```

9. To perform an encrypted backup of a tree "/usr" of filesystem "alpha.hp.com" with the label "work", using the `encode aes256` option, to the backup device "ENC1", execute:

```
omnib -filesystem alpha.hp.com:/work -device ENC1 -tree /usr -encode aes256 -
mode full
```

10. To back up a volume "/vol/vol1" of the Celerra NDMP Server "alpha.hp.com" using the NVB backup type option, to the backup device "DAT", execute:

```
omnib -filesystem alpha.hp.com:/vol/vol1 /vol/vol1 -device DAT -ndmp Celerra -
ndmp_bkptype nvb
```

11. To start a full backup using a Microsoft SharePoint Server 2013 backup specification named "myBackup", execute:

```
omnib -mssharepoint_list myBackup -barmode full
```

12. To perform a full backup of the Data Protector Internal Database (IDB) using the backup specification named "idb_weekly" and omit session monitoring in the command output, execute:

```
omnib -idb_list idb_weekly -no_monitor
```

13. To perform an incremental backup of the Data Protector Internal Database (IDB) using the backup specification named "idb_daily", with as little impact on the network traffic during the session as possible, execute:

```
omnib -idb_list idb_daily -barmode incr -load low
```

14. To start an incremental backup of Microsoft Hyper-V virtual machines using the backup specification named "hyperv_host_4" and disable session monitoring, execute:

```
omnib -veagent_list hyperv_sys_4 -barmode incr -no_monitor
```

15. To start a transaction log backup of your MySQL instance using the backup specification named "mysql_instance_core_sys", execute:

```
omnib -integ MySQL mysql_instance_core_sys -barmode trans
```

SEE ALSO

omnikeytool(1M), omniobjconsolidate(1), omniobjcopy(1), omniobjverify(1), omnir(1)

omnicc(1)

omnicc — handles the Data Protector licensing, reports the number of configured and available Data Protector licenses, installs the licenses, imports and exports Data Protector clients, manages access to secured clients, enables encrypted control communication, and creates a template for the user_ restrictions file
(this command is available on systems with any Data Protector component installed)

SYNOPSIS

omnicc -version|-help

omnicc -redistribute

omnicc -import_host *ClientName* [-virtual] [-encr_disable] [-encr_recreate_cert]

omnicc -import_host *ClientName* [-virtual]

omnicc -import_iap *ClientName* -port *Port* -user *UserName* -passwd *Password* -cert_mode {UseKnown|Download} -cert_name *CertificateName*

omnicc -import_ndmp *ClientName* -type *NdmpType* -port *Port* -user *UserName* -passwd *Password*

omnicc -import_is *ClientName*

omnicc -export_is *ClientName*

omnicc -update_host *ClientName*

omnicc -update_all [-force_cs]

omnicc -export_host *ClientName*

omnicc -list_authorities *ClientName*

omnicc -secure_client *ClientName* -authorities *ClientName1* [*ClientName2*...]

omnicc -unsecure_client *ClientName*

omnicc -install_license *password*

omnicc -password_info

omnicc -add_certificate *CertificateName PathOfCertificateFile*

omnicc -get_certificate *CertificateName*

omnicc -list_certificates

omnicc -confirm_mom_clients

omnicc -update_mom_server

omnicc -check_licenses [-detail]

omnicc -cbl_detail

omnicc [-query]

omnicc -create_userrestrictions_tmpl

omnicc -gre_license_info

omnicc -impersonation -add_user -user {*User@Domain* | *Domain\User*} {-host *ClientName* [-host *ClientName*...] | -all} {-passwd *Password* | -passwdfile *PasswordFile*}

omnicc -impersonation -modify_user -user {*User@Domain* | *Domain\User*} {-host *ClientName* [-host *ClientName*...] | -all} {-passwd *Password* | -passwdfile *FileName*} {-old_passwd *OldPassword* | -old_passwdfile *OldFileName*}

omnicc -impersonation -delete_user -user {*User@Domain* | *Domain\User*} {-host *ClientName* [-host *ClientName*...] | -all} {-passwd *Password* | -passwdfile *FileName*}

omnicc -encryption -enable {*ClientName1* [*ClientName2* ...] | -all} [-recreate_cert | [-cert *Cert* [-key *Key*]] [-trust *TrustedCert*]]

omnicc -encryption -enable_mom {*CSHostname1* [*CSHostName2* ...] | -all} [-recreate_cert]

omnicc -encryption-disable {*ClientName1* [*ClientName2* ...] | -all}

omnicc -encryption-disable_mom {*CSHostname1* [*CSHostName2* ...] | -all}

omnicc -encryption-status {*ClientName1* [*ClientName2* ...] | -all}

omnicc -encryption -update_trust {*Hostname1* [*HostName2* ...] | -all} -trust TrustedCerts [-replace]

omnicc -encryption -encr_param {*Hostname1* [*HostName2* ...] | -all} [-tls_min TLSvMin] [-tls_max TLSvMax]

omnicc -encryption -list_exceptions

omnicc -encryption -add_exception *ClientName1* [*ClientName2* ...]

omnicc -encryption -remove_exception *ClientName1* [*ClientName2* ...]

omnicc -encryption -status {*ClientName1* [*ClientName2* ...] | -all}

omnicc -import_esx *ClientName* -port *Port* -user *UserName* -passwd *Password* -web_root *WebRoot* -integrated_sec {0 | 1}

omnicc -import_vcenter *ClientName* -port *Port* -user *UserName* -passwd *Password* -web_root *WebRoot* -integrated_sec {0 | 1} -register_greplugin {0 | 1}

omnicc -import_hyperv *ClientName* -user *UserName* -passwd *Password*

omnicc -import_vcd *ClientName* -user *UserName* -passwd *Password*

omnicc -migrate_devfilter [HostName] [-delete_old_devfilter]

NdmpType

 Generic | NetApp | Celerra | BlueArc | Hitachi | HPX9000

DESCRIPTION

The omnicc command is used for licensing, importing and exporting clients, managing secured clients, enabling encrypted control communication, and creating a template for the user_restrictions file.

OPTIONS

`-version`

Displays the version of the `omnicc` command.

`-help`

Displays the usage synopsis for the `omnicc` command.

`-redistribute`

Displays licensing information for multicell environments. The first part shows the number of allocated licenses and the second shows the number of licenses actually used per server.

`-import_host`*ClientName* `[-virtual][-encr_disable][-encr_recreate_cert]`

Imports the specified server into the cell. If the client has multiple names, import each additional name with the `-virtual` option

If `-encr_disable` is specified, encrypted communication is not enabled on the client.

If `-encr_recreate_cert` is specified, certificates are recreated and the existing ones are overwritten.

`-import_host` *ClientName* `[-virtual]`

Imports the specified client into a cell. This allows you to move a client between two cells without reinstalling the Data Protector modules.

When you import the next one among multiple network names (clusters, service guards), use the `-virtual` option. This way you keep Data Protector from assigning licenses to all the network names of the same system.

`-import_ndmp` *ClientName*

Imports the specified NDMP server into the cell.

`-type` *NdmpType*

Sets the NDMP data format when importing an NDMP server into a cell.

`-port` *Port*

Sets the TCP/IP port number of the NDMP server when importing an NDMP server into a cell.

`-user` *UserName*

Sets the username that is used by Data Protector to establish the connection to the NDMP server when importing an NDMP server into a cell.

`-passwd` *Password*

Sets the password for the username that is used by Data Protector to establish the connection to the NDMP server when importing an NDMP server into a cell.

`-import_iap` *ClientName*

Imports the specified IAP Server into the cell.

`-cert_mode` {UseKnown | Download}

Sets the certificate mode when importing the IAP Server. If the `UseKnown` option is used, the certificate already uploaded to the Cell Manager will be used for connecting to the IAP appliance. If the `Download` option is used, the certificate for connecting to the IAP appliance will be downloaded from the IAP Server at the time of the first login. This certificate will be stored as *IAP_server_name*`.cert` and used for future connections.

-cert_name *CertificateName*

Specifies the name of the certificate that will be used for connecting to the IAP appliance.

-import_is *ClientName*

Imports an already installed Installation Server into the cell.

-export_is *ClientName*

Exports an already installed Installation Server from the cell.

-update_host *ClientName*

Updates the version information and installed components information in the Cell Manager configuration file for the specified client. You can use this option in circumstances when new remote installation packages for particular components exist on the client, but the component upgrade has failed.

-update_all [-force_cs]

Updates the version information and installed components information in the Cell Manager configuration file for all clients in the cell. You can use this option in circumstances when new remote installation packages for particular components exist on some clients, but the component upgrade processes have failed.

If the -force_cs option is specified, it checks if any clients have been improperly added to the current cell. If such clients exist, the command properly imports them into the cell before updating the information on the Cell Manager.

-export_host *ClientName*

Exports the specified client from the cell. This enables you to remove a client from the cell without uninstalling its Data Protector modules.

If the host is a vCenter system with the Advanced GRE Web Plug-in installed, then this command unregisters the Advanced GRE Web Plug-in from the vCenter and then exports the vCenter client from the cell server.

-list_authorities *ClientName*

Lists systems from which the specified client accepts requests on the Data Protector port (by default 5555).

-secure_client *ClientName*

Specifies the client to be secured.

-authorities *ClientName* [*ClientName2*...]

Specifies systems from which the specified client accepts requests on the Data Protector port (by default 5555). Consequently, other computers will not be able to access this client. For tasks like backup and restore, starting pre- or post-execution scripts, or importing and exporting clients, the client checks whether the computer which triggers one of these tasks via the Data Protector port is

allowed to do so. This security mechanism instructs the client to accept such actions only from the systems specified by this option.

-unsecure_client *ClientName*

Specifies the client from which you want to remove security. Such a client will enable access to all systems in the cell.

-install_license *password*

Installs an encrypted Data Protector license. The password must be formatted as a single line and must not contain any embedded carriage returns. The password must be in quotes. If the password includes also a description in quotes, the quotes in this description must be preceded with backslashes.

-password_info

Displays information about installed license passwords.

-add_certificate *CertificateName PathOfCertificateFile*

Adds a certificate to the Cell Manager.

-get_certificate *CertificateName*

Downloads the certificate from the Cell Manager and displays its content.

-list_certificates

Lists certificates uploaded to the Cell Manager.

-confirm_mom_clients

Collects the `cell_info` files from MoM clients (*Data_Protector_program_data*\Config\Server\cell\mom_info on Windows clients or /etc/opt/omni/server/cell/mom_info on UNIX clients) and stores them on the MoM Manager into the directory *Data_Protector_program_data*\Config\Server\mom\cell_info (Windows systems) or /etc/opt/omni/server/mom/cell_info (UNIX systems) under client Cell Manager name. Use this command when switching MoM clients to CMMDB mode. The `omnicc` command with this option specified has to be executed on the MoM Manager.

-update_mom_server

Pushes the `mom_info` file located in the directory *Data_Protector_program_data*\Config\Server\cell (Windows systems) or /etc/opt/omni/server/cell (UNIX systems) to MoM and CMMDB server to MoM into the directory *Data_Protector_program_data*\Config\Server\mom\cell_info (Windows systems) or /etc/opt/omni/server/mom/cell_info (UNIX systems) under client Cell Manager name. Use this command when switching to CMMDB mode. The `omnicc` command with this option specified has to be executed on the client Cell Manager.

-check_licenses [-detail]

Reports licensing related information from the cell.

If the -detail option is specified, a detailed report is produced.

If the -detail option is not specified, the command returns information on whether the Data Protector licensing is covered or not. The following information is returned: the time when the report was generated, the licensing mode, the license server and the total TB of data under protection.

Traditional licensing model: The license checker returns the following information for every license in the cell: license name, licenses installed, licenses in use, total TB of data under protection, and additional licenses required.

Note that in a traditional licensing model for drive extension licenses-to-use, the license checker also returns information about configured drives and recommended additional licenses. You need as many licenses as there are drives in use at any point in time. This is typically the total number of configured drives to allow all drives to be used simultaneously.

Capacity based licensing model: The license checker returns the following information: the license name and the capacity of installed licenses.

In a MoM environment with the CMMDB configured, when producing a license report for the items that are subject to libraries and devices related licenses, such as media (including advanced file device media), backup devices, drives and slots, the `omnicc` command must be executed on the Cell Manager with the CMMDB installed.

In a MoM environment, only the data specific to this Cell Manager is reported, not for all the cells in the MoM environment.

`-cbl_detail`

Displays a report which contains the backup type, object ID, backup name and object size of objects that are used to calculate the total protected data capacity.

`-query`

Displays information about the number of available licenses.

`-create_userrestrictions_tmpl`

Creates the `user_restrictions_tmpl` file which is a template for the `user_restrictions` file, populated by names of all systems of the Data Protector cell and names of all configured user groups other than *admin* and *operator*.

To put the template into use, change its contents as desired, and rename it to `user_restrictions`.

`-gre_license_info`

Reports Granular Recovery licensing related information from the cell. The following information is returned: the database server name, the application type, the time when the license was used for restore, the time when the license will be released for the next restore from another database server, and the number of days remaining until the license release.

`-impersonation -add_user -user` {*User@Domain*|*Domain\User*}{`-host` *ClientName* [`-host` *ClientName*...]|`-all`}{`-passwd` *Password*|`-passwdfile` *FileName*}

Sets up a user account for the Data Protector Inet service user impersonation on one or more specified clients, by specifying the user name and the password directly or by saving the user name and the password into the specified file.

To enable user impersonation on all clients in the cell, specify the `-all` option.

`-impersonation -modify_user -user` {*User@Domain*|*Domain\User*}{`-host` *ClientName* [`-host` *ClientName*...]|`-all`}{`-passwd` *Password*|`-passwdfile` *FileName*}{`-old_passwd` *OldPassword*|`-old_passwdfile` *OldFileName*}

Modifies a user account for the Data Protector Inet service user impersonation on one or more specified clients, by specifying the user name and the new password directly or by saving the user

name and the new password into the specified file and by specifying the user's old password directly or in the specified file.

To modify user impersonation on all clients in the cell, specify the `-all` option.

`-impersonation -delete_user -user {`*User@Domain*`|`*Domain*`\`*User*`} {-host` *ClientName* `[-host` *ClientName*`...] | -all} {-passwd` *Password* `| -passwdfile` *FileName*`}`

Deletes a user account for the Data Protector Inet service user impersonation on one or more specified clients.

To remove user impersonation from all clients in the cell, specify the `-all` option.

`-encryption -enable {`*ClientName1* `[`*ClientName2* `...] | -all} [-recreate_cert | [-cert` *Cert* `[-key` *Key*`]] [-trust` *TrustedCerts*`]]`

Enables encrypted control communication on one or more specified clients. If specified clients were listed in the Cell Manager's exception list, they are removed from it.

To enable encrypted communication on all clients in the cell, specify the `-all` option.

`-encryption -enable_mom {`*CSHostname1* `[`*CSHostName2* `...] | -all} [-recreate_cert]`

Enables encrypted control communication in the MoM environment in the specified Cell Managers.

To enable encrypted communication on all Cell Managers in the MoM environment, specify the `-all` option.

`-encryption -disable {`*ClientName1* `[`*ClientName2* `...] | -all}`

Disables encrypted control communication on specified clients and adds them to the Cell Manager's exception list.

To disable encrypted communication on all clients in the cell, specify the `-all` option.

`-encryption -disable_mom {`*CSHostname1* `[`*CSHostName2* `...] | -all}`

Disables encrypted control communication in the MoM environment in the specified Cell Managers.

To disable encrypted communication on all Cell Managers in the MoM environment, specify the `-all` option.

`-recreate_cert`

Recreates the certificates and overwrites existing certificates.

If no names are provided for the options cert, key, and trust, then certificates are generated for specified clients. If you require to use `hpdpcert.pem`, then you can do so by specifying the options cert, key and trust certificate as `hpdpcert.pem`.

> **Note**: The earlier versions of Data Protector did not create certificates automatically. The administrator had to create the certificates and point Data Protector to the certificate files. The `-cert`, `-key` and `-trust` options should be used with this earlier method of managing certificates. Note that it is not possible to specify the certificates using the GUI. However, it is possible from the CLI as in earlier versions.

`-cert` *Cert*

Remotely installs certificate on a selected client.

The default certificates file `hpdpcert.pem` is created during the installation or upgrade on the Cell Manager in the default server configuration directory under`certificates`.

-key *Key*

Remotely installs private key on a selected client.

The default key file `hpdpcert.pem` is created during the installation or upgrade on the Cell Manager in the default server configuration directory under `certificates`.

-trust*TrustedCerts*

Remotely installs trusted certificate that is used for peer certificate verification from the Cell Manager on a selected client.

The default trusted certificates file `hpdpcert.pem` is created during the installation or upgrade on the Cell Manager in the default server configuration directory under `certificates`.

-update_trust

This option updates the trusted certificates on specified clients. The specified trusted (`-trust`) certificate file is used to update the trusted certificate list. If the option `-replace` is also used, then the list is replaced with the provided certificates. If the option `-replace` is not used, then the new certificates are added to the existing certificates on specified clients by concatenating the two trusted certificate files.

-encr_param {*Hostname1* [*HostName2* ...] | -all} [-tls_min TLSvMin] [-tls_max TLSvMax]

This option specifies the minimum and/or maximum versions of TLS for host.

The values for `TLSvMin` and `TLSvMax` can be specified as:

TLSv1, TLSv1.1, TLSv1.2, (or) 1, 1.1, 1.2

-add_exception *ClientName1* [*ClientName2* ...]

Adds exception to the exception list on the Cell Manager.

-remove_exception *ClientName1* [*ClientName2* ...]

Removes exception from the exception list on the Cell Manager.

-list_exceptions

Lists exceptions from the exception list on the Cell Manager.

-status {*ClientName1* [*ClientName2* ...] | -all}

Checks whether encrypted control communication is enabled or disabled on specified clients. This option is useful for verification and troubleshooting.

If the `-all` option is specified, the command verifies the status of all clients in the cell.

-import_esx *ClientName*

This is a VMware specific option.

Specifies the VMware ESX(i) client to import.

-import_vcenter *ClientName*

This is a VMware specific option.

Specifies the VMware vCenter client to import.

`-import_hyperv` *ClientName*

> This is a Hyper-V specific option.

> Specifies the Hyper-V client to import.

`-import_cs` *HostName*

`-export_cs`*HostName*

> Imports or exports the remote Cell Manager.

`-import_vcd` *ClientName*

> This is a VMware specific option.

> Specifies the VMware vCloud Director client to import.

`-migrate_devfilter` [HostName] [-delete_old_devfilter]

> This option is used to migrate existing OMNIRC based device filter tags from Data Protector clients to a centralized Cell Manager file in the Cell Manager.

> # UNIX: `/etc/opt/omni/server/cell/hosttags`

> # Windows: `<Data_Protector_home>\Config\server\cell\hosttags`

> If the `HostName` option is specified, the device filter tag from the hostname is printed to the console in the following format:

> `<HostName> <tag>`

> If the `-delete_old_devfilter` option is specified, the OMNIRC variable `OB2DEVICEFILTER` and its value are removed from the host(s).

> If the `hosttags` file is already present and you run this option, the `hosttags_tmp` file is created in the same location as that of the `hosttags` file. You need to manually merge the `hosttags_tmp` file with the `hosttags` file.

`-port` *Port*

> This is a VMware specific option.

> Specifies the port to connect to (for example, 443).

`-user` *UserName*

> Specifies an operating system user account for the connection.

`-passwd` *Password*

> Specifies the user's password.

`-web_root` *WebRoot*

> This is a VMware specific option.

> Specifies the web service entry point URI (for example, /sdk).

`-integrated_sec` {0|1}

> This is a VMware specific option.

> Specifies the security mode.

If the 0 option is specified, you have to specify all login credentials manually (standard security).

If the 1 option is specified, Data Protector connects to the VMware vCenter Server system with the user account under which the Data Protector Inet service on the backup host is running (integrated security). Ensure this user account has appropriate rights to connect to the VMware vCenter Server system.

`-register_greplugin {0|1}`

This is a VMware specific option.

It registers the Advanced GRE Web Plug-in into the vCenter client.

If the 0 option is specified, Data Protector does not register the plug-in.

If the 1 option is specified, Data Protector registers the plug-in into the vCenter.

NOTES

If you change your licensing model from the traditional to the capacity based, the information about previously distributed traditional licenses will be overwritten.

EXAMPLES

The following examples illustrate how the omnicc command works.

1. To check if the licensing is covered within a Data Protector cell, execute:

   `omnicc -check_licenses`

2. To get information about configured drives and recommended additional drive extension licenses-to-use when you are using the traditional licensing model, execute:

   `omnicc -check_licenses -detail`

3. To get information about licenses capacity installed when you are using the capacity based licensing model, execute:

   `omnicc -check_licenses -detail`

4. To get information about the amount of data under protection in TB, execute:

   `omnicc -check_licenses -detail`

5. To get information about the objects that are used to calculate the amount of data under protection in TB, execute:

   `omnicc -cbl_detail`

6. To get information about used GRE licenses, execute:

   `omnicc -gre_license_info`

7. To enable encrypted control communication and remotely install default certificate "hpdpcert.pem", default private key "hpdpcert.pem", and default trusted certificate "hpdpcert.pem" from the Cell Manager on clients named "computer1.company.com" and "computer2.company.com", execute:

   `omnicc -encryption -enable computer1.company.com computer2.company.com -cert hpdpcert.pem -key hpdpcert.pem -trust hpdpcert.pem`

8. To enable encrypted control communication for the whole cell from the Cell Manager, execute:

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

9. To enable encrypted control communication and to recreate all certificates for the whole cell from the Cell Manager, execute:

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

10. To enable TLS version 1.2 with newly recreated certificates for the whole cell from the Cell Manager, execute the following command. When a connection is made from client to server or client to client, the highest common version is used. For systems updated to Data Protector 9.07 or higher, TLS version 1.2 is used.

    ```
    omnicc -encryption -encr_param -all -tls_min 1.0 -tls_max 1.2
    ```

11. To enable encrypted control communication and recreate certificates from the Cell Manager on clients named "computer1.company.com" and "computer2.company.com", execute:

    ```
    omnicc -encryption -enable computer1.company.com computer2.company.com -
    recreate
    ```

12. If the encryption was enabled with "hpdcert.pem", which was default prior to 9.03, the "hpdpcert.pem" has to be specified explicitly in the command line. To enable encrypted control communication from the Cell Manager on clients named "computer1.company.com" and "computer2.company.com", execute:

    ```
    omnicc -encryption -enable computer1.company.com computer2.company.com -cert
    hpdpcert.pem -key hpdpcert.pem -trust hpdpcert.pem
    ```

13. To add a client named "computer.company.com" to the exception list on the Cell Manager and allow plain (non-encrypted) communication, execute:

    ```
    omnicc -encryption -add_exception computer.company.com
    ```

14. To configure a Microsoft SharePoint 2007/2010 farm administrator which will be used for backup or restore on a medium farm (two web front ends, one application and one sql server), execute:

    ```
    omnicc -impersonation -add_user web1.domain.com web2.domain.com
    indexapp.domain.com sql.domain.com -user MyDomain\MyUser -passwd MyPassword
    ```

15. To import an "HPE X9000" NDMP server into a cell, execute:

    ```
    omnicc -import_ndmp lxdprnd5.ind.hp.com -type "HP X9000" -port 10000 -user root
    -passwd MyPassword
    ```

16. To check the value of the `DailyMaintenanceTime` option in the output of a debug text file named "CHECK", execute:

    ```
    omnicc -debug 20 CHECK.txt
    ```

SEE ALSO

omnicellinfo(1), omnicheck(1M), omnidlc(1M), omnisv(1M)

omnicellinfo(1)

omnicellinfo — displays configuration information about the Data Protector cell
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

```
omnicellinfo -version|-help
```

```
omnicellinfo -servers
```

```
omnicellinfo -group
```

```
omnicellinfo -object [schedule|no_schedule] [-group Group]
```

```
omnicellinfo -db
```

```
omnicellinfo -prepostinfocheck
```

```
omnicellinfo {-mm|-dev} [-detail]
```

```
omnicellinfo { -dlinfo [-group Group]} | -cell [brief] {-schinfo [Backup_Specification |-
days NumberDays |-group Group]} | {-dlobj [-group Group]} | {-trees [-group Group]} | -allbdf |-
acl
```

DESCRIPTION

The `omnicellinfo` command displays information about data objects, media pools, devices, clients, database, backup specifications and backup specification groups in the cell. It can be also used to display the cell managers in multicell environments.

Some options recognized by `omnicellinfo` are intended primarily for generating reports by shell/awk/perl scripts. Information produced is formatted in records with a newline as field separator and a blank line as record separator. Those options are: `-dlinfo`, `-schinfo`, `-dlobj`, `-trees` and `-allbdf`.

OPTIONS

`-version`

   Displays the version of the `omnicellinfo` command.

`-help`

   Displays the usage synopsis for the `omnicellinfo` command.

`-servers`

   Displays the list of cell managers that are included in the multicell environment.

`-group`

   Displays the backup specification groups that contain backup specifications. Note that the backup

specification group named `Default` is not displayed.

`-object[schedule |no_schedule]`

Displays information about objects (filesystems, databases and disk images) in the cell. The report shows: Object (object type, client name, and mountpoint), Label, and Next Scheduled Backup Date. When you use the `schedule` option, the report only shows those objects which are scheduled for backup. When you use the `-no_schedule` option, the report only shows those objects which are not scheduled for backup. By default, all objects (scheduled and unscheduled) are listed.

`-mm`

Displays information about the media and media pools in the cell. The report shows for each pool: the Pool Name, Media Class, Media Usage Policy, Media Allocation Policy, and Amount of Free Space in the pool.

`-dev`

Displays information about the backup devices in the cell. The report shows for each device: the Device Name, Client Name, Device Type and Media Pool.

`-db`

Displays information about the Data Protector Internal Database (IDB). The database is divided in logical structures, for each of these structures the report shows: Disk Space Used, Records Used and Records Total.

`-prepostinfocheck`

Searches all worklists configured in the Data Protector cell and checks for security compliance of the commands executed during the session. The rules are defined in the "Pre and Post-Exec Commands for a Backup Specification" in *HPE Data Protector Help* and enforced by Inet. You can execute this option after the patch installation to quickly figure out the Group or Name or command that is not proper.

> **Note:** PrePostInfoCheck validation includes file system path checking for defined pre and post scripts in data/bar lists. This is platform specific and it is recommended to execute checks on both Windows and Linux clients to check for errors that are reported for that platform.

`-cell`

Displays information about the configured clients in the cell. The report shows for each client: client name, operating system, cell console version, Disk Agent version, Media Agent version, GUI version, and all installed Data Protector integration versions. There is also a short summary which shows the total number of clients and, if the `brief` option was not specified, all possible Data Protector software components, together with the total number of every software component in the cell. If the `brief` option was specified, only the installed Data Protector software components together with the total number of every software component in the cell is listed.

The VADP feature introduced in Data Protector 8.14 provides enhanced reports for VMs. 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. If not, only the VM name is displayed.

`-detail`

The `-detail` option can be used in combination with the `-dev` and `-mm` options to produce a more detailed report.

-dlinfo

Shows information about backup specifications. For each backup specifications it lists the name of the backup specification, session owner, pre-exec and post-exec script. Session owner is in format *USER.GROUP@CLIENT*.

-schinfo [*Backup_Specification* | -days *NumberDays*]

Shows information about backup specification scheduling. If *Backup_Specification* and `-days` option are not specified, the command displays the next schedule time for each backup specification. If backup specification is specified the command lists all schedules in the next year for the specified backup specification. Option `-days` can be used to display schedules of all backup specifications for a specified number of days.

-dlobj

Shows information about all objects in backup specifications. For each object it lists object type, object name (in format *ClientName:PathName*), description, and the name of the backup specification. After this, the device and poolname fields are listed for each device used in the backup specification making the size of the records variable.

-trees

Shows information about all defined trees in backup specifications. For each tree, it lists filesystem name (in format *ClientName:Pathname*), tree, description, backup device, media pool and name of the backup specification.

-acl

Displays all Data Protector access permissions that the user running the command has.

-group *Group*

This option allows you to limit the output of the command to single backup specification group. The following options support this: `-dlinfo`, `-schinfo`, `-dlobj`, `-trees` and `-object`.

EXAMPLES

The following examples illustrate how the `omnicellinfo` command works.

1. To list detailed information about the selected objects, execute:

   `omnicellinfo -object schedule`

2. To list detailed information about the configured devices, execute:

   `omnicellinfo -dev -detail`

3. To display all virtual machines configured on ESXi servers or vCenters and imported to the Data Protector Cell Manager, execute:

   `omnicellinfo –cell brief`

SEE ALSO

omnicc(1), omnicheck(1M), omnidlc(1M), omnisv(1M)

omniclus(1)

omniclus — manages load balancing in a cluster environment in the event of an application (Data Protector or other) failover
(this command is available on systems with the Data Protector `MS Cluster Support` component installed (Windows systems) and on the Data Protector Cell Manager (UNIX systems))

SYNOPSIS

omniclus -version|-help

omniclus -clus *cluster_name* -session {*|*backup_specification*} -abortsess [-abortid {==| !=} *application_id*]

omniclus -clus *cluster_name* -inhibit {*|0|*minutes*}

omniclus -clus *cluster_name* -session {*|*backup_specification*} -symlink {split|active}

NOTE: On UNIX systems, replace the wildcard (*) with the string `'*'`.

NOTE: On Windows systems, the `-noclus` option can be specified directly after `-clus` to prevent loading of the cluster dynamic library.

DESCRIPTION

The `omniclus` command, which is common to all platforms (Windows and UNIX systems), allows you to communicate the Data Protector Cell Manager special events that in certain way control its behavior and behavior of the backup sessions in a cluster environment. `omniclus` allows load balancing by offering additional (CLI) control of the Cell Manager in cluster environments:

- Aborting sessions
- Temporarily disabling the Cell Manager for backup sessions
- Specifying the state of the EMC Symmetrix links after an application failover

Note that the system specified as the *cluster_name* argument of the `-clus` option must be a cluster-aware Data Protector Cell Manager.

OPTIONS

-version

   Displays the version of the `omniclus` command

-help

   Displays the usage synopsis for the `omniclus` command.

-clus *cluster_name*

   Specifies the cluster-aware Cell Manager.

`-session {*|`*`backup_specification`*`}`

Specifies the session(s) to which the abort message should be sent.

`-abortsess`

Specifies the abort session command.

`-abortid {==|!=}` *`application_id`*

Specifies the application identification.

`-inhibit {*|0|`*`minutes`*`}`

Specifies the number of minutes for Cell Manager backup inactivity, where `*` means forever and `0` means activate now.

`-symlink {active|split}`

Specifies the state of the EMC/Symmetrix links upon application failover if a backup is running.

NOTES

The `omniclus` command can only be used in cluster environments.

EXAMPLES

The following examples illustrate how the `omniclus` command works.

1. To abort all running sessions, execute:

   `omniclus -clus cluster.domain.com -session * -abortsess`

   NOTE: On UNIX systems, replace the wildcard (*) with the string `'*'`.

   The utility will connect to all running sessions and will send them abort messages. The state of the sessions can be then checked with the Data Protector `omnistat` utility.

2. To abort specific running sessions, execute:

   `omniclus -clus cluster.domain.com -session mybackup -abortsess`

   The utility will connect to backup session managers issuing abort messages and sending them additional information - the backup specification name. Each backup session manager checks whether the command addresses it and if this is the case it aborts.

3. To abort sessions (all or specific) with application identifications, execute:

   `omniclus -clus obvs.domain.com -session * -abortsess -abortid != 10`

   NOTE: On UNIX systems, replace the wildcard (*) with the string `'*'`.

   This way the user can define groups of sessions and abort only the ones that are actually related to the application that failed over. For example a backup session that performs a normal filesystem backup of a remote client is not aborted because an application server switches, while the application server backup can be aborted.

4. Temporarily disabling the Data Protector cell

   The following command will inhibit backup sessions for twenty minutes:

   `omniclus -clus cluster.domain.com -inhibit 20`

The following command will inhibit backup sessions forever:

```
omniclus -clus cluster.domain.com -inhibit *
```

NOTE: On UNIX systems, replace the wildcard (*) with the string '*'.

The following command will re-activate backup sessions immediately:

```
omniclus -clus cluster.domain.com -inhibit 0
```

5. EMC/Symmetrix links

The following syntax will connect to specific (running) backup session managers and inform them to left the EMC/Symmetrix links split:

```
omniclus -clus cluster.domain.com -session * -symlink split
```

NOTE: On UNIX systems, replace the wildcard (*) with the string '*'.

The following syntax will connect to specific (running) backup session managers and inform them to left the EMC/Symmetrix links active (established):

```
omniclus -clus cluster.domain.com -session * -symlink active
```

NOTE: On UNIX systems, replace the wildcard (*) with the string '*'.

SEE ALSO

omnirsh(1M)

omnicreatedl(1)

omnicreatedl — creates a filesystem backup specification file (datalist); or an HPE P9000 XP Disk Array Family or HPE P6000 EVA Disk Array Family Microsoft Exchange Server 2003 ZDB backup specification file (datalist)
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnicreatedl -version | -help

FILESYSTEM BACKUP

omnicreatedl [-datalist *Name*] [-host *HostName1* [*HostName2*...]] [-device *BackupDevice*]

MICROSOFT EXCHANGE SERVER 2003 ZERO DOWNTIME BACKUP

omnicreatedl -ex2000 -datalist *Name* [-device *Name*] {*P9000_DISK_ARRAY_XP_OPTIONS* | *P6000_ENTERPRISE_VIRTUAL_ARRAY_OPTIONS*} *EXCHANGE_OPTIONS* [-force] [-virtualSrv *Name*]

P9000_DISK_ARRAY_XP_OPTIONS

1. ZDB-to-disk and ZDB-to-disk+tape sessions (HPE Business Copy P9000 XP configurations):

 -split_mirror -sse -local *app_sys bck_sys* [-mirrors *MU_numbers*] -instant_restore [-leave_enabled_bs] [-split | -establish]

2. ZDB-to-tape sessions (HPE Business Copy P9000 XP configurations):

 -split_mirror -sse -local *app_sys bck_sys* [-mirrors *MU_numbers*] [-keep_version [-leave_enabled_bs]] [-split | -establish]

3. ZDB-to-tape sessions (HPE Continuous Access P9000 XP or combined (HPE CA+BC P9000 XP) configurations):

 -split_mirror -sse {-remote *app_sys bck_sys* | -combined *app_sys bck_sys*} [-keep_version [-leave_enabled_bs]] [-split | -establish]

P6000_ENTERPRISE_VIRTUAL_ARRAY_OPTIONS

1. ZDB-to-disk sessions:

 -snapshot -smis *app_sys bck_sys* -instant_recovery [-snapshots *number*]

2. ZDB-to-disk+tape sessions:

 -snapshot -smis *app_sys bck_sys* -instant_recovery [-snapshots *number*] [-wait_clonecopy *number*]

3. ZDB-to-tape sessions:

 -snapshot -smis *app_sys bck_sys* -snapshot_type {standard | vsnap | clone [-wait_clonecopy *number*]} -snapshot_policy {strict | loose} -replica_conf {local | combined [-ca_failover_option {follow_replica_direction | maintain_replica_location}]}

EXCHANGE_OPTIONS

 `-annotation` {MIS|SRS|KMS}

{`-all_storage_groups`|`-storage_group` *Storage_Group_Name1* [`-store` *Store1* [*Store2*...]] [`-storage_group` *Storage_Group_Name2* [`-store` *Store1* [*Store2*...]]...]}

DESCRIPTION

FILESYSTEM BACKUP

The `omnicreatedl` command creates a filesystem backup specification file (datalist). It searches all specified clients for local mount points and puts them in the backup specification or on the `stdout` if no backup specification name is specified. If no client is specified, all clients in the Data Protector cell are searched.

MICROSOFT EXCHANGE SERVER 2003 ZERO DOWNTIME BACKUP

The `omnicreatedl` command is also used to create an Exchange Server ZDB backup specification file for disk arrays of the following disk array families:

HPE P9000 XP Disk Array Family

HPE P6000 EVA Disk Array Family

When creating an Exchange ZDB backup specification file, if the circular logging is disabled for any storage group, an Exchange ZDB transaction logs backup specification file for each such storage group specified in the Exchange ZDB backup specification file is additionally created.

An Exchange ZDB backup specification file includes the stop/quiesce the application and restart the application scripts (`omniEx2000.exe`) sections for dismounting/mounting backed up stores and checking their consistency. A backup specification can be edited later using the Data Protector GUI to modify backup devices, ZDB options, schedule, and so on.

For a Microsoft Exchange Server 2003 ZDB, the *final* decision on whether the created backup specification will start a ZDB-to-disk, ZDB-to-disk+tape or ZDB-to-tape session depends on the Data Protector `omnib` command options selection.

OPTIONS

`-version`

Displays the version of the `omnicreatedl` command.

`-help`

Displays the usage synopsis for the `omnicreatedl` command.

FILESYSTEM BACKUP

`-datalist` *Name*

Specifies the name of the backup specification file (datalist) for filesystem backup. The backup specification file is created on the Cell Manager in the default server configuration directory under `datalists`. If this option is not specified, backup specification objects are written to `stdout`.

`-host` *HostName1 [HostName2]*

List of all clients whose filesystems will be included in the backup specification. If this option is not specified, all clients from the cell are used.

-device *BackupDevice*

Specifies the backup device to be used for backup. If this option is not used, the backup device must be specified using the Data Protector GUI.

MICROSOFT EXCHANGE SERVER 2003 ZERO DOWNTIME BACKUP

-ex2000

Instructs the `omnicreatedl` command to create a Microsoft Exchange Server 2003 ZDB backup specification file and, if circular logging is disabled for any storage group specified, a Microsoft Exchange Server 2003 ZDB transaction logs backup specification file(s) for every such storage group.

-datalist *Name*

Specifies the name of the Microsoft Exchange Server 2003 ZDB backup specification file for the Microsoft Exchange Server 2003 ZDB. The file is created on the Cell Manager in the default server configuration directory under `datalists`.

The corresponding datalist for Microsoft Exchange Server 2003 logs for every storage group specified that has the circular logging disabled are also created in the same directory with the file name *Storage_Group_Name* (LOGS) *app_sys*.

If any of the thus created backup specification files (datalists) has a name that already exists, the `omnicreatedl` command issues a warning and, depending on whether the -force option is set or not, overwrites the existing backup specification files with the same name or aborts the action.

-force

Forces overwriting of an existing backup specification file with the same name.

-virtualSrv *Name*

The name of the Microsoft Exchange Server 2003 virtual server. This option is obligatory and used only in cluster configurations.

P9000_DISK_ARRAY_XP_OPTIONS

-split_mirror -sse

Instructs the `omnicreatedl` command to create an HPE P9000 XP Disk Array Family Microsoft Exchange Server 2003 ZDB backup specification.

-local *app_sys bck_sys*

Specifies the HPE Business Copy (BC) P9000 XP configuration, with the application system *app_sys* and the backup system *bck_sys*.

-remote *app_sys bck_sys*

Selects the HPE Continuous Access (CA) P9000 XP configuration, with the application system *app_sys* and the backup system *bck_sys*.

-combined *app_sys bck_sys*

Selects the Combined (HPE Continuous Access + Business Copy (CA+BC) P9000 XP) configuration, with the application system *app_sys* and the backup system *bck_sys*.

-mirrors *MU_numbers*

This option is only considered when the HPE Business Copy (BC) P9000 XP configuration is chosen.

Specify the mirror unit (MU) number(s) of a replica or a replica set from which the Data Protector HPEP9000 XP Agent, according to the replica set rotation, selects the replica to be used in the zero downtime backup session. The replica selection rule is described in the *HPE Data Protector Concepts Guide*. The maximum number of replicas that can be created for the same source volumes is different for mirror copies and snapshots. Both limitations are imposed by the HPE P9000 XP Disk Array Family storage system.

You can specify one or more non-negative integer numbers, one or more ascending ranges of such numbers, or any combination of both. Use a comma as the separator character. Examples:

5

7-9

4,0,2-3

When a sequence is specified, it does not define the order in which the replicas are used. If this option is not specified, the MU number 0 is used.

-instant_restore

When specified, this option enables ZDB to disk or ZDB to disk+tape. Consequently, instant recovery can be run using the created replica in the ZDB session. If the option is not specified, it is only possible to perform a ZDB to tape. However, this option does not influence the replica set rotation.

If this option is specified, the omnicreatedl command automatically sets the -keep_version option.

-keep_version

If configuring a ZDB to tape, specify this option to keep the replica on the disk array after the zero downtime backup session. The replica becomes part of a replica set (specify a value for the option -mirrors). Unless the additional option -instant_restore is specified, the replica is not available for instant recovery.

If this option is not specified, the replica is removed at the end of the session. In this case, it is also not possible to specify the -leave_enabled_bs option.

-leave_enabled_bs

To specify this option, the -keep_version option has to be specified.

By default, Data Protector dismounts the filesystems on the backup system after each ZDB session.

If this option is specified, the filesystems remain mounted after the backup. Thus, you can use the backup system for some data warehouse activity afterwards, but not for instant recovery.

-split

If this option is specified, the volumes of the replica selected for the current ZDB session are prepared for the zero downtime backup at the start of the current ZDB session: mirrors are resynchronized with the P-VOLs, and volumes to be used for snapshot storage are made empty.

If neither the `-split` option nor the `-establish` option is specified, Data Protector acts as if the `-establish` option was specified.

`-establish`

If this option is specified, if the volumes of the replica to be used in the next ZDB session are not ready for ZDB, they are prepared for ZDB at the end of the current ZDB session.

If neither the `-split` option nor the `-establish` option is specified, Data Protector acts as if the `-establish` option was specified.

P6000_ENTERPRISE_VIRTUAL_ARRAY_OPTIONS

`-snapshot -smis` *app_sys bck_sys*

Instructs the `omnicreatedl` command to create an HPE P6000 EVA Disk Array Family snapshot backup specification file and sets the application system *app_sys* and the backup system *bck_sys*.

`-instant_recovery`

This parameter is optional. Specify this option, if you want to perform either a ZDB to disk or a ZDB to disk+tape and leave the replica on a disk array (after the backup session) to use it in future for instant recovery. If this option is not set, it is not possible to perform instant recovery from the replica created in this backup session.

Note that when this option is selected, the options `-snapshot_type clone` and `-snapshot_policy strict` are automatically set by Data Protector. If the option `-snapshots` *number* is not specified, it is set to `1`.

`-snapshots` *number*

This parameter is optional. By default, Data Protector automatically sets this option to `1` if the `-instant_recovery` option is specified.

Specify this option if you wish to keep the replica on a disk array after a backup session is completed. With *number*, specify the number of replicas you want to keep on a disk array. During every backup session, Data Protector creates a new replica and leaves it on a disk array as long as the specified number is not reached. When the specified number is reached, Data Protector deletes the oldest replica and creates a new one.

The maximum number for vsnaps and standard snapshots is `7`. Data Protector does not limit the number of replicas rotated, but the session will fail if the limit is exceeded.

Note that this option sets the number of replicas in the replica set for a backup specification.

`-snapshot_type {standard | vsnap | clone}`

This option instructs Data Protector to create one of the three types of HPE P6000 EVA Disk Array Family snapshots during the backup session.

Setting `standard` creates snapshots with the pre-allocation of disk space.

Setting `vsnap` creates snapshots without the pre-allocation of disk space.

Setting `clone` creates a clone of an original virtual disk.

`-snapshot_policy {strict | loose}`

Specifies how Data Protector creates snapshots with regard to types of already existing snapshots for the same original virtual disk.

When `strict` is set, Data Protector attempts to create snapshots of the type selected by the `-snapshot_type` option. If some of the original virtual disks used in the backup session already have existing snapshots of different type, the selected type of snapshots cannot be used. Such a backup session will be aborted.

When `loose` is set, Data Protector creates snapshots of different type than specified by the `-snapshot_type` option, when this would help to make a successful session. For example, if you select standard snapshots to be created, but Data Protector detects that standard snapshots cannot be created because some vsnaps or snapclones of the source volumes already exist in a replica set, the following happens: with the loose option selected, Data Protector creates either vsnaps (if vsnaps already exist) or snapclones (if snapclones already exist) instead of standard snapshots. Note that Data Protector can use only one type of snapshots in the backup session. In case when some of the original virtual disks used in the backup session have existing standard snapshots and some of them existing vsnaps, the backup session will be aborted.

`-wait_clonecopy` *number*

This parameter is optional and can be selected only if the `-snapshot_type clone` option is selected.

In the case of a ZDB to tape or a ZDB to disk+tape, specify this option if you want to delay moving data to tape media until the cloning process is completed. By *number*, specify the maximum waiting time in minutes. After the specified number of minutes, the backup to tape will start even if the cloning process is not finished yet.

With this option, you prevent degradation of the application data access times during the phase of backup to tape.

`-replica_conf` {`local`|`combined`}

Select the P6000 EVA Array configuration. Specify `local` to configure a backup specification for ZDB in HPE Business Copy (BC) P6000 EVA environments. Specify `combined` to configure a backup specification for ZDB in combined HPE Continuous Access + Business Copy (CA+BC) P6000 EVA environments.

`-ca_failover_option` {`follow_replica_direction`|

`maintain_replica_location`}

This parameter is optional and is available only if the `combined` replica configuration is selected. Specify this option to control the replication direction after a failover.

Select `follow_replica_direction` to follow the replication direction and create replicas on the array remote to current source. A failover reverses the replication direction and the replicas are created on the array that was originally a source P6000 EVA Array.

Select `maintain_replica_location` to maintain the replica location and create replicas on the array remote to home. After a failover, replicas continue to be created on the destination array that has also become a source P6000 EVA Array.

Note that when `-ca_failover_option` option is selected, `follow_replica_direction` is set as default.

EXCHANGE_OPTIONS

`-annotation` {`MIS`|`SRS`|`KMS`}

This option specifies the possible Microsoft Exchange Server 2003 annotations: Microsoft

Information Store (`MIS`), Site Replication Service (`SRS`), and Key Management Service (`KMS`). `MIS` is the default setting and does not need to be specified in case when the `MIS` will be backed up.

`-all_storage_groups`

This option creates a backup specification for all databases relating to Microsoft Exchange Server 2003 Microsoft Information Store. It must be specified by the `-annotation` MIS parameter.

`-storage_group` *storage group name*

This option creates a backup specification for all stores relating to the specified storage group. Multiple declarations of the `-storage_group` parameter are possible to create a backup specification for the selected storage groups.

Logical storage group names can be obtained by using the Exchange System Administrator tool, which is a part of Microsoft Exchange Server 2003.

`-store` *Store1* [*Store2*...]

When the `-store` parameter is specified, backup specification is created only for specified store(s) inside the storage group. List of stores can be specified after the `-store` parameter to create a backup specification for many stores.

Store names can be obtained by using Exchange System Administrator tool, which is a part of Microsoft Exchange Server 2003.

EXAMPLES

The following examples show how the `omnicreatedl` command works:

1. To create an HPE P9000 XP Disk Array Family Microsoft Exchange Server 2003 ZDB-to-tape backup specification file named "Exchange_example" for a Microsoft Exchange Server 2003 running on client "computer1.company.com" with the backup system "computer2.company.com", to back up all storage groups relating to Microsoft Information Store, execute:

   `omnicreatedl -ex2000 -datalist Exchange_example -all_storage_groups -split_ mirror -sse -local computer1.company.com computer2.company.com`

   `The omnicreatedl command` creates the HPE P9000 XP Disk Array Family Microsoft Exchange Server 2003 ZDB-to-tape backup specification file named "Exchange_example" and additional HPE P9000 XP Disk Array Family Microsoft Exchange Server 2003 ZDB transaction logs backup specification files (in case they do not already exist) for each storage group with disabled circular logging option.

2. To create an HPE P9000 XP Disk Array Family Microsoft Exchange Server 2003 ZDB-to-tape backup specification file named "Exchange_example" for a Microsoft Exchange Server 2003 running on client "computer1.company.com" with the backup system "computer2.company.com", to back up entire First Storage Group and Test Storage Group (both have circular logging disabled), execute:

   `omnicreatedl -ex2000 -datalist Exchange_example -storage_group "First Storage Group" -storage group "Test Storage Group" -split_mirror -sse -local computer1.company.com computer2.company.com`

   `The omnicreatedl command` creates the HPE P9000 XP Disk Array Family Microsoft Exchange Server 2003 ZDB-to-tape backup specification file (datalist) named "Exchange_example" and two

additional HPE P9000 XP Disk Array Family Microsoft Exchange Server 2003 ZDB transaction logs backup specification files (if they do not already exist) named: "First Storage Group (LOGS) computer1.company.com" for First Storage Group log files backup and "Test Storage Group (LOGS) computer1.company.com" for Test Storage Group log files backup.

3. To create an HPE P9000 XP Disk Array Family Microsoft Exchange Server 2003 ZDB-to-tape backup specification file named "Exchange_example" for a Microsoft Exchange Server 2003 running on "computer1.company.com" with the backup system "computer2.company.com", overwriting the possible already existent backup specification files with the same name to back up First Mailbox Store, Public Folder Store, part of First Storage group and Test Mailbox Store, part of Test Storage Group, execute:

```
omnicreatedl -ex2000 -datalist Exchange_example -storage_group "First Storage
Group" -store "First Mailbox Store" "Public Folder Store" -storage group "Test
Storage Group" -store "Test Mailbox Store" -split_ mirror -sse -local
computer1.company.com computer2.company.com -force
```

The `omnicreatedl` command creates the HPE P9000 XP Disk Array Family Microsoft Exchange Server 2003 ZDB-to-tape backup specification file (datalist) "Exchange_example" and two additional HPE P9000 XP Disk Array Family Microsoft Exchange Server 2003 ZDB transaction logs backup specification files if circular logging option is disabled for a particular storage group: "First Storage Group (LOGS) computer1.company.com" for First Storage Group log files backup and "Test Storage Group (LOGS) computer1.company.com" for Test Storage Group log files backup. Any possible already existent backup specification file with the same name is overwritten.

4. To create an HPE P6000 EVA Disk Array Family Microsoft Exchange Server 2003 ZDB-to-tape backup specification file (datalist) "Exchange_example", to back up Site Replication Service on "dev1" device, using the vsnap type of snapshot and the strict snapshot policy, execute:

```
omnicreatedl -ex2000 -datalist Exchange_example -device dev1 -annotation SRS -
snapshot -smis computer1.company.com computer2.company.com -snapshot_type vsnap
-snapshot_policy strict
```

The `omnicreatedl` command creates an HPE P6000 EVA Disk Array Family Microsoft Exchange Server 2003 ZDB-to-tape backup specification file named "Exchange_example" and an HPE P6000 EVA Disk Array Family Microsoft Exchange Server 2003 ZDB transaction logs backup specification file in case it does not already exist: "SRS (LOGS) computer1.company.com" for Site Replication Service log files backup if the circular logging is disabled. When the `omnib` command or Data Protector GUI is used to start the created backup specification, Data Protector tries to create the vsnap type of snapshots if they cannot be created, the session aborts.

5. To create an HPE P6000 EVA Disk Array Family Microsoft Exchange Server 2003 ZDB-to-disk backup specification file (datalist) "Exchange_example", to back up Site Replication Service on the backup device "dev1", using the replica set with "5" replicas, execute:

```
omnicreatedl -ex2000 -datalist Exchange_example -device dev1 -snapshot -smis
computer1.company.com computer2.company.com -instant_recovery -snapshots 5 -
annotation SRS
```

In case it does not already exist, `omnicreatedl` creates an HPE P6000 EVA Disk Array Family Microsoft Exchange Server 2003 transaction logs backup specification file "SRS (LOGS) computer1.company.com" for Site Replication Service log files backup (the circular logging must be disabled). When the `omnib` command or Data Protector GUI is used to start the created backup specification, you must choose the ZDB-to-disk session. Data Protector tries to create the

snapclone type of snapshots; if they cannot be created, the session aborts. After the backup session, the created replica is retained on a disk array and can be used for instant recovery.

6. To create an HPE P6000 EVA Disk Array Family Microsoft Exchange Server 2003 ZDB-to-disk+tape backup specification file (datalist) "Exchange_example", to back up Site Replication Service on the backup device "dev1", using the replica set with "3" replicas and to delay the backup to tape for the maximum of "50" minutes, execute:

```
omnicreatedl -ex2000 -datalist Exchange_example -device dev1 -snapshot -smis
computer1.company.com computer2.company.com -instant_recovery -snapshots 3 -
wait_clonecopy 50 -annotation SRS
```

In case it does not already exist, `omnicreatedl` creates an HPE P6000 EVA Disk Array Family Microsoft Exchange Server 2003 transaction logs backup specification file "SRS (LOGS) computer1.company.com" for Site Replication Service log files backup (the circular logging must be disabled). When the `omnib` command or Data Protector GUI is used to start the created backup specification, you must choose the ZDB-to-disk+tape session. Data Protector tries to create the snapclone type of snapshots; if they cannot be created, the session aborts. The backup to tape will start after the snapclones are fully created or after 50 minutes. After the backup session, the created replica is retained on a disk array and can be used for instant recovery.

SEE ALSO

omnib(1), omniintconfig.pl(1M), util_cmd(1M), util_oracle8.pl(1M), vepa_util.exe(1M)

omnidb(1)

omnidb — queries the Data Protector Internal Database (IDB)
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnidb -version|-help

omnidb -session [-datalist *Datalist*] [-type {restore|backup|verification}] [-user *User*] [-since *Date* -until *Date*|-last *Number*|-latest|-wo *start duration*] [-detail]

omnidb -filesearch [-n *N*] *Client Directory FileName*

omnidb *Object* [-session *SessionID*] [-copyid *CopyID*] -listdir *Directory*

omnidb -list_folders -session *SessionID* [-mailbox *MailboxName...*]

omnidb -rpt [*SessionID*|-latest] [-detail]

omnidb -rpt [-wo *start duration*]

omnidb -session *SessionID* [-report *Report* [warning|minor|major|critical]|-detail|-encryptioninfo|-strip|-purge|-change_protection *Protection*|-change_catprotection *Protection*|-media [-detail]|-remove_msgs]

omnidb -object [-detail|-encryptioninfo]

omnidb [-noexpand] {-filesystem|-winfs|-vbfs} *Client:MountPointLabel* [-file *FileName*] [-detail|-encryptioninfo]

omnidb *Object* [[-since *Date*] [-until *Date*]|-last *NumberOfDays*] [-change_protection *Protection*|-change_catprotection *Protection*] [-noexpand]

omnidb *Object* {[-since *Date*] [-until *Date*]|-last *NumberOfDays*} [-latest] [-detail|-encryptioninfo]

omnidb *Object* -strip *NumberOfDays*

omnidb -strip

omnidb -change_protection *Protection*

omnidb -change_catprotection *Protection*

omnidb [-noexpand] {-filesystem|-winfs|-vbfs} *Client:MountPoint Label* -fileversions *FileName* [{-detail|-encryptioninfo}]

omnidb *Object* [-detail|-encryptioninfo]

omnidb *Object* [-noexpand] -session *SessionID* [-copyid *CopyID*] [-report [*Report*]|-catalog|-change_protection *Protection*|-change_catprotection *Protection*|-strip|-encryptioninfo]

omnidb *Object* -session *SessionID* [-copyid *CopyID*] -media [-detail]

omnidb *Object* -session *SessionID* [-copyid *CopyID*] -listcopies [-detail|-encryptioninfo]

```
omnidb -auditing {-timeframe [StartDate] [EndDate] | -since Date [-until Date] | -last
NumberOfDays} [-detail]
```

Object

```
{ -filesystem Client:MountPoint Label |

 -winfs Client:MountPoint Label |

 -vbfs Client:MountPoint Label |

 -rawdisk Client Label |

 -stream Client:Set |

 -sap Client:Set |

 -sapdb Client:Set |

 -oracle8  Client:Set |

 -mssql Client:Set |

 -msese Client:Set |

 -e2010 Client:Set |

 -mbx Client:Set |

 -informix Client:Set |

 -sybase Client:Set |

 -lotus Client:Set |

 -vss Client:Set |

 -db2 Client:Set |

 -mssharepoint Client:Set |

 -veagent Client:Set |

 -idb [Client:Set] |

 -integ {MySQL | PostgreSQL} [Client:Set]}
```

Protection

```
{none | days n | weeks n | until Date | permanent}
```

Report

```
 warning | minor | major | critical
```

Date

```
 [YY]YY/MM/DD (1969 < [YY]YY < 2038)
```

DESCRIPTION

The `omnidb` command is used to query the IDB Log database.

This command can be used to:

- list sessions and their summary reports

- list backed up objects and their details (for example: client name, mountpoint, label, object type, object status, backup type, and so on), message logs, and media location

- search for all occurrences of a pathname pattern

The `omnidb` command performs basic IDB queries.

OPTIONS

`-version`

Displays the version of the `omnidb` command.

`-help`

Displays the usage synopsis for the `omnidb` command.

`-datalist` *IntegrationName BackupSpecificationName*

Lists the sessions resulting from backup specification backups created using this *BackupSpecificationName*.

NOTE: For non-filesystem backup specification (Microsoft Exchange Server, Microsoft SQL Server, Informix Server, and so on).

IntegrationName must be specified in front of *BackupSpecificationName*. Both must be in double quotes.

`-type` {`restore`|`backup`|`verification`}

If no *SessionID* is specified, the command lists either backup, restore, or verification sessions. If *SessionID* is specified for backup sessions, the command lists the objects created for that backup session.

`-user` *User*

Lists only the sessions belonging to the specified user.

`-since` *Date*

Lists sessions since the given *Date*.

`-until` *Date*

Lists sessions until the given *Date*.

`-last` *n*

Lists sessions that occurred within the last *n* days.

`-latest`

Lists the last active Data Protector session.

`-wo` *start duration*

Lists the sessions that started within a specified timeframe. *Start* defines the start of the timeframe. *Duration* is the duration of the timeframe in seconds.

`-detail`

Displays detailed information about the selected query, such as backup type, protection, whether or not encrypted. The VADP feature introduced in Data Protector 8.14 provides enhanced reports for Virtual Machines. Reports on VMware virtual machines display VM objects in the command line output in the same way as regular Data Protector clients. The reports for each VM object displays additional information, such as VM name, VM path, VM UUID, ESXi server, and full object name.

-encryptioninfo

Displays detailed encryption information for objects meeting the query criteria.

-session *SessionID*

Displays session information. If no *SessionID* is specified, all sessions are shown. The report shows for each session: the ID, type, status and user (UNIX login, UNIX group and client). If a *sessionID* is specified, then objects that are backed up within this session are shown. This information includes: client name, mountpoint, label, object type and object status.

If the -detail option is specified, more information is shown, such as the backup type (full, incr,...), protection status, encryption status, and so on. For integration objects, the backup ID is also shown. For VADP clients, the object name must use the VM name as reported from omnicellinfo -cell brief command and this applies to all variations of omnidb -session, where object name is <hostname>:/<vCenter>/<path>/<vmname>[<UUID>]. Here, <hostname> is DNS name of the guest virtual machine. If the DNS name is unknown, the IP address or VM name is used.

If the -encryption option is specified, the encryption *KeyID-StoreID* is displayed for each encrypted object created during the specified session. *SessionID* is mandatory in this case.

-auditing

Lists auditing related information from the cell. The following information is listed for each backup session: name, specification, completion status, backup type, start time, end time, and owner.

If the -detail option is specified, the command also lists used media and objects.

-copyid *CopyID*

If several copies of the same object exist in one session as a result of the object copy or object mirror operation, this option is obligatory. It selects a specific copy.

-filesearch [-n *N*] *Client Directory FileName*

Lists all the backed up files and directories that match the selection criteria set by the *Client Directory FileName* parameters. Wildcards can be used. The list can be limited to a certain number of displayed objects by setting the -n option, where *N* is the number of objects to be displayed. The following information is displayed about each object: object type, object name, object description, pathname.

-listdir *Directory*

Lists all the backed up objects in the specified directory.

-list_folders

Microsoft Exchange Server Single Mailbox integration: displays a list of all single mailbox folders (including their subfolders) backed up within a particular session.

-mailbox *MailboxName*

Microsoft Exchange Server Single Mailbox integration: displays mailbox folders for a particular

mailbox only. If the option is not specified, folders of all backed up mailboxes are listed.

-listcopies

Lists details on all existing object or mirror copies of the specified object for the specified session. The session ID, the CopyID, the time and the status of object copy or mirror sessions for the specified object are listed.

-rpt *SessionID*

Displays session information in a form specially suited for further use of awk, grep or perl. Records are separated with blank lines and line feed is the field separator. If no *SessionID* is specified, all backup sessions are shown. Each record contains the following fields: the ID, backup specification name, status, start time in format *HH:MM* and duration in hours as a floating point number.

-report *Report*

Lists all messages (of specified report level and higher) which were generated by the specified session. Messages are classified (in ascending order) as: `warning`, `minor`, `major` and `critical`. For example, if `major` is selected, only `major` and `critical` messages are reported. By default, all messages are reported.

-object

Displays information on all data objects. The report shows the client name, label, and object type.

If the `-detail` option is specified, more detailed information is displayed for each object, such as each session for which object versions were created, together with protection status, encryption status, and so on.

If the `-encryptioninfo` option is specified, for each object, the encryption *KeyID-StoreID* is displayed for each session in which object versions were created.

-filesystem *Client:MountPoint Label*

Displays information on all filesystem objects (displays the *Client:MountPoint Label* string for every filesystem object in the IDB). If a *Client:MountPoint Label* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-winfs *Client:MountPoint Label*

Displays information on all Windows filesystem objects (displays the *Client:MountPoint Label* string for every Windows filesystem object in the IDB). If a *Client:MountPoint Label* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-vbfs *Client:MountPoint Label*

Displays information on all Windows filesystem objects (displays the *Client:MountPoint Label* string for every Windows filesystem object in the IDB). If a *Client:MountPoint Label* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-rawdisk *Client Label*

Displays information on disk image objects (displays the *Client Label* string for every object in

the IDB). If a *Client Label* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-stream *Client:Set*

Displays information on stream objects (displays the *Client:Set* string for every stream object in the IDB). If a *Client:Set* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-sap *Client:Set*

Displays information on SAP R/3 data objects (displays the *Client:Set* string for every SAP R/3 object in the IDB). If *Client:Set* is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-sapdb *Client:Set*

Displays information on SAP MaxDB data objects (displays the *Client:Set* string for every SAP MaxDB object in the IDB). If a *Client:Set* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-oracle8 *Client:Set*

Displays information on Oracle objects (displays the *Client:Set* string for every Oracle object in the IDB). If a *Client:Set* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the status, size of object, and the number of session errors reported.

-mssql *Client:Set*

Displays information on Microsoft SQL Server objects (displays the *Client:Set* string for every Microsoft SQL Server object in the IDB). If a *Client:Set* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-msese *Client:Set*

Displays information on Microsoft Exchange Server objects (displays the *Client:Set* string for every Microsoft Exchange Server object in the IDB). If a *Client:Set* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-e2010 *Client:Set*

Displays information on Microsoft Exchange Server 2010/2013 objects (displays the *Client:Set* string for every Microsoft Exchange Server 2010/2013 object in the IDB). If a *Client:Set* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-mbx *Client:Set*

Displays information on Microsoft Exchange Server objects - single mailboxes (displays the *Client:Set* string for every Microsoft Exchange Server object - single mailboxes in the IDB). If a *Client:Set* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-informix *Client:Set*

Displays information on Informix Server objects (displays the *Client:Set* string for every Informix Server object in the IDB). If an *Client:Set* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-sybase *Client:Set*

Displays information on Sybase objects (displays the *Client:Set* string for every Sybase object in the IDB). If a *Client:Set* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-lotus *Client:Set*

Displays information on Lotus Notes/Domino objects (displays the *Client:Set* string for every Lotus Notes/Domino object in the IDB). If a *Client:Set* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-vss *Client:Set*

Displays information on Microsoft Volume Shadow Copy (VSS) objects (displays the *Client:Set* string for every VSS object in the IDB). If a *Client:Set* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-db2 *Client:Set*

Displays information on IBM DB2 UDB objects (displays the *Client:Set* string for every IBM DB2 UDB object in the IDB). If a *Client:Set* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-mssharepoint *Client:Set*

Displays information on Microsoft SharePoint Server 2007/2010 objects (displays the *Client:Set* string for every Microsoft SharePoint Server 2007/2010 object in the IDB). If a *Client:Set* string is specified, the backup sessions containing the object specified by this string are listed. For each backup session, the report shows: the session ID, session start time, session duration, backup object status and size, and the number of session errors reported.

-veagent *Client:Set*

Displays information on virtual environment objects (displays the *Client:Set* string for every virtual environments object in the IDB). If a *Client:Set* string is specified, the backup sessions

containing the object specified by this string are listed. For each backup session, the report shows: the status, size of object, and the number of session errors reported.

Reports display VMware virtual machines in the same way as Data Protector clients called VADP clients. The new object name format is as follows:

`<hostname>:/<vCenter>/<path>/<vmname>[<UUID>]`

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.

-idb [`Client:Set`]

Displays information on backup objects of the Internal Database type that are referenced in the IDB. In the `omnidb` output, such objects are marked with the `IDB` string.

If the argument `Client:Set` is not specified, `omnidb` lists the `Client:Set` string for every Internal Database backup object. If the `Client:Set` argument is specified, `omnidb` lists for each corresponding backup session its session ID, start time, duration, statuses and sizes of the backed up objects, and the number of errors reported.

-integ {MySQL | PostgreSQL} `Client:Set`

Displays information on backup objects of the MySQL or PostgreSQL type that are referenced in the IDB. In the `omnidb` output, such objects are marked with the `MySQL or PostgreSQL` string.

If the argument `Client:Set` is not specified, `omnidb` lists the `Client:Set` string for every MySQL or PostgreSQL backup object. If the `Client:Set` argument is specified, `omnidb` lists for each corresponding backup session its session ID, start time, duration, statuses and sizes of the backed up objects, and the number of errors reported.

Note: In the `omnidb` command output, MySQL pseudo-backup objects with metadata are denoted with the suffix `:METADATA` in their names.

-strip

This option works in three different ways. If `SessionID` is specified it strips the Detail Catalog of all the objects from the session with a specified session ID. If both `SessionID` and `ObjectName` are specified it strips the Detail Catalog of the object identified by `ObjectName` from the session with specified session ID. If no option is specified, it strips the Detail Catalog of all data objects that are no longer protected.

-strip `NumberOfDays`

This option can be used with ObjectName to strip the Detail Catalog of all versions of the specified object that are older than `NumberDays` days.

-fileversions `FileName`

Displays information on all sessions which contain the filesystem with specified file `FileName`, session ID, mode, date modified, size and type.

-media

Shows list of the media used in the backup session. If object is also specified then it only shows list of media containing that object.

-user_location

This option changes the output of media related reports to print out user defined location instead of

physical location used by default.

-change_protection *Protection*

Changes the current protection of the object versions identified by ObjectName and/or SessionID to the new protection defined as Protection. If it is specified without any other option then it changes protection for all Failed/Aborted objects. Protection can be none, permanent, until a specific date, or for a time interval. When the protection is until a specified date or for a time interval, you must specify the value. The Date form is [YY]YY/MM/DD. In the first case the value is the date until which the data is protected. In the second case the time interval is the number of days (after today) during which the data cannot be overwritten.

> **Note:** From Data Protector 9.05 onwards, the protection for any virtual machine can be modified for a specific VEAgent object. This option allows you to change the protection for all the VEAgent Disk objects, when the VEAgent object is selected.

-change_catprotection *Protection*

Changes the current protection of the catalog retention time. *Protection* can be none, same_as_ data_protection, until a specific date, or for a time interval. same_as_data_protection means that catalog will stay until data is overwritten/exported. When the protection is until a specified date or for a time interval, you must specify the value. The Date form is [YY]YY/MM/DD. In the first case the value is the date until which the data is protected. In the second case the time interval is the number of days (after today) during which the data cannot be overwritten.

> **Note:** From Data Protector 9.05 onwards, the catalog protection for any virtual machine can be modified for a specific VEAgent object. This option allows you to change the catalog protection for all the VEAgent Disk objects, when the VEAgent object is selected.

-catalog

Displays the Detail Catalog of a specified object - session combination. Use an object option (for example -filesystem) to specify the object and use the -session (and *sessionID*) to specify the session.

-purge

This option removes the session from the session list. All objects within the session become unprotected. It is still possible to make a restore from this session.

-timeframe *StartDate EndDate*

Lists the sessions that started within a specified timeframe.

-vdiskuuid *diskUUID*

This option can be used with the "-catalog" or "-media" to list the associated virtual machine disk objects. These options are supported only for the VMware.

-list_vdisks

Lists the disk information (disk name and disk UUID) for a backed up virtual machine object.

> **Note:** Both *-vdiskuuid diskUUID* and *-list_vdisks* options are available only for objects backed using Data Protector 9.05 and later.

NOTES

With clustered objects, the `Client` argument must define name of the virtual host.

With virtual environment objects, the VM UUID displayed by the `omnidb` command refers to the instance UUID of the virtual machine.

The virtual machine objects and its associated disk objects constitute the VEAgent object size.

EXAMPLES

The following examples illustrate how the `omnidb` command works.

1. To see details for the backup sessions started by user "root" in last three days, execute:

   `omnidb -session -user root -last 3 -type backup -detail`

2. To see critical errors for the session with the sessionID "2013/05/14-17", execute:

   `omnidb -session 2013/05/14-17 -report critical`

3. To see all virtual machines used in backup as VEPA objects and its additional information, execute:

   `omnidb -session 2015/11/05-1 -detail`

4. To see all virtual machines from all VEPA backups as objects, execute:

   `omnidb –veagent`

5. To display session information about a single virtual machine, execute:

   `omnidb –veagent host.domain.name:/vcenter.domain.name/datacenter/path/example_host[c6a20393-159d-4b9a-8671-73a4490ab032]`

   where, `<hostname>` is the DNS name of the guest virtual machine or IP address.

6. To see all objects of the type filesystem, execute:

   `omnidb -filesystem`

7. To see encryption information for all Windows filesystem objects, execute:

   `omnidb -winfs -encryptioninfo`

8. To see encryption information for objects created in session "2013/03/23-2" execute:

   `omnidb -session 2013/03/23-2 -encryptioninfo`

9. To see details for the filesystem "hpuljum.company.com:/ Label44" in the latest session, execute:

   `omnidb -filesystem hpuljum.company.com:/ Label44 -latest -detail`

10. To see catalog for the filesystem "bob:/" in the session "2012/07/14-6", execute:

    `omnidb -filesystem bob:/ -session 2012/07/14-6 -catalog`

11. To see details of the sessions that used a Microsoft Exchange Server backup specification named "MSExchange test", execute:

    `omnidb -session -datalist "E2010 MSExchange test" -details`

12. To list all Microsoft Exchange Server mailbox folders in the mailbox "User 2", backed up in the session "2013/03/16-10", execute:

```
omnidb -mbx -list_folders -session 2013/03/16-10 -mailbox "User 2"
```

13. To see information on Lotus Notes/Domino Server objects, execute:

```
omnidb -lotus
```

14. To see which Lotus Notes/Domino Server files are contained in the Lotus Notes/Domino Server object "computer.company.com:DREAM::Databases:5" from the session "2012/08/26-2", execute:

```
omnidb -lotus computer.company.com:DREAM::Databases:5 -session 2012/08/26-2 -
catalog
```

15. To see information on the SAP MaxDB object "machine.company.com:/instance1/Config/1", execute:

```
omnidb -sapdb machine.company.com:/instance1/Config/1
```

16. To see information on the SAP HANA object "hanasys.company.com:H95:7", execute:

```
omnidb -sapdb machine.company.com:/instance1/Config/1
```

17. To see detailed information on media used for the Windows filesystem object "system.company.com:/C" with the label "DTS_T" in the session "2012/07/14-17", with CopyID "d5032390-baba-4b3f-8c67-1f5b9273b242/1015", execute:

```
omnidb -winfs system.company.com:/C "DTS_T" -session 2012/07/14-17 -copyid
d5032390-baba-4b3f-8c67-1f5b9273b242/1015 -media -detail
```

18. To see detailed information on all existing object or mirror copies of the Windows filesystem object "system.company.com:/D" with the label "D1" with the sessionID "2013/05/01-12", execute:

```
omnidb -winfs system.company.com:/D "D1" -session 2013/05/01-12 -listcopies -
detail
```

19. To see information on Microsoft SharePoint Server 2010 configuration database objects, execute:

```
omnidb -mssharepoint helios.company.com:SharePoint_Config/1:SharePoint_Config
```

20. To display information on backup sessions that backed up the Internal Database (more specifically, the set "DPSPECs:0") on the Cell Manager with the fully qualified domain name "cmsys.company.com", execute:

```
omnidb -idb cmsys.company.com:DPSPECs:0
```

21. To display information on MySQL backup objects backed up with Data Protector, execute:

```
omnidb -integ MySQL
```

SEE ALSO

omnidbcheck(1M), omnidbinit(1M), omnidbp4000(1), omnidbsmis(1), omnidbutil(1M), omnidbvss(1), omnidbxp(1), omnidbzdb(1), omniofflr(1M)

omnidbp4000(1)

omnidbp4000 — manages the configuration data which the Data Protector HPE P4000 VSS Agent uses to connect to the CIMOM providers
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnidbp4000 --version | --help

omnidbp4000 --ompasswd --add *ClientName* [--ssl] [--port *PortNumber*] [--user *Username*] [--passwd *Password*] [--check [--host *ClientName*]]

omnidbp4000 --ompasswd --remove *ClientName* [--port *PortNumber*] [--user *Username*]

omnidbp4000 --ompasswd [--list [*ClientName*]]

omnidbp4000 --ompasswd --check [--host *ClientName*]

DESCRIPTION

The `omnidbp4000` command enables you to manage configuration data which is used for connections between the Data Protector HPE P4000 VSS Agent and the chosen Common Information Model Object Manager (CIMOM) providers. Such connections must be properly configured before storage systems of the HPE P4000 SAN Solutions family can be used for zero downtime backup and instant recovery purposes. For an overview, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

Using `omnidbp4000`, you should configure the connection to the chosen CIMOM provider. Once configured, the connection configuration data corresponding to the chosen CIMOM provider is stored in a separate configuration file located on the Cell Manager in the directory:

Windows systems: *Data_Protector_program_data*\server\db80\smisdb\p4000\login

UNIX systems: /var/opt/omni/server/db80/smisdb/p4000/login

With `omnidbp4000`, you can also update or remove the connection configuration data, list the contents of the configuration files, and check if the connection to a particular CIMOM provider can be established. For these purposes, the `omnidbp4000` command provides the basic options `--add`, `--remove`, `--list`, and `--check`. The option `--add` can be used for configuring a connection anew as well as updating the configuration data for an already configured connection.

OPTIONS

--version

   Displays the version of the `omnidbp4000` command.

--help

Displays the usage synopsis for the `omnidbp4000` command.

`--ompasswd --add` *ClientName* `[--ssl]` `[--port` *PortNumber*`]` `[--user` *Username*`]` `[--passwd` *Password*`]` `[--check` `[--host` *ClientName*`]]`

Configures or reconfigures the data which the Data Protector HPE P4000 VSS Agent uses to establish connection to a CIMOM provider whose service is running on the system *ClientName*. For *ClientName* you can specify either fully qualified domain name, host name, or IP address of the system. Host names are automatically expanded to fully qualified domain names before they are stored to the configuration files. If no additional options are specified, `omnidbp4000` configures the connection as a non-SSL connection, using the port number `5988` as the CIMOM service listening port, and using `administrator` as the user name. In this case, `omnidbp4000` prompts you to enter the password interactively, and omits the initial connection check.

If the option `--ssl` is specified, the connection is configured to use SSL.

If the option `--port` is specified, the connection is configured to use the port number *PortNumber*. If not specified, the default port number is used: `5988` for connections not using SSL, `5989` for connections using SSL. HPE recommends you use the default port number.

If the option `--user` is specified, the connection is configured to use the user name specified in *Username*. In the opposite case, the default user name `administrator` is used. If the option `--password` is specified, the connection is configured to use the password *Password*. If not specified, `omnidbp4000` prompts you to enter the password interactively,

If the option `--check` is specified, `omnidbp4000` checks if the connection to the CIMOM provider can be established after storing the data to the connection configuration file. If the option `--host` is specified, the Data Protector HPE P4000 VSS Agent checking the connections is started on the system *ClientName*, otherwise one of the systems with the Data Protector HPE P4000 VSS Agent installed is chosen by Data Protector. For *ClientName* you can specify either fully qualified domain name or IP address of the system.

`--ompasswd --remove` *ClientName* `[--port` *PortNumber*`]` `[--user` *Username*`]`

Removes the connection configuration data, which has been added by `omnidbp4000`, for the CIMOM providers whose service is running on the system *ClientName*. For *ClientName* you can specify either fully qualified domain name or IP address of the system. If the option `--port`, the option `--user`, or both are specified in addition, only the configuration files corresponding to connections whose port number matches *PortNumber*, whose user name matches *Username*, or whose port number and user name both match the specified values are removed, respectively.

`--ompasswd` `[--list` `[`*ClientName*`]]`

Lists all existing connection configuration data for the CIMOM providers, which has been added by `omnidbp4000`. For each provider, the following information is displayed: the user name, the fully qualified domain name or IP address of the system hosting the CIMOM service, the port number of the CIMOM service listening port, and the indicator whether the connection uses SSL. You can narrow the output to only a particular system by specifying the argument *ClientName*. For *ClientName* you can specify either fully qualified domain name or IP address of the system.

`--ompasswd --check` `[--host` *ClientName*`]`

Triggers a check if the configured connections from the Data Protector HPE P4000 VSS Agent to the CIMOM providers can be established. If the option `--host` is specified, the Data Protector HPE P4000 VSS Agent checking the connections is started on the system *ClientName*, otherwise one of the systems with the Data ProtectorHPE P4000 VSS Agent installed is chosen by Data

Protector. For `ClientName` you can specify either fully qualified domain name or IP address of the system.

NOTES

The `omnidbp4000` command is available on Windows systems only.

EXAMPLES

The following examples illustrate how the `omnidbp4000` command works.

1. To configure a connection to the CIMOM provider hosted on the system "cimom_host1" in the local domain, so that the connection uses SSL, the CIMOM service port number "5989", the user name "administrator", and the password "secretstring" to connect to the CIMOM provider, execute:

   `omnidbp4000 --ompasswd --add cimom_host1 --ssl --password secretstring`

2. To update the configuration of the connection to the CIMOM provider hosted on the system "cimom_host3.company.com" that does not use SSL and uses the user name "storagesys_admin" to connect to the CIMOM provider, so that the Data Protector HPE P4000 VSS Agent uses the new password "newsecretstring" to connect, execute:

   `omnidbp4000 --ompasswd --add cimom_host3.company.com --password newsecretstring`

3. To remove configuration data for connections to the CIMOM providers hosted on the system with the fully qualified domain name "cimom_host2.company.com" and for which the user name "backup_admin" is used, execute:

   `omnidbp4000 --ompasswd --remove cimom_host2.company.com --user backup_admin`

4. To list connection configuration data for connections to the CIMOM providers hosted on the system with the IP address "16.57.73.10", execute:

   `omnidbp4000 --ompasswd --list 16.57.73.10`

5. To trigger a check if the configured connections to the CIMOM providers can be established, and use the Data ProtectorHPE P4000 VSS Agent installed on the system "p4000_host1.company.com" for checking, execute:

   `omnidbp4000 --ompasswd --check --host p4000_host1.company.com`

SEE ALSO

omnidb(1), omnidbcheck(1M), omnidbinit(1M), omnidbrestore(1M), omnidbsmis(1), omnidbutil(1M), omnidbvss(1), omnidbdxp(1), omnidbzdb(1), omniofflr(1M)

omnidbsmis(1)

omnidbsmis — executes administrative tasks on the ZDB database (SMISDB) and on a disk array of the HPE P6000 EVA Disk Array Family

(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnidbsmis -version|-help

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

omnidbsmis -ompasswd {-remove *ClientName* | -delete*ClientName*} [-port *PortNumber*] [-user *Username*]

omnidbsmis -ompasswd -list [*ClientName*]

omnidbsmis -ompasswd -check [-host *ClientName*]

omnidbsmis -dgrules {-init | -put *FileName* | -get *FileName* | -check *EVA_WWN DG_name*}

omnidbsmis -caconf {-init | -put *FileName* | -get *FileName* | -list *EVA_WWN* | -check *DR_Group_Name*}

omnidbsmis [-list] {-session [-ir] [-excluded] [-original] | -datalist}

omnidbsmis [-show] {-session *SessionID* | -datalist *DatalistName*}

omnidbsmis -list -purge

omnidbsmis -purge [-force] [-host *ClientName*]

omnidbsmis -delete {-session *SessionID* | -datalist *DatalistName*} [-reference] [-preview] [-force] [-host *ClientName*]

omnidbsmis -sync_check [-host *ClientName*] [-session *SessionID* | -datalist *DatalistName*]

omnidbsmis {-exclude | -include} -session *SessionID*

DESCRIPTION

Using the `omnidbsmis` command, you can perform various tasks related to the SMISDB and the HPE SMI-S P6000 EVA Array provider.

SETTING, DELETING, LISTING, AND CHECKING THE LOGIN INFORMATION FOR THE SMI-S P6000 EVA ARRAY PROVIDER

The `omnidbsmis` command can be used to set, delete, list, and check the login information for the SMI-S P6000 EVA Array provider. The systems with the SMI-S P6000 EVA Array provider installed are referred to as management systems.

The `omnidbsmis` options used for manipulating the login information for SMI-S P6000 EVA Array provider, which should be used together with the `-ompasswd` option, are: `-add`, `-remove`, `-delete`, `-list`, `-check`, `-ssl,`-port`, `-user`, and `-passwd`.

SETTING THE DISK GROUP PAIRS CONFIGURATION FILE

The `omnidbsmis` command can be used to manipulate the P6000 EVA disk group pairs configuration file.

By default, Data Protector creates snapclones in the same disk group as the source volumes they belong to, and it creates mirrorclones in the same disk group as the original volumes they belong to. However, you can customize the allocation of snapclones and mirrorclones so that they are created in any disk group that is configured on the disk array. Note that standard snapshots and vsnaps are always created in the disk group of their source volumes whether the latter are original volumes or mirrorclones.

The `omnidbsmis` options used for manipulating the P6000 EVA disk group pairs configuration files, which should be used together with the `-dgrules` option, are: `-init`, `-put`, `-get`, `-check`.

SETTING UP THE P6000 EVA HOME CONFIGURATION FILE

The `omnidbsmis` command can be used to manipulate the HOME configuration file for the P6000 EVA storage system. You can create a new HOME configuration file template and store it in its default configuration directory, download the file for editing, and upload it back to the SMISDB. You can also list the data replication (DR) groups with a specified P6000 EVA storage system acting as home and check if a specified DR group is part of an HPE CA+BC P6000 EVA configuration.

The `omnidbsmis` options used for manipulating the P6000 EVA HOME configuration file, which should be used together with the `-caconf` option, are: `-init`, `-get`, `-put`, `-list`, and `-check`.

QUERYING THE INFORMATION ON THE BACKUP OBJECTS

The `omnidbsmis` command can be used to query the SMISDB for the information on the zero downtime backup (ZDB) sessions (the product of every successful ZDB session is a replica) and the ZDB backup specifications (a group of replicas created using the same ZDB backup specification is a replica set).

Using the `omnidbsmis` command to query the SMISDB, you can:

1. Get detailed information on a specific ZDB session (replica).

2. Get detailed information on all ZDB sessions created using a specific ZDB backup specification (replica set).

3. Get a list of all ZDB sessions created using the same ZDB backup specification.

4. Get a list of all ZDB sessions available for instant recovery.

5. Get a list of all ZDB backup specifications that have a replica created.

6. Get a list of replicas to be deleted (marked with the purge flag).

7. Get a list of replicas that are excluded from use.

8. Get a list of replicas for each of which an instant recovery session was performed and the corresponding original volumes were preserved on the disk array after the session.

Note that session details are only displayed for the sessions that have the `Keep the replica after the backup` option selected in the ZDB backup specification. Information about ZDB-to-tape sessions without this option selected is deleted from the SMISDB after each such session.

Entries which denote automatic mirrorclone creation operations performed by the Data Protector HPE P6000 / HPE 3PAR SMI-S Agent are presented as pseudo-ZDB sessions, and are listed together with the associated "regular" ZDB sessions.

The `omnidbsmis` options used for querying the SMISDB are: `-list`, `-show`, `-session`, `-datalist`, `-ir`, `-excluded`, `-original`, and `-purge`.

PURGING THE SMISDB

The `omnidbsmis` command can be used to run the purge operation that checks the SMISDB for the virtual disks with the purge flag and, in case of finding such disks, attempts to delete these objects.

The `omnidbsmis` options used for purging replicas and their entries in the SMISDB, which should be used together with the `-purge` option, are: `-force` and `-host`.

DELETING SPECIFIC REPLICAS FROM THE DISK ARRAY AND FROM THE SMISDB

The `omnidbsmis` command can be used to delete volumes (replicas or replica sets) associated with specific ZDB sessions from the disk array and information about them from the SMISDB. It can perform deletion only for a specific ZDB session (a replica), identified by the session ID, or for all sessions based on a specific ZDB backup specification (a replica set), identified by the backup specification name. Additional option is to only delete information about the specific replicas from the SMISDB. Mirrorclones created by Data Protector and their SMISDB entries can also be deleted. A mirrorclone can only be deleted provided that no snapshots are attached to it.

Note that it is not possible to perform instant recovery using a deleted replica or replica set.

The `omnidbsmis` options used for deleting replicas and SMISDB entries, or only SMISDB entries, which should be used together with the `-delete` option, are: `-session`, `-datalist`, `-reference`, `-preview`, `-force`, and `-host`.

COMPARING THE SMISDB CONTENTS WITH THE CURRENT STATE OF THE DISK ARRAY

The `omnidbsmis` command can be used to compare persistent data in the SMISDB with the current state of the P6000 EVA storage system, as retrieved by the Data Protector HPE P6000 / HPE 3PAR SMI-S Agent, and list the differences. The `omnidbsmis` options used for the comparison, which should be used together with the `-sync_check` option, are: `-host`, `-session`, and `-datalist`.

CAUTION

In specific circumstances, the comparison triggered by `omnidbsmis -sync_check` may give incorrect information. Before taking any actions based on the comparison results, you should therefore double-check if the results reflect the actual P6000 EVA storage system state.

EXCLUDING REPLICAS FROM USE AND BRINGING REPLICAS BACK INTO USE (INCLUDING REPLICAS)

The `omnidbsmis` command can be used to exclude a replica that was created in the ZDB session identified by the session ID from use (replica set rotation, instant recovery capability, possibility to delete its session from the SMISDB) or bring it back into use (include it).

The `omnidbsmis` options to be used for excluding or including replicas are: `-exclude`, `-include`, and `-session`.

OPTIONS

`-version`

Displays the version of the `omnidbsmis` command.

`-help`

Displays the usage synopsis for the `omnidbsmis` command.

`-ompasswd -add` *ClientName*

Stores the login information for the system with the name *ClientName*, on which the SMI-S P6000 EVA Array provider is installed, in the SMISDB.

The `-ssl` option specifies that HPE SMI-S P6000 EVA Array provider is SSL-enabled. In this case, the HPE P6000 / HPE 3PAR SMI-S Agent uses an SSL-based client connection to communicate with the SMI-S P6000 EVA Array provider.

The `-port`*PortNumber* option specifies the port number on which SMI-S P6000 EVA Array provider listens to requests. The default port number for SMI-S P6000 EVA Array provider is 5988 (the `-ssl` option is not selected) or 5989 (the `-ssl` option is selected). If your SMI-S P6000 EVA Array provider is configured to use a different port number, set it using this option.

The `-user`*Username* option sets the user of SMI-S P6000 EVA Array provider. The default user is `administrator`. Names of user accounts that are part of a domain must be specified in the format *Username@Domain*.

The `-passwd`*Password* option sets the password that will be used for logging in to SMI-S P6000 EVA Array provider. If you omit this option, the command will ask for a password interactively.

`-ompasswd {-remove` *ClientName* `| -delete` *ClientName*`}`

This option removes the system with the SMI-S P6000 EVA Array provider installed, specified by *ClientName*, from the SMISDB. The login and port number information is also removed. The option `-delete` is an alias for the option `-remove`.

Used together with the `-port` *PortNumber* option, the command will only remove the entries for the specified port. Use this option if you have more than one port configured on the same system, and you want to delete only one port from the configuration.

If the `-user` *Username* option is specified, the command will only remove the entries for the specified user. Use this option if you have more than one user configured on the same system, and you want to delete only one user from the configuration.

`-ompasswd -list` *ClientName*

Lists all systems that have SMI-S P6000 EVA Array provider installed, together with the port numbers, on which SMI-S P6000 EVA Array providers listen to requests. The *ClientName* value is optional: if you enter a name of the host, only the SMI-S EVA CIMOMs, configured for a specified host, will be displayed.

Note that you will get the same output if you execute the `omnidbsmis -ompasswd` command without the `-list` option.

`-ompasswd -check` [`-host` *ClientName*]

Checks if the SMI-S EVA CIMOMs were configured properly in the Data Protector cell. It performs a health check of your environment, which may help identify such potential problems as wrong user name or password provided, a broken network connection, a DNS resolution problem, and so on. The `-host` option is optional: if you specify the name of a host, the command will be run on the specified host, otherwise it will be run on the local host. Note that HPE P6000 / HPE 3PAR SMI-S Agent must be installed on the specified host.

-dgrules -init

Creates a template for P6000 EVA disk group pairs configuration file or overwrites an existing configuration file with the template. Note that only rules for configured disk group pairs are overwritten.

-dgrules -put *FileName*

Sets the configuration file for P6000 EVA disk group pairs by reading the input file, checking syntax of its contents, and uploading the file to the SMISDB. If the syntax is incorrect, the file is not uploaded.

-dgrules -get *FileName*

Prepares the configuration file for P6000 EVA disk group pairs for editing by reading appropriate contents from the SMISDB and saving them to a file *FileName*.

-dgrules-check*EVA_WWNDG_name*

Provides information on the disk group that is in pair with the disk group identified by *EVA_WWN* and *DG_name*. The command returns information on `1st disk group name`, `2nd disk group name`, and the EVA WWN. If there is no rule configured for the specified disk group, the same name is displayed for both disk groups.

-caconf -init

Creates a template P6000 EVA HOME configuration file or overwrites an old one with the new template.

-caconf -put *FileName*

Uploads the edited P6000 EVA HOME configuration file to the SMISDB. If the syntax of the file is inaccurate, the file is not uploaded.

-caconf -get *FileName*

Prepares the P6000 EVA HOME configuration file for editing by reading the contents of the file from the SMISDB and saving it under *FileName*.

-caconf -list *EVA_WWN*

Provides the information on the DR groups with the P6000 EVA Array identified by *EVA_WWN* acting as home. The command returns the information on *EVA_WWW* and the DR groups belonging to this P6000 EVA Array.

-caconf -check *DR_Group_Name*

Checks if a DR group, identified by *DR_Group_Name*, is part of HPE CA+BC P6000 EVA configuration. The command returns the information on the DR group and WWN of a home P6000 EVA Array.

-show -session *SessionID*

Lists expanded details of a session (identified by the backup session ID). The output of the command is the information on all target volumes created in the specified backup session. The following is displayed:

– Name, ID, and WWN of the target volume created in the backup session

– Name and ID of the P6000 EVA storage system on which the target volume was created

– Snapshot type used for the replica (preceded with the string `Mirroclone` if mirrorclones were used as the snapshot source)

– ID of the source volume used in the backup session

– The backup session ID

– Time stamp of the target volume

– The IR flag (determines if the replica can be used for instant recovery: 0 – the replica cannot be used, 1 – the replica can be used)

– The exclusion flag (determines if the created replica was subsequently excluded from use: 0 – the replica was not excluded, 1 – the replica was excluded)

– The source disk version (determines if the source volumes were preserved on the disk array after a corresponding instant recovery session was performed: 0 – the source volumes were not preserved, 1 – the source volumes were preserved)

– Name of the backup specification used in the backup session

– Names of the application and backup systems involved in the backup session

Note that you will get the same output if you execute the `omnidbsmis-session`*SessionID* command without the `-show` option.

`-show -datalist `*DatalistName*

Lists all replicas that are part of the replica set, which is identified by the ZDB backup specification name. Replicas displayed are identified by their ZDB session IDs. Note that you will get the same output if you execute the `omnidbsmis -datalist `*DatalistName* command without the `-show` option.

`-list -session `[`-ir`] [`-excluded`] [`-original`]

Lists all zero downtime backup sessions that have been run in the cell and in which replicas were created on a disk array of the HPE P6000 EVA Disk Array Family. For each such session, the following information is displayed: the session ID, the IR flag, snapshot type used for the replica (with the "mirrorclone" snapshot source denoted by the string (`MC`)), the exclusion flag, and the ZDB backup specification name.

Note that you will get the same output if you execute the `omnidbsmis -session` command without specifying the `-list` option.

Additionally, each successful automatic mirrorclone creation is denoted by a separate entry which is marked as a session for which instant recovery is not possible, with `Mirrorclone` as the snapshot type, and as a session that is excluded from use. The suffix `_MC` is added to its ZDB specification name.

If the option `-ir` is specified, the command only lists the sessions marked for instant recovery (ZDB-to-disk and ZDB-to-disk+tape sessions).

If the option `-excluded` is specified, the command only lists the sessions that are excluded from use. Excluded sessions do not participate in replica set rotation and do not offer a possibility to perform instant recovery.

If the option `-original` is specified, the command only lists the sessions for each of which the original volumes were preserved on the disk array after a corresponding instant recovery session was performed.

`-list -datalist`

Lists all ZDB backup specifications that were used to create replicas which are part of replica sets with existing members and which use a disk array of the HPE P6000 EVA Disk Array Family for replica storage.

Note that you will get the same output if you execute the `omnidbsmis -datalist` command without specifying the `-list` option.

Additionally, each successful automatic mirrorclone creation is denoted by a separate entry which has the suffix `_MC` added to the ZDB backup specification name.

`-list -purge`

Lists all virtual disks (source volumes or target volumes) that are marked for purging in the SMISDB.

`-purge`

Runs HPE P6000 / HPE 3PAR SMI-S Agent to perform the SMISDB purge operation that attempts to remove the virtual disks (source or target volumes) that could not be deleted although they should be. These elements are marked for purging, and the information about them is stored in the SMISDB.

Used together with the `-force` option, the command forces removal of the elements marked for deletion even if they are presented to the hosts.

If the `-host` *ClientName* option is specified, you can choose another location to start the SMISDB purge operation. Use this option if the systems, involved in a backup session, are no longer available, thus allowing redirection to another systems that have the HPE P6000 / HPE 3PAR SMI-S Agent installed.

`-delete -session` *SessionID* [`-reference`] [`-preview`] [`-force`] [`-host` *ClientName*]

Deletes a replica associated with a specific ZDB session identified by the session ID from the disk array, and deletes information about the replica from the SMISDB. These actions can only be performed for sessions which have not been excluded from use.

If the option `-reference` is specified, only information about the replica in the SMISDB is deleted. Use this option to remove entries that point to replicas that no longer exist on the disk array, or to make existing replicas independent from the Data Protector operation.

If the option `-preview` is specified, the `omnidbsmis` command does not delete anything, but lists the replica or replica set that would be deleted if `-preview` was not specified.

If the option `-force` is specified, the deletion actions are performed also for replicas that are presented to some system.

If the option `-host` *ClientName* is specified, it changes the location of the delete actions. Use this option to redirect deletion to another system, in circumstances when the system which was involved in the backup session(s) is no longer available. The specified system must have the Data ProtectorHPE P6000 / HPE 3PAR SMI-S Agent component installed.

The option combination `-delete -session` cannot be used for deletion of mirrorclones.

`-delete -datalist` *DatalistName*

Deletes a replica set associated with all ZDB sessions based on a specific ZDB backup specification from the disk array, and deletes information about the replica set from the SMISDB.

These actions can only be performed for sessions which have not been excluded from use.

To delete mirrorclones that were automatically created in all ZDB sessions based on a specific ZDB backup specification, specify *DatalistName*_MC instead of *DatalistName*. No mirrorclone snapshots should exist on the disk array for this operation to succeed.

Alternatively, to delete a replica set associated with all ZDB sessions based on a specific ZDB backup specification as well as the mirrorclones that were automatically created in these ZDB sessions, specify *DatalistName** instead of *DatalistName*.

CAUTION: The asterisk (*) is a wildcard character. If other ZDB backup specifications have names similar to the chosen ZDB backup specification name, the replica sets and the ZDB sessions based on these specifications may be affected as well.

The meaning of options `-reference`, `-preview`, `-force`, and `-host`*ClientName* is the same as when used in combination with the options `-delete -session`*SessionID*.

`-sync_check [-host `*ClientName*`] [-session `*SessionID*` | -datalist `*DatalistName*`]`

Compares persistent data in the SMISDB with the current state of the P6000 EVA storage system, and lists the differences. Entries which should be purged are also compared. Note that this option does not check whether configuration of the P6000 EVA storage system is correct, it only compares saved data against the actual setup. Before using the results for actual modifications, verify the configuration first.

If the option `-host` is specified, it changes the location of the comparison. Use this option to redirect the comparison to another system, in circumstances when the system which was involved in the backup session(s) is no longer available. The specified system must have the Data Protector HPE P6000 / HPE 3PAR SMI-S Agent component installed.

If the option `-session` is specified, only the entries related to the specified session are checked.

If the option `-datalist` is specified, only the entries related to the specified backup specification are checked.

`{-exclude|-include} -session `*SessionID*

Excludes a replica from use or brings it back into use (includes it). An excluded replica cannot participate in replica set rotation, cannot be used for instant recovery, and information about its ZDB session cannot be deleted from the SMISDB. To involve an excluded replica in replica set rotation, make it available for instant recovery, or enable deletion of information of its ZDB session from the SMISDB, bring it back into use.

EXAMPLES

1. To list all configured management systems together with the port numbers, on which the SMI-S P6000 EVA Array providers listen to requests, execute:

   ```
   omnidbsmis -ompasswd -list
   ```

2. To remove a management system with the hostname "system1", together with its login and port number information, from the SMISDB, execute:

   ```
   omnidbsmis -ompasswd -remove system1
   ```

3. To store the login information for the SMI-S P6000 EVA Array provider, installed and running on

the management system with the hostname "system1", in the SMISDB, execute:

```
omnidbsmis -ompasswd -add system1
```

You can also set optional parameters, such as the port number and username. If you omit these parameters, the command will take the default values.

4. To perform a health check of you environment on the local system, execute:

```
omnidbsmis -ompasswd -check
```

5. To create and set the disk group pairs configuration file or edit it, folow the steps below on the application system or the backup system:

a) To create a template disk group pairs configuration file or overwrite an old one with the template, execute:

```
omnidbsmis -dgrules -init
```

b) To get the file for editing and to save it as "c:\tmp\dgrules.txt", execute:

```
omnidbsmis -dgrules -get c:\tmp\dgrules.txt
```

The command reads the disk group pairs configuration file from the SMISDB and saves it in the "c:\tmp" directory on a local system under "dgrules.txt".

c) Edit the "dgrules.txt" file residing in the "c:\tmp" directory and save it. Note that the order of defining disk group names is ignored. This means that if a source volume is found in "disk group 1", its snapclone will be created in "disk group 2", and the other way round. Note that a certain disk group can be a member of only one disk group pair.

d) To upload the "dgrules.txt" file to the server, execute:

```
omnidbsmis -dgrules -put c:\tmp\dgrules.txt
```

The command reads the contents of the file, checks its syntax, and copies the file to its location on the Cell Manager.

6. To get the information on the disk group that is the pair of a disk group named `original_disk_group` configured on the P6000 EVA storage system with the WWN `50001FE15007CA90`, execute:

```
omnidbsmis -dgrules -check EVA1 original_disk_group
```

The following is the output of the command:

```
Configured disk group pair:
1st disk group name:"original_disk_group"
2nd disk group name:"paired_disk_group_name"
P6000 EVA Array Family name:"50001FE15007CA90"
```

If there is no rule for the specified disk group, the first and second disk groups are the same.

7. To create the P6000 EVA HOME configuration file or edit it, follow the steps below on the application system or the backup system:

a) To create a template P6000 EVA HOME configuration file or overwrite an old one with the template, execute:

```
omnidbsmis -caconf -init
```

b) To get the file for editing and to save it as "c:\tmp\cahome.txt", execute:

```
omnidbsmis -caconf -get c:\tmp\cahome.txt
```

The command reads the P6000 EVA HOME configuration file from the SMISDB and saves it in the "c:\tmp" directory on a local system under "cahome.txt".

c) Edit the "cahome.txt" file residing in the "c:\tmp" directory and save it.

d) To copy the "cahome.txt" file to its original place, execute:

```
omnidbsmis -caconf -put c:\tmp\cahome.txt
```

The command reads the contents of the file, checks its syntax, and copies the file back to the SMISDB.

8. To check if a DR group named `DR Group 1` is part of an HPE CA+BC P6000 EVA configuration, execute:

```
omnidbsmis -caconf -check DR_Group_1
```

The command reports the following:

```
DR Group :"DR_Group_1"
```

```
Home P6000 EVA Array :"EVA_WWN"
```

9. To list all existing ZDB sessions, together with their session IDs and ZDB backup specification names, execute:

```
omnidbsmis -session
```

10. To find out the name, ID, WWW, type, and time stamp of the target volumes created in the ZDB session with the session ID "2012/11/8-2", execute:

```
omnidbsmis -session 2012/11/8-2
```

11. To delete the replica created in the ZDB session with the session ID "2012/11/13-3" from the disk array and the associated information from the SMISDB, execute:

```
omnidbsmis -delete -session 2012/11/13-3
```

The command will not remove the mirrorclones that may have been automatically created in the ZDB session.

12. To delete the mirrorclones created in the ZDB sessions based on the ZDB backup specification "ZDB_mirrorclone_disk_C" from the disk array and the associated information from the SMISDB, execute:

```
omnidbsmis -delete -datalist ZDB_mirrorclone_disk_C_MC
```

The operation will only succeed if no snapshots of such mirrorclones exist on the disk array.

13. To delete the replicas created in the ZDB sessions based on the ZDB backup specification "ZDB_mirrorclone_disk_D" from the disk array and the associated information from the SMISDB, including the mirrorclones that were automatically created in these sessions, execute:

```
omnidbsmis -delete -datalist ZDB_mirrorclone_disk_D*
```

CAUTION: The asterisk (*) is a wildcard character. If other ZDB backup specifications have names similar to the specified ZDB backup specification name, the replica sets and the ZDB sessions based on these ZDB backup specifications may be affected by this operation as well.

14. To run a comparison of the SMISDB with the current state of the P6000 EVA storage system from the system "computer", execute:

```
omnidbsmis -sync_check -host computer
```

SEE ALSO

omnidb(1M), omnidbcheck(1M), omnidbinit(1M), omnidbp4000(1), omnidbutil(1M), omnidbvss(1), omnidbxp(1), omnidbzdb(1), omniofflr(1M), upgrade_cm_from_evaa(1M)

omnidbvss(1)

omnidbvss — queries the VSS database (VSSDB); browses, lists, saves, removes, and manages the items of the VSSDB
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnidbvss -version | -help

omnidbvss -init

omnidbvss -list session [-barlist *BackupSpecName*]

omnidbvss -list session_persistent [-barlist *BackupSpecName*] [-older_than *YYYY/MM/DD*]

omnidbvss -get session {*SessionKey* [*SessionKey* ...] | -barlist *BackupSpecName* | -all} [-detail] [-save_metadata *Directory*]

omnidbvss -get session_persistent {*SessionKey* [*SessionKey*...] | -barlist *BackupSpecName* [-older_than *YYYY/MM/DD*] | -all [-older_than *YYYY/MM/DD*]} [-save_metadata *Directory*]

omnidbvss -remove session {*SessionKey* [*SessionKey*...] | -barlist *BackupSpecName* | -all} -force -reference

omnidbvss -remove session_persistent {*SessionKey* [*SessionKey*...] | -barlist *BackupSpecName* -older_than *YYYY/MM/DD* | -all -older_than *YYYY/MM/DD* }

omnidbvss -disable session {*SessionKey* [*SessionKey* ...] | -barlist *BackupSpecName* | -all} [-force [-backhost *AlternativeBackupSystem*]]

omnidbvss -enable session {*SessionKey* [*SessionKey* ...] | -barlist *BackupSpecName* | -all} -backhost *BackupSystem* -mnt_target *MountPoint* [-mnt_sessionid_apphostname | -mnt_sessionid | -mnt_apphostname_sessionid | -mnt_apphostname | -mnt_direct] [-mnt_readwrite]

omnidbvss -resolve {-apphost *ApplicationSystem* | -all}

omnidbvss -resolve -session *SessionID*

SessionKey = *SessionID* [:*ClientName*]

DESCRIPTION

The `omnidbvss` command is used to query the VSSDB.

This command can be used to:

- list all available backup sessions (ZDB-to-disk, ZDB-to-disk+tape, and ZDB-to-tape)
- view information about a specific or all available backup sessions
- view details about a specific or all available ZDB-to-disk and ZDB-to-disk+tape sessions
- save backup components and writer metadata documents

- remove a specific or all available ZDB-to-disk and ZDB-to-disk+tape sessions, together with theirs replicas, from the VSSDB and from the disk array

- remove a reference to a specific or to all available backup sessions from the VSSDB

- disables the specified or all ZDB-to-disk or ZDB-to-disk+tape sessions

- enables the specified or all ZDB-to-disk, ZDB-to-disk+tape sessions

- initialize the VSSDB

- resolve the application systems in the Data Protector VSS integration cell.

OPTIONS

-version

Displays the version of the omnidbvss command.

-help

Displays the usage synopsis for the omnidbvss command.

-init

Initializes the VSSDB.

IMPORTANT: All data including sessions and created replicas is deleted from the VSSDB.

-list session [-barlist *barlist*]

Queries the VSSDB and lists all ZDB-to-disk and ZDB-to-disk+tape session IDs. If -barlist is specified, only the IDs of the ZDB-to-disk and ZDB-to-disk+tape sessions that were created using the backup specification are listed.

-list session_persistent [-barlist *barlist*] [-older_than *YYYY/MM/DD*]

Queries the VSSDB and lists all available backup session (ZDB to disk, ZDB to disk+tape, and ZDB to tape) IDs.

If -barlist is specified, only the IDs of the sessions that were created using the backup specification are listed.

If -older_than is specified, only the sessions IDs that were created before the specified date are listed.

-get session {*SessionKey* [*SessionKey...*] | -barlist *BackupSpecName* | -all}

[-detail] [-save_metadata *Directory*]

Displays information about the ZDB-to-disk and ZDB-to-disk+tape sessions.

By specifying *SessionKey*, the -barlist, or the -all option, information about the backup components and disks about the sessions that match the given criteria will be displayed.

–detail displays detailed information (components, disks) about the specified session.

-save_metadata saves the backup components document (Backup Components Document.xml) and writer metadata document (*writer_name*.xml) to the specified directory.

-get session_persistent {*SessionKey* [*SessionKey...*] | -barlist *BackupSpecName*

[-older_than *YYYY/MM/DD*] | -all [-older_than *YYYY/MM/DD*]} [-save_metadata *Directory*]

Displays information about any backup session created using VSS software or the hardware provider.

By specifying *SessionKey*, the -barlist, or the -all option, information about the sessions that match the given criteria will be displayed.

-older_than displays information about the backup sessions, specified with the -barlist option, or all sessions that were created before the specified date.

-save_metadata saves the backup components document (Backup Components Document.xml) and writer metadata document (*writer_name*.xml) to the specified *directory*.

-remove session {*SessionKey* [*SessionKey*...] | -barlist *BackupSpecName* | -all}

[-force] [-reference]

Removes the specified ZDB-to-disk or ZDB-to-disk+tape sessions and their replicas from the VSSDB (non-persistent metadata) and disk array.

By specifying the *SessionKey*, -barlist, or -all option, the information about the sessions that match the given criteria will be deleted from the VSSDB and the session's replicas will be deleted from the disk array.

If -reference is specified, only the reference information about the specified sessions and their replicas will be removed from the database. This option can be used to remove an entry that points to a replica that no longer exists.

The removing operation fails in the following cases, unless the -force option is used:

— If you have changed the disks' configuration manually after the backup (for example, the sessions target volumes were presented manually to some other system).

— If the target volumes cannot be dismounted because of a lock by some other process.

-remove session_persistent {*SessionKey* [*SessionKey*...] | -barlist

BackupSpecName [-older_than *YYYY/MM/DD*] | -all [-older_than *YYYY/MM/DD*]}

Removes the reference information about the specified sessions from the VSSDB (persistent metadata). It does not remove the session's replicas from the disk array.

By specifying the *SessionKey*, -barlist, or -all option, information about the sessions that match the given criteria will be removed.

-older_than removes the reference information about the backup sessions, specified by the -barlist option, or all sessions, that were created before the specified date.

-disable session {*SessionKey* [*SessionKey*...] | -barlist *BackupSpecName* | -all}

[-force [-backhost *AlternativeBackupSystem*]]

Disables the specified ZDB-disk or ZDB-to-disk+tape sessions (if *SessionKey* is used), sessions created by the specified backup specification (if -barlist is used), or all sessions (if -all is used). Disabling means that the replicas (target volumes) created in the specified sessions or using the specified backup specification are dismounted and unpresented from the backup system.

The disabling operation fails in the following cases, unless the -force option is used:

— If you have changed the disks' configuration manually after the backup (for example, the sessions target volumes were presented manually to some other system).

— If the target volumes cannot be dismounted because of a lock by some other process.

Use `-force -backhost` *AlternativeBackupSystem* if the backup system from which you want to disable a backup session is not available. *AlternativeBackupSystem* specifies an alternative client system (any client in the Data Protector cell hat has the VSS integration component installed), from which the session will be disabled by force.

`-enable session` {*SessionKey* [*SessionKey*...] | `-barlist` *BackupSpecName* | `-all`}

`-backhost` *BackupSystem* `-mnt_target` *MountPoint* [`-mnt_sessionid_apphostname` | `-mnt_sessionid` | `-mnt_apphostname_sessionid` | `-mnt_apphostname`] [`-mnt_direct`] [`-mnt_readwrite`] [`-force`]

Enables the specified ZDB-disk or ZDB-to-disk+tape sessions (if *SessionKey* is used), or sessions created by the specified backup specification (if `-barlist` is used), or all sessions (if `-all` is used). Enabling means that the replicas (target volumes) created in the specified sessions or using the specified backup specification are presented and mounted to the specified backup system.

`-backhost` specifies the target client system where you want the target volumes to be presented.

`-mnt_target` specifies the directory on the *BackupSystem* where you want the target volumes to be mounted. By default, a new directory with the session ID is created in the specified directory and the disks are mounted there. If `-mnt_direct` is used, the disks are mounted to the specified directory. Use `-mnt_direct` only when mounting disks from only one backup session.

You can select the suffix of the mount directory by selecting one of the following options:

`-mnt_sessionid_apphostname` The name of the application system follows the session ID.

`-mnt_sessionid` Only the sessionID is used.

`-mnt_apphostname_sessionid` The sessionID follows the application system name.

`-mnt_apphostname` Only the application system name is used.

The enabling operation fails in the following cases, unless the `-force` option is used:

— When `-mnt_direct` is used and another disk is already mounted in the specified directory.

— If the session to be enabled is already enabled on another backup system.

— If you have changed the disks' configuration manually after the backup.

If you use the `-force` option to enable disks on your specified system even if they are already specified on some other system, note that the disks will not be dismounted on the other system and you will need to clean the environment manually.

By default, the disks are mounted in read-only mode. If `-mnt_readwrite` is specified, the disks will be mounted in read/write mode.

`-resolve` {`-apphost` *ApplicationSystem* | `-all`}

Resolves the specified application system (if `-apphost` is used) or all application systems (if `-all` is used) in the Data Protector cell.

The command applies only to instant recovery-enabled backup sessions and must be run always after:

— installing or upgrading Data Protector

— your source volumes configuration on the application system has changed (for example, you have modified the existing source volumes or you have presented new source volumes)

— you have added a new storage object (for example, a Microsoft Exchange Server storage group)

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

`-resolve –session `*`SessionID`*

Resolves the target volumes created in the specified backup session.

The command applies only to instant recovery-enabled backup sessions and must be run always after:

— installing or upgrading Data Protector

— your source volumes configuration on the application system has changed (for example, you have modified the existing source volumes or you have presented new source volumes)

— you have added a new storage object (for example, a Microsoft Exchange Server storage group)

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

EXAMPLES

1. To list the instant recovery–enabled sessions (ZDB to disk or ZDB to disk + tape) created using the backup specification "Backup1", execute:

   ```
   omnidbvss -list session -barlist Backup1
   ```

2. To get information about the backup components of all backup sessions created before February 1st, 2013, and to save information about the backup components and writer metadata to the directory "C:\metadata", execute:

   ```
   omnidbvss -get session_persistent -all older_than 2013/02/01 -save_metadata
   C:\metadata
   ```

   Note that the specified directory must exist before you execute the command.

3. To remove the reference information about the sessions "2013/02/11-1" and "2013/02/11-2" from the VSSDB and to remove the associated replicas from the disk array, execute:

   ```
   omnidbvss -remove session 2013/02/11-1 2013/02/11-2
   ```

4. To mount the target volumes from the session "2013/02/15-1" in the directory "C:\mnt\", present them on the client system "backupsys", and to leave the volumes mounted in read/write mode, execute:

   ```
   omnidbvss -enable session 2013/02/15-1 -backhost backupsys -mnttarget C:\mnt -
   readwrite
   ```

   Note that a new directory with the session ID is created and that the target volumes are mounted in the directory "C:\mnt\2013-02-15-1".

SEE ALSO

omnidb(1M), omnidbcheck(1M), omnidbinit(1M), omnidbp4000(1), omnidbsmis(1), omnidbutil(1M), omnidbxp(1), omnidbzdb(1), omniofflr(1M)

omnidbxp(1)

omnidbxp — queries the ZDB database (XPDB), manipulates the P9000 XP LDEV exclude file, configures the HPE P9000 XP Disk Array Family command devices usage, and manages the user authentication data which the Data Protector HPE P9000 XP Agent uses to connect to specific disk arrays
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnidbxp -version | -help

omnidbxp -exclude {-put *filename* | -get *filename* | -check *SEQ LDEV* | -init | -delete}

omnidbxp [-ir] -session {-list | -show *sessionID*}

omnidbxp [-ir] -ldev {-list | -show *SEQ LDEV*}

omnidbxp -cm {-add *serial* {*CU:LDEV* | *LDEV*} *hostname* [*instance*] | -update *serial* {*CU:LDEV* | *LDEV*} *hostname* [*instance*]}

omnidbxp -cm -remove {all | *serial* [{*CU:LDEV* | *LDEV*} [*hostname*]]}

omnidbxp -cm -list

omnidbxp -user -add *SEQ* -username *Username* [-password *Password*]

omnidbxp -user -check *SEQ* -host *ClientName*

omnidbxp -user -update *SEQ* -username *Username* [-password *Password*]

omnidbxp -user -list *SEQ*

omnidbxp -user -remove *SEQ*

DESCRIPTION

Using the `omnidbxp` command, you can perform various tasks related to the XPDB and your HPE P9000 XP Disk Array Family storage system.

QUERYING THE INFORMATION ON BACKUP OBJECTS AND MANIPULATING THE LDEV EXCLUDE FILE

The `omnidbxp` command can be used to query the information stored in the ZDB database (XPDB), which stores the information about the configured LDEVs pairs (for both S-VOL types: mirror and snapshot) that is used during the Data Protector HPE P9000 XP Disk Array Family backup and restore sessions. The XPDB is a set of plain text files stored on the Cell Manager in the default Data Protector ZDB database directory. The XPDB records contain data about the P-VOL – S-VOL pairs which have been put in the SUSPENDED state by the Data Protector HPE P9000 XP Disk Array Family integration: the mirrors that have been split and the snapshots that have been created on the disk array. The XPDB is written to whenever a mirror is split or a snapshot is created. A pair is deleted from the

XPDB whenever the Data Protector HPE P9000 XP Agent resynchronizes a mirror (if the S-VOL is a mirror copy) or empties the volume that stores snapshot data (if the S-VOL is a snapshot).

The `omnidbxp` command can also be used to manipulate the P9000 XP LDEV exclude file. The P9000 XP LDEV exclude file enables disabling of using certain LDEVs on the backup system (S-VOL LDEVs) by Data Protector. Thus, it is possible to reserve certain LDEVs for other purposes than Data Protector backup and restore. The disabled LDEVs are, if used in a Data Protector session, ignored by Data Protector and such a session fails with critical error. The list of disabled LDEVs is kept in the P9000 XP LDEV exclude file on the Cell Manager: *Data_Protector_program_data*`\server\db80\exclude\XPexclude` (Windows systems) or `/var/opt/omni/server/db80/xpdb/exclude/XPexclude` (UNIX systems).

The `omnidbxp` options to be used for querying the XPDB and manipulating the P9000 XP LDEV exclude file are: `-exclude, -put, -get, -check, -init, -delete, -session, -list, -show, -ir, -ldev, -show`.

HPE P9000 XP Disk Array Family COMMAND DEVICE HANDLING

An HPE P9000 XP Disk Array Family command device is needed by any process that needs access to a disk array of the HPE P9000 XP Disk Array Family. The information about HPE P9000 XP Disk Array Family command devices is kept in the XPDB for the purpose of eliminating duplicate instance usage and overallocation. Data Protector provides the following mechanism to prevent duplicate instance usage and overallocation:

1. Whenever a session is started, Data Protector queries the XPDB for a list of command devices. If none is found in the XPDB (default behavior when the first session is started), Data Protector identifies command devices and generates a list of command devices in the XPDB as connected to every application system and every backup system in the cell.

2. Every command device is assigned an instance number (starting from 301) and the system (hostname) having access to it. If a command device can be accessed from more than one system, the hostname identifier enables Data Protector to be aware of the fact that the command device is already meant to be used by some other system; next available instance number is assigned to such a command device–hostname combination.

3. When the list is created, every disk array of the HPE P9000 XP Disk Array Family which is attached to an application system or a backup system has a list of its command devices and systems having access to them (together with an instance number) assigned.

4. During a session, whenever an application system or a backup system needs access to a P9000 XP Array, it uses the first assigned command device with the instance number from the list. If the command device fails, the next command device from the list assigned to a particular system is used. If all of them fail, the session fails. The successful command device is used by a particular system until the end of the session and the list of command devices is used for all consecutive sessions.

Using the `omnidbxp` command, it is possible to:

1. Specify a particular command device (identified by the serial number of a P9000 XP Array and the LDEV number) to be used by a particular system. Optionally, an instance number can be assigned too. If the instance number is not specified, Data Protector assigns the lowest not yet assigned instance number. The entire information is written in the XPDB.

2. List all command devices in the XPDB.

3. Remove a specific or all command devices from the XPDB or update the information about a specific command device in the XPDB.

The `omnidbxp` option combinations to be used for command device handling begin with the `-cm` option. The options that can follow are: `-add`, `-update`, `-remove`, `-list`.

CONFIGURING THE USER AUTHENTICATION DATA

You can use the `omnidbxp` command to add user credentials of disk array user accounts to the XPDB and to manage stored credentials. For each particular disk array serial number, you can add user credentials of a single user account. The credentials are used by the HPE P9000 XP Agent when it attempts to connect to a disk array through a command device which has the user authentication mode enabled. They must match those configured on the P9000 XP Array. The user credentials are required for the following types of Data Protector sessions:

— zero downtime backup and instant recovery sessions (involving only the Data Protector HPE P9000 XP Agent)

— VSS instant recovery sessions (involving the Data Protector HPE P9000 XP Agent and the Data Protector Microsoft Volume Shadow Copy Service integration)

Before running a particular Data Protector session, you can use the `omnidbxp` command to verify that the HPE P9000 XP Agent can actually connect to the disk array using the preconfigured user credentials from the XPDB.

The `omnidbxp` option combinations to be used for configuring the user authentication data begin with the `-user` option. The options that can follow are: `-add`, `-username`, `-password`, `-check`, `-host`, `-update`, `-list`, `-remove`.

OPTIONS

`-version`

> Displays the version of the `omnidbxp` command

`-help`

> Displays the usage synopsis for the `omnidbxp` command.

`-exclude -put` *filename*

> Sets the list of excluded LDEVs by reading the contents of the *filename*, checking its syntax and if the syntax is correct, copying the file to its position on the Cell Manager. If the syntax is not correct, the file is not copied.

`-exclude -get` *filename*

> Prepares the P9000 XP LDEV exclude file for editing by reading the contents of the P9000 XP LDEV exclude file on the Cell Manager and saving it under the *filename*.

`-exclude -check` *SEQ LDEV*

> Checks whether the specified LDEV, identified by its backup system disk array serial number (*SEQ*) and LDEV number (*LDEV*) is specified in the P9000 XP LDEV exclude file on the Cell Manager. The LDEV number must be specified as the CU#:LDEV in decimal format. If the queried LDEV is specified in the P9000 XP LDEV exclude file, the command returns the string `YES!`. If the queried LDEV is not specified in the P9000 XP LDEV exclude file, the command returns the string `NO!`.

`-exclude -init`

> Overwrites the current P9000 XP LDEV exclude file on the Cell Manager with the template P9000

XP LDEV exclude file.

`-exclude -delete`

Deletes the contents of the P9000 XP LDEV exclude file on the Cell Manager.

`-ir`

Specifies that the current `omnidbxp` command is executed only for the LDEVs pairs marked for the instant recovery in the XPDB. If this option is not specified, the current command is executed for all LDEVs pairs in the XPDB.

`-session -list`

Lists all available sessions in the XPDB.

`-session -show` *sessionID*

Lists all backup system S-VOL LDEVs that were involved in the session with the *sessionID*.

`-ldev -list`

Lists all S-VOL LDEVs in the XPDB together with their corresponding backup session ID.

`-ldev -show` *SEQ LDEV*

Lists all available XPDB information about the S-VOL specified by its *SEQ* and *LDEV* identifiers. The following information is listed: session ID, timestamp (date and time), CRC data, instant recovery flag, the *SEQ* and *LDEV* identifiers and port number of the corresponding primary volume (P-VOL), mirror type, mirror unit (MU) number, fully qualified domain name (FQDN) of the application system name, and FQDN of the backup system.

`-cm -add` *serial* {*CU:LDEV* | *LDEV*} *hostname* [*instance*]

Adds the command device identified by the serial number of a P9000 XP Array (*serial*) and serial number of command device in the hexadecimal or decimal format (*CU:LDEV* or *LDEV*) to the XPDB, and assigns it the hostname of the system accessing it (*hostname*) and optionally the instance number (*instance*). If the instance number is not specified, Data Protector assigns the lowest not yet assigned instance number.

The instance number must be any number in the range between 301 and 399.

The command does not check whether the specified command device or system exist, it only checks if the optional instance number specified is within the correct range and if the command device together with the instance number is not already assigned to be used by some other system. If checks fail, the command fails with an appropriate error message.

`-cm -update` *serial* {*CU:LDEV* | *LDEV*} *hostname* [*instance*]

Updates the XPDB information about the command device identified by the serial number of a P9000 XP Array (*serial*), serial number of command device in the hexadecimal or decimal format (*CU:LDEV* or *LDEV*) and the specified hostname of the system accessing it (*hostname*), by assigning the newly specified instance number (*instance*) to the P9000 XP Array serial number, serial number of command device and hostname combination. If the instance number is not specified, Data Protector assigns the lowest not yet assigned instance number.

The instance number must be any number in the range between 301 and 399.

The command does not check whether the specified command device or system exist, it only checks if the optional instance number specified is within the correct range and if the command

device together with the instance number is not already assigned to be used by some other system. If the checks fail, the command fails with an error message.

`-cm -remove all`

Removes the information about all command devices from the XPDB.

`-cm -remove` *serial* [{*CU:LDEV* | *LDEV*} [*hostname*]]

If only the *serial* argument is specified, the command removes the information about command devices within a specific P9000 XP Array identified by the serial number of this P9000 XP Array (*serial*) from the XPDB.

If the *CU:LDEV* | *LDEV* and optionally *hostname* arguments are specified as well, the command removes the information about the command device identified by the serial number of the P9000 XP Array (*serial*), serial number of command device in the hexadecimal or decimal format (*CU:LDEV* or *LDEV*) and optionally by the hostname of the system (*hostname*).

When removing the information about the command device without specifying the system (*hostname*), the command deletes all entries for the specified command device, regardless of the system(s) assigned to it.

`-cm -list`

Lists all command devices in the XPDB.

`-user -add` *SEQ* `-username` *Username* [`-password` *Password*]

Adds user authentication data for the disk array with the serial number *SEQ* to the XPDB. For each particular disk array serial number, the XPDB can only contain authentication data of a single disk array user account.

If the option `-password` is specified, `omnidbxp` uses the password specified in the command line instead of prompting for it to be entered interactively.

`-user -check` *SEQ* `-host` *ClientName*

Checks if the HPE P9000 XP Agent can connect to the disk array with the serial number *SEQ* from the system *ClientName* using the user authentication data configured for this disk array. *ClientName* can be a name of either an application system or the backup system. This action actually checks if the user name and password configured in the XPDB for this disk array match any of the user accounts that are configured on the disk array. If successful, `omnidbxp` reports the command device and the instance number that were used for the connection.

`-user -update` *SEQ* `-username` *Username* [`-password` *Password*]

Updates the user authentication data for the disk array with the serial number *SEQ* in the XPDB.

If the option `-password` is specified, `omnidbxp` uses the password specified in the command line instead of prompting for it to be entered interactively.

`-user -list` [*SEQ*]

Lists the user authentication records that are stored in the XPDB, in the form of serial number–user name pairs.

If the argument *SEQ* is specified, `omnidbxp` only lists the records that belong to the disk array with this particular serial number.

`-user -remove` *SEQ*

Removes the user authentication data for the disk array with the serial number *SEQ* from the XPDB.

EXAMPLES

1. To set or change the P9000 XP LDEV exclude file:

   a.) Use the following command on the application or backup system:

   `omnidbxp -exclude -get c:\tmp\filename.txt`

   This command reads the P9000 XP LDEV exclude file from the Cell Manager and saves it in the "c:\tmp\filename.txt" file.

   b.) Edit the `c:\tmp\filename.txt` file and save it when you are done editing.

   c.) Use the following command on the application or backup system:

   This command reads the contents of the "c:\tmp\filename.txt", checks its syntax and if the syntax is correct, copies the file to its position on the Cell Manager.

   `omnidbxp -exclude -put c:\tmp\filename.txt`

2. To check whether the LDEV identified by the serial number "12345" and the LDEV number "123" is specified in the P9000 XP LDEV exclude file, execute the following command:

   `omnidbxp -exclude -check 12345 2864`

3. To list all backup system LDEVs, regardless of they being marked for instant recovery or not, that were involved in the session with the sessionID "2013/05/18-22", execute:

   `omnidbxp -session -show 2013/05/18-22`

4. To list all command devices in the XPDB, execute:

   `omnidbxp -cm -list`

5. To add the command device identified by the P9000 XP Array serial number "00035371" and command device serial number "103" to the XPDB and assign it to be used on the "computer.company.com" system by the instance number "303", execute:

   `omnidbxp -cm -add 00035371 103 computer.company.com 303`

6. To remove the information about all command devices from the XPDB, execute:

   `omnidbxp -cm -remove all`

7. To add the user name "data_protector_admin_3" and the password "3drowssap2xelpmoc1ym" as the user authentication data for the disk array with the serial number "80134" to the XPDB, execute:

   `omnidbxp -user -add 80134 -username data_protector_admin_3 -password 3drowssap2xelpmoc1ym`

8. To check if the HPE P9000 XP Agent installed on the system "p9500_bkp_sys.company.com" can connect to the disk array with the serial number "80134" using the user authentication data configured in the XPDB, execute:

   `omnidbxp -user -check 80134 -host p9500_bkp_sys.company.com`

9. To update the user authentication data that is configured in the XPDB for the disk array with the serial number "80134" with the user name "data_protector_admin_5" and a password that you will enter interactively, execute:

   `omnidbxp -user -update 80134 -username data_protector_admin_5`

SEE ALSO

omnidb(1), omnidbcheck(1M), omnidbinit(1M), omnidbp4000(1), omnidbsmis(1), omnidbutil(1M), omnidbvss(1), omnidbzdb(1), omniofflr(1M)

omnidbzdb(1)

omnidbzdb — executes administrative tasks on the HPE 3PAR StoreServe Storage, NetApp Storage, EMC VNX Storage, or EMC VMAX Storage and manages the configuration data which the integration agents use to connect to the CIMOM providers or to the storage systems (this command is available on systems with the Data Protector `User Interface` component installed).

SYNOPSIS

omnidbzdb --version | --help

omnidbzdb --diskarray *ArrayFamily* --ompasswd --add *ClientName* [--ssl] [--port *PortNumber*] [--namespace *Namespace*] [--user *Username*] [--passwd *Password*]

omnidbzdb --diskarray *ArrayFamily* --ompasswd --remove *ClientName* [--port *PortNumber*] [--user *Username*]

omnidbzdb --diskarray *ArrayFamily* --ompasswd [--list [*ClientName*]]

omnidbzdb --diskarray *ArrayFamily* --ompasswd --check [--host *ClientName*]

omnidbzdb --diskarray *ArrayFamily* --list {--session [--ir] [--excluded] [--original] | --datalist}

omnidbzdb --diskarray *ArrayFamily* --show {--session SessionID |-datalist BackupSpecName}

omnidbzdb --list --purge

omnidbzdb --purge [--force] [--host *ClientName*]

omnidbzdb --delete {--session *SessionID* | --datalist *BackupSpecName*} [--reference] [--preview] [--force] [--host *ClientName*]

omnidbzdb --sync_check [--host *ClientName*] [--session *SessionID* | --datalist *BackupSpec*]

omnidbzdb {--exclude | --include} --session *SessionID*

DESCRIPTION

The `omnidbzdb` command enables you to manage configuration data used for connections between a Data Protector HPE 3PAR StoreServ Storage integration agent and the chosen Common Information Model Object Manager (CIMOM) providers, for connection between Data Protector NetApp Storage integration agent and the chosen NetApp storage system, for connection between EMC VNX Storage and the chosen EMC VNX system, or for connection between EMC VMAX Storage and the chosen EMC VMAX system. Such connections must be properly configured before a storage system can be used for zero downtime backup and instant recovery purposes. To integrate with these storage system families, Data Protector uses different integration agents, depending on the operating system running on the application and backup systems:

- **Data Protector HPE 3PAR StoreServ Storage integration:** HPE 3PAR VSS Agent for Windows systems and HPE P6000 / HPE 3PAR SMI-S Agent for Windows, Linux, and HP-UX systems.

- **Data Protector NetApp Storage integration:** NetApp Storage Provider, which is a plug-in to the Data Protector SMI-S Agent

- **Data Protector EMC VNX Storage Family integration:** EMC VNX Storage Provider, which is a plug-in to the Data Protector SMI-S Agent

- **Data Protector EMC VMAX Storage Family integration:** EMC VMAX Storage Provider, which is a plug-in to the Data Protector SMI-S Agent

For an overview, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

Use `omnidbzdb` to configure the connection to the chosen CIMOM provider or NetApp storage system. Once configured, the connection configuration data for the chosen system is stored in a separate configuration file located on the Cell Manager in the directory:

- **Data Protector HPE 3PAR StoreServ Storage integration:**

  Data_Protector_program_data`\server\db80\smisdb\p10000\login` (Windows) and `/var/opt/omni/server/db80/smisdb/p10000/login` (UNIX)

- **Data Protector NetApp Storage integration:**

  Data_Protector_program_data`\server\db80\smisdb\netapp\login` (Windows) and `/var/opt/omni/server/db80/smisdb/netapp/login` (UNIX)

- **Data Protector EMC VNX Storage integration:**

  Data_Protector_program_data`\server\db80\smisdb\emcvnx\login` (Windows) and `/var/opt/omni/server/db80/smisdb/emcvnx/login` (UNIX)

- **Data Protector EMC VMAX Storage integration:**

  Data_Protector_program_data`\server\db80\smisdb\emcvmax\login` (Windows) and `/var/opt/omni/server/db80/smisdb/emcvmax/login` (UNIX)

With `omnidbzdb`, you can also update or remove the connection configuration data, list the contents of the configuration files, and check if the connection to a particular CIMOM provider or storage system can be established. For these purposes, the `omnidbzdb` command provides the basic options `--add`, `--remove`, `--list`, and `--check`. The option `--add` can be used for configuring a new connection and updating the configuration data for an already configured connection.

OPTIONS

`--version`

    Displays the version of the `omnidbzdb` command.

`--help`

    Displays the usage synopsis for the `omnidbzdb` command.

`--diskarray` *ArrayFamily*

    Selects the disk array family on which to perform configuration data management. `omnidbzdb` of the current Data Protector product version supports HPE 3PAR StoreServ Storage, NetApp Storage, EMC VNX Storage, and EMC VMAX Storage. You can select one of them by specifying the corresponding value for the *ArrayFamily* argument (`P10000` or `3PAR` for the HPE 3PAR StoreServ Storage, `NetApp` for the NetApp Storage, `EmcVnx` for the EMC VNX Storage, or `EmcVmax` for the EMC VMAX Storage). In an `omnidbzdb` command line, this option must precede all other options

and option combinations.

`--ompasswd --add` *ClientName* [`--ssl`] [`--port` *PortNumber*] [`--namespace` *Namespace*] [`--user` *Username*] [`--passwd` *Password*]

> Configures or reconfigures the data, which the appropriate Data Protector ZDB integration agent uses to establish connection to a CIMOM provider running on the system *ClientName*, or to a storage residing on this system. For *ClientName* you can specify either fully qualified domain name, host name, or IP address of the system. Host names are automatically expanded to fully qualified domain names before they are stored to the configuration files. If no additional options are specified, `omnidbzdb` configures the connection as non-SSL, using the port number `5988` as a system listening port, and using `administrator` as the user name. In this case, `omnidbzdb` prompts you to enter the password interactively and omits the initial connection check.

> If the option `--ssl` is specified, the connection is configured to use SSL.

> If the option `--port` is specified for a CIMOM provider, the connection is configured to use the port number *PortNumber*. If not specified, the default port number is used: `5988` for connections not using SSL, `5989` for connections using SSL. HPE recommends you use the default port number.

> If the option `--user` is specified, the connection is configured to use the user name specified in *Username*. If not specified, the default user name `administrator` is used. If the option `--password` is specified, the connection is configured to use the password *Password*. If not specified, `omnidbzdb` prompts you to enter the password interactively.

`--ompasswd --remove` *ClientName* [`--port` *PortNumber*] [`--user` *Username*]

> Removes the connection configuration data, which was added with `omnidbzdb`. For *ClientName* you can specify either fully qualified domain name or IP address of the system with running CIMOM providers services, or of the storage system. If you additionally specify the options `--port`, `--user`, or both, only those configuration files are removed where connection values match the specified ones.

`--ompasswd` [`--list` [*ClientName*]]

> Lists all existing connection configuration data for the CIMOM providers or storage systems, which was added with `omnidbzdb`. For each system, the following information is displayed: the user name, the fully qualified domain name or IP address of the system hosting the CIMOM service or the storage, the port number of the listening port, and the indicator whether the connection uses SSL. You can narrow the output by specifying the argument *ClientName*. For *ClientName* you can specify either fully qualified domain name or IP address of the system.

`--ompasswd --check` [`--host` *ClientName*]

> Triggers a check if the configured connections from the Data Protector integration agent to the CIMOM providers or to the storage systems can be established. If the option `--host` is specified, the connections check starts on system *ClientName*, otherwise it starts on one of the systems with the appropriate integration agent. For *ClientName* you can specify either fully qualified domain name or IP address of the system.

`--list` {`-session` [`--ir`] [`--excluded`] [`--original`] | `--datalist`}

> Lists all zero downtime backup sessions running in the cell that matches the specified criteria, or the backup specifications that were used to create replicas. Specify the `--ir` option to list only sessions for which the "Track the replica for instant recovery" option was selected. To list excluded sessions, specify `--excluded` option. Specify the `--original` option to list only the sessions with

the original volumes preserved on the disk array after a corresponding instant recovery session was performed. The `--datalist` option lists all ZDB backup specifications which were used to create the replicas that are part of replica sets with existing members.

`-show {-session `*SessionID*` | -datalist `*BackupSpecName*`}`

When used with `--session` option, the command lists expanded details of a session. When used with `--datalist`, the command lists replicas that are a part of replica set identified by the backup specification name.

`--list --purge`

Lists virtual disks marked for purging.

`--purge [--force] [--host `*ClientName*`]`

Removes virtual disks marked for purge. The `--force` option removes elements marked for deletion even if they are presented to clients. Use the `--host` option to change location to start the SMISDB purge operation.

`--delete {--session `*SessionID*` | --datalist `*BackupSpecName*`} [--reference] [--preview] [--force] [--host `*ClientName*`]`

Deletes information about session or backup specification from the SMISDB. Specify the `--session` option to delete information about the session. Specify `--datalist` to delete replicas associated with the specified backup specification and the linked information from SMISDB. Specify the `--reference` option to only delete information about the replica from the SMISDB. Use this option to remove entries that point to replicas that no longer exist on the disk array, or to make existing replicas independent from the Data Protector operation. The `--preview` option lists the replicas that will be deleted, but does not delete them nor does it delete the information from the SMISDB. Specify the `--force` option to force deletion even if replicas are presented to other hosts. Use the `--host` option to change the location of the deleted actions when the system from the backup session is no longer available.

`--sync_check [--host `*ClientName*`] [--session `*SessionID*` | --datalist `*BackupSpec*`]`

Compares persistent data in SMISDB with the current state of the storage system and lists the differences for all ZDB sessions. In specific circumstances, the comparison output might be incorrect, so double check whether the results reflect the actual storage system state before taking any action based on the comparison results. The `--host` option changes the location of comparison. Use the `--datalist` to check for entries related to the specified backup specification or the `--session` option to lists the differences for the specified session.

`{--exclude | --include} --session `*SessionID*

Excludes or includes a replica for use.

EXAMPLES

The following examples illustrate how the `omnidbzdb` command works.

1. To configure a connection to the CIMOM provider available to the Windows system "cimom_host3" in the local domain, so that the connection uses the default CIMOM service port number "5989", the user name "Cimomuser", and the password "drowssapelpmis" to connect to the CIMOM provider, execute:

```
omnidbzdb --diskarray 3PAR --ompasswd --add cimom_host3 --user Cimomuser --
password drowssapelpmis
```

2. To configure a connection to the NetApp Storage residing on "netapp_box" in the local domain, so that the connection uses the default port number "5989", the user name "NetApp_admin", and the password "netapppwd4" to connect to the NetApp Storage, execute:

```
omnidbzdb --diskarray NetApp --ompasswd --add netapp_box --user NetApp_admin --
password netapppwd4
```

3. To update the configuration of the connection to the CIMOM provider available to the system "cimom_host5.company.com" that uses SSL and the user name "administrator" to connect to the CIMOM provider, so that the Data Protector HPE 3PAR VSS Agent and the Data Protector HPE P6000 / HPE 3PAR SMI-S Agent use the new password "drowssapregnorts" to connect, execute:

```
omnidbzdb --diskarray 3PAR --ompasswd --add cimom_host5.company.com --ssl --
password drowssapregnorts
```

4. To remove configuration data for connections to the NetApp storage systems available to the Windows system with the fully qualified domain name "netapp_storage5.company.com" and for which the user name "backup_operator" is used, execute:

```
omnidbzdb --diskarray NetApp --ompasswd --remove netapp_storage5.company.com --
user backup_operator
```

5. To list connection configuration data for connections to the CIMOM providers available to the system with the IP address "19.105.89.43", execute:

```
omnidbzdb --diskarray 3PAR --ompasswd --list 19.105.89.43
```

6. To verify that the Data Protector NetApp Storage integration agent can connect to the NetApp storage system using the configured user authentication data, execute:

```
omnidbzdb --diskarray NetApp --ompasswd --check
```

7. To list all ZDB sessions where the replica is tracked for instant recovery, execute:

```
omnidbzdb --diskarray 3PAR -session --ir
```

8. To compare the information in the SMISDB with the current state of the HPE 3PAR StoreServe Storage system from the system "Computer", execute:

```
omnidbzdb --diskarray 3PAR --sync_check --host Computer
```

9. To list replicas that are to be deleted from a session 2012/12/1-2, execute:

```
omnidbzdb --diskarray 3PAR --delete --session 2012/12/1-2 -preview
```

10. To exclude a replica from use, execute:

```
omnidbzdb --diskarray 3PAR --exclude --session 2013/12/1-2
```

SEE ALSO

omnidb(1M), omnidbcheck(1M), omnidbinit(1M), omnidbp4000(1), omnidbsmis(1), omnidbutil(1M), omnidbvss(1), omnidbxp(1), omnioffir(1M)

omnidownload(1)

omnidownload — downloads information about a backup device and a library from the Data Protector Internal Database (IDB)

(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnidownload -version|-help

omnidownload -list_devices [-detail]

omnidownload -dev_info

omnidownload -device *BackupDevice* [-file *FileName*]

omnidownload -list_libraries [-detail]

omnidownload -library *Library* [-file *FileName*]

DESCRIPTION

Allows the user to display information about backup devices or download the configuration of the specified backup device to an ASCII file. Used together with the `omniupload` utility, this command enables you to create and maintain backup devices using the command-line interface.

OPTIONS

-version

Displays the version of the `omnidownload` command.

-help

Displays the usage synopsis for the `omnidownload` command.

-device *BackupDevice*

Specifies the name of the backup device you want to download to an ASCII file.

-library *Library*

Specifies the name of the library you want to download to an ASCII file.

-file *FileName*

Specifies the name of the target ASCII file for the backup device. By default, the file is created in the local directory. If this option is omitted, the data is sent to the standard output.

-list_devices

Displays information about the Data Protector backup devices. The report includes the following information for each device: device name, client, device type and pool.

`-dev_info`

> Same as `-list_devices` option. Used only for compatibility with old Data Protector releases.

`-list_libraries`

> Displays information about the Data Protector libraries. The report includes the following information for each device: library name, client and library type.

`-detail`

> This option can be used in combination with the `-list_devices` and `-list_libraries` options to display more detailed information about the Data Protector backup devices or libraries.

EXAMPLES

The following examples illustrate how the `omnidownload` command works.

1. To download backup device "DAT1" into file "/tmp/DAT1", execute:

   `omnidownload -device DAT1 -file /tmp/DAT1`

2. To review the information about a virtual tape library named "VTL" in ASCII format that will be saved as the file "libVTL.txt" to the directory "C:\Temp", execute:

   `omnidownload -library VTL -file C:\Temp\libVTL.txt`

SEE ALSO

omniamo(1), omnib2dinfo(1M), omnimcopy(1), omniminit(1), omnimlist(1), omnimm(1), omnimnt(1), omnimver(1), omniupload(1), sanconf(1M), uma(1M)

omniiso(1)

omniiso — primarily serves as a pre-exec script to prepare the ISO image file for One Button Disaster Recovery (OBDR); can also be used as a standalone command to automate your backup and disaster recovery process
(this command is available on systems with the Data Protector `Automatic Disaster Recovery` component installed)

SYNOPSIS

`omniiso -version`|`-help`

`omniiso` [`-session` *FSSessionID* [*IDBSessionID*]] [`-host` *ClientName* [`-remotehost` *ClientName*]] [`-cd`|`-net`] [`-out` *ISOImagePath*] [`-srd` *SRDPath*] [`-rset` *P1SPath ImgPath*] [`-autoinject`] [`-waik` *WAIKPath*] [`-inject_drivers` *DriverPath_1 DriverPath_2*...] [`-use_raw_object`] [`-move_to` *Path*] [`-unique_name`] [`-exec_script` *ScriptFilePath*] [`-password` [*Passwd*]] [`-anyobj`]

DESCRIPTION

The `omniiso` command can be used as a:

STANDALONE COMMAND

Although all functionality of the command is also available through the Disaster Recovery Wizard in the Data Protector GUI, it can also be used as a standalone command to automate your backup and disaster recovery process.

The command merges

- the recovery set (the data required for temporary DR OS installation and configuration that is created during a full client backup),
- the SRD file (a file that contains all required backup and restore object information to perform the restore),
- and the P1S file (a file that contains information on how to format and partition all disks installed in the system)

with disaster recovery installation into a disaster recovery ISO image or creates a network recovery image and saves the created image to a file on disk. By default, the DR OS image files are created in the Data Protector temporary files directory and are used to perform disaster recovery.

Such DR OS image can also be created using the OBDR Wizard in the Data Protector GUI instead of this command (recommended).

PRE-EXEC SCRIPT

If the command is used as a pre-exec script in the OBDR Wizard in the Data Protector GUI to prepare the disaster recovery ISO image, you do not have to specify parameters as their values are automatically obtained from the environment.

> **Tip:** You cannot use `omniiso` in a pre-exec or post-exec script to create ISO image on the Cell

Manager because the IDB is backed up in a separate session.

OPTIONS

`-version`

Displays the version of the `omniiso` command.

`-help`

Displays the usage synopsis for the `omniiso` command.

`-session` *FSSessionID* [*IDBSessionID*]

Specifies IDs of the backup sessions that serve as the basis for updating the DR OS image file. All object backed up in the specified sessions and included in the SRD file are used for the update.

If you are updating the DR OS image file for a Data Protector client, specify the *FSSessionID* argument for the most recent full or incremental filesystem backup session that involves the entire client. If you are updating the DR OS image file for the Data Protector Cell Manager, additionally specify the *IDBSessionID* argument for an appropriate full or incremental Data Protector Internal Database backup session.

CAUTION: The specified Data Protector Internal Database backup session must be a session that was run after the specified filesystem backup session had completed. To ensure the highest consistency of the included data, the time frame between both sessions' start times should be minimal.

`-host` *ClientName*

Specifies the client system for which the DR OS image is created. If not specified, the local system (the system on which the command is executed) is used.

`-remotehost` *ClientName*

Specifies the client system where the DR OS image is created. If not specified, the system specified with `-host` is used.

`-cd`

If this option is specified, `omniiso` creates an ISO file that can be written to a CD-ROM. If this option is not specified, the command creates disaster recovery ISO file to be written on a backup medium.

`-net`

If this option is specified, `omniiso` creates a network recovery image file that can be then used to boot the target system over the network. If this option is not specified, the command creates disaster recovery ISO file to be written on a backup medium.

`-out` *ISOPath*

Specifies the location where the DR OS image is created. If this option is not specified, the DR OS image file is created in the Data Protector temporary files directory.

`-srd` *SRDPath*

Specifies the path to the SRD file. If the `-srd` option is not specified, the command creates a SRD

file on the system, where `omniiso` is running and uses it to create the disaster recovery ISO image. If the `-remotehost` option is specified, the SRD file is created on the remote client.

`-rset` *P1SPath ImgPath*

Specifies the full path to the P1S file and the recovery set. If this option is not specified, the command creates the P1S file and the recovery set on the system, where `omniiso` is running and uses them to create the disaster recovery ISO image. If the `-remotehost` option is specified, the P1S and recovery set parameters specify the path on the remote client.

`-autoinject`

Automatically injects drivers into the DR OS image.

This option is available only for Windows Vista and later releases.

`-waik` *WAIKPath*

Specifies the Windows Automated Installation Kit (WAIK) or Assessment and Deployment Kit (ADK) installation directory. If the `-remotehost` option is specified, the path is searched on the remote client.

This option is available only for Windows Vista and later releases.

`-inject_drivers` *DriverPath_1 DriverPath_2 ...*

Injects additional drivers into the DR OS image. You must specify a full path to the driver. A maximum of 50 paths can be specified. If the `-remotehost` option is specified, the paths are searched on the remote client.

This option is available only for Windows Vista and later releases.

`-use_raw_object`

If the specified backup session contains both filesystem and disk image backup objects for the same volume, this option specifies that a disk image backup object should be used. If this option is not specified, filesystem backup objects have a higher priority. If only one backup object for the same volume is present in the specified backup session, this option is ignored.

`-anyobj`

Enables you to create a recovery image even if the specified backup session does not contain all client volumes. Note that all host critical volumes must be part of the specified backup session:

- the boot and system volumes
- the Data Protector installation volume
- the volume where the CONFIGURATION object is located
- the Active Directory database volume (in case of an Active Directory controller)
- the quorum volume (in case of a Microsoft Cluster)

`-password` [*Passwd*]

Specifies the password that is used during the creation of the recovery media. By using a password you can prevent unauthorized use of the recovery media after boot. If you only specify the `-password` option without a password, the command will prompt you to provide one at the start of the image creation process.

`-move_to` *Path*

Moves the DR OS image file to the specified location.

`-unique_name`

Renames the ISO image file to a unique name consisting of the platform type (Windows, Linux), followed by the client name, platform type (amd64, ia64), the date and time when the image was created, and the system BIOS UUID. All name components are separated by the '#' character. For example:

`windows#computer.company.com#amd64#2013-04-25-09-03#844D978B-1D69-BC7D-EB0D-3B93628059A1.iso`

`-exec_script ` *ScriptFilePath*

Executes the specified script after the DR OS file is created. The executed script receives the full path of the DR OS recovery image. You can use scripts to automate various image post-processing tasks.

> **Note**: Some options, like `move_to Path`, `-exec_Script ScriptFilePath`, and `-unique_name`, are not applicable to remote ISO generation, but applicable to only local ISO generation.

NOTES

- The `omniiso` command is available on Windows and Linux systems only.
- If the BTRFS volume is detected, you get the following **Warning** message:

  Warning: BTRFS volume detected. Make sure that you have included all the BTRFS sub volumes in the specified version.

EXAMPLES

The following examples illustrate how the `omniiso` command works.

1. To create and save a disaster recovery ISO file for a Data Protector client (Windows Server 2003) in the CD-ROM-ready format at "C:\iso\dr\omnidr.iso" on the local system, containing objects backed up in the session with the session ID "2013/05/16-23", using information stored in the SRD and P1S files stored in "C:\iso\dr\srd\machine101.company.com" and "C:\iso\dr\p1s\machine101.company.com", using the recovery set stored in "C:\iso\dr\img\machine101.company.com.img", execute:

   `omniiso -session 2013/05/16-23 -cd`

   `-out c:\iso\dr\omnidr.iso`

   `-srd C:\iso\dr\srd\machine101.company.com`

   `-rset C:\iso\dr\p1s\machine101.company.com`

   `C:\iso\dr\img\machine101.company.com.img`

2. To create and save a disaster recovery ISO file for a Data Protector Windows client (Windows Vista and later releases) in the CD-ROM-ready format at "C:\iso\dr\omnidr.iso" on the local system, containing objects backed up in the session with the session ID "2013/05/22-23", using information stored in the SRD and P1S files stored in "C:\iso\dr\srd\machine102.company.com" and "C:\iso\dr\p1s\machine102.company.com", using the recovery set stored in "C:\iso\dr\img\machine102.company.com.img" where the drivers are automatically injected,

execute:

```
omniiso -session 2013/05/22-23 -cd

-out C:\iso\dr\omnidr.iso

-srd C:\iso\dr\srd\machine102.company.com

-rset C:\iso\dr\p1s\machine102.company.com

C:\iso\dr\img\machine102.company.com.img -autoinject
```

3. To create and save a disaster recovery ISO file for a Data Protector Linux client in the CD-ROM-ready format at "/data/iso/dr/omnidr.iso" on the local system, containing objects backed up in the session with the session ID "2013/04/12-35", using information stored in the SRD and P1S files stored in "/etc/opt/omni/server/dr/srd/machine106.company.com" and "/etc/opt/omni/server/dr/p1s/machine106.company.com", using the recovery set stored in "/etc/opt/omni/server/dr/p1s/machine106.company.com.img", execute:

```
omniiso -session 2013/04/12-35 -cd

-out /tmp/omnidr.iso

-srd /etc/opt/omni/server/dr/srd/machine106.company.com

-rset /etc/opt/omni/server/dr/p1s/machine106.company.com

/etc/opt/omni/server/dr/p1s/machine106.company.com.img
```

4. To create and save a disaster recovery network image for the Data Protector Cell Manager system "machine202.company.com" (Windows Vista and later releases) at "C:\iso\dr\omnidr.iso", containing objects backed up in the sessions with the session IDs "2013/04/12-43" and "2013/04/12-44", using information stored in the SRD and P1S files stored in "C:\iso\dr\srd\machine102.company.com" and "C:\iso\dr\p1s\machine102.company.com", using the recovery set stored in "C:\iso\dr\img\machine102.company.com.img" where the drivers are automatically injected, and with the image protected by a password which must be provided from the command prompt, execute:

```
omniiso -session 2013/04/12-43 2013/04/12-44 -net

-host machine202.company.com

-out C:\iso\dr\omnidr.iso

-srd C:\iso\dr\srd\machine102.company.com

-rset C:\iso\dr\p1s\machine102.company.com

C:\iso\dr\img\machine102.company.com.img -autoinject -password
```

SEE ALSO

omnidr(1M), omniofflr(1M), omnisrdupdate(1M), omniusb(1)

omnimcopy(1)

omnimcopy — makes a copy of a Data Protector medium using Data Protector backup devices as the source and destination
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnimcopy -version|-help

omnimcopy -copy *BackupDevice* [-slot *Slot*...] -from *BackupDevice* [-src_slot *Slot*...]
[*BasicOptions*] [*LabelOptions*]

BasicOptions

 -pool *PoolName*

 -location *Location*

 -force

 -size *SpecSize*

 -encrypt

 -eject

 -permanent|-until *Date*

Date = [YY]YY/MM/DD (1969 < [YY]YY < 2038)

LabelOptions

 -label *UserLabel* [-no_barcode_as_label]|-autolabel|-[no_]barcode_as_label

DESCRIPTION

The `omnimcopy` copies a Data Protector medium. It reads data from the input medium and writes the data to the output medium. Note that the output medium is first initialized. During initialization, a medium is assigned a:

- Data Protector Medium Label: Depending on the selected options, the media labels can be user defined or generated automatically. By default, Data Protector automatically generates media labels from the media pool names, unless the `Use barcode as media label during initialization` option is selected in the library properties. This behavior can be changed during the initialization of media using `-barcode_as_label`, `-no_barcode_as_label` and `MediumLabel` options.

- Medium ID (system-assigned)

- Location

The physical devices used for the input and output must be the same device type and have the same block size. This copy functionality allows the user to use multiple tapes in order to implement vaulting with Data Protector. This copy function is a separate function within Data Protector and cannot be done

automatically during backup. Main advantage of this implementation is that all devices can be used during backup (better performance).

The source and destination devices are backup devices which means they can be located everywhere in the Data Protector cell. During the copy the destination tape will be initialized before all data from the source tape is copied.

The writing destination tape will ignore the early end of tape mark and will write until the physical EOT is reached. If the space on the destination is not sufficient to keep the whole original tape the copy has to be restarted with a new tape.

After a copy operation both media are tracked in the media management database.

This enables also a listing of the copies for an original media as well as the listing of the original tape for a copy. If a mount request is issued during a restore session all tapes which contains the data will be listed (original and copies).

After the operation copy both tapes become nonappendable.

A copy of a copy is not possible.

If the original media get obsolete in the database, which means it is overwritten or it is exported from the cell, the first copy becomes automatically the original tape.

OPTIONS

-version

Displays the version of the `omnimcopy` command

-help

Displays the usage synopsis for the `omnimcopy` command

-copy *BackupDevice* [-slot *Slot*...]

Specifies the output backup device - the device used to create a copy of the medium (target medium).

-from *BackupDevice* [-src_slot *Slot*...]

Specifies the input backup device — the device which is used as a source. You can specify only one slot. The `-src_slot` parameter takes the barcode value of the source tape.

-pool *PoolName*

Specify the poolname to which the copy of the medium is added. By default the medium is added to the source media poolname.

-location *Location*

Specifies the location of the media, when you keep them out of the library. Used for the vaulting purposes.

-force

Overwrites the data on the target medium even if this data is still protected by the Data Protector media management system. Note that this option must be used with an unprotected medium as well.

-size *SpecSize*

This option specifies the size of the target medium.

-encrypt

This option turns on hardware encryption on all destination drives.

-eject

Ejects the target medium from the drive after the medium is copied.

-permanent

This backup protection option provides permanent protection of backup media. This means that the data is permanently protected from being overwritten.

-until *Date*

This backup protection option provides protection until a date of your choice. This means that the data on the medium cannot be overwritten until the specified date. Protection for the data stops at noon on the selected day.

-label *UserLabel*

Manually specify the medium label for the copy of the medium. A description can have a maximum of 80 characters, including any keyboard character or space. If the Use barcode as medium label on initialization option is selected in the library properties, you have to specify also the -no_barcode_as_label option.

-autolabel

If this option is specified, the medium is labeled automatically by the Data Protector media management system.

-barcode_as_label

Data Protector uses barcode as a medium label during the initialization of the medium instead of generating media labels based on the media pools names. This option is supported only on library devices with enabled barcode support.

-no_barcode_as_label

Data Protector does not use barcodes as a medium label during the initialization of the medium, but generates media labels based on the media pools names. This option can be used to override the Use barcode as medium label on initialization option (if it is selected) in the library properties in the Data Protector GUI.

SEE ALSO

omniamo(1), omnib2dinfo(1M), omnidownload(1), omniminit(1), omnimlist(1), omnimm(1), omnimnt(1), omnimver(1), omniupload(1), sanconf(1M), uma(1M)

omniminit(1)

omniminit — initializes a Data Protector medium
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omniminit -version | -help

omniminit -init *BackupDevice* [*MediumLabel*] [*BasicOptions*] [*SlotOptions*] [-no_barcode_as_
label]

omniminit -init *BackupDevice* [*BasicOptions*] [*SlotOptions*] [-barcode_as_label]

omniminit -init_magazine *BackupDevice* [*MagazineDescription*] [*BasicOptions*]

omniminit -init_mag_medium *BackupDevice MagazineDescription* [*BasicOptions*] [*SlotOptions*]

omniminit -preerase *BackupDevice* [*SlotOptions*] [-eject]

BasicOptions

 -force

 -pool *PoolName*

 -size *n*

 -location *OffLineLoc*

 -wipe_on_init

 -eject

SlotOptions

 -slot *SlotID* [*Side*]

DESCRIPTION

The `omniminit` command initializes a backup medium. During initialization, a medium is assigned a:

- Data Protector Medium Label: Depending on the selected options, the media labels can be user defined or generated automatically. By default, Data Protector automatically generates media labels from the media pool names, unless the `Use barcode as media label during initialization` option is selected in the library properties. This behavior can be changed during the initialization of media using `-barcode_as_label`, `-no_barcode_as_label` and `MediumLabel` options.
- Medium ID (system-assigned)
- Location

This information is added to the Data Protector Internal Database (IDB) and the medium is added to a Data Protector media pool. Medium ID is its unique identifier. The Medium Label does not necessarily

have to be unique, but it is recommended. The medium location is optional, and can be used to define an offline location for the medium.

OPTIONS

`-version`

Displays the version of the `omniminit` command.

`-help`

Displays the usage synopsis for the `omniminit` command.

`-init` *BackupDevice* [*MediumLabel*]

Specifies two items: the name of the *BackupDevice* where the medium is mounted and the *MediumLabel* which is assigned to the medium by Data Protector after initialization. The *MediumLabel* can be up to 32 characters long. Any printable character, including spaces, can be used. The text must be enclosed in quotation marks.

`-init_magazine` *BackupDevice* [*MagazineDescription*]

Specifies two items: the name of the *BackupDevice* where the magazine is mounted and the *MagazineDescription* (optional) which is assigned to the magazine. Note that the *MagazineDescription* must be unique for each magazine. The description is also used for assigning *MediumLabel* to each medium.

`-init_mag_medium` *BackupDevice* *MagazineDescription*

Initializes single medium from magazine. *BackupDevice* specifies the device where the magazine is mounted. *MagazineDescription* must also be specified to identify the magazine. Note that single medium from the magazine can be initialized only if the magazine has been initialized before and therefore has a unique *MagazineDescription*.

`-preerase` *BackupDevice*

Pre-erases the optical disk. Pre-erasing a medium enables backups which are twice as fast. This is because the pre-erase step is removed from the backup process. For best performance, optical disks should be pre-erased before each backup.

`-force`

Overrides the initialization safety checks. By default, a medium containing protected data or being in a non-Data Protector format cannot be initialized.

`-pool` *PoolName*

Specifies the name of the media pool to which this medium will be added. If no *PoolName* is specified, the medium is added to the default pool for the specified backup device.

`-slot` *SlotID* [*Side*]

Specifies the *SlotID* of the exchanger backup device where the medium is mounted. This option is only valid for this backup device type, but must be given for magneto-optical devices. To specify the side of the platter in this slot, use the additional *Side* parameter. Values of *Side* are A or B.

`-size` *n*

Specifies the medium capacity in MB. If not specified, Data Protector uses the standard capacity of

the media class used with the backup device selected for initialization. The size is later used to calculate the free space remaining on the medium. (FreeSpace = Size - SpaceUsed)

-location *OffSiteLoc*

Specifies the location of the medium. This information is useful if media is stored off-site. The location can have maximum 32 characters. Any printable character, including spaces, can be used. The text must be enclosed in quotation marks.

-wipe_on_init

Wipes the data on medium after it has been initialized. This is done by overwriting the data on medium so it is impossible to restore the original data on medium after it has been wiped.

-barcode_as_label

Data Protector uses barcode as a medium label during the initialization of the medium instead of generating media labels based on the media pools names. This option is supported only on library devices with enabled barcode support.

-no_barcode_as_label

Data Protector does not use barcodes as a medium label during the initialization of the medium, but generates media labels based on the media pools names. This option can be used to override the Use barcode as medium label on initializationoption (if it is selected) in the library properties in the Data Protector GUI.

EXAMPLES

The following examples illustrate how the omniminit command works.

1. To initialize slot "4" of backup device "ADIC" with medium label "Label4", in location "Backup Room", execute:

   ```
   omniminit -init ADIC Label4 -slot 4 -location "Backup Room"
   ```

2. To preerase slot "8" side "A" of magneto-optical tape library unit "MO_Changer", execute:

   ```
   omniminit -preerase MO_Changer -slot 8 A
   ```

SEE ALSO

omniamo(1), omnib2dinfo(1M), omnidownload(1), omnimcopy(1), omnimlist(1), omnimm(1), omnimnt(1), omnimver(1), omniupload(1), sanconf(1M), uma(1M)

omnimlist(1)

omnimlist — lists the contents of a Data Protector medium
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnimlist -version | -help

omnimlist -device *BackupDevice* [-slot *SlotID* [*Side*]] [-monitor] [-detail]

omnimlist -device *BackupDevice* [-slot *SlotID* [*Side*]] [-header] [-monitor]

omnimlist -device *BackupDevice* -session [-slot *SlotID* [*Side*]] [-monitor] [-detail]

omnimlist -device *BackupDevice* -session *SessionID* [-slot *SlotID* [*Side*]] [-monitor] [-detail]

omnimlist -device *BackupDevice* -catalog [-slot *SlotID* [*Side*]] [-monitor]

omnimlist -device *BackupDevice* -catalog *DiskAgentID* [-slot *SlotID* [*Side*]] [-monitor]

DESCRIPTION

The `omnimlist` command lists the contents of a Data Protector medium. The command scans the catalog (index) of the medium and shows all objects and sessions on the medium.

The command can also be used to display the Data Protector medium tape header. If used for such purpose, the command reads the first block of the tape and then displays the information.

OPTIONS

-version

> Displays the version of the `omnimlist` command.

-help

> Displays the usage synopsis for the `omnimlist` command.

-device *BackupDevice*

> Specifies the *BackupDevice* where the medium is mounted. If no other option is specified the command lists all sessions and all their objects.

-slot *SlotID* [*Side*]

> Specifies the *SlotID* of the library where the medium is mounted. This option is only valid for this backup device type, but must be given for magneto-optical devices. To specify the side of the platter in this slot, use the additional *Side* parameter. Values of *Side* are A or B.

-session [*SessionID*]

> Displays information about the sessions on the medium. If no *SessionID* is specified, all sessions

are shown. This reports shows for each session: the *SessionID*, Session Type, Session Status. For the user who initiated the session it shows: the UNIX Login, UNIX Group, and ClientName. If a *SessionID* is specified, the objects for that session are shown. The session report shows for each object: the Client, Mountpoint, Object Label, Disk Agent ID and Object Status.

`-catalog` [*DiskAgentID*]

Displays the Detail Catalog for single or multiple objects. The catalog shows file information for all the files included in the backup of the object in that session. The *DiskAgentID* is used to uniquely identify the backup object-session combination. If not specified all found objects are processed.

`-monitor`

Displays information about the Medium (Pool, Medium ID, Medium Label, Location, and Initialization date/time), the Session (Session ID, Owner, Datalist used, and Start date/time), Objects (Type, Start date/time, Backup Mode), and Session (Client, Mountpoint, Object Label, Disk Agent ID, and Object Status).

`-header`

The command first checks if the media header is in Data Protector format and if it is corrupted. If the media header is not in Data Protector format or if it is corrupted, an appropriate message is displayed. Otherwise the following information from the media header is displayed: medium ID, medium label, medium location, initialization date, last access date, last write date, last overwrite date, number of writes, number of overwrites, pool label, device information, device capacity, tape format version, medium ID from original tape (for replicated media only), medium data format type and medium data format subtype. For random access media, date and time information (last access date, last write date and last overwrite date) is updated every time the medium is accessed/written/overwritten. For all other media, header information is not updated except when initializing the medium.

`-detail`

Displays detailed information about the selected query.

NOTES

For the `-header` option, the following limitation applies: the command displays the header information stored on the medium, ignoring possible updates in the Data Protector Internal Database (IDB).

EXAMPLES

The following examples show how the `omnimlist` command works.

1. To list sessions and corresponding Disk Agents from device "DAT2", execute:

   `omnimlist -device DAT2 -monitor`

2. To list sessions on slot "43" side "B" of a magneto-optical tape library unit "MO_Changer", execute:

   `omnimlist -device MO_Changer -slot 43 B -session`

3. To list all Disk Agents for the session "2013/05/13-23" on the device "Exa8500", execute:

   `omnimlist -device Exa8500 -session 2013/05/13-23`

4. To list the catalog for the object-session combination with the DiskAgentID "774226832", from the medium located in slot "7" of device "Herstal2", execute:

```
omnimlist -device Herstal2 -slot 7 -catalog 774226832
```

5. To display media header for the medium in the backup device named "dev_1", execute:

```
omnimlist -device dev_1 -header
```

SEE ALSO

omniamo(1), omnib2dinfo(1M), omnidownload(1), omnimcopy(1), omniminit(1), omnimm(1), omnimnt (1), omnimver(1), omniupload(1), sanconf(1M), uma(1M)

omnimm(1)

omnimm — provides media management for Data Protector
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnimm -version|-help

omnimm -create_pool *PoolName MediaType* [*Policy AgeLimit MaxOverWrites*] [-[no_]alloc_uninit_first] [-[no_]free_pool [*FreePoolName*]] [-[no_]move_free_media]

omnimm -modify_pool *PoolNameNewPoolName* [*Policy AgeLimit MaxOverWrites*] [-[no_]alloc_uninit_first] [-[no_]free_pool [*FreePoolName*]] [-[no_]move_free_media]

omnimm -create_free_pool *PoolName MediaType* [*AgeLimit MaxOverWrites*]

omnimm -modify_free_pool *PoolName NewPoolName* [*AgeLimit MaxOverWrites*]

omnimm -create_mag_pool *PoolName MediaType* [*Policy AgeLimit MaxOverWrites*]

omnimm -modify_mag_pool *PoolName NewPoolName* [*Policy AgeLimit MaxOverWrites*]

omnimm -remove_pool *PoolName*

omnimm -remove_mag_pool *PoolName*

omnimm -show_pools [*PoolName*]

omnimm -move_medium *Medium ToPoolName*

omnimm -move_magazine *MagazineDescription NewPoolName*

omnimm -modify_medium *Medium NewMediumLabel NewLocation*

omnimm -modify_magazine *MagazineDescription NewLocation* [*NewMagazineDescription*]

omnimm -reset_poor_medium *Medium*

omnimm -reset_wp_medium *Medium*

omnimm -list_pool [*PoolName*] [-detail]

omnimm -show_pool_alloc *PoolName*

omnimm -list_scratch_media *PoolName* [-detail]

omnimm -show_repository_alloc *Library PoolName* [-detail]

omnimm -list_media *Medium* [-detail|-encryptioninfo]

omnimm -list_appendable_media *PoolName*

omnimm -list_copy *Medium*

omnimm -media_info *Medium* [-detail|-encryptioninfo]

omnimm -list_magazines_of_pool *PoolName* [-detail]

omnimm -list_media_magazine *MagazineDescription* [-detail]

omnimm -catalog *Medium*

omnimm -check_protection *Medium*

omnimm -recycle *Medium*

omnimm -recycle_magazine *MagazineDescription*

omnimm -export *Medium*

omnimm -export_magazine *MagazineDescription*

omnimm -copy_to_mcf {*Medium1* [*Medium2...*]} [-output_directory *Pathname*]

omnimm -import *LogicalDevice* [-slot *SlotID* [*Side*]] [-no_log | -log_dirs | -log_file | -log] [-pool *PoolName*] [-import_as_original]

omnimm -import_catalog *LogicalDevice* [-slot *SlotID* [*Side*]] [-no_log | -log_dirs | -log_file | -log]

omnimm -import_magazine *LogicalDevice* [*MagazineDescription*] [-slot *SlotID* [*Side*]] [-no_log | -log_dirs | -log_file | -log] [-pool *PoolName*] [-import_as_original]

omnimm -import_from_mcf {*File...*} [[-pool_prefix *PoolPrefix*] | [-no_pool_prefix]] [-[no_] orig_pool] [-import_as_original]

omnimm -disable_lockname *LockName*

omnimm -enable_lockname *LockName*

omnimm -disable_device *DeviceName* [-ignore_lockname]

omnimm -enable_device *DeviceName* [-ignore_lockname]

omnimm -repository *LibraryName*

omnimm -repository_barcode_scan *LibraryName*

omnimm -repository_update *DriveName* [-slot *SlotID* [*Side*]]

omnimm -add_slots *LibraryName* {*Slot...* | *FromSlot-ToSlot...*}

omnimm -remove_slots *LibraryName* {*Slot...* | *FromSlot-ToSlot...*}

omnimm -silo_query *LibraryName* [-range *FromSlot-ToSlot*]

omnimm -silo_enter *LibraryName* -cap *CapID*

omnimm -silo_eject *LibraryNme* {*Volser...* | *FromVolser-ToVolser...*} -cap *CapID* [-location *Location*]

omnimm -enter *LibraryName* {*Slot...* | *FromSlot-ToSlot...*}

omnimm -eject *LibraryName* {*Slot...* | *FromSlot-ToSlot...*} [-location *Location*]

omnimm -group *PoolName MagazineDescription Medium...*

omnimm -ungroup *MagazineDescription*

omnimm -reload_serial_number *DeviceName*

omnimm -show_locked_devs [-all]

omnimm -delete_unprotected_media [Library | -all] [-force]

omnimm -merge_library sourceLibrary destinationLibrary [-nofallback]

```
Policy =
 Loose|
 Strict|
 App+Loose|
 App+Strict|
 AppIncr+Loose|
 AppIncr+Strict


Medium =
 Medium_Label|
 Barcode|
 Medium ID


Basic Options =
 -force
 -pool PoolName
 -size n
 -location OffLineLoc
 -eject
```

DESCRIPTION

The main purpose of *media management* is to protect valuable user data.

To achieve this goal Data Protector provides the following functionality: protecting data from being overwritten, detecting and tracking bad or old media, utilizing and reporting space in auto changers, use of media within pools, drive cleaning, detecting standard tape and magneto-optical format. All this information is stored into the Data Protector Internal Database.

The `omnimm` command manages media pools, checks the protection of a medium, maintains and updates the contents of the repository in the library.

Protecting data is more than just stopping Data Protector from overwriting the tape. The detection of an old and poor media informs the administrator before data loss so that he can react before he needs to restore the data and tape will never be used for backups again. This means protection of data which are on Data Protector tapes and protection for data which is still on the system and needs to be backed up.

For the list of supported media classes, see the *HPE Data Protector Product Announcements, Software Notes, and References*.

Data Protector has the concept of `media pools` to manage large numbers of cartridges. Pools are logical collection of cartridges with same common media or data properties. One pool can only contain media of one type. Data Protector support several media `pool policies`:

- *Loose* (loose, non-appendable); When Data Protector prompts for a medium and loose policy is selected, any medium in the defined pool will be accepted.

- *Strict* (strict, non-appendable); Data Protector decides which medium must be inserted for backup and only this medium will be accepted.

- *App+Loose* (loose, appendable);

- *App+Strict* (strict,appendable);

- *AppIncr+Loose* (loose, appendable for incrementals);

- *AppIncr+Strict* (strict, appendable for incrementals).

OPTIONS

-version

> Displays the version of the omnimm command

-help

> Displays the usage synopsis for the omnimm command

-create_pool *PoolName MediaType* [*Policy AgeLimit MaxOverWrites*]

> Creates a new pool with *PoolName* for the medium of *MediaType* with the policy defined by *Policy*. For the list of supported media classes, see the *HPE Data Protector Product Announcements, Software Notes, and References*. Supported policies are: Loose, Strict, App+Loose, App+Strict, AppIncr+Loose and AppIncr+Strict. *AgeLimit* is set in months. The *MaxOverWrites* is the maximum number of times that the medium can be overwritten. The default is 250 overwrites.

-[no_]alloc_uninit_first

> Option -noalloc_uninit_first sets/resets "Use uninitialized media first" pool policy. This option can be used with *Loose* policy only.

-[no_]free_pool [*FreePoolName*]

> If -free_pool is set, the pool is linked to the free pool specified with FreePoolName in order to share free media. Condition factors are inherited from the free pool. If the -no_free_pool is set, the pool is not linked. The default setting is -no_free_pool.

-[no_]move_free_media

> The -move_free_pool option can only be set if the -free_pool option was set. If -move_free_media is set, de-allocation of free media from a regular to a free pool is done automatically. If -no_move_free_media is set, there is no automatic de-allocation of free media. The default setting is -no_move_free_media.

-modify_pool *PoolName NewPoolName* [*Policy AgeLimit MaxOverWrites*]

> Renames the pool *PoolName* into *NewPoolName*. The *Policy*, *AgeLimit* and *MaxOverWrites* can also be changed. Supported policies are: Loose, Strict, App+Loose, App+Strict, AppIncr+Loose and AppIncr+Strict. *AgeLimit* is set in months. The *MaxOverWrites* is the maximum number of times that the medium can be overwritten. The default is 250 overwrites.

-create_free_pool *PoolName MediaType* [*AgeLimit MaxOverWrites*]

> Creates a new free pool with *PoolName* for the medium of *MediaType*. The *MaxOverWrites* is the

maximum number of times that the medium can be overwritten. The default is 250 overwrites.

-modify_free_pool *PoolName* *NewPoolName* [*AgeLimit* *MaxOverWrites*]

Renames the free pool *PoolName* into *NewPoolName*. The *AgeLimit* and *MaxOverWrites* can also be changed. *AgeLimit* is set in months. The *MaxOverWrites* is the maximum number of times that the medium can be overwritten. The default is 250 overwrites.

-create_mag_pool *PoolName* *MediaType* [*Policy* *AgeLimit* *MaxOverWrites*]

Creates pool *PoolName* with magazine support.

-modify_mag_pool *PoolName* *NewPoolName* [*Policy* *AgeLimit* *MaxOverWrites*]

Renames the magazine pool *PoolName* into *NewPoolName*. The *Policy*, *AgeLimit* and *MaxOverWrites* can also be changed. *AgeLimit* is set in months. The *MaxOverWrites* is the maximum number of times that the medium can be overwritten. The default is 250 overwrites.

-remove_pool *PoolName*

Removes the pool specified by *PoolName*.

-remove_mag_pool *PoolName*

Removes the magazine pool specified by *PoolName*.

-show_pools [*PoolName*]

Shows media from the specified *PoolName* pool or from all pools if *PoolName* is omitted.

-move_medium *Medium* *ToPoolName*

Moves medium from the current pool to the pool specified by *ToPoolName*.

-move_magazine *MagazineDescription* *NewPoolName*

Moves magazine *MagazineDescription* from the current pool to the pool specified by *NewPoolName*.

-modify_medium *Medium* *NewMediumLabel* *NewLocation*

Modifies medium with the specified *Medium*. Note that you should always enter the medium label *NewMediumLabel* and location *NewLocation* in that sequence.

-modify_magazine *MagDescription* *NewLocation* [*NewMagDescription*]

Changes the location of the magazine *MagDescription* to *NewLocation*. If *NewMagDescription* is specified, it is assigned to the magazine as a new MagazineDescription. Note that each magazine must have a unique MagazineDescription.

-reset_poor_medium *Medium*

Resets the media condition factors. Once the medium has expired (its maximum usage criteria), it is marked as poor and can no longer be used for backup. This option resets the medium quality status, thus enabling it to be used for backup. You have to be very cautious using this option, because a backup stored on an expired medium might not be recoverable.

-reset_wp_medium *Medium*

Removes the write-protected flag for the specified medium from the MMDB, thus making the medium available for writing.

-list_pool [*PoolName*] [-detail]

Displays all the media from pool *PoolName*. The report shows: medium label, status, location, appendability and protection. Appendability is shown under item FULL. If displayed status under FULL is "YES" then medium is unappendable, otherwise it is appendable. If *PoolName* is not specified, the command lists all the configured media pools. This report shows: pool name, status, media class, the number of media and free space in pool.

-detail

Displays information in a more detailed format.

-show_pool_alloc *PoolName*

Displays the sequence in which the media from the specified pool will be used for backup. The report shows: sequence, medium label and location.

-list_scratch_media *PoolName*

Displays media from the specified pool which are not protected and can be used for backup. The report shows sequence, medium label and location.

-show_repository_alloc *Library PoolName*

Displays the order in which the media in the repository of the specified *Library* will be used. The report shows: sequence, medium label, location and slot number.

-list_media *Medium*

Displays all the objects, their type and their protection status for the medium you specified.

The VADP feature introduced in Data Protector 8.14 provides enhanced reports for Virtual Machines. The reports display VMware virtual machines in the same way as Data Protector clients called VADP clients. The object name must use the VM name as reported from `omnicellinfo -cell brief` command, where object name is `<hostname>:/<vCenter>/<path>/<vmname> [<UUID>]`. Here, `<hostname>` is DNS name of the guest virtual machine. If the DNS name is unknown, the IP address or VM name is used.

-list_appendable_media *PoolName*

Displays all appendable media from the specified media pool.

-list_copy *Medium*

List all copies of the given medium.

-media_info *Medium*

Displays information on the given medium.

-encryptioninfo

Displays detailed encryption information for objects on the specified medium.

-list_magazines_of_pool *PoolName*

Lists magazines of the pool *PoolName*.

-list_media_magazine *MagazineDescription*

Lists all the media in specified magazine.

-catalog *Medium*

Lists catalog for all object versions located on the specified medium. Only files located on this

medium are displayed.

-recycle *Medium*

Resets the protection of data on medium. The present data can now be overwritten and medium can be used to store new data.

-recycle_magazine *MagazineDescription*

Recycles all media of specified magazine.

-export *Medium*

Purges from the database all data associated with the medium and the object versions it contains. This option is used when the medium will no longer be used for backup in this cell. A medium containing protected data cannot be exported.

-export_magazine *MagazineDescription*

Exports all media of specified magazine.

-copy_to_mcf *Medium*

Copies media-related catalog data into media container format (MCF) files, which you can transfer to another Cell Manager, thus enabling you to import all media-related information on another Cell Manager where it is then available for browsing. The media-related catalog data are not removed from the original Cell Manager. You can specify one or more media by either medium ID or medium label.

-output_directory *Pathname*

Specifies the directory where MCF files are stored. You must specify a full path to the files. If not specified, the files are by default copied on the Cell Manager to the directory *Data_Protector_program_data*\Config\Server\export\mcf (Windows systems) or /var/opt/omni/server/export/mcf (UNIX systems).

-import *LogicalDevice*

Imports a medium from a different cell. The medium is put in the default pool of the specified backup device. Information about the new medium is added to the database. Slot side must be specified for magneto-optical devices.

-no_log

Used with the -import option, this option omits the detail part of the catalog from the import.

-log

Used with the -import option, this option logs all detailed information of the backed up directory such as versions, numbers, and attributes.

-log_dirs

Used with the -import option, this option imports only the detail part of the directories.

-pool *PoolName*

Specifies the name of the pool.

-import_catalog *LogicalDevice*

Rereads the Detail Catalog from the specified device into the database, in case this information has

been deleted. If the Detail Catalog for the specified medium already exists in the database, import will fail.

-import_magazine *LogicalDevice* [*MagazineDescription*]

Imports a magazine from a different cell. The magazine is put in the default pool of the specified backup device. Information about the new magazine and its media is added to the database.

-import_from_mcf *File*...

Imports one or more MCF files that contain copies of media-related catalog data from the original Cell Manager. You must specify full paths to the files on the current Cell Manager.

-pool_prefix

Specifies an optional prefix for a media pool to which you want to import MCF files with media-related catalog data copies. If this option is not specified, the default prefix IMPORTED is used.

If the -no_pool_prefix option is set, no prefix is generated for a pool.

-[no_]orig_pool

Specifies a media pool for import. By default, the -orig_pool option is set.

It can be disabled with the -[no_]orig_pool option.

-import_as_original

Imports a medium copy or a medium-related catalog data copy as original if an original medium does not exist in a database.

-disable_lockname *LockName*

Disables devices with the *LockName* for any operation. The *LockName* must be defined using the Data Protector GUI or using the omniupload command.

-enable_lockname *LockName*

Enables devices with the *LockName*. The *LockName* must be defined using the Data Protector GUI or using the omniupload command.

-disable_device *DeviceName* [-ignore_lockname]

Disables the device with the *DeviceName* for any operation. The *DeviceName* must be defined using the Data Protector GUI or using the omniupload command. Unless the option -ignore_lockname is specified, if the device has a lockname defined, all devices with the same lockname are also disabled.

-enable_device *DeviceName* [-ignore_lockname]

Enables the device with the *DeviceName*. The *DeviceName* must be defined using the Data Protector GUI or using the omniupload command. Unless the option -ignore_lockname is specified, if the device has a lockname defined, all devices with the same lockname are also enabled.

-repository *LibraryName*

This option is used to specify the repository backup device that you want to check. This information is then used to update the database.

-repository_barcode_scan *LibraryName*

If this option is used then barcode reader is used to update the database. This option should be used

only with devices that have enabled barcode reader.

-repository_update *DriveName*

Updates the database by reading all the slots (loads media in drive) in the device repository. If you additionally specify the slot number of the slot that is defined for a CL cartridge, then a cleaning operation is performed on the specified drive.

-slot *SlotID* [*Side*]

Specifies the *SlotID* of the library where the medium is mounted. This option is only valid for this backup device type. To specify the side of the platter in this slot, use the additional *Side* parameter. Slot *SlotID* must be specified for magneto-optical devices. Values of *Side* are A or B.

-add_slots *LibraryName* {*Slot*... | *FromSlot-ToSlot*...}

Adds slots to the selected library. With ADIC/GRAU DAS or StorageTek ACS libraries, this option adds volsers to the selected library. 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.

-remove_slots *LibraryName* {*Slot*... | *FromSlot-ToSlot*...}

Removes slots from the selected library.

-silo_query *LibraryName*

Queries ACS/DAS server for the list of currently resident volsers and updates the Data Protector repository of specified library. This option is not recommended to be used with an ACS/DAS Server when querying logical libraries configured for the same physical library. In such a case, use the -add_slots option to add volsers manually.

With DAS Server, however, when logical libraries are not configured using Data Protector, but using the DAS utilities, the Data Protector query operation can safely be used on such libraries instead of adding volsers manually.

-silo_enter *LibraryName*

Moves ACS/DAS media from the CAP (ACS) or insert/eject area (DAS) to the repository.

-cap *CapID*

ID of Control Access Port of ACS or insert/eject area of DAS library.

-silo_eject *LibraryName* {*Volser*... | *FromVolser-ToVolser*...}

Moves media from the ACS/DAS repository into the CAP.

-location *Location*

Specifies the new location for the ejected media. Only media with barcode will be updated.

-enter *LibraryName* {*Slot*... | *FromSlot-ToSlot*...}

Moves media from the mail slots to the repository slots. This option is available only for SCSI libraries.

-eject *LibraryName* {*Slot*... | *FromSlot-ToSlot*...}

Moves media from the repository slots into the mail slots. This option is available only for SCSI libraries.

-group *PoolName MagazineDescription Medium*...

Creates a magazine *MagazineDescription* out of the specified non-magazine media. Note that all specified media must be resident in the same SCSI library at the time. The magazine is added to the pool *PoolName* which must be configured to support magazines.

`-ungroup` *MagazineDescription*

Splits the magazine *MagazineDescription* so that the magazine media become non-magazine media.

`-reload_serial_number` *DeviceName*

Reloads the device serial number and overwrites the serial number stored in the Internal Database. A physical device can therefore be replaced without changing the logical device properties.

`-show_locked_devs` [`-all`]

Lists all locked devices, target volumes, media, and slots in the Data Protector cell. The `-all` option applies only when you execute the command on a MoM system, in which case locked devices, target volumes, media, and slots from all cells are listed.

`-delete_unprotected_media` [`Library`|`-all`][`-force`]

This option is used for deleting unprotected media. This process is automatically carried out by Data Protector during its maintenance window. To trigger it when required, use the command with the options shown here.

Library - This option deletes all unprotected media for a specific B2D library.

all - This option deletes all unprotected media belonging to all the B2D devices.

force - This option makes Data Protector delete a medium from the IDB, even if the deletion process reports an error for that entry.

`-merge_library` *sourceLibrary destinationLibrary* [`-nofallback`]

This option is used to merge gateways of StoreOnce Backup system or Data Domain Boost devices (configured using FC identifiers) to IP- based devices. This enables the gateways to support multi-interface access.

After migration is complete, all gateways from the source library are, by default, set to use the FC paths, with an option to fall back to an IP path.

Note that you can *only* merge from DDBoost to DDBoost or StoreOnce to StoreOnce and only from FC to IP. Also, the media pools for the migrated gateways are not changed. If required, you can change the gateway pools using the DP GUI.

sourceLibrary - The library configured with an FC interface (DDBoost or StoreOnce System).

destinationLibrary - The library configured with an IP interface (DDBoost or StoreOnce System).

nofallback - This option can be specified to prevent migrated gateways from falling back to IP paths in case of errors.

RETURN VALUES

See the man page `omniintro` for return values.

Additional return values of the `omnimm` command are:

1 - Program failed, user error.

2 - Program failed, environmental malfunction.

3 - Program failed, internal malfunction.

4 - Program failed, reason unknown.

NOTES

Make sure that the *PoolName* does not exceed a limit of 32 characters.

The omnimm command displays the virtual machine object and its associated disk objects. Other commands will not display the same.

EXAMPLES

The following examples illustrate how the omnimm command works.

1. To create pool "DDS_Pool" of the class "DDS", with policy "App+Loose". Media in the pool will be usable for 12 months or for 100 overwrites., execute:

   ```
   omnimm -create_pool DDS_Pool "DDS" App+Loose 12 100
   ```

2. To modify the medium with label "Label23" changing the label to "LABEL23" and location to "Backup Room", execute:

   ```
   omnimm -modify_medium Label23 LABEL23 "Backup Room"
   ```

3. To list detailed information for medium "dat1", execute:

   ```
   omnimm -list_media dat1 -detail
   ```

4. To list the virtual machine objects and disk objects for medium "medId", execute:

   ```
   omnimm -list_media medId -detail
   ```

5. To list encryption information for medium "MediaPool1_10", execute:

   ```
   omnimm -list_media MediaPool1_10 -encryptioninfo
   ```

6. To import a medium in the backup device "Pool1" into pool "Default DDS", execute:

   ```
   omnimm -import Pool1 -pool "Default DDS"
   ```

7. To copy media catalogs of media "DefaultFile_1" and "MyDLT_35" to the MCF directory on a UNIX system, execute:

   ```
   omnimm -copy_to_mcf "DefaultFile_1" "MyDLT_35" -output_directory /tmp/mcf
   ```

8. To import media-related catalog data copies "2401110a_47d7f516_0aa0_0001.mcf" and "2401110a_47e26bc2_0a74_0002.mcf" from the default MCF directory on Windows Server 2003 into a new media pool with prefix "MCF_" located on another Cell Manager, execute:

   ```
   omnimm -import_from_mcf "C:\Program Files\OmniBack\Config\
   Server\import\mcf\2401110a_47d7f516_0aa0_0001.mcf" "C:\Program
   Files\OmniBack\Config\Server\import\mcf\ 2401110a_47e26bc2_0a74_0002.mcf" -
   pool_prefix "MCF_" -no_orig_pool
   ```

SEE ALSO

omniamo(1), omnib2dinfo(1M), omnidownload(1), omnimcopy(1), omniminit(1), omnimlist(1), omnimnt (1), omnimver(1), omniupload(1), sanconf(1M), uma(1M)

omnimnt(1)

omnimnt — responds to a Data Protector mount requests for a medium
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnimnt -version | -help

omnimnt -device *BackupDevice* -session *SessionID* [-cancel]

DESCRIPTION

The `omnimnt` command satisfies or aborts a Data Protector mount request. A mount request is issued by a backup device once it has filled all the available media. A mount request is a prompt to mount a new medium. Once the requested medium is inserted in the device drive, the `omnimnt` command should be used to confirm that the correct medium is inserted. The mount request can also be canceled which is done by canceling device. If you cancel device, all data objects associated with the backup device that issued the mount request will not be processed any further. To view information on currently active sessions, use the `omnistat` command.

OPTIONS

-version

　　Displays the version of the `omnimnt` command

-help

　　Displays the usage synopsis for the `omnimnt` command

-cancel

　　Cancels the device. This will terminate processing of all objects that are associated with the backup device which issued the request.

-device *BackupDevice*

　　References the backup device *BackupDevice* which issued the mount request, in order to confirm mount request or cancel the device.

-session *SessionID*

　　Specifies the session using the backup device which issued the mount request.

EXAMPLES

The following examples illustrate how the `omnimnt` command works.

1. To satisfy a mount request issued by device "DAT1" in a session with SessionID "R-2013/05/05-275", execute:

   ```
   omnimnt -device DAT1 -session R-2013/05/05-275
   ```

2. To cancel device for the backup device "Juke" in the session with SessionID "R-2013/05/25-3", execute:

   ```
   omnimnt -device Juke -session R-2013/05/25-3 -cancel
   ```

SEE ALSO

omniamo(1), omnib2dinfo(1M), omnidownload(1), omnimcopy(1), omniminit(1), omnimlist(1), omnimm (1), omnimver(1), omniupload(1), sanconf(1M), uma(1M)

omnimver(1)

omnimver — verifies data on a medium
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnimver -version|-help

omnimver -device *BackupDevice* [-slot *SlotID* [*Side*]] [-eject]

DESCRIPTION

The `omnimver` command is used to verify the contents of a Data Protector backup medium. It reads the data and verifies that data is written in the Data Protector format. If the `-crc` option was used to back up the data, it also checks the CRC for each block.

OPTIONS

-version

   Displays the version of the `omnimver` command

-help

   Displays an extended usage synopsis for the `omnimver` command

-device *BackupDevice*

   Specifies the backup device where medium is located.

-slot *SlotID* [*Side*]

   Specifies the *SlotID* of the Exchanger backup device where the medium is mounted. This option is only valid for this backup device type. To specify the side of the platter in this slot, use the additional *Side* parameter. Slot *Side* must be specified for magneto-optical devices.

-eject

   Ejects the medium from the drive after the verification.

EXAMPLES

1. To verify slot 32 of backup device "Spectra60", execute:

   ```
   omnimver -device Spectra60 -slot 32
   ```

SEE ALSO

omniamo(1), omnib2dinfo(1M), omnidownload(1), omnimcopy(1), omniminit(1), omnimlist(1), omnimm
(1), omnimnt(1), omniupload(1), sanconf(1M), uma(1M)

omniobjconsolidate(1)

omniobjconsolidate — consolidates Data Protector backup objects into synthetic full backups
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omniobjconsolidate -version|-help

omniobjconsolidate -consolidationlist *ConsolidationSpecificationName* -scheduled
[*GeneralOptions*]

omniobjconsolidate -consolidationlist *ConsolidationSpecificationName* -postbackup -
session *SessionID* [*GeneralOptions*]

omniobjconsolidate [*GeneralOptions*] [*Device*...] *Object* [*Object*...]

GeneralOptions

 [-dynamic *min max*]

 [-protect {none|weeks *n*|days *n*|until *Date*|permanent}]

 [-keepcatalog {weeks *n*|days *n*|until *Date*|same_as_data_protection}]

 [-[no_]log|-log_dirs|-log_file]

 [-recycle]

 [-locationpriority *MediumLocation* [*MediumLocation*...]]

 [-no_monitor]

 [-priority *NumValue*]

MediumLocation

 = "=*MediumLocation*"|"<*MediumLocation*"

Device

 -targetdevice *LogicalDevice* [*DeviceOptions*]

DeviceOptions

 [-concurrency *ConcurrencyNumber*]

 [-crc]

 [-encrypt]

 [-pool *PoolName*]

 [-prealloc *MediumID* [*MediumID*...]]

Object

 {-filesystem|-winfs} *Client:MountPoint Label*

```
-session SessionID

[-copy CopyID]

[-sourcedevice BackupDevice]

-consolidationdevice LogicalDevice

[-targetdevice LogicalDevice]

[-protect {none|weeks n|days n|until Date|permanent}]

[-keepcatalog {weeks n|days n|until Date|same_as_data_protection}]

[-[no_]log|-log_dirs|-log_file]

[-[no_]recycle]

OtherOptions

Date= [YY]YY/MM/DD (1969 < [YY]YY < 2038)
```

DESCRIPTION

The `omniobjconsolidate` command creates synthetic full backups from full and incremental backups. It can be used to:

- consolidate objects that you specify
- start a post-backup object consolidation specification
- start a scheduled object consolidation specification

To consolidate an object to a specific point in time, specify only the incremental version of that point in time. The restore chain is retrieved automatically.

To obtain the information about all backed up objects or sessions containing the objects you want to consolidate, use the `omnidb` command.

OPTIONS

`-version`

Displays the version of the `omniobjconsolidate` command.

`-help`

Displays the usage synopsis of the `omniobjconsolidate` command.

`-consolidationlist ConsolidationSpecificationName`

Specifies the object consolidation specification identified by `ConsolidationSpecificationName` for object consolidation.

`-scheduled`

Immediately starts a scheduled object consolidation specification.

`-postbackup`

Immediately starts a post-backup object consolidation specification specified by the `-session`

SessionID option.

-session *SessionID*

> If specified with the -postbackup option, provides the session ID for the post-backup object consolidation session.

> If specified as part of the object definition, selects the point in time for object consolidation.

-dynamic *min max*

> Specifies how many devices are locked prior to starting a session. Devices that are specified per object through the -targetdevice option are locked in any case. The *max* value is increased by Data Protector if the number of statically assigned devices is higher than the specified *max* value.

> *Min* specifies the minimum number of available devices (devices that are not being used by another Data Protector session and have the license to be started) required for starting the session. If fewer devices are available than specified here, the session will queue. The default is 1.

> *Max* specifies the maximum number of available devices that Data Protector will use in the session. The highest number you can specify is 32. The default is 5. Data Protector will lock the number of devices that you specify using this parameter if so many devices are available. If this option is not specified, the default value for *max* is the number of specified devices.

-protect {none | weeks *n* | days *n* | until *Date* | permanent}

> Sets a period of protection for the consolidated data on the backup medium to prevent the data from being overwritten. If this option is not specified, the data protection of the consolidated objects is the same as the protection of the full backup of the objects. If a relative period of protection was set for the full backup, such as *n* days or weeks, the same protection period is counted from the creation time of the synthetic full backup.

-keepcatalog {weeks *n* | days *n* | until *Date* | same_as_data_protection}

> Specifies file catalog retention time. If you do not want to save the file catalog, use the -no_log option. If this option is not specified, the catalog protection of the consolidated objects is the same as the catalog protection of the full backup of the objects. If a relative period of catalog protection was set for the full backup, such as *n* days or weeks, the same protection period is counted from the creation time of the synthetic full backup.

-log

> Specifies the logging level of the object consolidation session. All detailed information about backed up files and directories (filenames, file versions, and attributes) are logged to the IDB. You can browse directories and files before restoring and in addition look at file attributes. Data Protector can fast position on the tape when restoring a specific file or directory.

> If the logging level is not specified, the logging level of the source object is used.

-no_log

> Specifies the logging level of the object consolidation session. No information about backed up files and directories is logged to the IDB. You will not be able to search and browse files and directories before restoring.

-log_dirs

> Specifies the logging level of the object consolidation session. All detailed information about backed up directories (names, versions, and attributes) is logged to the IDB. You can browse only

directories before restoring.

-log_file

Specifies the logging level of the object consolidation session. All detailed information about backed up files and directories (filenames and file versions) is logged to the IDB. You can browse directories and files before restoring, and Data Protector can fast position on the tape when restoring a specific file or directory. The information does not occupy much space, since not all file details (file attributes) are logged to the database.

-[no_]recycle

The -recycle option removes data and catalog protection of the objects on the source media. When there are no more protected objects on the media, the media can be overwritten. The -no_recycle option is available as part of the object definition if the -recycle option is specified as part of *GENERAL_OPTIONS*.

IMPORTANT: If you recycle data protection of source objects, the recycled points in time will no longer be available. Unless copies of these points in time exist, you will be able to restore only to the latest (consolidated) point in time.

-locationpriority *MediumLocation* [*MediumLocation*]

The order in which media are selected for object consolidation in case copies of the same object version exist in more than one location. By default, Data Protector automatically selects the most appropriate media set. Media location priority is considered if more than one media set equally matches the conditions of the media set selection algorithm.

The priority must be specified in the form "=*MediumLocation*" (equal to) or "<*MediumLocation*" (lower priority than).

If you specify -locationpriority "=Loc1" "<Loc2" "=Loc3" "<Loc4", than Loc1 has the highest priority, Loc2 and Loc3 have a lower priority, and Loc4 has the lowest priority.

-no_monitor

If this option is used, the command displays only the session ID. By default, the command monitors the session and displays all messages.

-priority *NumValue*

In case multiple running sessions request access to a specific device at the same time, this option determines the order in which the sessions will be queued. The *NumValue* can be any value from 1 (the highest priority) to 6000 (the lowest priority). In case the option is not specified, the default value of 3000 is set. If a low priority session is running when a high priority session starts queuing, the currently running session is allowed to finish. When more sessions request access to a device with the same priority, any of these sessions might acquire access first.

-filesystem *Client:MountPoint Label*

Selects the filesystem identified with *Client:MountPoint Label* for object consolidation.

-winfs *Client:MountPoint Label*

Selects the Windows filesystem identified with *Client:MountPoint Label* for object consolidation.

-copy *CopyID*

Selects the copy identified with *CopyID*. If not specified, Data Protector automatically selects the

most appropriate copy as the source for object consolidation.

-sourcedevice *LogicalDevice*

Specifies a logical device to be used for reading full object versions from the source media. If this option is not specified, Data Protector uses the logical device that was used for writing the objects.

-consolidationdevice *LogicalDevice*

Specifies a logical device that will read incremental object versions and perform object consolidation.

-targetdevice *LogicalDevice*

Specifies a logical device that will be used for writing consolidated object versions to the target media. If specified as a part of *GeneralOptions*, the device is used for all objects. In this case, you can also specify device options. If you specify several devices, the devices will be dynamically assigned to objects.

If specified as part of *Object*, the device is used only for this object.

You can combine static and dynamic assignment of devices by specifying some devices as part of *GeneralOptions*, and for some objects, specifying a device per object.

-concurrency *ConcurrencyNumber*

Specifies the number of restore Media Agents that can send data to a device concurrently.

The maximum concurrency value is 32.

-crc

The CRC check is an enhanced checksum function. When this option is selected, cyclic redundancy check sums (CRC) are written to the media during object consolidation. The CRC checks enables you to verify the media after the operation. Data Protector re-calculates the CRC during a restore and compares it to the CRC on the medium. It is also used while verifying and copying the media.

-encrypt

If this option is used, the backup Media Agent enables hardware encryption on the device. Consolidated data is encrypted and written to media.

-pool *PoolName*

Selects a specific media pool for object consolidation. If not defined, a default media pool from the device definition will be used.

-prealloc *MediumID* [*MediumID*]...

Defines the prealloc list. This is a subset of media used for object consolidation in the specified sequence.

When using the prealloc list and the strict media allocation policy with the backup device, Data Protector expects the sequence of the media in the device to correspond with that specified in the prealloc list. If the media are not available in this sequence, Data Protector issues a mount request. If no media are specified in this list, the Data Protector allocation procedure is used to allocate media.

NOTES

All options specified before the first *Object* are applied to all objects. Options specified as a part of an *Object* are applied only to that object and may override general options.

RETURN VALUES

See the man page omniintro for return values.

Additional return values of the `omniobjconsolidate` command are:

`10` - There was an error while consolidating some files. All agents completed successfully.

`11` - One or more agents failed, or there was a database error.

`12` - None of the agents completed the operation.

`13` - Session was aborted.

EXAMPLES

1. To start an object consolidation session that consolidates the WinFS object versions for "OBJECT1" on the host "system1.company.com" to the point in time defined with the session ID "2013/05/06-1", using the device "LTO3" as the source device and the file library "FILEDEV1" as the consolidation device, and writes the consolidated objects to the device "LTO4", use:

   ```
   omniobjconsolidate -winfs system1.company.com:/C 'OBJECT1' -session 2013/05/06-
   1 -sourcedevice 'LTO3' -consolidationdevice 'FILEDEV1' -targetdevice 'LTO4'
   ```

2. To start an interactive object consolidation session for the filesystem object "system1.company.com:/ 'Label42'" from the session "2013/05/01-2", using the device "DEV1" to read the source object and the device "DEV2" to consolidate the object, and write the consolidated object to the device "DEV3", use:

   ```
   omniobjconsolidate -filesystem system1.company.com:/ 'Label42' -session
   2013/05/01-2 -sourcedevice 'DEV1' -consolidationdevice 'DEV2' -targetdevice
   'DEV3'
   ```

3. To immediately start a post-backup object consolidation specification named "post_BU1" for the session "2013/05/03-1", execute:

   ```
   omniobjconsolidate -consolidationlist post_BU1 -postbackup -session 2013/05/03-
   1
   ```

4. To immediately start a scheduled object consolidation specification named "Consolidation_16_Spec", execute:

   ```
   omniobjconsolidate -consolidationlist Consolidation_16_Spec -scheduled
   ```

SEE ALSO

omnib(1), omnikeytool(1M), omniobjcopy(1), omniobjverify(1), omnir(1)

omniobjcopy(1)

omniobjcopy — creates additional copies of objects backed up with Data Protector on a different media set
(this command is available on systems with the Data Protector User Interface component installed)

SYNOPSIS

omniobjcopy -version | -help

omniobjcopy -copylist *CopySpecificationName* -scheduled [*[GENERAL_OPTIONS]*]

omniobjcopy -copylist *CopySpecificationName* -postbackup -session *SessionID* [*[GENERAL_ OPTIONS]*]

omniobjcopy -replist *ReplicationSpecificationName* -scheduled [*[GENERAL_OPTIONS]*]

omniobjcopy -replist *ReplicationSpecificationName* -postbackup -session *SessionID* [*[GENERAL_OPTIONS]*]

omniobjcopy -restart *SessionID*

omniobjcopy [*GENERAL_OPTIONS*] *Device ... Object* [*Object*]...

GENERAL_OPTIONS

[-replication]

[-dynamic *min max*]

[-targetprotect {none | weeks *n* | days *n* | until *Date* | permanent}]

[-keepcatalog {weeks *n* | days *n* | until *Date* | same_as_data_protection}]

[-[no_]log | -log_dirs | -log_file]

[-sourceprotect {none | weeks *n* | days *n* | until *Date* | permanent}]

[-locationpriority *MediumLocation* [*MediumLocation*...]]

[-no_monitor]

[-no_auto_device_selection]

[-session Session ID]

      -replication

      -targetcs HostName

      -targetcsdevice DeviceName

[-priority*NumValue*]

MediumLocation

 = "=*MediumLocation*" | "<*MediumLocation*"

Device

 =-targetdevice *LogicalDevice* [*DeviceOptions*]

DeviceOptions

 [-concurrency *ConcurrencyNumber*]

 [-crc]

 [-encrypt]

 [-pool *PoolName*]

 [-prealloc *MediumID* [*MediumID*...]]

Object

 {-filesystem|-winfs} *Client:MountPoint Label*

 -session *SessionID*

 [-copyid *N* [-fixedcopy...]]

 [-sourcedevice *LogicalDevice*]

 [-targetdevice *LogicalDevice*]

 [-targetprotect {none|weeks *n*|days *n*|until *Date*|permanent}]

 [-keepcatalog {weeks *n*|days *n*|until *Date*|same_as_data_protection}]

 [[-no_]log|-log_dirs|-log_file]

 [-sourceprotect {none|weeks *n*|days *n*|until *Date*|permanent|keep}]

 [-full]

Object

 -rawdisk *Client Label*

 -session *SessionID*

 [-copyid *N* [-fixedcopy...]]

 [-sourcedevice *LogicalDevice*]

 [-targetdevice *LogicalDevice*]

 [-targetprotect {none|weeks *n*|days *n*|until *Date*|permanent}]

 [-sourceprotect {none|weeks *n*|days *n*|until *Date*|permanent|keep}]

Object

 {-sap|-oracle8|-integ {MySQL | PostgreSQL}|-informix|-msese|-e2010|-mssql|-lotus|
 -mbx|-sapdb|-saphana|-msvssw|-db2|-sybase|-mssharepoint|-veagent|-idb} *Client:Set*

 -session *SessionID*

 [-copyid *N* [-fixedcopy...]]

 [-sourcedevice *LogicalDevice*]

 [-targetdevice *LogicalDevice*]

 [-targetprotect {none|weeks *n*|days *n*|until *Date*|permanent}]

[-sourceprotect {none|weeks *n*|days *n*|until *Date*|permanent|keep}]

OtherOptions

Date= [YY]YY/MM/DD (1969 < [YY]YY < 2038)

DESCRIPTION

The omniobjcopy command creates additional copies of objects backed up using Data Protector. You can use the omniobjcopy command to copy objects such as filesystems (UNIX or Windows), very big file systems, disk image sections, and Data Protector Internal Database (IDB) to an additional media set. The command can be also used for copying the integration objects (SAP R/3, Oracle, Informix Server, VMware vSphere, Microsoft Hyper-V, Microsoft Exchange Server 2007, Microsoft Exchange Server 2010/2013, Microsoft Exchange Server single mailboxes, Microsoft SharePoint Server 2007/2010/2013, Microsoft SQL Server, Lotus, Sybase, DB2, Microsoft Volume Shadow Copy Service, SAP MaxDB, and SAP HANA Appliance).

To obtain the information about all backed up objects or sessions containing the objects you want to copy, use the omnidb command.

This command starts an interactive or automated object copy session. Use this command to immediately start an automated (scheduled or post-backup) object copy specification.

From Data Protector 9.05 onwards for VMware integrations, the virtual machine disks are considered as objects that run in parallel. When the virtual machine object is selected for object copy operations, its associated disk objects are also considered for these copy operations.

OPTIONS

-version

Displays the version of the omniobjcopy command.

-help

Displays the usage synopsis for the omniobjcopy command.

-copylist *CopySpecificationName*

Specifies the name of the object copy specification identified by *CopySpecificationName* for object copying.

-replist *ReplicationSpecificationName*

Specifies the name of the replication specification identified by *ReplicationSpecificationName* for replication.

-scheduled

Immediately starts a scheduled object copy specification.

-postbackup

Immediately starts a post-backup object copy specification specified by the -session *SessionID* option.

-replication

Enables replication for supported B2D devices in interactive sessions.

-session *SessionID*

Selects the session ID for the -postbackup option or for the object definition.

-restart *SessionID*

Tries to restart a failed non-interactive object copy session, specified by its session ID.

-dynamic *min max*

Specifies how many devices are locked prior to starting a session. Devices that are specified per object through the -targetdevice option are locked in any case. The *max* value is increased by Data Protector if the number of statically assigned devices is higher than the specified *max* value.

Min specifies the minimum number of available devices (devices that are not being used by another Data Protector session and have the license to be started) required for starting the session. If fewer devices are available than specified here, the session will queue. The default is 1.

Max specifies the maximum number of available devices that Data Protector will use in the session. The highest number you can specify is 32. The default is 5. Data Protector will lock the number of devices that you specify using this parameter if so many devices are available. If this option is not specified, the default value for *max* is the number of specified devices.

-targetprotect {none|weeks *n*|days *n*|until *Date*|permanent}

Sets the level of protection for the copy object. The media containing this object copy session cannot be overwritten until the protection expires. By default (if this option is not specified), the protection is the same as the original protection for the source object.

-keepcatalog {weeks *n*|days *n*|until *Date*|same_as_data_protection}

Specifies file catalog retention time. If you do not want to save the file catalog at all, use the -no_ log option. By default (if this option is not specified), the protection is the same as for the source object.

-log

Specifies the logging level of the object copy session. All detailed information about backed up files and directories (filenames, file versions, and attributes) are logged to the Data Protector Internal Database (IDB). This allows you to browse directories and files before restore and in addition look at the file attributes. Data Protector can fast position on the tape when restoring a specific file.

If the logging level is not specified, it is set to the same logging level as for the source object.

-no_log

Specifies the logging level of the object copy session. Disables the logging of backed up files to the catalog database. By default, the filename and backup history of each backed up file is written to the catalog database.

-log_dirs

Specifies the logging level of the object copy session. If this option is specified, only the directories are logged into the database. By default, the filename and backup history of each backed up file is written to the catalog database.

-log_file

Specifies the logging level of the object copy session. All detailed information about backed up files

and directories (filenames and file versions) is logged to the Data Protector Internal Database (IDB). This information allows you to search for backed up files and allows Data Protector to fast position the tape. It also does not take much space since some information on file details (file attributes) is not logged to the database.

-sourceprotect {none|weeks *n*|days *n*|until *Date*|permanent|keep}

Sets the level of protection for the source object after a successful copy. The media containing this source object cannot be overwritten until the protection expires. By default (if this option is not specified), the protection is not changed.

The none option specifies that protection is removed from the source object immediately, allowing recycling.

The keep option can only be specified at the object level and specifies that the protection for that source object should not be changed.

-locationpriority *MediumLocation* [*MediumLocation*]

The order in which media are selected for the object copy in case that the same object version exist in more than one location. By default, Data Protector automatically selects the most appropriate media set. Media location priority is considered if more than one media set equally matches the conditions of the media set selection algorithm.

The priority must be specified in the form "=*MediumLocation*" (equal to) or "<*MediumLocation*" (lower priority than).

If you specify -locationpriority "=Loc1" "<Loc2" "=Loc3" "<Loc4", than Loc1 has the highest priority, Loc2 and Loc3 have a lower priority, and Loc4 has the lowest priority.

-no_monitor

If this option is used, the command displays only the session ID. By default, the command monitors the session and displays all messages.

-no_auto_device_selection

If this option is specified, Data Protector does not automatically replace unavailable devices with available devices of the same device tag.

-priority *NumValue*

In case multiple running sessions request access to a specific device at the same time, this option determines the order in which the sessions will be queued. The *NumValue* can be any value from 1 (the highest priority) to 6000 (the lowest priority). In case the option is not specified, the default value of 3000 is set. If a low priority session is running when a high priority session starts queuing, the currently running session is allowed to finish. When more sessions request access to a device with the same priority, any of these sessions might acquire access first.

-concurrency *ConcurrencyNumber*

Specifies the number of restore Media Agents that can send data to a device concurrently.

The maximum concurrency value is 32.

-crc

The CRC check is an enhanced checksum function. When this option is selected, cyclic redundancy check sums (CRC) are written to the media during an object copy. The CRC checks enables you to verify the media after the operation. Data Protector re-calculates the CRC during a

restore and compares it to the CRC on the medium. It is also used while verifying and the media.

-encrypt

If this option is used, the backup Media Agent enables hardware encryption on the device. Data is encrypted and copied.

-pool *PoolName*

Selects a specific media pool for object copy. If not defined, a default media pool from the device definition will be used.

-prealloc *MediumID* [*MediumID*]…

Defines the prealloc list. This is a subset of media used for object copy in the specified sequence.

When using the prealloc list and the strict media allocation policy with the backup device, Data Protector expects the sequence of the media in the device to correspond with that specified in the prealloc list. If the media are not available in this sequence, Data Protector issues a mount request. If no media are specified in this list, the Data Protector allocation procedure is used to allocate media.

-filesystem *Client:MountPoint Label*

Selects the filesystem identified by the *Client:MountPoint Label* string for object copying.

-winfs *Client:MountPoint Label*

Selects the Windows filesystem identified by the *Client:MountPoint Label* string for object copying.

-copyid *N* [-fixedcopy]

Selects the specified object copy as a source for object copying.

If -fixedcopy option is not specified, Data Protector selects the needed media set automatically. If several copies of the same object exist in one session as a result of the object copy or object mirror operation, this option is obligatory.

-sourcedevice *LogicalDevice*

Specifies a logical device different from the one used for the backup to be used for reading backed up objects from the source media. By default (if this option is not specified), the same backup device is used for backing up and reading backed up objects from the source media.

-targetdevice *LogicalDevice*

Specifies a backup device that will be used for writing object copies to the target media.

-full

Selects the whole restore chain of full and incremental backups for the object copy operation. This option is not supported for Data Protector application integrations.

-sap *Client:Set*

Selects the SAP R/3 object identified by the *Client:Set* string for object copying.

-informix *Client:Set*

Selects the Informix Server object identified by the *Client:Set* string for object copying.

-msese *Client:Set*

Selects the Microsoft Exchange Server 2007 object identified by the `Client:Set` string for object copying.

-e2010 `Client:Set`

Selects the Microsoft Exchange Server 2010/2013 object identified by the `Client:Set` string for object copying.

-mssql `Client:Set`

Selects the Microsoft SQL Server object identified by the `Client:Set` string for object copying.

-lotus `Client:Set`

Selects the Lotus Notes/Domino Server object identified by the `Client:Set` string for object copying.

-mbx `Client:Set`

Selects the Microsoft Exchange Server single mailbox object identified by the `Client:Set` string for object copying.

-sapdb `Client:Set`

Selects the SAP MaxDB object identified by the `Client:Set` string for object copying.

-msvssw `Client:Set`

Selects the Microsoft Volume Shadow Copy Service object identified by the `Client:Set` string for object copying.

-db2 `Client:Set`

Selects the DB2 object identified by the `Client:Set` string for object copying.

-sybase `Client:Set`

Selects the Sybase object identified by the `Client:Set` string for object copying.

-mssharepoint `Client:Set`

Selects the Microsoft SharePoint Server 2007/2010 object identified by the `Client:Set` for object copying.

-veagent `Client:Set`

Selects the virtual environment object identified by the `Client:Set` string for object copying.

-idb `Client:Set`

Selects the Internal Database backup object identified by the `Client:Set` string for object copying.

-saphana `Client:Set`

Selects the SAP HANA backup object identified by the `Client:Set` string for object copying.

-integ {MySQL | PostgreSQL} `Client:Set`
Selects the MySQL or PostgreSQL backup object identified by the `Client:Set` string for object copying.

RETURN VALUES

For common return values, see the `omniintro` man page.

Additional return values of the `omniobjcopy` command are:

`10` - There was an error while copying some files. All agents completed successfully.

`11` - One or more agents failed, or there was a database error.

`12` - None of the agents completed the operation.

`13` - Session was aborted.

EXAMPLES

1. To start an interactive object copy session for copying two WinFS objects "system.company.com:/C 'Object1'" and "system.company.com:/C 'Object1'" from two different sessions to the device "DEV1", so that the source object version for "Object1" is then recycled, execute:

   ```
   omniobjcopy -winfs system.company.com:/C 'Object1' -session 2013/04/01-3 -
   targetdevice 'DEV1' -recycle -winfs systems.company.com:/C 'Object2' -session
   2013/04/25-9 -targetdevice 'DEV1'
   ```

2. To start an interactive object copy session for copying the whole restore chain of full and incremental backups for the filesystem object "system1.company.com:/ 'Label42'" from the session "2013/05/01-2", using the device "DEV1" to read the source objects and the device "DEV2" copy the objects, execute:

   ```
   omniobjcopy -filesystem system1.company.com:/ 'Label42' -session 2013/05/01-2 -
   sourcedevice 'DEV1' -targetdevice 'DEV2' -full
   ```

3. To start an interactive replication session for copying the whole restore chain of full and incremental backups from the session "2013/05/01-2", using the device "B2D1" as the source and the device "B2D2" as the target device, execute:

   ```
   omniobjcopy -replication -session 2013/05/01-2 -sourcedevice 'B2D1' -
   targetdevice 'B2D2' -full
   ```

4. To immediately start a post-backup object copy specification named "post_BU1" for the session "2013/05/03-1", and to make the devices available to this session with the highest priority in case of resource conflicts, execute:

   ```
   omniobjcopy -copylist post_BU1 -postbackup -session 2013/05/03-1 -priority 1
   ```

5. To immediately start a scheduled object copy specification named "CopySpec", use:

   ```
   omniobjcopy -copylist CopySpec -scheduled
   ```

6. To immediately start a scheduled replication specification named "ReplicSpec", use:

   ```
   omniobjcopy -replist ReplicSpec -scheduled
   ```

7. To restart a failed post-backup object copy specification "2013/03/16-10", use:

   ```
   omniobjcopy -restart "2013/03/16-10"
   ```

8. To start an interactive object copying session for the MySQL object "mysqlsys.company.com:MYSQL56.1408541577" from the session "2014/11/20-21", setting the level of protection for the object copy to five weeks, and suppressing messages from the command output, execute:

   ```
   omniobjcopy -no_monitor -integ MySQL mysqlsys.company.com:MYSQL56.1408541577 -
   session 2014/11/20-21 -targetproctect weeks 5
   ```

SEE ALSO

omnib(1), omnikeytool(1M), omniobjconsolidate(1), omniobjverify(1), omnir(1)

omniobjverify(1)

omniobjverify — verifies Data Protector backup objects, either interactively or using pre-configured post-backup, or scheduled verification specifications
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omniobjverify -version|-help

omniobjverify -verificationlist *VerificationSpecificationName* -scheduled [*GENERAL_OPTIONS*]

omniobjverify -verificationlist *VerificationSpecificationName* -postbackup -session *SessionID* [*GENERAL_OPTIONS*]

omniobjverify [*GENERAL_OPTIONS*] *Object[[Object]...]*

GENERAL_OPTIONS

[-verify_on_source|-verify_on_mahost|-verify_on_host *hostname*]

[-locationpriority *MediumLocation* [*MediumLocation*]...]

[-no_monitor]

[-priority*NumValue*]

MediumLocation

= "=*MediumLocation*"|"<*MediumLocation*"

Object

{-filesystem|-winfs|-rawdisk}*Client:ObjectName Label*

-session *SessionID*

[-copyid *N* [-fixedcopy...]]

[-sourcedevice *LogicalDevice*]

Object

{-sap|-oracle8|-integ {MySQL | PostgreSQL}|-informix|-msese|-e2010|-mssql|-lotus|
-mbx|-sapdb|-saphana|-msvssw|-db2|-sybase|-mssharepoint|-veagent|-idb}
Client:ObjectName

-session *SessionID*

[-copyid *N* [-fixedcopy...]]

[-sourcedevice *LogicalDevice*]

DESCRIPTION

The omniobjverify command verifies backup objects that have been created by Data Protector backup, object copy, or object consolidation sessions. You can use the omniobjverify command to verify objects such as filesystems (UNIX or Windows), very big file systems, disk image sections, and the Data Protector Internal Database (IDB).

The command can be also used to verify integration objects (SAP R/3, Oracle, MySQL, PostgreSQL, Informix Server, VMware vSphere, Microsoft Hyper-V, Microsoft Exchange Server 2007, Microsoft Exchange Server 2010/2013, Microsoft Exchange Server single mailboxes, Microsoft SharePoint Server 2007/2010/2013, Microsoft SQL Server, Lotus, Sybase, DB2, Microsoft Volume Shadow Copy Service, SAP MaxDB, and SAP HANA Appliance). It verifies the data integrity of the objects and the ability of Data Protector to deliver them to the application integration, not the application integration's ability to restore them.

To obtain the information about all backed up objects or sessions containing the objects you want to verify, use the omnidb command.

This command can be used to start an interactive object verification session or immediately start an automated (scheduled or post-backup) object verification specification.

From Data Protector 9.05 onwards for VMware integrations, the virtual machine disks are considered as objects that run in parallel. When the virtual machine object is selected for object copy operations, its associated disk objects are also considered for these copy operations.

OPTIONS

-version

Displays the version of the omniobjverify command.

-help

Displays the usage synopsis for the omniobjverify command.

-verificationlist *VerificationSpecificationName*

Specifies the name of the verification specification, identified by *VerificationSpecificationName*, for object verification.

-scheduled

Immediately starts a scheduled verification specification.

-postbackup

Immediately starts a post-backup verification specification specified by the -session*SessionID* option.

-session *SessionID*

Selects the session ID for the -postbackup option or for the object definition.

-verify_on_source

Specifies the original backup object source host as the host on which the object verification process

will be performed.

-verify_on_mahost

Specifies the Media Agent host as the host on which the object verification process will be performed.

-verify_on_host *hostname*

Specifies the host identified by *hostname* as the host on which the object verification process will be performed.

-locationpriority *MediumLocation* [*MediumLocation*]

The order in which media are selected for object verification if the same object version exists in more than one location. By default, Data Protector automatically selects the most appropriate media set. Media location priority is considered if more than one media set equally match the conditions of the media set selection algorithm.

The priority must be specified in the form "=*MediumLocation*" (equal to) or "<*MediumLocation*" (lower priority than).

If you specify -locationpriority "=Loc1" "<Loc2" "=Loc3" "<Loc4", then Loc1 has the highest priority, Loc2 and Loc3 have a lower priority, and Loc4 has the lowest priority.

-no_monitor

If this option is used, the command displays only the session ID. By default, the command monitors the session and displays all messages.

-priority *NumValue*

In case multiple running sessions request access to a specific device at the same time, this option determines the order in which the sessions will be queued. The *NumValue* can be any value from 1 (the highest priority) to 6000 (the lowest priority). In case the option is not specified, the default value of 3000 is set. If a low priority session is running when a high priority session starts queuing, the currently running session is allowed to finish. When more sessions request access to a device with the same priority, any of these sessions might acquire access first.

-filesystem *Client:ObjectName Label*

Selects the filesystem identified by *Client:ObjectName Label* for object verification.

-winfs *Client:ObjectName Label*

Selects the Windows filesystem identified by the *Client:ObjectName Label* string for object verification.

-rawdisk *Client:ObjectName Label*

Selects the disk image identified by the *Client:Label* string for object verification. *ObjectName* is blank in this case

-copyid *N* [-fixedcopy]

Selects the specified copy of an object version as a source for object verification.

If -fixedcopy option is not specified, Data Protector selects the needed media set automatically. If several copies of the same object version exist in one session as a result of the object copy or object mirror operation, this option is obligatory.

-sourcedevice *LogicalDevice*

Specifies a logical device different from the one used for the backup to be used for reading backed up objects from the source media. By default (if this option is not specified), the original backup device is used for reading backed-up objects from the source media.

-sap *Client:ObjectName*

Selects the SAP R/3 object identified by the *Client:ObjectName* string for object verification.

-oracle8 *Client:ObjectName*

Selects the Oracle object identified by the *Client:ObjectName* string for object verification.

-informix *Client:ObjectName*

Selects the Informix Server object identified by the *Client:ObjectName* string for object verification.

-msese *Client:ObjectName*

Selects the Microsoft Exchange Server 2007 object identified by the *Client:ObjectName* string for object verification.

-e2010 *Client:ObjectName*

Selects the Microsoft Exchange Server 2010/2013 object identified by the *Client:ObjectName* string for object verification.

-mssql *Client:ObjectName*

Selects the Microsoft SQL Server object identified by the *Client:ObjectName* string for object verification.

-lotus *Client:ObjectName*

Selects the Lotus Notes/Domino Server object identified by the *Client:ObjectName* string for object verification.

-mbx *Client:ObjectName*

Selects the Microsoft Exchange Server single mailbox object identified by the *Client:ObjectName* string for object verification.

-sapdb *Client:ObjectName*

Selects the SAP MaxDB object identified by the *Client:ObjectName* string for object copying.

-msvssw *Client:ObjectName*

Selects the Microsoft Volume Shadow Copy Service object identified by the *Client:ObjectName* string for object verification.

-db2 *Client:ObjectName*

Selects the DB2 object identified by the *Client:ObjectName* string for object verification.

-sybase *Client:ObjectName*

Selects the Sybase object identified by the *Client:ObjectName* string for object verification.

-mssps *Client:ObjectName*

Selects the Microsoft SharePoint Portal Server object identified by the *Client:ObjectName* string for object verification.

-mssharepoint *Client:ObjectName*

    Selects the Microsoft SharePoint 2007/2010 Server object identified by *Client:ObjectName* string for object verification.

-veagent *Client:ObjectName*

    Selects the virtual environment object identified by the *Client:ObjectName* string for object verification.

-idb *Client:Set*

    Selects the Internal Database backup object identified by the *Client:Set* string for object verification.

-saphana *Client:Set*

    Selects the SAP HANA backup object identified by the *Client:Set* string for object verification.

-integ {MySQL | PostgreSQL} *Client:Set*

    Selects the MySQL or PostgreSQL backup object identified by the *Client:Set* string for object verification.

RETURN VALUES

For common return values, see the omniintro man page.

Additional return values of the omniobjverify command are:

10 - There was an error while copying some files. All agents completed successfully.

11 - One or more agents failed, or there was a database error.

12 - None of the agents completed the operation.

13 - Session was aborted.

EXAMPLES

1. To start an interactive object verification session for verifying one WinFS object "system.company.com:/C 'Object1'" from session 2013/02/06-1, using the original host as the verification host, execute:

   ```
   omniobjverify -winfs system.company.com:/C 'Object1' -session 2013/02/06-1
   ```

2. To start an interactive verification session for verifying two filesystem objects "system1.company.com:/ 'Label1'" and "system1.company.com:/ 'Label2'" from session 2013/03/01-2, on host "system2.company.com", execute:

   ```
   omniobjverify -verify_on_host system2.company.com -filesystem
   system1.company.com:/ 'Label1' -session 2013/03/01-2 -filesystem
   system1.company.com:/ 'Label2' -session 2013/03/01-2
   ```

3. To immediately start a post-backup verification specification named "post_bu_verify1" for the session "2013/01/03-1", execute:

```
omniobjverify -verificationlist post_bu_verify1 -postbackup -session
2013/01/03-1
```

4. To immediately start a scheduled verification specification named "sched_verify1", execute:

```
omniobjverify -verificationlist sched_verify1 -scheduled
```

5. To start an object verification session for the MySQL object "
mysqlsys.company.com:MYSQL58.1411047241" from the session "2014/11/21-5", using the
backup device "FastestDriveOfAll ",execute:

```
omniobjverify -integ MySQL mysqlsys.company.com:MYSQL58.1411047241 -session
2014/08/21-5 -sourcedevice FastestDriveOfAll
```

SEE ALSO

omnib(1), omnidb(1), omnikeytool(1M), omniobjconsolidate(1), omniobjcopy(1), omnir(1)

omnir(1)

omnir — restores filesystems, disk images, the Data Protector Internal Database (IDB), Microsoft Exchange Server single mailboxes and Public Folders, Microsoft Exchange Server 2007, Microsoft Exchange Server 2010/2013, Microsoft SQL Server, Microsoft SharePoint Server 2007/2010/2013, MySQL, PostgreSQL, SAP R/3, SAP MaxDB, Informix Server, VMware vSphere, VMware vCloud Director, Microsoft Hyper-V, Lotus, IBM DB2 UDB, and NDMP objects backed up with Data Protector. The command is also used to start the instant recovery process. To restore a Sybase database, see the `syb_tool` man page.
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnir *SESSION_OPTIONS* [-noexpand] *Object* [*Object* ...]

 SESSION_OPTIONS

 -[no_]preview

 -report {warning|minor|major|critical}

omnir -resume *SessionID* [-no_monitor]

FILESYSTEM RESTORE

Object

 {-filesystem|-winfs} *Client:MountPoint Label*

 -session *SessionID* [-copyid *CopyID*]

 -tree *TreeName*...

 [*DATA_OPTIONS*]

 [*FILESYSTEM_OPTIONS*]

 [*SPLIT_MIRROR_OPTIONS*]

 [*GENERAL_OPTIONS*]

Object

 {-filesystem|-winfs} *Client:MountPoint Label*

 -full [-session *SessionID*]

 -tree *TreeName*...

 [*DATA_OPTIONS*]

 [*FILESYSTEM_OPTIONS*]

[*SPLIT_MIRROR_OPTIONS*]

[*GENERAL_OPTIONS*]

Object

{-filesystem|-winfs}*Client:MountPoint Label*

-omit_deleted_files [-session *SessionID* [-copyid *CopyID*]]

-overwrite

-tree*TreeName*...

[*DATA_OPTIONS*]

[*FILESYSTEM_OPTIONS*]

[*SPLIT_MIRROR_OPTIONS*]

[*GENERAL_OPTIONS*]

Object

{-filesystem|-winfs}*Client:MountPoint Label*

-tree *TreeName*...

MEDIUM_OPTIONS

[*DATA_OPTIONS*]

[*FILESYSTEM_OPTIONS*]

[*GENERAL_OPTIONS*]

Object

-host *Clientname*

-session *SessionID*

[-full|-omit_deleted_files -overwrite]

[*FILESYSTEM_OPTIONS*]

[*GENERAL_OPTIONS*]

DISK IMAGE RESTORE

Object

-rawdisk *Host Label*

-session *SessionID* [-copyid *CopyID*]

-section [*ToSection1=*]*Section1* [-section *ToSection2=Section2*...]

[*SPLIT_MIRROR_OPTIONS*]

[*GENERAL_OPTIONS*]

Object

-rawdisk *Host Label*

-section [*ToSection1=*]*Section1* [-section *ToSection2=Section2*...]

MEDIUM_OPTIONS

[*GENERAL_OPTIONS*]

INSTANT RECOVERY

omnir -host *ClientName*

 -session *SessionID*

 -instant_restore

[*P9000_DISK_ARRAY_XP_OPTIONS* | *P6000_ENTERPRISE_VIRTUAL_ARRAY_OPTIONS* | *3PAR_DISK_ARRAY_OPTIONS*]

[*ORACLE_SPECIFIC_OPTIONS*]

[*SAP_SPECIFIC_OPTIONS*]

P9000_DISK_ARRAY_XP_OPTIONS

 -keep_version

 -check_config

P6000_ENTERPRISE_VIRTUAL_ARRAY_OPTIONS

{-copyback wait_clonecopy *Minutes* | -switch}

{-leave_source | -no_leave_source}

{-check_config | -no_check_config}

[-force_prp_replica]

3PAR_DISK_ARRAY_OPTIONS

{-copyback wait_clonecopy *Minutes*}

{-check_config | -no_check_config}

[-force_prp_replica]

[-force_restore_volset]

SAP_SPECIFIC_OPTIONS

 -sap

 -user *UserName* -group *GroupName*

 -recover {now | time *MM/DD/YY hh:mm:ss* | logseq *LogSeqNumber* thread *ThreadNumber* | SCN *Number*} [-open [-resetlogs]]

 -appname *ApplicationDatabaseName*

ORACLE_SPECIFIC_OPTIONS

 `-oracle`

 `-user` *UserName* `-group` *GroupName*

 `-recover` {`now`|`time` *MM/DD/YY hh:mm:ss*|`logseq` *LogSeqNum* `thread` *ThreadNum*|`SCN` *Number*} [`-open` [`-resetlogs`]]

 `-appname` *ApplicationDatabaseName*

 `-parallelism` *Number*

NDMP RESTORE

Object

 `-filesystem` *Host:MountPoint Label*

 `-full` [`-session` *SessionID*]

 `-device` *BackupDevice*

 `-tree` *TreeName*...

 [*NDMP_DATA_OPTIONS*]

 [*NDMP_GENERAL_OPTIONS*]

Object

 `-filesystem` *Host:MountPoint Label*

 `-session` *SessionID* [`-full`]

 `-device` *BackupDevice*

 `-tree` *TreeName*...

 [*NDMP_DATA_OPTIONS*]

 [*NDMP_GENERAL_OPTIONS*]

NDMP_DATA_OPTIONS

 `-into` *PathName*

 `-ndmp_env` *FileName*

 `-ndmp_user` *UserName*

 `-ndmp_passwd` *Password*

NDMP_GENERAL_OPTIONS

 `-server` *ServerName*

 `-no_monitor`

 `-variable` *VariableName VariableValue*

SAP R/3 FILE RESTORE

Object

 -sap *Client:Set*

 -session *SessionID* [-copyid *CopyID*]

 -tree *FileName*...

 [*DATA_OPTIONS*]

 [*FILESYSTEM_OPTIONS*]

 [*GENERAL_OPTIONS*]

VIRTUAL ENVIRONMENT RESTORE

omnir -veagent

 -virtual-environment {vmware | hyperv | vcd}

 -barhost *BackupHost*

 -apphost *OriginalAppHost*

 -instance *OriginalDatacenter*

 -method {vStorageImage | vCDvStorageImage | vStorageImageOpenStack}

 [-session *BackupID*|

 -fromsession *BackupID* -untilsession *BackupID*]

 VirtualMachine [*VirtualMachine* ...]

 [*NewInstance* | *Directory* | *NewOrganization*]

 [*RESTORE_OPTIONS*]

 VirtualMachine

 -vm *vmpath* -instanceUUID *vmInstanceUUID* [-versionID *VersionID* -new_name *NewVirtualMachineName*] [-disk *DiskName* ...]

 NewInstance

 -newinstance *TargetDatacenter*

 [-store *TargetDatastore*]

 [-network_name *TargetNetwork*]

 [-destination *RestoreClient*]

 [-host/cluster *HostOrCluster*]

 [-resourcePool *ResourcePool*]

 [-specificHost *SpecificHost*]

[-targetstoragepath *TargetStoragePathOfAllHyper-V-VMs*]

NewOrganization

-neworganization *TargetOrganization*

[[-virtual_datacenter_path|-virtual_datacenter_uuid] *TargetVDC*]

[[-vapp_path|-vapp_uuid] *TargetVApp*]

[[-vcenter_path|-vcenter_uuid] *TargetVCenter*]

[[-network_name|-network_uuid] *TargetNetwork*]

Directory

-directory *RestoreDirectory*

[-overwrite|-skip|-latest]

RESTORE_OPTIONS

[-consolidate]

[-register]

[-poweron]

[-deletebefore|-deleteafter|-skip|-keep_for_forensics]

[-removeSnapshots]

SAP MAXDB RESTORE

omnir -sapdb

 -barhost *ClientName*

 -instance *InstanceName*

[-destination *ClientName*]

[-newinstance *DestinationInstanceName*]

[-session *BackupID*]

[-recover [-endlogs|-time:*YYYY-MM-DD.hh.mm.ss*] [-from_disk]]

[-nochain]

INFORMIX SERVER RESTORE

omnir -informix

 -barhost *ClientName*

 -barcmnd *PathName*

```
-user User:Group

-appname ApplicationDatabaseName

-bararg OnBarRestoreArguments

[SESSION_OPTIONS]

[GENERAL_OPTIONS]


SESSION_OPTIONS

-report {warning|minor|major|critical}

-load {low|medium|high}

-no_monitor
```

MICROSOFT EXCHANGE SERVER 2007 RESTORE

```
omnir -msese

 -barhost ClientName

[-destination ClientName]

 -appname full_application_name

{-base DBName -session BackupID}...

 -logpath Path

[-last [-mount] [-consistent]]

[GENERAL_OPTIONS]
```

MICROSOFT EXCHANGE SERVER 2010/2013 RESTORE

```
STANDARD RESTORE
omnir -e2010

 -barhost ClientName

 Database [Database ...]

[-user User:Domain]

[VSS_EXCHANGE_SPECIFIC_OPTIONS]

[GENERAL_OPTIONS]


INSTANT RECOVERY
omnir -e2010

 -barhost ClientName
```

```
-instant_restore

Database [Database ...]

[-user User:Domain]

[VSS_INSTANT_RECOVERY_OPTIONS]

[VSS_EXCHANGE_SPECIFIC_OPTIONS]

[GENERAL_OPTIONS]


Database

{-db_name SourceDatabaseName | -db_guid SourceDatabaseGUID}

[-source SourceClientName]

{-repair | -latest | -pit | -new | -temp} E2010_METHOD_OPTIONS


E2010_REPAIR_METHOD_OPTIONS

[-no_resume_replication]


E2010_LATEST_METHOD_OPTIONS

[-node TargetNode... | -all]

[-no_resume_replication]

[-no_recover]

[-no_mount]

[E2010_IR_SPECIFIC_OPTIONS]


E2010_PIT_METHOD_OPTIONS

-session ID

[-node TargetNode... | -all]

[-no_resume_replication]

[-no_recover]

[-no_mount]

[E2010_IR_SPECIFIC_OPTIONS]


E2010_NEW_METHOD_OPTIONS

-session ID

-client TargetClientName

-location TargetDatabasePath
```

```
-name TargetDatabaseName

[-recoverydb]

[-no_recover]

[-no_mount]

[E2010_IR_SPECIFIC_OPTIONS]


E2010_TEMP_METHOD_OPTIONS

-session ID

-client TargetClientName

-location TargetDatabasePath

[-no_chain]

[-edb_only]

[-no_recover]

[E2010_IR_SPECIFIC_OPTIONS]


E2010_IR_SPECIFIC_OPTIONS

[-from_session SessionID]
```

MICROSOFT EXCHANGE SINGLE MAILBOX RESTORE

```
omnir -mbx
 -barhost HostName
[-destination HostName]
-mailbox MailboxName -session BackupID [MAILBOX_OPTIONS]...
-public -session BackupID [PUBLIC_FOLDERS_OPTIONS]
[GENERAL_OPTIONS]


MAILBOX_OPTIONS

-folder FolderName

-exclude FolderName

-originalfolder {-keep_msg | -overwrite_msg}

-destmailbox DestMailboxName

-chain


PUBLIC_FOLDERS_OPTIONS
```

    -folder *FolderName*

    -exclude *FolderName*

    -originalfolder {-keep_msg | -overwrite_msg}

    -chain

MICROSOFT SQL SERVER RESTORE

omnir -mssql

 -barhost *ClientName*

[-destination *ClientName*]

[-instance *SourceInstanceName*]

[-destinstance *DestinationInstanceName*]

{ -base *DBName* -session *BackupID* [*MSSQL_OPTIONS*]... | -base *DBName* -datafile
GroupName/DataFileName -session *BackupID* [*DATAFILE_OPTIONS*...]}

[*GENERAL_OPTIONS*]

MSSQL_OPTIONS

 -asbase *NewDBName* {-file *LogicalFileName1 PhysicalFileName1* [-file *LogicalFileName2
PhysicalFileName2* ...]}

 -replace

 -singleuser

 -nochain

 -recovery {rec | norec}

 -stopat *yyyy/mm/dd.hh:mm:ss*

 -standby *File*

 -tail_log *BackupSpecificationName*

DATAFILE_OPTIONS

 -replace

 -singleuser

 -nochain

 -recovery {rec | norec}

MICROSOFT SHAREPOINT SERVER 2007/2010/2013 RESTORE

omnir -mssharepoint

 -barhost *HostName*

[-destination *RestoreClientName*]

-user *User*:*Group*

[-session *BackupID*]

[-replace]

[-byserver *ServerName* [-byserver *ServerName*...]]

-farmname *FarmName*

[*Component* [*Component*...]]

[*GENERAL_OPTIONS*]

Component

 -configdb |

 -webapplication *WebApplicationName* [WEB_APPLICATION_OPTIONS] [*ContentDatabase*
[*ContentDatabase*...]] |

 -ssp *SSPName* [*SSP_OPTIONS*] [-index *INDEX_OPTIONS*] [*Database* [*Database*...]]

 [-webapp *WebApplicationName* [*WEB_APPLICATION_OPTIONS*] [*ContentDatabase*
[*ContentDatabase*...]]]

 -wsssearch [*Database*] |

 -ssodb [*DB_OPTIONS*]

ContentDatabase

 -db *DBName* -host *DBHostName* [-unlink] [*DB_OPTIONS*]

Database

 -db *DBName* -host *DBHostName* [*DB_OPTIONS*]

WEB_APPLICATION_OPTIONS

 -as *WebApplicationName*

 -url *WebApplicationURL*

 -poolusername *Username* [-poolpassword *Password*]

 -replace

DB_OPTIONS

 -sqllogin *Username* [-sqlpassword *Password*]

 -instance *SourceInstanceName*

 -as *NewDBName*

    -tohost *DBHostName*

    -newinstance *DestinationInstanceName*

    -todir *NewDirectoryName*

    -replace

    SSP_OPTIONS

    -ssplogin *Username* [-sspassword *Password*]

    -as *SSPName*

    -mysiteurl *MySiteWebAppUrl*

    INDEX_OPTIONS

    -tohost *IndexServerHostName*

    -todir *NewDirectoryName*

LOTUS RESTORE

omnir -lotus

 -barhost *ClientName*

[-user *User:Group*]

[-destination *ClientName*]

[-parallelism *n*]

 -domino_server *srv_name*

 -appname

-db *db1* [-db *db2*...]

[-NSF] [-NTF] [-BOX] [-ALL]

[-direx *direx1* [-direx *direx2*...]]

[-r_dest *restore_dir*]

[-recover | recovery_time *yyyy/mm/dd.hh:mm:ss*]

[-reset_replica]

[-session *BackupID*]

MICROSOFT VOLUME SHADOW COPY SERVICE RESTORE

STANDARD RESTORE

omnir -vss

    -barhost *ClientName*

    -session *BackupID1* {*Tree* [*Tree*...]}

    [-session *BackupID2* {*Tree* [*Tree*...]}...]

    [-no_recovery]

    [-into *PathName*]

    [-destination *ClientName*]

    [*VSS_EXCHANGE_SPECIFIC_OPTIONS*]

    [*GENERAL_OPTIONS*]

    INSTANT RECOVERY

    omnir -vss

    -instant_restore

    -barhost *ClientName*

    -session *SessionID1* {*Tree* [*Tree*...]}

    [-session *SessionID2* {*Tree* [*Tree*...]}...]

    [-no_recovery]

    [-destination *ClientName*]

    [*VSS_INSTANT_RECOVERY_OPTIONS*]

    [*VSS_EXCHANGE_SPECIFIC_OPTIONS*]

    [*GENERAL_OPTIONS*]

    Tree

    -tree *TreeName* [*VSS_EXCHANGE_2007_SPECIFIC_OPTIONS*]

    VSS_INSTANT_RECOVERY_OPTIONS

    [-conf_check {strict|non-strict|disabled}]

    [-no_retain_source]

    [-use_vds|-use_vss|*VSS_P6000_ENTERPRISE_VIRTUAL_ARRAY_OPTIONS*|*VSS_P9000_DISK_ARRAY_XP_OPTIONS*|*VSS_P10000_OPTIONS*|*VSS_P4000_OPTIONS*]

    VSS_P6000_ENTERPRISE_VIRTUAL_ARRAY_OPTIONS

    [-no_copy_back|-copy_back [-diskarray_wait *Minutes*|-no_diskarray_wait]]

    [-no_retain_source]

    VSS_P9000_DISK_ARRAY_XP_OPTIONS

```
-copy_back -no_retain_source [-no_diskarray_wait]
```

VSS_P10000_OPTIONS

```
-copy_back
```

VSS_P4000_OPTIONS

```
-copy_back
```

VSS_EXCHANGE_SPECIFIC_OPTIONS

```
[-exch_check [-exch_throttle Value] | -exch_checklogs]
```

VSS_EXCHANGE_2007_SPECIFIC_OPTIONS

```
[[-target_tree TargetStoreName | -exch_RSG LinkedStoreName] -target_dir Directory]
```

DB2 RESTORE

```
omnir -db2
 -barhost ClientName
 -instance InstName
 {[-dbname DBName [-session BackupID] [-newdbname NewDBName...]] [-tsname DBName*TSName [-
session BackupID] [-offline...]] [-logfile DBName*LogFileName [-session BackupID...]]}
 [DB2_OPTIONS]
```

DB2_OPTIONS

```
 -destination ClientName
 -rollforward [-time YYYY-MM-DD.hh.mm.ss]
 -frominstance InstName
```

MYSQL RESTORE

```
omnir SESSION_OPTIONS
-integ MySQL
 -barhost TargetMySQLHostname
 -appname TargetInstanceName
 -user Username:GroupName
 -options MYSQL_OPTIONS
 [GENERAL_OPTIONS]
```

SESSION_OPTIONS

[-report {warning | minor | major | critical}]

MYSQL_OPTIONS

-source_client *SourceMySQLHostname*

-source_instance *SourceInstanceName*

-database *DATABASE_OPTIONS* | -binary_log *BINARY_LOG_OPTIONS*

DATABASE_OPTIONS

–session*SessionID*

{-staging [*CustomStagePath*]

[–copy_back [–target_dir *NonOriginalTargetPath*]

–import] |

–inplace [–target_dir *NonOriginalTargetPath*]}

[–include {*DatabaseName* | *DatabaseName.TableName*}] ...

[-roll_forward [YYYY–MM–DD hh:mm:ss]]

BINARY_LOG_OPTIONS

-include *BinaryLogFilename* [-include *BinaryLogFilename*] ...

[-target_dir *NonOriginalTargetPath*]

POSTGRESQL RESTORE

omnir *SESSION_OPTIONS*

 -integ *PostgreSQL*

 -barhost *TargetPostgreSQLHostname*

 -appname *TargetInstanceName*

 -user *Username:GroupName*

 -options *POSTGRESQL_OPTIONS*

[*GENERAL_OPTIONS*]

SESSION_OPTIONS

[-report {warning | minor | major | critical}]

POSTGRESQL_OPTIONS

 -source_client *SourcePostgreSQLHostname*

 -source_instance *SourceInstanceName*

 -target_dir *NonOriginalTargetPath*

[-session *SessionID*] |

[−roll_forward [YYYY−MM−DD hh:mm:ss]]

DATA_OPTIONS

 -exclude *PathName* ...

 -skip *MatchPattern* ...

 -only *MatchPattern* ...

 -as *Pathname*

 -into *Pathname*

MEDIUM_OPTIONS

 -device *BackupDevice*

 -medium *MediumID*

 -id *DiskAgentID*

[-slot *SlotID* [*Side*]]

FILESYSTEM_OPTIONS

 -touch

 -lock

 -no_protection

 -[no_]overwrite | -merge

 -catalog

 -sparse

 -move_busy

 -no_share[_info]

 -omit_unrequired_object_versions

 -[no_]resumable

IDB RESTORE

 -idb

-barhost *ClientName*

[-restoredb [*RESTORE_DB_OPTIONS*]]

[-restoreconf [*RESTORE_CONF_OPTIONS*]]

[-restoredcbf [*RESTORE_DCBF_OPTIONS*]]

[-client *SourceClientName*]

[-until *YYYY-MM-DD*[*hh.mm.ss*]]

[-pre *PathName*]

[-post *PathName*]

[*GENERAL_OPTIONS*]

RESTORE_DB_OPTIONS

-targetdir *TargetDataFolderPath*

-port *TargetDatabasePort*

[-nodbrecover | -nouseasnewidb]

RESTORE_CONF_OPTIONS

[-keeprecent | -nooverwrite | -overwrite]

[-session *SessionID*]

[-targetdir *TargetConfFolderPath*]

[-name *FileOrFolderName*...]

RESTORE_DCBF_OPTIONS

[-targetdir *TargetDCBFFolderPath*]

GENERAL_OPTIONS

-device *BackupDevice*

-no_auto_device_selection

-server *ServerName*

-target *Client*

-profile

-load {low | medium | high}

-pre_exec *PathName*

-post_exec *PathName*

-variable *VariableName VariableValue*

```
-no_monitor
```

```
[-priorityNumValue]
```

SPLIT_MIRROR_OPTIONS

```
-sse|-symmetrix
```

```
-remote ApplicationSystem BackupSystem|-local ApplicationSystem BackupSystem|-
combined ApplicationSystem BackupSystem
```

```
[-quiesce cmd]
```

```
[-restart cmd]
```

```
[-mirrors list]
```

```
[-discovery]
```

```
[-re_establish_links_before_restore]
```

```
[-disable_disks]
```

```
[-restore_links_after_restore]
```

DESCRIPTION

The `omnir` command restores objects backed up using Data Protector. You can use the `omnir` command to restore filesystems, very big file systems, disk image sections, NDMP objects, and Data Protector Internal Database (IDB) to their original or different location. It can also be used for restoring application integration objects (SAP R/3, Microsoft Exchange Server 2007, Microsoft Exchange Server 2010/2013, Microsoft Exchange Server single mailboxes, Microsoft SQL Server, Microsoft SharePoint Server 2007/2010/2013, Lotus, Informix Server, VMware vSphere, VMware vCloud Director, Microsoft Hyper-V, IBM DB2 UDB, MySQL, or SAP MaxDB), or to start the instant recovery process. To restore a Sybase database, see the `syb_tool` man pages.

If several copies of the same object version exist, you can let Data Protector select which media set will be used for the restore. You can also specify the media set from which you want to restore the data, except when restoring the IDB or an integration object. It is not possible to specify the media set created as a result of the media copy operation.

The `omnir` command also supports parallel restore. You can achieve this by specifying more than one object using the command line options. It is not possible to use the `-medium` option when performing a parallel restore. The number of objects for parallel restore is limited by the `MaxSessions` global option.

NOTE: It is not allowed to specify the same object more than once within the same `omnir` command. To differentiate options for the same object (for example, the `-tree` option) specify these options for the same object as many times as needed.

Information about all backed up objects can be obtained from the IDB by using `omnidb` command or, in the case of the instant recovery, from a ZDB database or VSS database by using the `omnidbxp`, `omnidbsmis`, or `omnidbvss` command. For more information on these commands, see the related man pages. For most restore actions you need to specify the *SessionID* of the session containing the object you want to restore, which can be obtained by the `omnidb` command.

NOTE: When restoring integration objects, provide the *SessionID* of the backup session. In case of object copies, do not use the object copy session ID, but the object's *BackupID*, which equals the original object's backup session ID. If imported backup media are used for restoring an object, do not specify the new session ID which is assigned to the imported backup session, but the object's *BackupID* which is the original backup session ID for that object.

To restore objects from a medium that is not in the IDB, use the `-medium` *MediumID* option, instead of the *SessionID*.

NOTE: The `-medium` option is not possible when performing a parallel restore.

To get the *MediumID* and *DiskAgentID* from the medium, use the `omnimlist` command to read the medium. See the `omnimlist` man page for more information on this command.

NOTE: When restoring a Microsoft SQL Server with the `-tail_log` option specified, a tail log backup session is performed before the actual restore session starts.

OPTIONS

`-version`

Displays the version of the `omnir` command.

`-help`

Displays the usage synopsis of the `omnir` command.

`-resume` *SessionID*

Starts a new session that continues with the restore from where the failed session *SessionID* left off, using the same options as used in the failed session. This functionality is supported for filesystem restore sessions and Data Protector Oracle Server integration restore sessions.

FILESYSTEM RESTORE

`-filesystem` *Client:MountPoint Label*

Selects the filesystem identified with *Client:MountPoint Label* for restore.

`-winfs` *Client:MountPoint Label*

Selects the Windows filesystem identified with *Client:MountPoint Label* for restore.

`-session` *SessionID*

Specifies the session to be used for restore.

`-copyid` *CopyID*

If several copies of the same object exist in one session as a result of the object copy or object mirror operation, this option identifies the specific object copy (object mirror or object copy) to be used for restore. By default (if this option is not specified), Data Protector selects the media set to restore from automatically. When using this option, it is necessary to specify both the object and the session.

`-tree` *TreeName*

Specifies the file, component, or tree to restore. Note that when specifying trees on UNIX systems,

complete trees must be specified including the mount points, whereas on Windows systems, trees must be specified without volumes (drives). For example: `-tree /usr/temp` (UNIX systems) and `-tree /temp/Filesystem/E` (Windows systems).

`-full`

Specifies that the selected object will be restored from the last full backup and all incremental backups related to this full backup.

`-omit_deleted_files`

This option can only be used in combination with the `-overwrite` option. 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 specified, Data Protector recreates the state of the backed up directory tree at the time of the chosen 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 not restored.

If this option is not specified, 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.

If you use this option in combination with the `-as` or `-into` option, carefully choose the restore location to prevent accidental removal of existing files.

`-host` *ClientName*

Restores all objects of the specified client that were backed up in the specified session. This option is only valid for the filesystem restore.

DISK IMAGE RESTORE

`-rawdisk` *Client Label*

Selects the disk image identified by *Client* and *Label* for restore.

`-session` *SessionID*

Specifies the session to be used for restore.

`-copyid` *CopyID*

If several copies of the same object exist in one session as a result of the object copy or object mirror operation, this option identifies the specific object copy (object mirror or object copy) to be used for restore. By default (if this option is not specified), Data Protector selects the media set to restore from automatically. When using this option, it is necessary to specify both the object and the session.

`-section` [*ToSection*=]*Section*

Specifies the disk image section to be restored. To restore the section to a new section, include both the source and destination section.

NDMP RESTORE

-full

> Specifies that the selected object will be restored from the last full backup and all incremental backups related to this full backup.

-filesystem *Client:MountPoint Label*

> Selects the filesystem identified with `Client:MountPoint Label` for restore.

-session *SessionID*

> Specifies the session to be used for restore.

-tree *TreeName*

> Specifies the file, component, or tree to restore. Note that when specifying trees on UNIX systems, complete trees must be specified including the mount points, whereas on Windows systems, trees must be specified without volumes (drives). For example: `-tree /usr/temp` (UNIX system) and `-tree /temp/Filesystem/E` (Windows system).

-into *Pathname*

> Restores the selected fileset into the given directory.

-ndmp_user *UserName*

> Sets the username that is used by Data Protector to establish the connection to the NDMP server.

-ndmp_passwd *Password*

> Sets the password for the username that is used by Data Protector to establish the connection to the NDMP server.

-ndmp_env *FileName*

> Specifies the filename of file with NDMP environment variables for specific NDMP implementations.

SAP R/3 FILE RESTORE

-sap *Client:Set*

> Selects the SAP R/3 object identified by `Client:Set` for restore.

-session *SessionID*

> Specifies the session to be used for restore.

-copyid *CopyID*

> If several copies of the same object exist in one session as a result of the object copy or object mirror operation, this option identifies the specific object copy (object mirror or object copy) to be used for restore. By default (if this option is not specified), Data Protector selects the media set to restore from automatically. When using this option, it is necessary to specify both the object and the session.

-tree *TreeName*

> Specifies the file, component, or tree to restore. Note that when specifying trees on UNIX systems, complete trees must be specified including the mount points, whereas on Windows systems, trees must be specified without volumes (drives). For example: `-tree /usr/temp` (UNIX system) and -

```
tree /temp/Filesystem/E
```
(Windows system).

INFORMIX SERVER RESTORE

`-informix`

Selects the Informix Server object for restore.

`-barhost` *ClientName*

Specifies the Informix Server client from which the data was backed up.

`-barcmnd` *PathName*

The value of the `barcmnd` option has to be set to `ob2onbar.pl`.

`-user` *UserName:GroupName*

Specifies *Username* and *GroupName* that started the script specified by the `-barcmnd` option.

`-appname` *ApplicationDatabaseName*

Specifies the database server name of Informix Server to be restored.

`-bararg` *OnBarRestoreArguments*

Specifies the onbar restore arguments. Each onbar restore argument has to be put in double quotes.

MICROSOFT EXCHANGE SERVER 2007 RESTORE

`-msese`

Selects the Microsoft Exchange Server object for restore.

`-barhost` *ClientName*

Specifies the Microsoft Exchange Server client from which the data was backed up.

`-destination` *ClientName*

Specifies the target client for restore. Use this option only when you restore to some other instance than the one that was backed up.

`-appname` *full_application_name*

Specifies a Microsoft Exchange Server Information Store, Site Replication Service or Key Management Service for the restore. The name of the Store/Service (*full_application_name*) must be provided in double quotes as follows:

- For the Information Store: `Microsoft Exchange Server (Microsoft Information Store)`
- For the Site Replication Service: `Microsoft Exchange Server (Microsoft Site Replication Service)`
- For the Key Management Service: `Microsoft Exchange Server (Microsoft Key Management Service)`

`-base` *DBName*

Specifies the Microsoft Exchange Server store or logs for restore.

`-session` *BackupID*

Specifies from which backup data (`BackupID`) to restore, for example, `2011/10/09-2`.

A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

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

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

This option must be set for every `-base` option specified.

`-logpath` *path*

Specifying this option, you set the temporary directory for the Microsoft Exchange Server log files. Data Protector restores the log files to this directory. Using this directory, the Microsoft Exchange Server then recovers the database - this operation is referred to as hard recovery.

`-last`

Hard recovery is performed after the restore of the Microsoft Exchange Server object. Use this option if you are restoring the last set of files. If you do not set this option, you have to start the recovery manually by running the `eseutil /cc /t` utility from the directory for temporary log files. If this option is not specified, soft recovery is performed after the restore.

`-mount`

The restored Microsoft Exchange Server databases will be automatically mounted after the soft or hard recovery.

`-consistent`

Restores the database to its last consistent state. The latest log files, created after backup, are applied to the restored database during recovery.

MICROSOFT EXCHANGE SERVER 2010/2013 RESTORE

`-e2010`

Selects the Microsoft Exchange Server 2010/2013 object for restore.

`-barhost` *ClientName*

Specifies on which client to start the Data Protector Microsoft Exchange Server 2010 integration agent (`e2010_bar.exe`). This can be any client that has the `MS Exchange Server 2010+ Integration` component installed.

`-instant_restore`

Performs an instant recovery.

`-user` *User:Domain*

Specifies which Windows domain user account to use to start the restore session. Ensure that the

specified user has appropriate Microsoft Exchange Server permissions, is added to the Data Protector `admin` or `operator` user group, and is saved to a Windows Registry on the Microsoft Exchange Server client on which the integration agent (`e2010_bar.exe`) will be started (see the Data Protector `omnicc` command).

If this option is not specified, the restore session is started under the user account under which the `Data Protector Inet` service is running.

{-db_name *SourceDatabaseName* |-db_guid *SourceDatabaseGUID*}

Specifies which database to restore. If the database no longer exists, use the `-db_guid` option.

-source *SourceClientName*

Specifies from which client the database was backed up. For databases that are part of a DAG, specify the DAG virtual system (host). If this option is not specified, Data Protector assumes that the database was backed up from the client specified with the `-barhost` option.

{-repair |-latest |-pit |-new |-temp}

Specifies which restore method to use:

`repair`: Available only for databases that are part of a Microsoft Exchange Server Database Availability Group (DAG). Automatically restores all the corrupt passive copies (copies with the status `Failed` or `FailedAndSuspended`).

`latest`: Restores a corrupt database to the latest possible point in time.

`pit`: Restores an existing database to a specific point in time.

`new`: Restores files to a different database, either because the original database no longer exists or in order to move the data elsewhere.

`temp`: Restores files to a location of your choice.

-no_resume_replication

Specifies that the replication between the active and passive copies should not be resumed after the restore session completes.

-node *TargetNode* ... |-all

Specifies which clients (that is, database copies) to restore.

-no_recover

Specifies that logs should not be applied to the database file after the restore completes.

-no_mount

Specifies that the database should not be mounted after the database recovery completes.

-session {*BackupID* |*SessionID*}

Specifies from which backup data to restore, for example, `2012/10/09-2`.

For standard restore, specify *BackupID*. A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup

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

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

If a differential backup session is selected, the `.log` files backed up in the selected differential backup session are restored.

If an incremental backup session is selected, the `.log` files backed up in all subsequent incremental backup sessions, up to the selected incremental backup session, are restored.

For instant recovery, specify *SessionID* of a ZDB-to-disk or ZDB-to-disk+tape session.

-client *TargetClientName*

Specifies to which client to restore.

-location *TargetDatabasePath*

Specifies to which directory to restore.

-name *TargetDatabaseName*

Specifies which name to use for the new database. If another database with the same name already exists, the restore is not performed.

-recoverydb

Restores files to a Microsoft Exchange Server recovery database.

Although multiple recovery databases can exist in parallel, only one recovery database can be mounted to the Microsoft Exchange Server at a time.

-no_chain

Restores only the files backed up in the selected session.

By default, the complete chain is restored.

-edb_only

Restores only the database file (`.edb`). Logs (`.log`) and checkpoint files (`.chk`) are not restored.

-from_session

An instant recovery specific option that specifies which full or copy ZDB session to use as a starting session in a restore chain.

Use this option if the session that you specified for instant recovery is an incremental or a differential session. If you do not use it, the integration agent uses the last full or copy session as the starting point in a restore chain for instant recovery.

MICROSOFT EXCHANGE SINGLE MAILBOX RESTORE

-mbx

Selects Microsoft Exchange Server single mailboxes and Public Folders for restore.

-barhost *ClientName*

Specifies the Microsoft Exchange Server client from which the data was backed up.

-destination *ClientName*

Specifies the target client for restore. Use this option only when you restore to some other instance than the one that was backed up.

-mailbox *MailboxName*

Specifies the Microsoft Exchange Server single mailboxes for restore.

-session *BackupID*

Specifies from which backup data (*BackupID*) to restore, for example, 2012/10/09-2.

A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

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

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

-public

Specifies the Microsoft Exchange Server Public Folders for restore (as part of the Microsoft Exchange Server single mailbox restore).

-folder *FolderName*

Specifies folders to be restored. Note that the subfolders are also restored. If this option is not specified, all backed up folders are restored.

-exclude *FolderName*

Specifies the folders to be excluded from restore.

-originalfolder {-keep_msg | -overwrite_msg}

If this option is selected, Data Protector restores Exchange Server items to the same folders in which they were when the backup was performed.

If -keep_msg is selected, the messages in the mailbox or Public Folders are not restored, even if they are different from their backed up version.

If -overwrite_msg is selected, all messages are restored, replacing their current versions (if they exist). If different versions of the same message exist in the mailbox or Public Folders (for example, if you have a copy of the message), only one is replaced with the backed up version and all other versions remain intact.

The messages in the mailbox that were not backed up in the specified backup session (or the restore chain of backup sessions) always remain intact.

If -originalfolder is not specified, Data Protector creates a new folder in the root of the mailbox or in the root of All Public Folders and restores Exchange items into it. For a mailbox restore, the folder is named Data Protector *BackupDate BackupTime*, and for a Public Folders restore, it is named Data Protector *BackupDate BackupTime* - public folder. If you restore a mailbox or Public Folders from the same backup several times, a number is appended to the folder name. For example, in the second restore session of a mailbox, the folder Data Protector *BackupDate BackupTime* (1) is created.

-destmailbox *DestMailboxName*

Specifies the destination mailbox, into which data will be restored. The destination mailbox must exist on the target Microsoft Exchange Server. If this option is not specified, data is restored to the original mailbox.

-chain

If this option is specified, data is restored not only from the specified backup session, but also from the latest full, the latest incremental1 (if exists), and all incremental backups from the last incremental1 up to the specified version.

LOTUS RESTORE

-lotus

Selects the Lotus Notes/Domino Server object for restore.

-barhost *ClientName*

Specifies the Lotus Notes/Domino Server client from which the data was backed up.

-destination *ClientName*

Specifies the target client for restore. Use this option only when you restore to some other instance than the one that was backed up.

-parallelism *n*

Sets the number of restore streams, running in parallel. The default is 1.

-domino_server *srv_name*

Sets the name of the Lotus Notes/Domino Server which you want to restore.

-appname

Specifies the Lotus Notes/Domino Server instance source.

-db *db*

Sets the restore of an individual Lotus Notes/Domino Server database.

-NSF

Sets the restore of all NSF (Notes Storage Facility) databases.

-NTF

Sets the restore of all NTF (Notes Templates Facility) files.

-BOX

Sets the restore of all BOX files.

-ALL

Sets the restore of all objects, NSF databases, NTF files and BOX files.

-dir *dir*

Sets the Lotus Notes/Domino data directories that you want to include in the restore. Enter their relative pathnames to the Lotus Notes/Domino data directory.

-direx *direx*

Sets the Lotus Notes/Domino data directories that you want to exclude from the restore. Enter their relative pathname to the Lotus Notes/Domino data directory.

-r_dest *restore_dir*

Sets the relative pathname to the restored database directory.

-recover

Specify this option to perform the recovery of the restored database to the last possible consistent state.

-recovery_time *yyyy/mm/dd.hh:mm:ss*

Sets a point in time to which you want the database to be recovered.

-reset_replica

This option should be used only when restoring to the last possible consistent state. If the option is specified, each restored storage database (NSF database) is assigned a new replica ID.

-session *BackupID*

Specifies from which backup data (*BackupID*) to restore, for example, 2012/10/09-2.

A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

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

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

DB2 RESTORE

-db2

Selects the IBM DB2 UDB object to restore.

-barhost *ClientName*

Specifies the IBM DB2 UDB client from which the data was backed up.

-instance *InstName*

Sets the name of the database instance that was backed up.

-dbname *DBName*

Sets the name of the DB2 database that you want to restore.

-newdbname *NewDBName*

Specify this option if you want to restore the whole DB2 database into a new database.

-tsname *DBName*TSName*

Sets the name of the DB2 table space that you want to restore. To specify the table space you would like to restore, write the name of the database, then the "*" character and finally the name of the table space (without spaces).

-logfile *DBName*LogFileName*

Sets the name of the DB2 Log file that you want to restore. It should not be used with the -rollforward option. To specify the Log file you would like to restore, write the name of the database, then the "*" character and finally the name of the Log file (without spaces).

-offline

Specify this option if you want to restore a table space offline.

-destination *ClientName*

Specifies the target client for restore. Use this option only when you restore to some other instance than the one that was backed up.

-rollforward [time:*YYYY-MM-DD.hh.mm.ss*]

Specify the point in time when you want a rollforward to be performed to. The rollforward point in time *must* be entered in local time (as it is set on the DB2 target server) and not in coordinated universal time (UTC). If you specify a rollforward option without time argument, a rollforward will be performed to the end of the logs.

-frominstance *InstName*

Sets the name of the DB2 instance from which you want to restore the data.

MICROSOFT VOLUME SHADOW COPY SERVICE RESTORE

-vss

Selects the VSS object for restore.

-barhost *ClientName*

Specifies the system on which the backup session was originally performed.

-session {*BackupID* | *SessionID*}

Specifies from which backup data to restore, for example, 2012/10/09-2.

For standard restore, specify *BackupID*. A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

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

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

For instant recovery, specify *SessionID* of a ZDB-to-disk or ZDB-to-disk+tape session.

-tree *TreeName*

Specifies the file, component, or tree to restore. For example, to specify a component, you can use:
`-tree "/Microsoft Exchange Writer(Exchange Information Store)/Microsoft Information Store/First Storage Group/StoreOne"`

When specifying trees, the trees must be specified without the drive letter.

-into *Pathname*

Restores the selected files, component, or tree into the given directory.

-destination *ClientName*

Specifies the target client for restore. Use this option only when you restore to some other instance than the one that was backed up. If not specified, the components are always restored to the server from where they were backed up. Note that all objects in one restore session must be restored to the same system.

-instant_restore

Selects instant recovery for ZDB and VSS integrations.

-conf_check {strict|non-strict|disabled}

Defines the configuration check mode. If this option is specified, Data Protector checks whether the individual components can be selectively restored using the instant recovery functionality. The check detects whether there is more than one component on the volume or there is any data besides the component's data on the volume. If the check fails, the instant recovery session will fail. Specify the strict mode to check each file or folder. Specify the non-strict mode to check each folder. Disable configuration check only if instant recovery cannot be performed with an enabled configuration check and only after you make sure that this will not result in a loss of data. In case of a data loss, the data that does not belong to a component, but resides on the same volume, will be lost.

-no_recovery

Leaves the application database in the recovery mode after completion of the restore session, enabling you to manually apply transaction logs to the database.

This option is available only for the SQL Server writer and Microsoft Exchange Server 2007 writer. It is not supported for Microsoft Exchange Server 2003 writer, where the transaction logs are always applied when the store is mounted.

-use_vds

    Switches a replica from the specified backup session with the source volume. Once switched, the replica is not available for another instant recovery session and also information about this replica is deleted from the database (VSSDB). Does not use a ZDB array specific options or agents.

    With disk arrays of the HPE P9000 XP Disk Array Family, this option must be used after the backup created with the P9000 XP Array provider in the VSS compliant mode.

-use_vss

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

VSS_P6000_ENTERPRISE_VIRTUAL_ARRAY_OPTIONS

-no_copy_back

    If this option is specified, a replica from the specified backup session is switched with the source volume. Once used, the replica is not available for another instant recovery session.

-copy_back

    If this option is specified, copy back is performed. This is also the default behavior when neither -no_copy_back nor -copy_back is specified.

-diskarray_wait *Minutes*

    If this option is specified, there is a delay before the background processes can run. The duration of the delay (in minutes) is determined by *Minutes*. This is also the default behavior when neither -diskarray_wait nor -no_diskarray_wait are specified, in which case there is a 60–minute delay.

-no_diskarray_wait

    If this option is specified, the background processes, such as integrity check, will not stop during the copy creation. This may cause a slowdown of the copy process.

-no_retain_source

    Deletes the source volume during restore. If this option is used with -copy_back, the disk is overwritten during restore. Failure during such restore will cause the source volume data to be lost. If used with -no_copy_back, the disk is deleted after successful restore.

VSS_P9000_DISK_ARRAY_XP_OPTIONS

-copy_back

    Performs resynchronization of the disk pair, copying data from the target volume (backup disk) to the source volume. This option must be specified if the data was backed up with VSS provider in the resync mode.

-no_retain_source

    Deletes the source volume during restore. This option must be specified if the data was backed up with VSS provider in the resync mode since there is no possibility to retain the source during re-synchronization of replica and source disk.

-no_diskarray_wait

If this option is specified, the source volume is immediately available while the synchronization or copy process is running in the background (quick restore). The SSE Agent does not wait for the synchronization or copy process to complete. If this option is not specified, there is a 60–minute delay before the background processes can run.

VSS_P10000_OPTIONS

`-copy_back`

Performs a restore of snapshot data to the source volume.

VSS_P4000_OPTIONS

`-copy_back`

Performs a restore of snapshot data to the source volume.

NOTE: All snapshots dependent on the snapshot being used for restore are deleted.

VSS_EXCHANGE_SPECIFIC_OPTIONS

`-exch_check`

Performs the consistency check of the Microsoft Exchange Server database replicated datafiles. The Microsoft Exchange Server database backup is considered as successful only if the consistency check succeeds. Use this option if consistency check was not performed during backup.

`-exch_throttle` *Value*

Throttles down the consistency check to lessen impact on restore performance. Set the number of input/output operations, after which the check is stopped for one second.

`-exch_checklogs`

Performs the consistency check of the log files only, which is enough for Microsoft Exchange Server to guarantee backup data consistency.

VSS_EXCHANGE_2007_SPECIFIC_OPTIONS

`-target_tree` *TargetStoreName*

Specifies the target component to which the source component will be restored and enables you to restore a subcomponent to a different component than the one from which it was backed up. This option can be used only once for each `-tree` option and cannot be specified together with `-exch_RSG`.

TreeName and its *TargetStoreName* pair must always be fully expanded subcomponents representing an Exchange store or logs. See also the Exchange 2007 examples. To get a list of available targets on a specific host, execute the command:

`vssbar -appsrv:`*HostName* `-perfom:browse -all`

Potential targets can be identified by the string "RESTOREMODE = 1".

NOTE: You cannot restore only a store without logs to a different location. If you specify a target store for an original store, you must also specify logs with an additional `-tree` *TreeName* `-target_tree` *TargetStoreName* pair.

The option must be specified together with `-target_dir`.

`-exch_RSG` *LinkedStoreName*

Creates a new Recovery Storage Group (RSG) and links it to *LinkedStoreName*. This option can be used only once for each `-tree` option and cannot be specified together with `-target_tree`. Only one storage group per session can be restored with this option due to an Exchange limitation. *LinkedStoreName* and its *TreeName* pair must always be fully expanded subcomponents, representing an Exchange store or logs. See also the Exchange 2007 examples.

IMPORTANT: If the RSG already exists, it is removed and a new one is created. Any existing data in it will be lost. NOTE: You cannot restore only a store without logs to a different location. If you specify a target store for an original store, you must also specify logs with an additional `-tree` *TreeName* `-target_tree` *TargetStoreName* pair.

The option must be specified together with `-target_dir`.

`-target_dir` *Directory*

During an instant recovery session, the replica will be mounted to *Directory*. The target directory for one session must always be the same, for example, you cannot specify one target directory for the store(s) and another one for the logs.

SAP MAXDB RESTORE

`-sapdb`

Selects the SAP MaxDB object for restore.

`-barhost` *ClientName*

Specifies the SAP MaxDB client from which the data was backed up.

`-instance` *InstName*

Sets the name of the database instance that was backed up.

`-destination` *ClientName*

Specifies the target client for restore. Use this option only when you restore to some other instance than the one that was backed up.

`-newinstance` *DestinationInstanceName*

Performs a restore to the SAP MaxDB instance with the instance name *DestinationInstanceName*. This option is to be used only when a restore to an instance other than the one that was backed up is to be performed. Note that the specified instance must already exist and must be configured for use with Data Protector. This option does not create a new instance.

`-session` *BackupID*

Specifies from which backup data (*BackupID*) to restore, for example, `2012/10/09-2`.

A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

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

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

If this option is not specified, backup data created in the last backup session is restored regardless of the -endlogs or the -time option selection.

-recover [-endlogs | -time:*YYYY-MM-DD.hh.mm.ss*]

Specify this option to recover the restored SAP MaxDB database by applying the restored (if the -from_disk option is not specified) or client-resident logs (if the -from_disk option is specified) to the last available log (the default behavior, or if the -endlogs option is specified), or to the specified point in time (if the -time: option is specified).

Make sure that the backup session selected by the -session option will restore enough data for the integration to apply the redo logs until the last available log or until the specified point in time.

When this option is not specified, the following happens after the restore:

- If archive logs are not restored (if restore from a full backup session is performed), the database remains in the Admin mode after the restore.

- If archive logs are restored, the database is, if the restored archive logs allow it, switched to the Online mode. If the database, however, cannot be switched to the Online mode (because the restored archive logs do not allow it), it remains in the Admin mode.

-endlogs

Specify this option to recover the database until the last log. This is the default option.

-time: *YYYY-MM-DD.hh.mm.ss*

Specify the -time: option to recover the database until the point specified by the *YYYY-MM-DD.hh.mm.ss* argument.

Note that the specified time is the system time on the system running the Data Protector CLI. If the system to be recovered is not in the same time zone as the system running the Data Protector CLI, the point of recovery is adjusted to the local time setting on the system to be restored.

-from_disk

Specify this option to apply the existing archive logs on the SAP MaxDB Server to SAP MaxDB Server redo logs.

If this option is not specified, the backed up archive logs on backup media are applied to the redo logs (if trans backup session is restored), or the redo logs are left intact together with the existing archive logs on the SAP MaxDB Server (if full or diff backup session is restored).

When a transactional backup session is selected for restore or when it is a part of the needed restore chain, and the this option is specified at the same time, the archive logs from Data Protector media are applied to the redo logs. Thereafter, the archive logs on the SAP MaxDB Server are applied to redo logs.

This option is ignored in case of SAP MaxDB migration, thus allowing only for the restore of redo logs from the backed up archive logs on backup media (if trans backup session is restored).

-nochain

This option instructs the command to restore only the selected or last backup session; the integration does not restore the whole restore chain of full, differential, and transactional backups.

MICROSOFT SQL SERVER RESTORE

`-mssql`

Selects the Microsoft SQL Server object, identified with *DBName*, for restore.

`-barhost` *ClientName*

Specifies the Microsoft SQL Server client from which the data was backed up.

`-destination` *ClientName*

Specifies the target client for restore. Use this option only when you restore to some other instance than the one that was backed up.

`-instance` *SourceInstanceName*

Sets the name of the Microsoft SQL Server instance to be restored. `omnir` takes the (`DEFAULT`) instance by default.

The *SourceInstanceName* is case-sensitive; it has to be the same as the name of the SQL Server instance that you specified in the backup specification.

`-destinstance` *DestinationInstanceName*

Specify this option to determine an Microsoft SQL Server instance into which the data will be restored. `omnir` takes the (DEFAULT) instance by default.

`-base` *DBName*

Specifies the SQL Server database for restore. The database name is case–sensitive.

`-session` *BackupID*

Specifies from which backup data (*BackupID*) to restore, for example, `2012/10/09-2`.

A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

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

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

`-datafile` *GroupName*/*DataFileName*

Specifies an SQL Server data file for restore. *GroupName* is the name of the group the data file belongs to.

`-asbase` *NewDBName* {`-file` *LogicalFileName1 PhysicalFileName1*

   [-file *LogicalFileName2 PhysicalFileName2*]...}

This option can only be used for database restore.

Enables restore of the Microsoft SQL Server database under a new name and restore of files to a new location. If the `-asbase` option is used, all logical and physical filenames have to be specified with the `-file` option.

`-replace`

Specify this option if a database with the same name but a different internal structure already exists at the target Microsoft SQL Server instance.

If this option is not specified, the Microsoft SQL Server does not let you overwrite the existing database - the restore will fail.

If you are restoring a data file from the PRIMARY group to an existing database, you must specify the option at the data file level.

When using this option, ensure that the most recent logs are backed up before the restore.

`-singleuser`

Disconnects all users that are connected to the target Microsoft SQL Server database and puts the database in the single user mode. Note that if the database is not in the simple recovery mode, the `-replace` option should also be specified.

`-nochain`

Microsoft SQL Server integration: Restores only the data identified by the `-session` option. If the option `-session` is not specified, backup data created in the latest backup session is restored.

`-recovery {rec | norec}`

Specifies the state (recovered, nonrecovered) of the Microsoft SQL Server database after the restore. The default value for this option is `rec`.

`-stopat` *yyyy/mm/dd.hh:mm:ss*

This option is only available for database objects.

Specifies the exact time when the rollforward of transactions will be stopped. Therefore, to enable database recovery to a particular point in time, the backup you restore from must be a transaction log backup.

You cannot use this option with norecovery or standby. If you specify a stop at time that is after the end of the restore log operation, the database is left in a non-recovered state (as if the restore log is run with norecovery).

`-standby` *File*

This option can only be used for database restore.

Specifies the standby state of the Microsoft SQL Server database after the restore.

`-tail_log` *BackupSpecificationName*

Specify this option to perform a tail log backup session before the actual restore session starts.

MYSQL RESTORE

`-integ MySQL`

Selects a MySQL backup object for restore.

-barhost *TargetMySQLHostname*

> Specifies the Data Protector client to restore data to. You can specify any client that hosts MySQL database management system and has the Data Protector `MySQL Integration` component installed. On this client, the Data Protector MySQL integration agent is started at the beginning of the restore session.

-appname *TargetInstanceName*

> Specifies the name of the MySQL instance you want to restore data to. If the instance does not exist yet, Data Protector automatically creates and registers it at the end of the restore session.

-user *Username:GroupName*

> Specifies the username of the operating system user account to use for the restore session. The chosen account must be granted appropriate privileges as a MySQL database administrator and be a Data Protector user with the proper user rights for the restore scenario (`Start restore`, `Restore from other users`, `Restore to other clients`, and so on). If no value is specified, username of the Data Protector Inet account on the target client is used.

-source_client *SourceMySQLHostname*

> Specifies the Data Protector client from which MySQL data was backed up.

-source_instance *SourceInstanceName*

> Specifies the name of the original MySQL instance whose data was backed up.

-database

> Enables Data Protector to primarily restore MySQL databases, database tables, or both, as opposed to restoring MySQL binary log files only. The MySQL integration agent is used in this process. If the `-roll_forward` option is also specified, Data Protector also restores and applies all the required binary log files according to the chosen time period. In this case, the `Disk Agent` is additionally used to restore binary log files.

-binary_log

> Enables Data Protector to restore one or more MySQL binary log files only, as opposed to MySQL databases, database tables, or both. The `Disk Agent` is used in this process.

-session *SessionID*

> This option can only be used in combination with the `-database` option.

> Enables Data Protector to process the restore chain from its beginning up to end including the MySQL backup session with the specified session ID. Ensure the session ID belongs to a valid backup session.

-staging [*CustomStagePath*]

> This option can only be used in combination with the `-database` option.

> If the additional option `-copy_back` is not specified, performs the first phase of a staged restore. In this scenario, data from valid backup images of the restore chain is placed to an intermediate location on the target client leaving MySQL production data intact.

> If the additional option `-copy_back` is specified, performs a complete staged restore.

If the additional option `-import` is specified, performs data migration of the database and/or database tables.

Specify the *CustomStagePath* parameter to use a custom folder for staging the restored data, instead of the Data Protector default folder for temporary files.

`-copy_back`

This option can only be used in combination with the `-database` and `-staging` options.

Performs the complete staged restore. Data from backup images of the restore chain is placed to an intermediate location on the target client first. Afterwards, this data is copied to the target location. The binary log is filtered and only the content applicable to the selected tables is recovered. The system tablespace is always restored regardless of the restore scope.

This restore method requires more storage space from an in-place restore, but can better prevent potential data inconsistency in the event something goes wrong.

Use the `-target_dir` option to redirect restore to a location that differs from the original one

`-target_dir` *NonOriginalTargetPath*

This option can only be used in combination with the `-copy_back`, `-inplace`, or `-include` option.

Redirects restore of databases, database tables, or binary log files to a location that differs from the original one.

`-import`

This option can only be used in combination with the `-database` and `-staging` options.

Imports the selected MySQL databases, database tables, or both to the target MySQL instance. The database tables with the same name should not exist on the target instance. The target MySQL instance should be offline during the restore session.

The binary log is filtered and only the content applicable to the selected tables is recovered.

This options is supported with MySQL 5.6.6 and later versions.

`-inplace` [*CustomStagePath*]

This option can only be used in combination with the `-database` option.

Performs an in-place restore of MySQL data, as opposed to one or both phases of a staged restore. In this scenario, data from the backup images of the restore chain overwrites the MySQL production data (if it exists). Such restore process requires less storage space from a complete staged restore, but is more prone to potential data inconsistency in the event something goes wrong.

> **Note:** With this option selected, you can restore only the entire backup image to the target location regardless of the restore scope.

Use the `-target_dir` option to redirect restore to a location that differs from the original one.

`-include {DatabaseName | DatabaseName.TableName}`

This option can only be used in combination with the `-database` and `-copy_back` or `-import` options.

Narrows the scope of the restore process to the specified databases or database tables. Other MySQL entities are not restored even if they exist in the backup images of the restore chain. You can specify the `-include` option and a corresponding parameter more than once.

If this option is not specified, all backup data is included in the restore process.

`-roll_forward [EndDateTime]`

This option can only be used in combination with the `-database` option.

Recovers the restored MySQL entity (instance, database, or database table) by rolling it forward using transactions from the corresponding binary log files that Data Protector restores as needed. If you want Data Protector to bring the entity to a certain point in time (not the latest available state), specify the *EndDateTime* parameter to stop the rollforward at that particular date and time. For *EndDateTime*, use local time on the source client, not the coordinated universal time (UTC).

`-include BinaryLogFilename`

This option can only be used in combination with the `-binary_log` option.

Specifies the name of the binary log file which you want to restore. You can specify the `-include` option and a corresponding parameter more than once, thus restoring multiple binary log files in the same session.

VIRTUAL ENVIRONMENT RESTORE

`-veagent`

Selects the virtual environment objects for restore.

`-virtual-environment` {vmware | hyperv | vcd}

Specifies the virtual environment type.

`-barhost BackupHost`

Specifies the client with the Virtual Environment Integration component installed to control the restore session.

`-apphost OriginalAppHost`

Specifies the client that the virtual machine objects were backed up from.

`-instance OriginalDatacenter`

Specifies the instance from which the virtual machines were backed up.

`-method` {vStorageImage | vCDvStorageImage | vStorageImageOpenStack}

This is a VMware specific option.

Specifies the method that was used for backup.

`-session BackupID`

Specifies from which backup data (*BackupID*) to restore, for example, `2012/10/09-2`.

A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

Mirrored objects and objects created in an object copy session have the same backup ID as objects created in the original backup session. Suppose the media set created in the original backup

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

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

If you specify the session ID of an incremental or differential backup session, all backup data from the corresponding backup chain is restored as well.

-fromsession *BackupID1* -untilsession *BackupID2*

Restores from the backup data created in the time interval between *BackupID1* and *BackupID2*.

-vm *vmpath*

For VMware virtual machines, *vmpath* is the complete virtual machine path (for example, `/MyVirtualMachines/VMname`).

For Microsoft Hyper-V virtual machines, *vmpath* is the GUID (for example, `991B483A-C177-4EB0-9DBE-998E96692783`).

-instanceUUID *vmInstanceUUID*

This is a VMware specific option.

Specifies the instanceUUID of virtual machine for restore.

> **Note:** You should not specify the instanceUUID parameter for restore while restoring the virtual machine backed up from Data Protector 8.1 and below.

-versionID *VersionID*

This is a VMware specific option.

Specifies the version of a backed up object selected for the restore.

-new_name *NewVirtualMachineName*

This is a VMware specific option.

Restores a virtual machine under a new name.

-disk *DiskName*

For VMware virtual machines, `DiskName` refers to the name of the disk. For example, `scsi0:0`.

For Microsoft Hyper-V virtual machines, `DiskName` refers to the path of the disk. For example, `c:\Disk1.vhdx`. If the path of the disk contains special characters, the path must be enclosed in a single quote.

-newinstance *TargetDatacenter*

This is a VMware specific option.

Specifies the datacenter that the virtual machines are restored to. If this option is not specified, the virtual machines are restored to the original datacenter.

-store *TargetDatastore*

This is a VMware specific option.

Specifies the datastore to which the virtual machines should be restored. You can choose among all datastores that are accessible by the specified restore target host. If this option is not specified, the virtual machines are restored to the original datastore.

`-destination` *DifferentAppHost*

Specifies the client that the virtual machines are restored to. If you specify an ESX(i) Server system, the virtual machines are registered in and restored to it. If you specify a vCenter Server system, the virtual machines are registered in the vCenter Server but restored to one of its ESX(i) Server systems.

If this option is not specified, the virtual machines are restored to the original client where they were backed up from.

`-host/cluster` *HostOrCluster*

Specifies the ESX(i) Server system or the cluster that the virtual machines are restored to. If this option is not specified, the virtual machines are restored to the original ESX(i) Server system or cluster.

`-resourcePool` *ResourcePool*

Specifies the resource pool on the ESX(i) Server system or the cluster that the virtual machines are restored to. If this option is not specified, the virtual machines are restored to the original resource pool.

`-specificHost` *SpecificHost*

Specifies the specific ESX(i) Server system in the cluster that the virtual machines are restored to. If this option is not specified, the virtual machines are restored to the original ESX(i) Server system.

`-targetstoragepath` *TargetStoragePathOfAllHyper-V-VMs*

This is a Hyper-V specific option.

Specifies the complete path for a different location where the virtual machines should be restored to. The original path is appended to the specified path. For example, if the original path is `C:\VMStorage` and the target path is `D:\Restore`, the virtual machines will be restored to `D:\Restore\C\VMStorage`.

`-neworganization` *TargetOrganization*

This is a VMware specific option.

Specifies the organization in vCloudDirector to which the virtual machines should be restored. If this option is not specified, the virtual machines are restored to the original organization.

`-virtual_datacenter_path` *TargetVDC*

This is a VMware specific option.

Specifies the path of the vDatacenter to which the virtual machines should be restored. If this option is not specified, the virtual machines are restored to the original vDatacenter.

`-virtual_datacenter_uuid` *TargetVDC*

This is a VMware specific option.

Specifies the UUID of the vDatacenter to which the virtual machines should be restored.

`-vapp_path` *TargetVApp*

This is a VMware specific option.

Specifies the path of the vApp to which the virtual machines should be restored. If this option is not specified, the virtual machines are restored to the original vApp.

Note that the virtual machines are restored as a new vApp if the original vApp is no longer available or all virtual machines of the selected vApp are restored.

-vapp_uuid *TargetVApp*

This is a VMware specific option.

Specifies the UUID of the vApp to which the virtual machines should be restored.

-vcenter_path *TargetVCenter*

This is a VMware specific option.

Specifies the path of the vCenter to which the virtual machines should be restored.

-vcenter_uuid *TargetVCenter*

This is a VMware specific option.

Specifies the UUID of the vCenter to which the virtual machines should be restored.

-network_name *TargetNetwork*

This is a VMware specific option.

VMware vSphere behavior:

Specifies the name of the network that enables virtual machines communication.

The target network can be selected for all virtual machines specified in the restore session.

If an individual virtual machine does not have a network adapter, no action is taken.

If an individual virtual machine has multiple network adapters, the first in the list is selected.

If you leave this option empty, the virtual machine is connected to the network available at the time of backup even though it might not be available anymore.

VMware vCloud Director behavior:

Specifies the name of the network that enables virtual machines communication.

If an individual virtual machine is restored into an existing vApp, the vApp network is specified.

If all virtual machines of the selected vApp are restored, the Organization network is specified.

-network_uuid *TargetNetwork*

This is a VMware specific option.

Specifies the UUID of the network that enables virtual machines communication.

If an individual virtual machine is restored into an existing vApp, the vApp network is specified.

If all virtual machines of the selected vApp are restored, the Organization network is specified.

-consolidate

This is a VMware specific option.

Commits all snapshots (including non-Data Protector ones) to the virtual machine base once a
virtual machine is restored.

-register

This is a VMware specific option.

Registers the virtual machines once they are restored. If this option is not specified, you need to
manually recover the restored virtual machines. By default, the option is selected.

-poweron

Puts the newly restored virtual machines online once they are restored.

[-deletebefore|-deleteafter|-skip|-keep_for_forensics]

VMware behavior:

The -deletebefore option deletes an existing virtual machine before it is restored, even if it resides
in a different datacenter than your target datacenter, and then restores it from new. This is the space
efficient option, but is less secure, since the old virtual machine is not available if the restore fails.
Therefore, it should be selected with caution.

The -deleteafter option deletes an existing virtual machine after it is restored, even if it resides in
a different datacenter than your target datacenter. If the restore fails, the existing virtual machine is
not deleted.

The -skip option skips the restore of an existing virtual machine. This allows you to restore missing
virtual machines without affecting existing ones.

The -keep_for_forensics option marks an existing virtual machine with a timestamp. The virtual
machine which is kept for forensics is powered off after the restore and remains at the original
location. It does not affect consecutive backups of the original virtual machine.

If none of these options are specified, an existing virtual machine is deleted after the restore
completes. If the restore fails, the existing virtual machine is not deleted.

Hyper-V behavior:

The -deletebefore option deletes an existing virtual machine before it is restored and then
restores it from new.

The -skip option skips the restore of an existing virtual machine. When restoring multiple virtual
machines, selecting this option enables you to restore only the virtual machines that do not exist at
restore time.

If none of these options are specified, the behavior is the same as with the -deletebefore option
(an existing virtual machine is deleted before the restore by default).

[-removeSnapshots]

This is an Hyper-V specific option.

The -removeSnapshots option consolidates all snapshots before disk restore.

-directory *RestoreDirectory*

Restores virtual-machine files to a directory on the backup host. After such a restore, the virtual
machines are not functional.

[-overwrite|-skip|-latest]

These are VMware specific options.

The -overwrite option overwrites existing files with those from the backup. By default, this option is used.

The -skip option leaves an existing file intact if it is more recent than the one from the backup. Otherwise, it overwrites the file with the one from the backup.

The -latest option preserves an existing file (the file is not restored from the backup).

MICROSOFT SHAREPOINT SERVER 2007/2010/2013 RESTORE

-mssharepoint

Selects the Microsoft SharePoint Server object for restore.

-barhost *HostName*

Specifies the front-end Web server system that was used during backup.

-destination

Specifies the client on which the Data Protector Microsoft SharePoint Server integration agent should be started. It also specifies to which farm the components are restored.

-user

Specifies the Windows domain user under which the Data Protector Microsoft SharePoint Server integration agent should run. This user must be a farm administrator.

-webapplication

Specifies a Web application for restore. Shows the original Web application name.

-session *BackupID*

Specifies from which backup data (*BackupID*) to restore, for example, 2012/10/09-2.

A backup ID is a point in time. All objects (backup data) created in a backup session have the same backup ID, which is the same as the session ID of the backup session.

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

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

-db

Specifies different options for different databases.

-ssodb

Specifies the Microsoft SharePoint Server single sign-on database for restore.

-ssp

Specifies the Shared Services Provider (SSP) for restore.

-wsshelpsearch

Specifies the Windows SharePoint Services (WSS) Help Search for restore.

-tohost *Client*

Specifies the client to restore to. When you restore Microsoft SQL Server databases, the client must be an SQL Server system.

-instance *SourceInstanceName*

Specifies the original Microsoft SQL Server instance name.

-newinstance *DestinationInstanceName*

Specifies the Microsoft SQL Server instance to which the database should be restored.

-as *NewDBName*

Specifies the name under which the database should be restored. By default, the Microsoft SQL Server databases are restored under the original name. You can restore the Microsoft SQL Server database under a different name.

-todir *NewDirectoryName*

Specifies the path to the directory to which the files (database files, index files) should be restored. By default, index files are restored to their original directories.

-replace

Overwrites any existing database. Overwrites all the existing redirection options specified for the selected component. A restore to the original location is performed.

INSTANT RECOVERY

-instant_restore

Restores data on a disk array using instant recovery.

-host *ClientName*

Restores all objects of the specified client that were backed up in the specified session.

-session *SessionID*

Specifies the session to be used for restore.

P6000_ENTERPRISE_VIRTUAL_ARRAY_OPTIONS and *3PAR_DISK_ARRAY_OPTIONS*

-copyback [wait_clonecopy *Minutes*]

If this option is specified, the instant recovery method of copying replica data (the "copy-back" method) is used in the instant recovery session. With this method, volumes of the replica are copied to the disk group of the current source volumes. If mirrorclones were used in the corresponding zero downtime backup session, volumes of the replica are copied to the disk group of the original volumes, not mirrorclones.

Before the actual data copy operation, storage for the replica to be restored is allocated. Although the copy of the replica is only virtual at that time, it is immediately available for use. In the

background, however, a process is still copying data from the replica to the source location (the replica normalization process). The copy process may degrade the disk array performance, and indirectly the application system performance as well. To reduce a potential degradation of the application system performance, specify the option `wait_clonecopy` *Minutes* to make Data Protector wait for the copy to complete before the session continues. If the copy process completes before the delay expires, the session continues immediately. Additionally, you can control the copy process by setting appropriate omnirc options.

`-switch`

If this option is specified, the instant recovery method of switching the disks (the "switch" method) is used in the instant recovery session. With this method, volumes of the replica replace the source volumes.

Note that if this option is specified, and the target volumes to be used in the instant recovery session are standard snapshots or vsnaps, the session automatically uses the instant recovery method of copying replica data instead. In such a case, Data Protector does not wait for the copy to complete, and the instant recovery session continues or finishes immediately.

`{-leave_source|-no_leave_source}`

These options determine whether original data from the source volumes is preserved on the disk array after instant recovery or not. For example, you can specify the option `-leave_source` to investigate why the original data got corrupted.

If the `-no_leave_source` option is specified, the source volumes are either overwritten with data from the replica (with the "copy-back" instant recovery method) or deleted (with the "switch" instant recovery method) during the instant recovery session. In case of the "copy-back" instant recovery method in which the replica used consists of snapclones, the source volumes are converted into containers before being overwritten, provided that the source and target volumes match in size, redundancy level, and belong to the same P6000 EVA disk group.

CAUTION

If you decide to perform instant recovery by copying replica data and not to preserve source volumes after the session (the options `-copyback` and `-no_leave_source` are specified), and the instant recovery session fails, a data loss on the source volumes may occur.

`{-check_config|-no_check_config}`

These options determine whether a sanity check and a comparison of current volume group configuration of the volume groups participating in the instant recovery session and the volume group configuration information kept in the SMISDB after the corresponding zero downtime backup session are performed or not. If the sanity check fails or the volume group configuration has changed since the zero downtime backup session, the instant recovery session aborts.

In an HPE Serviceguard cluster, when performing instant recovery to some other node than the one from which data was backed up, you must specify the `-check_config` option. In such circumstances, the current volume group configuration on the node to which data is to be restored differs from the volume group configuration kept in the SMISDB. Consequently, the SMISDB volume group configuration data is replaced by the current volume group configuration data on the node to which data is to be restored, and the instant recovery session succeeds.

`-force_prp_replica`

If this option is specified and any target volume containing data to be restored is presented to a

system other than the backup system, the HPE P6000 / HPE 3PAR SMI-S Agent removes such presentation. If the option is not specified, the instant recovery session fails in such circumstances.

If this option is specified and a target volume containing data to be restored is presented to the backup system, but cannot be dismounted in an operating system-compliant way, the HPE P6000 / HPE 3PAR SMI-S Agent performs a forced dismount. If the option is not specified, the instant recovery session fails in such circumstances.

-force_restore_volset

If this option is specified and a source volume (a member of the volume set) is exported to the application host using volume set, the HPE P6000 / HPE 3PAR SMI-S Agent removes all volumes that are part of the volume set presentation during instant recovery and adds them back after the restore completes. If the option is not specified, the instant recovery session fails in such circumstances.

Note that if this option is selected during remove presentation, none of the volumes part of the volume set can be accessed.

P9000_DISK_ARRAY_XP_OPTIONS

-keep_version

If this option is specified, the LDEV pairs involved in the current instant recovery session are split and left in the SUSPENDED state after the restore of data is complete. In the opposite case, the LDEV pairs are left in the PAIR state.

Even if the instant recovery is successful, it is recommended to keep the replica until the next ZDB session.

On Linux systems, you must specify this option if the replica set consists of more than a single replica.

-check_config

If this option is specified, the current configuration of the participating volume groups is compared with the volume group configuration as it was during the ZDB session and which is stored in the XPDB. If the configuration has changed since the ZDB session, the instant recovery session aborts. Additionally, the CRC check information for the selected LDEV pairs stored in the XPDB is compared to the current CRC check information. If the items compared do not match, the session aborts. A RAID Manager Library flag, which is set whenever the selected mirror LDEV is accessed/changed by any process (including non-Data Protector processes) is checked. If the flag is set, the session fails with an appropriate warning.

In HPE Serviceguard clusters, if instant recovery is performed to some other node than the one from where the volumes were backed up, the current volume group configuration on the target node is different from the volume group configuration kept in the XPDB. In such a case, the XPDB volume group configuration data is replaced by the current volume group configuration data on the target node, and the session does not abort. When performing instant recovery to some other node than the one that was backed up, specify this option.

ORACLE/SAP_SPECIFIC_OPTIONS

-oracle

Selects the Oracle options for instant recovery.

-sap

Selects the SAP R/3 options for instant recovery.

`-recover {now|time Time|logseq LogSeqNumber`

`thread ThreadNumber | SCN Number}`

Selects the point in time to which the database is recovered. The following options are available:

`now`

All existing archive logs are applied.

`time MM/DD/YY hh:mm:ss`

Specifies an incomplete recovery. Archive logs are applied only to a specific point in time.

`logseq LogSeqNumber thread ThreadNumber`

Specifies an incomplete recovery. Archive logs are applied only to the specified redo log sequence and thread number.

`SCN Number`

Specifies an incomplete recovery. The archive logs are applied only to the specified SCN number.

`-open`

Opens the database after recovery.

`-resetlogs`

Resets the logs after the database is opened. Available only if the `-open` option is specified. This option is not available if the `-recovery` option is set to `now`.

The following are recommendations on when to reset the logs.

Always reset the logs:

- After an incomplete recovery, that is if not all archive redo logs will be applied.

- If a backup of a control file is used in recovery.

Do not reset the logs:

- After a complete recovery where a backup of a control file is not used in recovery.

- If the archive logs are used for a standby database. If you must reset the archive logs, then you have to recreate the standby database.

`-user UserName -group GroupName`

Specifies the username and group name of the account under which Data Protector starts instant recovery. Required only for UNIX clients.

`-appname ApplicationDatabaseName`

Name of the backed up database.

ORACLE_SPECIFIC_OPTIONS

`-parallelism Number`

Selects the parallelism for the restore of archive logs and restore from incremental backups.

DATA_OPTIONS

`-exclude` *TreeName*

Excludes the specified tree from the restore. This option is not supported with the Data Protector NDMP server integration.

`-skip` *MatchPattern*

Excludes files matching *MatchPattern* from restore. This option is not supported with Data Protector NDMP server integration.

`-only` *MatchPattern*

Restores only files that match the given *MatchPattern*. This option is not supported with Data Protector NDMP server integration.

`-as` *Pathname*

Restores the selected fileset as the specified tree.

`-into` *Pathname*

Restores the selected fileset into the given directory.

SESSION_OPTIONS

`-preview`

Checks the restore parameters without performing the actual restore.

Restore preview is not available for Internal Database restore sessions.

`-report` {`warning`|`minor`|`major`|`critical`}

Sets the level of error notification for the session. Errors are classified (in ascending order) as: `warning`, `minor`, `major` and `critical`. When you select a level, errors of this level and higher are displayed in the Monitor window. For example, if `major` is selected, only `major` and `critical` errors are reported. By default, all errors are reported.

MEDIUM_OPTIONS

`-device` *BackupDevice*

Specifies the backup device where the backup medium is mounted.

`-medium` *MediumID*

Specifies the medium from which data will be restored.

This option is not possible when performing a parallel restore.

`-slot` *SlotID* [*Side*]

Specifies the *SlotID* of the tape library unit where the medium is mounted. This option is only valid for this backup device type. To specify the side of the platter in this slot, use the additional *Side* parameter. Slot *Side* must be specified for magneto-optical devices. Values for side are `A` or `B`.

`-id` *DiskAgentID*

Specifies the ID of the Disk Agent which should be used for restore.

FILESYSTEM_OPTIONS

`-touch`

Updates the access date/time of the file during the restore. By default the access date/time of the backup version is used.

`-lock`

When performing a restore of a file, the Disk Agent tries to lock the file. By default the file is not locked.

`-no_protection`

Do not restore protection of the backed up files, instead use the default protection settings.

`-overwrite`

Overwrites files with the same name in the specified fileset on the disk.

`-no_overwrite`

Does not overwrite existing files with the same name.

`-merge`

This option merges files from the backup medium to the target directory and replaces older versions that exist in the directory with newer (if they exist on the medium) files. Existing files are overwritten if the version on the medium is newer than version on disk. No existing directory is deleted.

If a directory or file doesn't exist on disk (but is on the backup medium) it is restored (created).

`-catalog`

Displays the restored files and directories.

`-sparse`

Restores sparse files in their original form.

`-move_busy`

This option is useful only in case the option `-overwrite` is specified. A problem can occur if, for example, a file to be overwritten cannot be deleted because it is currently in use. Setting this option causes busy files to be moved to a filename starting with #. The original file can thus be deleted as the lock is transferred to the corresponding file starting with # sign. For example, `/tmp/DIR1/DIR2/FILE` would be moved to `/tmp/DIR1/DIR2/#FILE`.

`-no_share[_info]`

If this option is specified, share information for directories on Windows is not restored. If a directory was shared on the network when a backup was run with the `Backup share information for directories` option set (by default), it will be automatically shared after restore, unless this option is selected for restore.

`-omit_unrequired_object_versions`

This option applies if you select directories for restore and the backup was performed with the

logging level `-log` or `-log_files`. If specified, 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 not specified, each backup in the restore chain is read, even if there was no change since the previous backup. To restore empty directories, do not specify this option.

`-[no_]resumable`

By default, Data Protector creates checkpoint files during the restore session. The checkpoint files are needed if the restore session fails and you want to restart the failed session, using the Data Protector resume session functionality. If you specify the option `-no_resumable`, the checkpoint files are not created.

If you have changed the default using the global option `ResumableRestoreDefault`, specify the option `-resumable` if you want checkpoint files to be created.

IDB RESTORE

`-idb`

Selects the Internal Database backup object for restore.

`-barhost ClientName`

Specifies the Cell Manager system to which the Internal Database (IDB) should be restored to, in case of a Cell Manager migration. For `ClientName` you can specify either fully qualified domain name, host name, or IP address.

`-restoredb`

Instructs Data Protector to restore the basic IDB parts: the Catalog Database (CDB), the Media Management Database (MMDB), and the Session Messages Binary Files (SMBF).

If no additional options `-nodbrecover` and `-nouseasnewidb` are specified, after a successful restore Data Protector starts the Internal Database Service, performs recovery of the basic IDB parts using both the backed up and the not yet backed up IDB archived log files, and finally starts using the recovered IDB as the new Internal Database of the cell.

However, if the restored database is not used as a new Internal Database (`-nouseasnewidb` option), then along with restore of Internal Database files (files in PG,IDB and JCE folder), all backed up Session Messages Binary files (SMBF) and all backed up Data Protector IDB specific files (DPSPEC) will be restored to the temporary location (the specified Restore location).

DPSPEC files are all Data Protector Internal Database specific files and these are not Postgres related files, DCBFs, SMBFs and Configuration files. These are usually: Auditing files, Data Protector logs, keystore, log files, meta, reportdb, smisdb, sqldb, sysdb, vssdb, and xpdb files.

If needed, this restored database can be used as an Internal Database. However, before switching over to the new Internal Database, all SMBF and DPSPEC files should be copied from the temporary location to the original location. This is required for the Cell manager functionality.

`-restoreconf`

Instructs Data Protector to restore the Cell Manager configuration data. A prerequisite for this operation is a successful restore of the basic Internal Database part in the same session (if the latter is also selected for restore).

-restoredcbf

Instructs Data Protector to restore the Detail Catalog Binary Files (DCBF) part of the IDB. A prerequisite for this operation is a successful restore of the basic Internal Database part in the same session (if the latter is also selected for restore).

-client *SourceClientName*

Specifies the Cell Manager system from where the Internal Database (IDB) was backed up. This system should be running on the same operating system version as the original Cell Manager system. For *SoruceClientName* specify the fully qualified domain name.

-until *YYYY-MM-DD*[*hh.mm.ss*]

Specifies that a point-in-time restore should be performed, returning the IDB to the state it was in at the specified date (and optional time).

If this option is not specified, the restore process creates a copy of the IDB in the latest backed up state. Additionally, in this case, the not yet backed up IDB archived log files are copied from the original IDB location to the target restore location.

Important: 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 -until option.

-pre *PathName*

Specifies the path name of the pre-exec command or script on the Cell Manager system. This command is invoked on the Cell Manager before the IDB restore process is initiated.

-post *PathName*

Specifies the path name of the post-exec command or script on the Cell Manager system. This command is invoked on the Cell Manager after the IDB restore process is completed.

RESTORE_DB_OPTIONS

-targetdir *TargetDataFolderPath*

Specifies the target directory on the Cell Manager where the basic IDB parts (CDB, MMDB, SMBF) should be restored to. Before invoking the restore, make sure this directory is empty and provides enough free storage space. Note that the *TargetDataFolderPath* length should not exceed 80 characters.

Important: Do not reuse the original IDB directory as the target directory.

-port

Specifies the number of the port that is temporarily used for the Internal Database Service during the restore process. After the process completes, this service is restarted on the original port

defined during Data Protector Cell Manager installation.

Important: Do not reuse the original Internal Database Service port as the temporary port. HPE recommends to use the port 7114 for this purpose.

-nodbrecover

If specified, this option instructs Data Protector not to start the Internal Database Service after a successful restore. Thus, recovery of the basic IDB parts (CDB, MMDB, SMBF) using the IDB archived log files is not performed.

-nouseasnewidb

This option can only be specified if the -nodbrecover option is not specified.

If specified, this option instructs Data Protector not to use the recovered IDB as the new Data Protector Internal Database in the cell.

RESTORE_CONF_OPTIONS

-keeprecent

Instructs Data Protector to keep the most recent version of each Cell Manager configuration file: the existing version on the Cell Manager system (when newer from the version in the IDB backup image) or the backed up version (when newer from the version that already exists on the Cell Manager system). This is the default behavior when neither -keeprecent, nor -nooverwrite, nor the -overwrite option is specified.

-nooverwrite

Instructs Data Protector to preserve each existing Cell Manager configuration file even when its counterpart is present in the IDB backup image.

-overwrite

Instructs Data Protector to unconditionally overwrite each existing Cell Manager configuration file with its counterpart from the IDB backup image. You can use this selection in the event that only a few configuration files are missing on the Cell Manager.

-session *SessionID*

This option must be specified if the -restoredb option is not specified.

If the -session option is specified, Data Protector processes the restore chain of the IDB backup session with the specified session ID. Ensure your session ID belongs to a valid backup session.

If the -session option is not specified, Data Protector automatically selects and processes the restore chain that suits your restore chain selection for the basic IDB parts (CDB, MMDB, SMBF).

-targetdir *TargetConfFolderPath*

Specifies the target directory on the Cell Manager where the Cell Manager configuration data should be restored to. Before invoking the restore, make sure this directory is empty and provides enough free storage space.

If this option is not specified, the original Cell Manager configuration data location is used for the restore session.

-name *FileOrFolderName*...

If specified, narrows the scope of the Cell Manager configuration data restore to the specified files or folders.

RESTORE_DCBF_OPTIONS

`-targetdir` *TargetDCBFFolderPath*

Specifies the target directory on the Cell Manager where the DCBF part of the IDB should be restored to. Before invoking the restore, make sure this directory is empty and provides enough free storage space.

If this option is not specified, the original DCBF location is used for the restore session.

SPLIT_MIRROR_OPTIONS

`-sse`

Selects the HPE P9000 XP Disk Array Family split mirror restore.

`-symmetrix`

Selects the EMC Symmetrix split mirror restore.

`-remote` *ApplicationSystem BackupSystem*

If the `-symmetrix` option is specified, this option selects the EMC Symmetrix Remote Data Facility (SRDF) split mirror configuration.

If the `-sse` option is specified, this option selects the HPE Continuous Access (CA) P9000 XP configuration.

`-local` *ApplicationSystem BackupSystem*

If the `-symmetrix` option is specified, this option selects the EMC Symmetrix Time Finder split mirror configuration.

If the `-sse` option is specified, this option selects the HPE Business Copy (BC) P9000 XP configuration.

`-combined` *ApplicationSystem BackupSystem*

If the `-symmetrix` option is specified, this option selects the EMC Symmetrix combined (SRDF & Time Finder) split mirror configuration.

If the `-sse` option is specified, this option selects the combined HPE Continuous Access+Business Copy (CA+BC) P9000 XP configuration.

`-mirrors` *list*

Specifies the mirror unit (MU) number of a specific replica to be used in the restore session, or the MU numbers of a range or sequence of replicas which define a replica set from which the integration, according to the replica set rotation, selects one replica to be used in the restore session. If this option is not specified, the MU number 0 is used.

`-quiesce` *cmd*

Specifies the command/script to be run before the LDEV pairs are split (put into the SUSPENDED state). The command/script must reside on the application system in the default Data Protector administrative commands directory. It can be used, for example, for stopping the application,

dismounting the file systems not to be restored in the active session, but belong to the same volume group or disk, or preparing the volume group for deactivation.

If this command/script fails, the command/script specified with the option `-restart` is not executed. Therefore, you need to implement a cleanup procedure in this command/script. Note that if the omnirc option `ZDB_ALWAYS_POST_SCRIPT` is set to `1`, the command/script specified with the option `-restart` is always executed. For details, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

`-restart` *cmd*

Specifies the command/script to be run immediately after the LDEV pairs are resynchronized (put into the PAIR state). The command/script must reside on the application system in the default Data Protector administrative commands directory. It can be used, for example, for restarting the application or mounting the filesystems.

`-discovery`

This option can only be specified for the EMC Symmetrix split mirror restore sessions.

Directs the Data Protector EMC Symmetrix Agent to build or re-build the Data Protector Symmetrix database on both the application system and the backup system. Its effect is the same as that of the command `syma -init`. For details, see the *HPE Data Protector Zero Downtime Backup Administrator's Guide*.

`-re-establish_links_before_restore`

Directs the Data Protector disk array agent to synchronize the LDEV pairs, that is, to copy the application data to the disks which store backup data. This is necessary to prepare the disks for restore and to enable consistent data restore. If the paired LDEVs have been split (put into the SUSPENDED state) before the restore, and only some files need to be restored, then this option updates the backup system. This will ensure that the correct data is resynchronized to the application system. If this option is not specified, the synchronization is not performed.

`-disable_disks`

Directs the Data Protector disk array agent to disable disks on the application system, that is, dismount the filesystems and deactivate the volume groups. This is performed before the LDEV pairs are split. The disks are enabled after the links are restored. Note that only filesystems selected for restore are dismounted. If other filesystems exist on the volumes of the volume group or on the disk, appropriate commands/scripts must be used to dismount these filesystems (specified with the options `-quiesce` and `-restart`). You must always select this option for restore when you want to copy data from the backup system to the application system, that is, to incrementally restore links. The application system disks have to be disabled to provide data integrity after the links are restored, that is, data is copied.

`-restore_links_after_restore`

Directs the Data Protector disk array agent to incrementally restore the links for the LDEVs that Data Protector has successfully restored to the backup system. The HPEP9000 XP Agent also incrementally re-establishes links for the LDEVs for which the Data Protector restore failed.

GENERAL_OPTIONS

`-device` *BackupDevice*

Specifies the backup device where the backup medium is mounted.

-no_auto_device_selection

If this option is specified, Data Protector does not automatically replace unavailable devices with available devices of the same device tag.

-server *ServerName*

Selects the Cell Manager with the client name *ServerName* as the Cell Manager. Use this option to perform a restore to a client that is not in the current Data Protector cell.

-target *Client*

Restores the selected fileset to the specified client.

-profile

Displays restore statistics.

-load {low|medium| high}

Specifies the level of network traffic generated by a session during a time period. High level generates as much traffic as allowed by the network, resulting in a faster restore. A low level has less impact on network performance, but results in a slower restore. By default, this option is set to high.

-pre_exec *PathName*

Instructs the Disk Agent to execute this command before restoring the data object. The complete pathname of the command should be specified.

-post_exec *PathName*

Instructs the Disk Agent to execute this command after restoring the data object. The complete pathname of the command should be specified.

-variable *VariableName VariableValue*

This option lets you specify a variable name and its value for proper operation of some platforms and integrations. Setting user definable variables (a variable name and its value) enables flexible operation on some platforms and integrations with Data Protector. The list of variables and their values that are configurable with Data Protector is dynamic and comes with Data Protector patches.

-no_monitor

By default the command monitors the session and displays all messages. If this option is used, the command displays only the session ID.

-priority *NumValue*

In case multiple running sessions request access to a specific device at the same time, this option determines the order in which the sessions will be queued. The *NumValue* can be any value from 1 (the highest priority) to 6000 (the lowest priority). In case the option is not specified, the default value of 3000 is set. If a low priority session is running when a high priority session starts queuing, the currently running session is allowed to finish. When more sessions request access to a device with the same priority, any of these sessions might acquire access first.

RETURN VALUES

See the man page `omniintro` for return values.

Additional return values of the `omnir` command are:

`10` - There was an error while copying some files. All agents completed successfully.

`11` - One or more agents failed, or there was a database error.

`12` - None of the agents completed the operation.

`13` - Session was aborted.

EXAMPLES

The following examples illustrate how the `omnir` command works.

1. To restore trees "/tree1" and "/tree2" of the root filesystem on "fs", with the label "lb1", from data created in the session "2013/05/12-33", as the trees "/tmp/tree1" and "/tmp/tree2", skipping ".xyz" files, execute:

   ```
   omnir -filesystem fs:/ lb1 -session 2013/05/12-33 -tree /tree1 -as /tmp/tree1 -
   tree /tree2 -as /tmp/tree2 -skip *.xyz
   ```

2. To perform a full restore of tree "/ac" on filesystem "bb:/", with no label, from data created in the session "2013/05/12-2, execute":

   ```
   omnir -filesystem bb:/ -full -session 2013/05/12-2 -tree /ac
   ```

3. To perform restore of the section "/dev/rdsk/c201d6s0" of the disk image labeled "RawRoot" on the client "machine" from data created in the backup session "2013/05/23-12", execute:

   ```
   omnir -rawdisk machine "RawRoot" -section /dev/rdsk/c201d6s0 -session
   2013/05/23-12
   ```

4. To use parallel restore for restoring two objects, execute:

   ```
   omnir -filesystem client1:/ -session 2013/04/17-2 -tree /users -into /tmp -
   filesystem client2:/opt -session 2013/04/17-3 -tree /opt -into /tmp
   ```

5. To perform an instant recovery to the system named "machine" from data created in the backup session "2013/03/08-1", keeping the replica on the disk array, execute:

   ```
   omnir -host machine -session 2013/03/08-1 -instant_restore -keep_version
   ```

6. To perform an instant recovery of filesystem backup data on a disk array of the HPE P9000 XP Disk Array Family to the system named "computer" from data created in the backup session "2013/05/02-1", keeping the replica on the disk array, execute:

   ```
   omnir -host computer -session 2013/05/02-1 -instant_restore -keep_version
   ```

7. To perform an instant recovery of data residing on a disk array of the HPE P6000 EVA Disk Array Family to the system named "computer" from data created in the filesystem backup session "2013/01/08-1" by copying replica data, preserve source volumes on the disk array, and perform volume group configuration check in advance, execute:

   ```
   omnir -host computer -session 2013/01/08-1 -instant_restore -copyback -leave_
   source -check_config
   ```

8. To perform an instant recovery of data residing on a disk array of the HPE P6000 EVA Disk Array
Family to the system named "computer" from data created in the filesystem backup session
"2013/01/08-1" by switching disks, preserve source volumes on the disk array, and not perform
volume group configuration check in advance, execute:

```
omnir -host computer -session 2013/01/08-2 -instant_restore -switch -leave_
source -no_check_config
```

9. To perform a point in time recovery of the database "dbase.nsf" and all Lotus Notes/Domino
Server NTF files of the Lotus Notes/Domino Server "BLUE" from the system "computer", to the
original location with parallelism 4, execute:

```
omnir -lotus -barhost computer -domino_server BLUE -parallelism 4 -db dbase.nsf
-NTF -recovery_time 2012/08/15.15:00:00
```

10. To perform an Informix Server restore of the database server "ol_computer" on the UNIX system
"computer" with the bar argument "-r rootdbs", and to make the devices available to this session
with the highest priority in case of resource conflicts, execute:

```
omnir -informix -barhost computer -barcmnd ob2onbar.pl -user informix:informix
-bararg "-r rootdbs" -appname ol_computer -priority 1
```

11. The Microsoft Information Store with the "/First Storage Group/STORE/Public Folder Store" store
and "/First Storage Group/LOGS/Logs" logs is to be restored to the system called
"computer.company.com" (where it was backed up), from data created in the backup session
"2013/05/07-13". The Microsoft Exchange Server log files are to be restored to "c:\temp" directory,
the hard recovery is to be performed after the restore has finished. The database is to be mounted
after the hard recovery. Execute:

```
omnir -msese -barhost computer.company.com -appname "Microsoft Exchange Server
(Microsoft Information Store)" -base "/First Storage Group/LOGS/Logs" -session
"2013/05/07-13" -base "/First Storage Group/STORE/Public Folder Store" -session
"2013/05/07-13" -logpath c:\temp -last -mount
```

12. Microsoft Exchange Server 2010/2013 restore: Suppose you want to restore the backup of the
database "DB1" to a recovery database that should be created on the client
"exchange2.company.com" and named "Recovery1", with the files in the "C:\Recovery1Folder"
directory. Suppose the database "DB1" was backed up in the session "2013/5/14-1" from a DAG
whose virtual system name was "dag0.company.com". To also ensure that the integration agent
(e2010_bar.exe) is started on the client "exchange1.company.com", execute:

```
omnir -e2010 -barhost exchange1.company.com -db_name DB1 -source
dag0.company.com -new -session 2013/5/14-1 -client exchange2.company.com -
location C:\Recovery1Folder -name Recovery1 -recoverydb
```

13. Microsoft Exchange Server 2010/2013 restore (instant recovery): Suppose you want to restore the
corrupt standalone database "DB1", which resides on the client "exchange1.company.com". The
database was backed up in the ZDB session "2013/05/20-3". To ensure that the integration agent
(e2010_bar.exe) is started on the client "exchange1.company.com", and that the database is
restored to the latest state, using the copy-back instant recovery method, execute:

```
omnir -e2010 -barhost exchange1.company.com -instant_restore -copy_back -db_
name DB1 -latest
```

14. Virtual Environment (VMware vSphere) restore: Suppose you want to restore the virtual machine
"/vm/machineA" and the individual disks ("scsi0:0" and "scsi0:1") of the virtual machine
"/vm/machineB". At the time of backup, the virtual machines were running on the ESX Server

systems that belonged to the datacenter "/MyDatacenter" managed by the vCenter Server system "vcenter.company.com". The virtual machines were backed up with the "vStorageImage" backup method.

To restore them to the original location, using the backup session "2013/01/11-1" and to ensure that the newly restored virtual machines are put online when the session completes, execute:

```
omnir -veagent -virtual-environment vmware -barhost backuphost.company.com -
apphost vcenter.company.com -instance /MyDatacenter -method vStorageImage -
session 2013/1/11-1 -vm -instanceUUID vmInstanceUUID /MyDataCenter/vm/machineA
-vm MyDataCenter/vm/machineB -disk scsi0:0 -disk scsi0:1 -register -poweron
```

15. Virtual Environment (VMware vSphere) restore: Suppose the virtual machines "/MyVirtualMachines/machineA" and "/MyVirtualMachines/machineB" were backed up in the session "2013/02/12-5" from the datacenter "/MyDatacenter" that is managed by the vCenter Server system "vcenter.company.com", using the "vStorageImage" backup method. To restore the virtual machines outside the datacenter, to the directory "C:\tmp" on the backup host "backuphost.company.com", execute:

```
omnir -veagent -virtual-environment vmware -barhost backuphost.company.com -
apphost vcenter.company.com -instance /MyDatacenter -method vStorageImage -
session 2013/2/12-5 -vm -instanceUUID vmInstanceUUID
/MyVirtualMachines/machineA -vm /MyVirtualMachines/machineB -directory c:\tmp
```

16. Virtual Environment (Restoring virtual machines to a Microsoft Hyper-V system): Suppose you want to restore the virtual machines "VM1" with the GUID "62BD6C3C-D4BE-44F4-88D6-E439C96C4B0C" and "VM2" with the GUID "54C22930-E3B9-43AA-AFCD-1E90BB99F130". At the time of backup, the virtual machines were running on the Microsoft Hyper-V system "hyperv1.company.com". The virtual machines were backed up with the "Hyper-V Image" backup method.

To restore the virtual machines to the Microsoft Hyper-V system "hyperv2.company.com" to the default location, using backup data created in the backup session "2013/01/11-1" and to power the newly restored virtual machines on when the session completes, execute:

```
omnir -veagent -virtual-environment hyperv -barhost backuphost.company.com -
apphost hyperv1.company.com -instance hyperv -session 2013/1/11-1 apphost
hyperv1.company.com -instance hyperv -session 2013/1/11-1 -vm 62BD6C3C-D4BE-
44F4-88D6-E439C96C4B0C -vm 54C22930-E3B9-43AA-AFCD-1E90BB99F130 -destination
hyperv2.company.com -poweron
```

17. Virtual Environment (Restoring virtual machines outside a Microsoft Hyper-V system): Suppose the virtual machines "VM1" with the GUID "54C22930-E3B9-43AA-AFCD-1E90BB99F130" was backed up in the session "2013/02/12-5" from the Microsoft Hyper-V system "hyperv.company.com", using the "Hyper-V Image" backup method. To restore the virtual machine outside the Microsoft Hyper-V system, to the directory "c:\tmp" on the restore client "client.company.com", execute:

```
omnir -veagent -virtual-environment hyperv -barhost client.company.com -apphost
hyperv.company.com -instance hyperv -session 2013/2/12-5 -vm 54C22930-E3B9-
43AA-AFCD-1E90BB99F130 -directory c:\tmp
```

18. Virtual Environment (Restoring individual virtual machine disks to a directory in a Microsoft Hyper-V system):

Suppose the virtual machine VM1 with the GUID "54C22930-E3B9-43AA-AFCD-1E90BB99F130" was backed up in the Microsoft Hyper-V system "hyperv.company.com" using the "Hyper-V Image" backup method in session "2016/02/02-5". To restore disks (DiskPath1 "C:\Hyper-V\Virtual Hard Disks\Disk1.vhdx" and DiskPath2 "c:\Disk2.vhdx") to the directory "C:\tmp" on the restore client "client.company.com", execute:

```
omnir -veagent -virtual-environment hyperv -barhost client.company.com -apphost
hyperv.company.com -instance hyperv -destination hyperv.company.com -session
2016/02/02-5 -vm 54C22930-E3B9-43AA-AFCD-1E90BB99F130 -disk C:\Hyper-V\Virtual
Hard Disks\Disk1.vhdx -disk c:\Disk2.vhdx -directory C:\tmp
```

19. Virtual Environment (Restoring individual virtual machine disks to original location in a Microsoft Hyper-V system)

    Suppose the virtual machine VM1 with the GUID "54C22930-E3B9-43AA-AFCD-1E90BB99F130" was backed up in the Microsoft Hyper-V system "hyperv.company.com" using the "Hyper-V Image" backup method in session "2016/02/02-5". To restore disks (DiskPath1 "C:\Hyper-V\Virtual Hard Disks\Disk1.vhdx" and DiskPath2 "c:\Disk2.vhdx") to its original location on the restore client "client.company.com", execute:

```
omnir -veagent -virtual-environment hyperv -barhost client.company.com -apphost
hyperv.company.com -instance hyperv -destination hyperv.company.com -session
2016/02/02-5 -removeSnapshots -vm 54C22930-E3B9-43AA-AFCD-1E90BB99F130 -disk
C:\Hyper-V\Virtual Hard Disks\Disk1.vhdx -disk c:\Disk2.vhdx
```

> **Note:** If the disk path (DiskPath1 and DiskPath2 in example) contains special characters, the path must be enclosed in a single quote.
>
> The -removeSnapshots option removes existing snapshots from the virtual machine.

20. To perform a VSS restore of the "Registry Writer" and "System Writer" trees from the backup session "2013/05/20-3" and the "Event Log Writer" tree from data created in the backup session "2013/05/27-1", which were both performed on the client "system1.company.com" to the client "system2.company.com" into the "c:\tmp directory", execute:

```
omnir -vss -barhost system1.company.com -session 2013/05/20-3 -tree /"Registry
Writer" -tree /"System Writer" -session 2013/05/27-1 -tree /"Event Log Writer"
-destination "system2.company.com" -into c:\tmp
```

21. To start an online restore of a DB2 database called "TEMP" from instance "DB2Inst" on the client "splendid" and roll it forward till the 16th March 2013, 9:15 a.m., execute:

```
omnir -db2 -barhost splendid -instance DB2Inst -dbname TEMP -rollforward -time
2013-03-16.09.15.00
```

22. To restore the contents of a mailbox called "FIRST" residing on an Microsoft Exchange Server system called "infinity.ipr.company.com" from data created in the backup session 2013/01/10-1, into the new mailbox called "TEMP", execute:

```
omnir -mbx -barhost infinity.ipr.company.com -mailbox FIRST -session
2013/01/10-1 -destmailbox TEMP
```

23. To restore all messages from the "Inbox" folder (and all subfolders) from the "User 1" mailbox residing on the Microsoft Exchange Server system called "exchange.hp.com", into the original location, from data created in the backup session "2013/03/10-18", without overwriting the messages, execute:

```
omnir -mbx -barhost exchange.hp.com -mailbox "User 1" -session 2013/03/10-18 -
folder Inbox -originalfolder -keep_msg
```

24. To restore all messages from the "User 2" mailbox residing on the Microsoft Exchange Server system called "exchange.hp.com", except for the messages in the folder "Deleted Items", into a new location, from data created in the backup session "2013/03/10-19" (for example, performed at 13:47:00), execute:

```
omnir -mbx -barhost exchange.hp.com -mailbox "User 2" -session 2013/03/10-19 -
exclude "Deleted Items"
```

The messages will be restored in the "Data Protector 03/10/13 13:47:00" mailbox on the "exchange.hp.com" Microsoft Exchange Server.

25. To start an online restore of an SAP MaxDB database called "TEMP" on the client "splendid" and roll it forward till the 10th January 2013, 9:15 a.m. from the archive logs already residing on the client, execute:

```
omnir -sapdb -barhost splendid -instance TEMP -recover -time: 2013-01-
10.09.15.00 -from_disk
```

26. With disk arrays of the HPE P9000 XP Disk Array Family, to recover an Oracle database "DB1" on the Windows client "san32" using the user account "sys" that belongs to the "sysgroup" user group, from data created in the backup session "2013/02/05-18", until the most recent time, to open the database after the recovery, to keep the replica on the disk array, and to use "1" as the parallelism setting, execute:

```
omnir -host san32 -session 2013/02/05-18 -instant_restore -keep_version -oracle
-user sys -group sysgroup -recover now -open -appname DB1 -parallelism 1
```

27. To perform restore of the section "/dev/rdsk/c201d6s0" of the disk image labeled "Raw" on the client "system1" from data created in the backup session "2013/05/23-12" using the media set containing the object copy with ID "d5032390-baba-4b3f-8c67-1f5b9273b242/1013", execute:

```
omnir -rawdisk system "Raw" -section /dev/rdsk/c201d6s0 -session 2013/05/23-12
-copyid d5032390-baba-4b3f-8c67-1f5b9273b242/1013
```

28. To start instant recovery of data on a disk array of the HPE P6000 EVA Disk Array Family on the system named "system1" from data created in the VSS backup session "2013/05/08-14" which copies the data from the replica to the source disk group overwriting the source volume, execute the following command:

```
omnir -vss -instant_restore -barhost system1 -session 2013/05/08-14 -copy_back
-no_retain_source
```

29. To restore the SqlServerWriter from the VSS backup session "2013/05/07-9" on the system named "system1" using the Microsoft Virtual Disk Service with the possibility to later apply transaction logs on the SQL Server, execute the following command:

```
omnir -vss -instant_restore -use_vds -barhost system1 -session 2013/05/07-9 -
tree "/SqlServerWriter(SQL Server 2005:SQLWriter)" -no_recovery
```

30. Exchange 2007 VSS restore to a different storage group:

To restore the Exchange 2007 Writer logs on the system "exch2007.company.com" from the storage group copy "Replicated Storage Group" created by LCR, from data created in the backup session "2013/04/08-12", to storage group "Original Storage Group", and with the files restored in the "C:\Omni" directory, execute the following command:

```
omnir -vss -instant_restore -use_vds -barhost exch2007.company.com -session
2013/04/08-12 -tree "/Microsoft Exchange Writer(Exchange Replication Service)
/Microsoft Information Store/Replicated Storage Group/Logs" -target_tree
"/Microsoft Exchange Writer(Exchange Information Store)/Microsoft Information
Store/Original Storage Group/Logs" -target_dir "C:\Omni"
```

31. Exchange 2007 VSS instant recovery to a non-Exchange location:

     To perform instant recovery of the Exchange 2007 Writer store "StoreOne" from the storage group
     "First Storage Group" from data created in the backup session "2013/04/08-9" on the system
     "exch2007.company.com", to the system "server2.company.com", and with the replicas mounted
     to "C:\Omni_Mnt", execute:

```
omnir -vss -instant_restore -use_vds -barhost exch2007.company.com -destination
server2.company.com -session 2013/04/08-9 -tree "/Microsoft Exchange Writer
(Exchange Information Store)/Microsoft Information Store/First Storage
Group/StoreOne" -target_dir "c:\mnt" -tree "/Microsoft Exchange Writer(Exchange
Information Store)/Microsoft Information Store/First Storage Group/Logs" -
target_dir "C:\Omni_Mnt"
```

32. Exchange 2007 VSS restore to a non-Exchange location and creating RSG:

     To restore the Exchange 2007 Writer store "Store One" from the storage group named "First
     Storage Group" from data created in the backup session "2013/04/10-9" that was performed on the
     system "exch2007.company.com", and to create the Recovery Storage Group "DP RSG" that
     links restored store to "Store Two" in storage group "Second Storage Group", and with the files
     restored in the "C:\Omni" directory, execute:

```
omnir -vss -instant_restore -use_vds -barhost exch2007.company.com -session
2013/04/10-9 -tree "/Microsoft Exchange Writer(Exchange Information Store)
/Microsoft Information Store/First Storage Group/Store One" -exch_RSG
"/Microsoft Exchange Writer(Exchange Information Store)/Microsoft Information
Store/Second Storage Group/Store Two/" -target_dir "c:\mount" -tree "/Microsoft
Exchange Writer(Exchange Information Store)/Microsoft Information Store/First
Storage Group/Logs" -exch_RSG "/Microsoft Exchange Writer(Exchange Information
Store)/Microsoft Information Store/Second Storage Group/Logs" -target_dir
"C:\Omni"
```

33. To perform a full restore of the tree "/vol/vol1" of the NDMP client alpha.hp.com, from data created
     in the backup session "2013/05/12-2", using the device "LTO" connected to the client beta,
     execute:

```
omnir -filesystem alpha.hp.com:/vol/vol1 /vol/vol1 -full -session 2013/05/12-2
-tree "/vol/vol1" -device LTO
```

34. To restore an entire Microsoft SharePoint Server 2007 server (moss.domain.com) from the latest
     session, execute:

```
omnir -mssharepoint -barhost wfe1.domain.com -server moss.domain.com
```

35. To restore a Microsoft SharePoint Server 2010 Web application content database from the latest
     session to the alternate location, changing a name, sql server, an instance and a data file path,
     execute:

```
omnir -mssharepoint -barhost wfe1.domain.com -webapplication "SharePoint -
2224" -db "WSS_Content_2224" -as "WSS_new_DB" -tohost mosssql2.domain.com -
newinstance moss1 -todir "f:\program files\SQL\data"
```

36. To restore the database "TEST1" on the Microsoft SQL Server instance "TEST_INSTANCE" and client "system1.company.com", and to perform a tail log backup session before the actual restore session starts, by using the backup specification "DB1_Backup", execute:

```
omnir -mssql -barhost system1.company.com -instance TEST_INSTANCE -base TEST1 -
tail_log DB1_Backup
```

37. To perform online restore of the entire Internal Database (IDB) and the Cell Manager configuration files backed up from the Windows system "cmsys-win.company.com", by restoring the basic IDB parts (CDB, MMDB, SMBF) to the path "D:\Data_Protector_temp\idb", by restoring the Cell Manager configuration files and the DCBF part of the IDB to their original location, by restoring to the latest state that was backed up before 24 May 2013, using the temporary port "7114", and without performing the IDB recovery, execute:

```
omnir -idb -barhost cmsys-win.company.com -restoredb -targetdir D:\Data_
Protector_temp\idb -port 7114 -until 2013-05-24 -nodbrecover -restoreconf -
restoredcbf
```

38. To perform online restore of the basic IDB parts (CDB, MMDB, SMBF) backed up from the UNIX system "cmsys-ux.company.com", by restoring the data to the path "/var/tmp/Data_Protector_temp/idb", by restoring to the latest backed up state, using the temporary port "7114", and by performing the IDB recovery without putting the recovered IDB into use as the new IDB in the cell, execute:

```
omnir -idb -barhost cmsys-ux.company.com -restoredb -targetdir /var/tmp/Data_
Protector_temp/idb -port 7114 -nouseasnewidb
```

39. To perform online restore of the DCBF part of the IDB backed up from the system "cmsysx.company.com" to its original location, and to the latest state that was backed up before 12 April 2013 at 16:00, execute:

```
omnir -idb -barhost cmsys.company.com -restoredcbf -until 2013-04-12.16.00.00
```

40. To support restore of object names with instanceUUID in its name, execute:

```
omnir.exe -veagent -virtual-environment vmware -barhost barHostName -apphost
appHostName -instance instanceName -method vStorageImage -session sessionID -vm
vmPath –instanceUUID vmInstanceUUID -register -poweron -deletebefore
```

41. To restore OpenStack Nova Instances backed up as virtual machines from the VMware vCenter, execute:

```
omnir.exe -veagent -virtual-environment vmware -barhost barHostName -apphost
appHostName -instance /Datacenter -method vStorageImageOpenStack -session
sessionID -vm vmPath -instanceUUID vmInstanceUUID -register -poweron -
deletebefore
```

42. To restore the MySQL database "db1" and database table "db2.table1" of the "MYSQL56" instance to the original system "winsys.company.com" as a new instance named "MYSQL56 _ NEW " using the user account "MYSQLDOMAIN\Administrator" in a complete staged restore session, to use the restore chain of the backup session with the ID "2014/11/22-13", and to roll the restored data forward until the last available state in the backup images of the corresponding binary log files, execute:

```
omnir -integ MySQL -barhost winsys.company.com -appname MYSQL56 _NEW -user
Administrator:MYSQLDOMAIN -options -source_client winsys.company.com -source_
instance MYSQL56 -database -session 2014/11/22-13 -staging -copy_back -include
db1 -include db2.table1 -roll_forward
```

43. To restore the MySQL binary log stored in the "mysql-bin.000001" file of the "MYSQL55" instance to the non-original target system "linuxsys2.company.com" and to the non-original path "C:\Users\MySQL\temp" using the user account with which the Data Protector `Inet` service is running, execute:

```
omnir -integ MySQL -barhost linuxsys2.company.com -appname MYSQL55 -options -
source_client linuxsys1.company.com -source_instance MYSQL55 -binary_log -
include mysql-bin.000001 -target_dir C:\Users\MySQL\temp
```

SEE ALSO

omnib(1), omnikeytool(1M), omniobjconsolidate(1), omniobjcopy(1), omniobjverify(1)

omnirpt(1)

omnirpt — generates various reports about the Data Protector environment, for example, about backup, object copy, object consolidation, and object verification sessions in a specific time frame, session specifications, media, Data Protector configuration, and single sessions
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnirpt -version | -help

omnirpt -report *ReportName REPORT_OPTIONS* [*METHOD_OPTIONS*] [*FORMAT_OPTIONS*] [-header] [-multicell] [-[no]_multiple]

omnirpt -rptgroup *ReportGroup*

FORMAT_OPTIONS

 -ascii |

 -html |

 -tab |

 -short

METHOD_OPTIONS

 -email *EmailAddress* ... |

 -smtp *EmailAddress* ... |

 -snmp *Hostname* ... |

 -broadcast *Hostname* ... |

 -log *Filename* ... |

 -external *CommandName* ...

ReportName

 list_sessions |

 session_flow |

 device_flow |

 used_media |

 used_media_extended |

 host_statistics |

 session_statistics |

 session_errors |

```
      dl_trees|

      obj_nobackup|

      obj_copies|

      obj_lastbackup|

      obj_avesize|

      fs_not_conf|

      dl_info|

      dl_sched|

      db_size|

      cell_info|

      hosts_unused|

      dev_unused|

      lookup_sch|

      hosts_not_conf|

      licensing|

      host|

      media_list|

      media_list_extended|

      media_statistics|

      pool_list|

      single_session|

      session_objects|

      session_hosts|

      session_devices|

      session_media|

      session_objcopies
  REPORT_OPTIONS

  SessionOption

   -session SessionID

  PoolOption

   -pool Poolname ...

  LabelOption

   -label Label

  LocationOption
```

  -location *Location* ...

  LibraryOption

  -[no_]library *Library* ...

  ProtectionOption

  -[no_]protection *NoOfDays*

  MediaClassOption

  -class *MediaClass*

  MediaStatusOption

  -status *MediaStatus*

  SpecificationOptions

  -datalist *BackupSpecificationName* ...

  -copylist_sch *ScheduledCopySpecificationName* ...

  -copylist_post *PostbackupCopySpecificationName* ...

  -verificationlist_sch *ScheduledVerificationSpecificationName* ...

  -verificationlist_post *PostbackupVerificationSpecificationName* ...

  -conslist_sch *ScheduledConsolidationSpecificationName* ...

  -conslist_post *PostbackupConsolidationSpecificationName* ...

  -no_datalist

  -no_copylist

  -no_verificationlist

  -no_conslist

  BackupSpecificationGroupOption

  -group*BackupSpecificationGroup*

  LookupSchedulesOption

  -schedule *NoOfdays*

  NetworkOption

  -network *IP_Address* ...

  HostsOption

  -hosts *Hostname* ...

  HostOption

  -host *Hostname*

  LevelOption

  -level *Level*

  ObjectCopiesOption

```
-num_copies {less|equal|more} NumberOfCopies
```

TimeframeOption

```
-timeframe {Start Duration|Day Hour Day Hour}
```

LatestObjectOption

```
-days NoOfdays
```

Level: {warning|minor|major|critical}

Day: *[YY]YY/MM/DD*

Hour: *HH:MM*

DESCRIPTION

The `omnirpt` command generates various reports about Data Protector environment: reports about backup, object copy, object consolidation, and object verification sessions in a specific time frame, about backup, object copy, object consolidation, and object verification specifications, media, Data Protector configuration and single sessions. Each report is defined by its name `-report` *ReportName* and a set of options that specify report parameters (described below). The reports are provided in four different formats: ASCII, HTML, tabulator separated format and short ASCII format. Each report is described in two parts: input (what you have to/may specify to configure a report) and output (what is the content of the report). Input items that are enclosed in square brackets ([ ]) are optional, while all others are required. The following *report categories* are available:

Sessions in Timeframe

"Sessions in Timeframe" reports provide reports about backup, object copy, object consolidation, and object verification activities in a certain past time period. This time period can either be defined in relative terms (such as last 24 hours) or absolute (15/03/12 00:00 - 16/03/12 00:00). Two other common report options for all "Sessions in Timeframe" reports are backup specification and backup specification group. These two limit the report to selected backup specifications. "Session in Timeframe" reports are:

- List of Sessions (`list_sessions`)
- Session Flow Report (`session_flow`)
- Device Flow Report (`device_flow`)
- Report on Used Media (`used_media`)
- Extended Report on Used Media (`used_media_extended`)
- Client Statistics (`host_statistics`)
- Session Statistics (`session_statistics`)
- Session Errors (`session_errors`)
- Object Copies Report (`obj_copies`)

Session Specifications

"Session Specifications" reports provide different configuration reports which are based on backup, object copy, object consolidation, and object verification specifications. By default, all backup, object

copy, object consolidation, and object verification specifications are used, but you may choose to limit a report to a certain session specification. Selection of a backup specification group is available only for backup specifications. "Session Specifications" reports are:

- Trees in Backup Specification (`dl_trees`)
- Objects Without Backup (`obj_nobackup`)
- Object's Latest Backup (`obj_lastbackup`)
- Average Backup Object Sizes (`obj_avesize`)
- Filesystems Not Configured for Backup (`fs_not_conf`)
- Session Specification Information (`dl_info`)
- Session Specification Schedule (`dl_sched`)

Internal Database

The "Internal Database" report provides information about Data Protector Internal Database (IDB) size. The "Internal Database" report is:

- Internal Database Size Report (`db_size`)

Configuration

"Configuration" reports provide various reports about Data Protector environment. "Configuration" reports are:

- Cell Information (`cell_info`)
- Configured Clients not Used by Data Protector (`hosts_unused`)
- Configured Devices not Used by Data Protector (`dev_unused`)
- Look up Schedule (`lookup_sch`)
- Clients not Configured for Data Protector (`hosts_not_conf`)
- Licensing report (`licensing`)
- Client Backup Report (`host`)

Pools and Media

"Pools and Media" reports provide four reports that search through Data Protector pools for media that match the search criteria. The default is to list all media or pools and each report option can then limit the search to a certain set of media. "Pools and Media" reports are:

- List of Pools (`media_list`)
- Extended List of Media (`media_list_extended`)
- Media Statistics (`media_statistics`)
- List of Media (`pool_list`)

Single Session

"Single session" reports provide various information about single Data Protector backup, object copy, object consolidation, or object verification sessions. These reports are mostly used as End of Session notification. In this case, Data Protector will use the session ID of the current session (the one that generated the End of Session event) to create the appropriate report. "Single session" reports are:

- Single Session Report (`single_session`)
- Session Objects Report (`session_objects`)

- Session per Client Report (`session_hosts`)
- Session Devices Report (`session_devices`)
- Session Media Report (`session_media`)
- Session Object Copies Report (`session_objcopies`)

OPTIONS

`-version`

Displays the version of the `omnirpt` command.

`-help`

Displays the usage synopsis for the `omnirpt` command.

`-header`

This option is not used for the reports that have no required or optional report options. If this option is set, the output of the report will display report options too. If it is not set, only the output of the report is displayed.

`-multicell`

This option is only used with Manager-of-Managers. If this option is specified, the report will be generated for all Cell Managers configured in the MoM environment (multicell report).

`-[no_]multiple`

This option is only used for enterprise reports (multicell) and for Session per Client reports. If this option is specified, the report will be divided into sections. For enterprise reports the report will be divided by Cell Manager and for Session per Client reports it will be divided by client.

`Report Names`

`list_sessions`

Lists all sessions in the specified time frame. The report is defined by set of options that specify report parameters. By default, the report is generated for all session specifications. Use the report filtering options to generate a report only for a specific backup, object copy, object consolidation, or object verification specification.

Report options are:

`-timeframe` {*Start Duration*|*Day Hour Day Hour*}

[`-datalist` *BackupSpecificationName* ...]

[`-group` *BackupSpecificationGroup*]

[`-copylist_sch` *ScheduledCopySpecificationName* ...]

[`-copylist_post` *PostbackupCopySpecificationName* ...]

[`-verificationlist_sch` *ScheduledVerificationSpecificationName* ...]

[`-verificationlist_post` *PostbackupVerificationSpecificationName* ...]

[`-conslist_sch` *ScheduledConsolidationSpecificationName* ...]

[-conslist_post *PostbackupConsolidationSpecificationName* ...]

Report filtering options are:

[-no_datalist]

[-no_copylist]

[-no_verificationlist]

[-no_conslist]

session_flow

Graphically presents duration of each session specified in certain time frame. Flow chart of the backup, object copy, object consolidation and object verification sessions matching search criteria is shown. By default, the report is generated for all session specifications. Use the report filtering options to generate a report only for a specific backup, object copy, object consolidation, or object verification specification.

Report options are:

-timeframe {*Start Duration* | *Day Hour Day Hour*}

[-datalist *BackupSpecificationName* ...]

[-group *BackupSpecificationGroup*]

[-copylist_sch *ScheduledCopySpecificationName* ...]

[-copylist_post *PostbackupCopySpecificationName* ...]

[-verificationlist_sch *ScheduledVerificationSpecificationName* ...]

[-verificationlist_post *PostbackupVerificationSpecificationName* ...]

[-conslist_sch *ScheduledConsolidationSpecificationName* ...]

[-conslist_post *PostbackupConsolidationSpecificationName* ...]

Report filtering options are:

[-no_datalist]

[-no_copylist]

[-no_verificationlist]

[-no_conslist]

device_flow

Graphically presents usage of each device. Flow chart of the backup, object copy, and object consolidation sessions matching search criteria is shown. If you set the RptShowPhysicalDeviceInDeviceFlowReport 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. By default, the report is generated for all session specifications. Use the report filtering options to generate a report only for a specific backup, object copy, or object consolidation specification.

Report options are:

-timeframe {*Start Duration* | *Day Hour Day Hour*}

[-datalist *BackupSpecificationName* ...]

[-group *BackupSpecificationGroup*]

[-copylist_sch *ScheduledCopySpecificationName* ...]

[-copylist_post *PostbackupCopySpecificationName* ...]

[-conslist_sch *ScheduledConsolidationSpecificationName* ...]

[-conslist_post *PostbackupConsolidationSpecificationName* ...]

Report filtering options are:

[-no_datalist]

[-no_copylist]

[-no_conslist]

used_media

Lists destination media that have been used by backup, object copy, and object consolidation sessions in the specific time frame together with their statistics. By default, the report is generated for all session specifications. Use the report filtering options to generate a report only for a specific backup, object copy, or object consolidation specification.

Report options are:

-timeframe {*Start Duration*|*Day Hour Day Hour*}

[-datalist *BackupSpecificationName* ...]

[-group *BackupSpecificationGroup*]

[-copylist_sch *ScheduledCopySpecificationName* ...]

[-copylist_post *PostbackupCopySpecificationName* ...]

[-conslist_sch *ScheduledConsolidationSpecificationName* ...]

[-conslist_post *PostbackupConsolidationSpecificationName* ...]

Report filtering options are:

[-no_datalist]

[-no_copylist]

[-no_conslist]

used_media_extended

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. By default, the report is generated for all session specifications. Use the report filtering options to generate a report only for a specific backup, object copy, or object consolidation specification.

Report options are:

-timeframe {*Start Duration*|*Day Hour Day Hour*}

[-datalist *BackupSpecificationName* ...]

[-group *BackupSpecificationGroup*]

[-copylist_sch *ScheduledCopySpecificationName* ...]

[-copylist_post *PostbackupCopySpecificationName* ...]

[-conslist_sch *ScheduledConsolidationSpecificationName* ...]

[-conslist_post *PostbackupConsolidationSpecificationName* ...]

Report filtering options are:

[-no_datalist]

[-no_copylist]

[-no_conslist]

host_statistics

Lists of clients and their backup status - only clients that were used by the backup sessions matching the search criteria are displayed.

Additionally, clients can be limited also with the -hosts report option.

The VADP feature introduced in Data Protector 8.14 provides enhanced reports for Virtual Machines. Reports display VMware virtual machines in the same way as Data Protector clients called VADP clients. The VM name is the client name.

Report options are:

-timeframe {*Start Duration* | *Day Hour Day Hour*}

[-datalist *BackupSpecificationName* ...]

[-group *BackupSpecificationGroup*]

[-hosts]

session_statistics

Shows statistics about backup, object copy, and object consolidation status in the selected time frame, limited to sessions matching the search criteria. By default, the report is generated for all session specifications. Use the report filtering options to generate a report only for a specific backup, object copy, or object consolidation specification.

Report options are:

-timeframe {*Start Duration* | *Day Hour Day Hour*}

[-datalist*BackupSpecificationName* ...]

[-group*BackupSpecificationGroup*]

[-copylist_sch*ScheduledCopySpecificationName* ...]

[-copylist_post*PostbackupCopySpecificationName* ...]

[-conslist_sch*ScheduledConsolidationSpecificationName* ...]

[-conslist_post*PostbackupConsolidationSpecificationName* ...]

Report filtering options are:

[-no_datalist]

[-no_copylist]

[-no_conslist]

session_errors

Shows list of messages that occur during backup, object copy, object consolidation, and object verification sessions in the specified time frame for selected session specifications. The messages are grouped by clients (for all selected clients). By default, the report is generated for all session specifications. Use the report filtering options to generate a report only for a specific backup, object copy, object consolidation, or object verification specification.

Report options are:

-timeframe {*Start Duration*|*Day Hour Day Hour*}

[-datalist *BackupSpecificationName* ...]

[-group *BackupSpecificationGroup*]

[-copylist_sch *ScheduledCopySpecificationName* ...]

[-copylist_post *PostbackupCopySpecificationName* ...]

[-conslist_sch *ScheduledConsolidationSpecificationName* ...]

[-conslist_post *PostbackupConsolidationSpecificationName* ...]

[-verificationlist_sch *ScheduledVerificationSpecificationName* ...]

[-verificationlist_post *PostbackupVerificationSpecificationName* ...]

[-hosts *Hostname* ...]

[-level *Level*]

Report filtering options are:

[-no_datalist]

[-no_copylist]

[-no_conslist]

[-no_verificationlist]

obj_copies

Lists object versions that are created in the specified time frame with the number of their valid copies. The number of copies includes the original object version. By default, the report is generated for all session specifications. Use the report filtering options to generate a report only for a specific backup, object copy, or object consolidation specification.

Reports display VMware virtual machines in the same way as Data Protector clients called VADP clients. The new object name format is as follows:

<hostname>:/<vCenter>/<path>/<vmname>[<UUID>]

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.

Report options are:

-timeframe {*Start Duration*|*Day Hour Day Hour*}

`-num_copies` {*less* | *equal* | *more*} *NumberOfCopies*

[`-datalist` *BackupSpecificationName* ...]

[`-group` *BackupSpecificationGroup*]

[`-copylist_sch` *ScheduledCopySpecificationName* ...]

[`-copylist_post` *PostbackupCopySpecificationName* ...]

[`-conslist_sch` *ScheduledConsolidationSpecificationName* ...]

[`-conslist_post` *PostbackupConsolidationSpecificationName* ...]

Report filtering options are:

[`-no_datalist`]

[`-no_copylist`]

[`-no_conslist`]

`dl_trees`

Lists all trees in the specified backup specification. It also shows names of drives and the name of a tree.

Reports display VMware virtual machines in the same way as Data Protector clients called VADP clients. The report displays all the VM names for VMware objects.

Report options are:

[`-datalist` *BackupSpecificationName* ...]

[`-group` *BackupSpecificationGroup*]

`obj_nobackup`

Lists all objects, specified for backup in selected backup specifications, which do not have a valid backup. A valid backup means that the backup completed successfully and its protection has not expired. For each object that does not have a valid protected full backup, the following items are shown: backup specification, an object type, an object name and a description. Only objects from the selected backup specification are used for the report. If HOST object is used: Host object is expanded (get disks) and report checks that expanded objects are in database. UNIX and Windows filesystems are supported. This option is not available for backup specifications for integrations.

Report options are:

[`-datalist` *BackupSpecificationName* ...]

[`-group` *BackupSpecificationGroup*]

[`-days` *NoOfDays*]

`obj_lastbackup`

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.

Objects of the Client System type (host backup) are expanded; it means that the information is listed for each volume separately. As for objects of the Filesystem type (filesystem objects), only the UNIX and Windows filesystems are supported.

Reports display VMware virtual machines in the same way as Data Protector clients called VADP clients. The new object name format is as follows:

`<hostname>:/<vCenter>/<path>/<vmname>[<UUID>]`

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.

You can narrow the scope of objects listed using the following report options:

`[-datalist` *BackupSpecificationName* ...`]`

`[-group` *BackupSpecificationGroup*`]`

`[-days` *NoOfDays*`]`

However, note the following:

– Filesystem objects that do not match the condition in the object creation time filter are listed anyway. However, in this case, the 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.

The above note is not applicable for objects of the Bar type (integration objects).

`obj_avesize`

Lists all objects, specified for backup in selected backup specifications, which have a valid backup. A valid backup means that the backup completed successfully and its protection has not expired. For each object average full and average incremental backup size is displayed. If HOST object is used: Host object is expanded (get disks) and report checks that expanded objects are in database. UNIX and Windows filesystems are supported.

Reports display VMware virtual machines in the same way as Data Protector clients called VADP clients.

The new object name format is as follows:

`<hostname>:/<vCenter>/<path>/<vmname> [<UUID>]`

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.

Report options are:

`[-datalist` *BackupSpecificationName* ...`]`

`[-group` *BackupSpecificationGroup*`]`

`[-days` *NoOfDays*`]`

`fs_not_conf`

Displays a list of mounted filesystems which are not in selected backup specifications. Output is a list of filesystems. If HOST object is used, the report will not report any disk from client as not configured (assuming that HOST backup will backup all disks). If HOST object is used, the report will not report any disk from client as not configured (assuming that HOST backup will backup all disks).

Report options are:

`[-datalist` *BackupSpecificationName* ...`]`

[-group *BackupSpecificationGroup*]

dl_info

Shows 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 & post exec commands. Host does not influence the report. By default, the report is generated for all session specifications. Use the report filtering options to generate a report only for a specific backup, object copy, object consolidation, or object verification specification.

Report options are:

[-datalist *BackupSpecificationName* ...]

[-group *BackupSpecificationGroup*]

[-copylist_sch *ScheduledCopySpecificationName* ...]

[-copylist_post *PostbackupCopySpecificationName* ...]

[-verificationlist_sch *ScheduledVerificationSpecificationName* ...]

[-verificationlist_post *PostbackupVerificationSpecificationName* ...]

[-conslist_sch *ScheduledConsolidationSpecificationName* ...]

[-conslist_post *PostbackupConsolidationSpecificationName* ...]

Report filtering options are:

[-no_datalist]

[-no_copylist]

[-no_verificationlist]

[-no_conslist]

dl_sched

Shows information about all selected backup, object copy, object consolidation, and object verification specifications and their next scheduled time up to one year in advance (type, session type, session specification name, group, next execution, and backup operation time). HOST does not influence report. By default, the report is generated for all session specifications. Use the report filtering options to generate a report only for a specific backup, object copy, object consolidation, or object verification specification.

Report options are:

[-datalist *BackupSpecificationName* ...]

[-group *BackupSpecificationGroup*]

[-copylist_sch *ScheduledCopySpecificationName* ...]

[-copylist_post *PostbackupCopySpecificationName* ...]

[-verificationlist_sch *ScheduledVerificationSpecificationName* ...]

[-verificationlist_post *PostbackupVerificationSpecificationName* ...]

[-conslist_sch *ScheduledConsolidationSpecificationName* ...]

[-conslist_post *PostbackupConsolidationSpecificationName* ...]

Report filtering options are:

[-no_datalist]

[-no_copylist]

[-no_verificationlist]

[-no_conslist]

db_size

Provides a table that contains information about the MMDB, CDB, archived log files, datafiles, and information for DCBF and SMBF.

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.

hosts_unused

Lists configured clients that are not used for backup and do not have any device configured.

dev_unused

Lists configured destination devices that are not used for backup, object copy, or object consolidation at all.

lookup_sch

List of backup, object copy, and object consolidation specifications that are scheduled to start in the next n number of days up to one year in advance (where n is the number of days specified by user).

Report option is:

[-schedule *NoOfDays*]

hosts_not_conf

List of clients in selected domain(s) that are not configured for Data Protector. Note that Data Protector will display also routers and other machines that have IP address in selected domain.

Report option is:

-network *IP_Address*...

licensing

Lists all licenses and the available number of licenses.

host

Report output is all end-user backup related information about specific client: list of filesystems not configured for selected clients, list of all objects configured in backup specifications for the selected client, list of all objects with a valid backup for specified client with times and average sizes.

Note that Client Backup reports do not include information about application integration backup objects and backup specifications.

Report option is:

-host*HostName*

media_list

List of all media matching the search criteria. The following information is provided for each medium: ID, label, location, status, protection, used and total MB, the time when media was last used, the media pool, and media class.

Report options are:

[-label *Label*]

[-location *Location* ...]

[-pool *PoolName* ...]

[-class *MediaClass*]

[-status *MediaStatus*]

[-[no_]protection *NoOfDays*]

[-timeframe {*Start Duration*|*Day Hour Day Hour*}]

[-[no_]library *Library* ...]

media_list_extended

List of all media matching the search criteria. The following information is provided for each medium: ID, label, location, status, protection, used and total MB, the time when media was last used, the media pool and media type, session specifications that have used this medium for backup, object copy, or object consolidation, as well as the session type and subtype. By default, the report is generated for all session specifications. Use the report filtering options to generate a report only for a specific backup, object copy, or object consolidation specification.

Report options are:

[-label *Label*]

[-location *Location* ...]

[-pool *Pool Name* ...]

[-class *MediaClass*]

[-status *MediaStatus*]

[-[no_]protection *NoOfDays*]

[-timeframe {*Start Duration*|*Day Hour Day Hour*}]

[-[no_]library*Library* ...]

[-datalist *BackupSpecificationName*...]

[-group *BackupSpecificationGroup*]

[-copylist_sch *ScheduledCopySpecificationName* ...]

[-copylist_post *PostbackupCopySpecificationName* ...]

[-conslist_sch *ScheduledConsolidationSpecificationName* ...]

[-conslist_post *PostbackupConsolidationSpecificationName* ...]

Report filtering options are:

[-no_datalist]

[-no_copylist]

[-no_conslist]

media_statistics

Reports the 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; and total, used, and free space on media.

Report options are:

[-label *Label*]

[-location *Location* ...]

[-pool *PoolName* ...]

[-class *MediaClass*]

[-status *MediaStatus*]

[-[no_]protection *NoOfDays*]

[-timeframe {*Start Duration* | *Day Hour Day Hour*}]

[-[no_]library *Library* ...]

pool_list

Lists all pools matching a specified search criteria. For each pool the following information is provided: 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.

Report options are:

[-pool *PoolName* ...]

[-location *Location* ...]

[-class *MediaClass*]

[-[no_]library *Library* ...]

[-timeframe {*Start Duration* | *Day Hour Day Hour*}]

single_session

Report displays all relevant information about single Data Protector backup, object copy, object consolidation, and object verification sessions.

Report option is:

-session *SessionID*

[-level *Level*]

session_objects

Returns all information about all backup, object copy, or object consolidation objects that took part

in a selected session.

Returns information about VM name and VM path for VMware virtual machines manifested as Data Protector clients called VADP clients.

Report option is:

-session *SessionID*

session_hosts

Provides information for each client that took part in the selected backup session: statistics about backup status for the client, list of objects and their related information for the client, error messages for the client.

All information is grouped for each client separately. Using the -multiple option, this report can be split into smaller reports, one for each client (see section Notifications for details).

Reports display VMware virtual machines in the same way as Data Protector clients called VADP clients.

The new object name format is as follows:

<hostname>:/<vCenter>/<path>/<vmname>[<UUID>]

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.

Report option is:

-session *SessionID*

[-level *Level*]

session_devices

Provides information about all devices that took part in a selected session.

Report option is:

-session *SessionID*

session_media

Provides information about all destination media that took part in a selected session.

Report option is:

-session *SessionID*

session_objcopies

Lists object versions that are created in the selected backup, object copy, and object consolidation session with the number of their valid copies.

Reports display VMware virtual machines in the same way as Data Protector clients called VADP clients.

The new object name format is as follows:

<hostname>:/<vCenter>/<path>/<vmname>[<UUID>]

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.

Report option is:

`-session` *SessionID*

`Method options`

`-email` *EmailAddress*

Sends the report to the specified *EmailAddress*.

On Windows systems, you need a configured MAPI profile. You can either use an existing mail profile or create a new one, named `Omniback`. To use an existing profile, edit the omnirc option `OB2_MAPIPROFILE`.

On UNIX systems, `/usr/bin/mail` is used for sending the e-mails.

`-smtp` *EmailAddress*

The recommended option for sending reports by e-mail. Sends the report to the specified *EmailAddress* using the SMTP protocol.

By default, the SMTP server address is set to the Cell Manager address. To change the SMTP server, edit the `SMTPServer` global option. The server must be accessible from the Cell Manager system, but does not need to be part of the Data Protector cell.

`-snmp` *Hostname*

Report is send as an SNMP (Simple Network Mailing Protocol) trap.

`-broadcast` *Hostname*

Report is broadcasted to the selected machine. NOTE: Only Windows machines can be specified as broadcast destination.

`-log` *Filename*

Report is saved in to the log file specified with *Filename*.

`-external` *CommandName*

Specifies a script which receives the report. Optionally the script can than parse the report and forward it to user configured recipient. Usually, TAB report format is used in combination with `-external` option.

`Report options`

`-rptgroup` *ReportGroup*

This option executes the specified *ReportGroup*.

`-session` *SessionID*

This option is used to specify the session ID.

`-pool` *Poolname* ...

This option is used to specify the media pool name.

`-label`*Label*

This option is used to specify the medium label.

`-location` *Location* ...

This option is used to specify the medium location.

-[no_]library *Library* ...

This option is used to specify the library. If it is set to `-no_library`, all libraries in the cell are selected for the report.

-[no_]protection *NoOfDays*

This option is used to specify the protection. The number of days in which the protection will expire can be specified. If it is set to `no_protection`, all media in the cell will be selected for the report.

-class *MediaClass*

This option is used to specify the media class. It can have one of the following values: `DDS`, `QIC`, `Exabyte`, `AIT`, `SAIT`, `T3480/T4890/T9490`, `Optical`, `File`, `T9840`, `Tape`, `DLT`, `SD-3`, `T3590`, `T3592`, `LTO-Ultrium`, `SuperDLT`, `VXA`, `DTF`, `T9940`, `T10000`, `StoreOnceSoftware`, `DataDomainBoost` or `ObjectStore`.

-status *MediaStatus*

This option is used to specify the media status. It can have one of the following values: `poor`, `fair`, or `good`.

-datalist *BackupSpecificationName* ...

This option is used to specify the backup specifications for the report. If you specify more than one backup specification, separate the specification names with spaces.

-copylist_sch *ScheduledCopySpecificationName* ...

This option is used to specify the scheduled object copy specifications for the report. If you specify more than one scheduled object copy specification, separate the specification names with spaces.

-copylist_post *PostbackupCopySpecificationName* ...

This option is used to specify the post-backup object copy specifications. If you specify more than one post-backup object copy specification, separate the specification names with spaces.

-verificationlist_sch *ScheduledVerificationSpecificationName* ...

This option is used to specify the scheduled object verification specifications for the report. If you specify more than one scheduled object verification specification, separate the specification names with spaces.

-verificationlist_post*PostbackupVerificationSpecificationName* ...

This option is used to specify the post-backup object verification specifications. If you specify more than one post-backup object verification specification, separate the specification names with spaces.

-conslist_sch *ScheduledConsolidationSpecificationName* ...

This option is used to specify the scheduled object consolidation specifications. If you specify more than one scheduled object consolidation specification, separate the specification names with spaces.

-conslist_post *PostbackupConsolidationSpecificationName* ...

This option is used to specify the post-backup object consolidation specifications. If you specify more than one post-backup object consolidation specification, separate the specification names

with spaces.

-group *BackupSpecificationGroup*

This option is used to specify backup specification group for the report.

-schedule *NoOfDays*

This option is used to specify the number of days for which to display the schedule information.

-network *IP_Address* ...

This option specifies one or more IP addresses. Valid IP address forms are:

- a.b.c.d – a complete IPv4 address (for example, 10.17.1.1)
- a.b.c – an IPv4 C-class network address (for example, 10.17.1)
- IPv6 addresses in any valid form (for example, ::1, td10::abba:1603, and so on)

You can specify more than one IP address by using spaces in between.

-hosts *Hostname* ...

Select the client systems for which you want to create the report.

-host *Hostname*

Select the client system for which you want to create the report.

-level *Level*

Select the level of warnings that should be included in the report. The levels are `warning`, `minor`, `major`, and `critical`.

-num_copies {*less* | *equal* | *more*} *NumberOfCopies*

This option is used to specify the number of valid object versions copies. Note that you can specify more than, equal to, or less than the selected number of copies.

-timeframe *Start Duration*

This option is used to specify a relative time frame. It is useful for recurrent reports, for example you can use `-timeframe 24 24` each day to set the time frame to last 24 hours.

-timeframe *Day Hour Day Hour*

This option is used to specify an absolute time frame.

-days *NoOfDays*

The report will filter objects that have been backed up recently. Specify the number of days.

Report filtering options

-no_datalist

This option is used to exclude all backup specifications from the report.

-no_copylist

This option is used to exclude all object copy specifications from the report.

-no_verificationlist

This option is used to exclude all object verification specifications from the report.

`-no_conslist`

This option is used to exclude all object consolidation specifications from the report.

`Report Formats`

`-ascii`

Specifies report format: ASCII

`-html`

Specifies report format: HTML

`-tab`

Specifies report format: TAB

`-short`

Specifies report format: SHORT

NOTES

The virtual machine objects and its associated disk objects constitute the VEAgent object size. The other fields in the omnirpt output for the VMware objects display the details of the virtual machine object, and not the disk objects.

EXAMPLES

1. To list all backup sessions that have started in the last 24 hours and display the report in the default ASCII format, execute:

   `omnirpt -report list_sessions -timeframe 24 24 -no_copylist -no_conslist -no_ verificationlist`

2. To list all objects from session "2012/11/16-1" in tabulator separated format, which is useful for additional parsing or can be used with other tools for analysis, execute:

   `omnirpt -report session_objects -session 2012/11/16-1 -tab`

3. To list all media of class DLT with location string "COMPANY", for which protection will expire in the next 5 days, execute:

   `omnirpt -report media_list -protection 5 -class DLT -location COMPANY`

   This report can be used as a base for the vaulting process, as it can list you media that need to be taken to the vault.

4. To send "Internal Database Size Report" in HTML format to the user "name@domain.com" using the SMTP protocol, execute:

   `omnirpt -report db_size -html -smtp name@domain.com`

5. To execute the report group named "MyReportGroup", execute:

   `omnirpt -rptgroup MyReportGroup`

6. To graphically present the usage of devices that were used for backup and object consolidation (but not object copy) sessions in the last 48 hours in HTML format that will be sent as the file

"session1.html" to the directory "C:\Temp", execute:

```
omnirpt -report device_flow -timeframe 48 48 -no_copylist -html
>C:\Temp\session1.html
```

7. To list all the media used only for object copy and object consolidation sessions, execute:

```
omnirpt -report media_list_extended -no_datalist
```

8. To list all object versions created in the last 72 hours that have less than 5 valid copies, execute:

```
omnirpt -report obj_copies -timeframe 72 72 -num_copies less 5
```

9. To list all destination media that were used only for scheduled object copy specification named "Alpha" in the last 2 days, execute:

```
omnirpt -report used_media -timeframe 48 48 -copylist_sch Alpha -no_datalist -
no_conslist
```

10. To show statistics about backup status (but not object copy, object consolidation, and object verification) in the last 24 hours, execute:

```
omnirpt -report session_statistics -timeframe 24 24 -no_copylist -no_conslist -
no_verificationlist
```

11. To graphically present duration of all object consolidation sessions in the last 24 hours in HTML format that will be sent as the file "session_flow1.html" to the directory "C:\Temp", execute:

```
omnirpt -report session_flow -timeframe 24 24 -no_datalist -no_copylist -no_
verificationlist -html >C:\Temp\session_flow1.html
```

12. To show all virtual machines selected in the VEPA backup specifications, execute:

```
omnirpt –report dl_trees
```

13. To show the backup information about vCenters, ESXi servers and virtual machines in the last 24 hours, execute:

```
omnirpt –report host_statistics –timeframe 24 24
```

14. To show copy IDs for backed up virtual machines in the last 24 hours, execute:

```
omnirpt –report obj_copies –timeframe 24 24
```

15. To list information about the latest backup for all objects including virtual machines, execute:

```
omnirpt –report obj_lastbackup
```

16. To list average object size for all objects including virtual machines, execute:

```
omnirpt –report obj_avesize
```

17. To list information about objects including virtual machines for a given session, execute:

```
omnirpt -report session_objects -session 2015/10/14-1
```

18. To list virtual machines as clients included in the selected backup session, execute:

```
omnirpt -report session_hosts -session 2015/10/14-1
```

19. To list object versions of backed up virtual machines for the selected session, execute:

```
omnirpt -report session_objcopies -session 2015/10/14-1
```

SEE ALSO

omnihealthcheck(1M), omnitrig(1M)

omnistat(1)

omnistat — displays the status of active Data Protector backup and restore sessions
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnistat -version | -help

omnistat -session *SessionID* [-status_only | -monitor | -detail]

omnistat [-user *Username*] [-mount] [-error] [-detail]

omnistat -previous [-user *Username*] [{-since *Date* | -until *Date*} | -last *Number*] [-failed]

Date

 [*YY*]*YY/MM/DD*

DESCRIPTION

The `omnistat` command displays information on active sessions. You can view all active sessions (default) or only details of a specific session. An active session is referenced by its *SessionID*.

OPTIONS

-version

Displays the version of the `omnistat` command.

-help

Displays the usage synopsis for the `omnistat` command.

-session *SessionID*

Displays detailed information on the single active session identified by this *SessionID*.

-monitor

`omnistat` connects to the specified active session and starts monitoring the progress of the session.

-status_only

Displays only the overall status of the active session.

-detail

Displays detailed information about all current sessions.

-user *Username*

Displays information on active sessions belonging to the specified user.

`-failed`

Displays information on sessions containing data objects that failed due to errors.

`-error`

Displays information on active sessions with the status "In Progress (errors)"

`-mount`

Displays all active sessions with mount requests pending.

`-previous`

Lists all sessions from the Data Protector Internal Database (IDB).

`-since` *Date*

Lists all sessions since the specified *Date*.

`-until` *Date*

Lists all sessions until the specified *Date*.

`-last` *n*

Lists all sessions within the last *n* days.

EXAMPLES

The following examples illustrate how some options of the `omnistat` command work.

1. To view sessions that are currently active and have mount requests pending, execute:

   `omnistat -mount`

2. To see detailed information for the session with the SessionID "2013/04/24-32", execute the following commands. The SessionID can be specified in two different formats. If the short format is used, the ID refers to the session that was run in the same day:

   `omnistat -detail -session 2013/04/24-32`

   `omnistat -detail -session 32`

3. To see an overview of the sessions that occurred in last 3 days and were run by user root, execute:

   `omnistat -previous -user root -last 3`

4. To see information regarding the sessions that occurred within the last 3 days and had objects that have failed, execute:

   `omnistat -previous -last 3 -failed`

5. To see only the status of session with this SessionID, execute:

   `omnistat -status_only -session 2`

6. To monitor the session with the SessionID "R-2013/05/13-8", execute:

   `omnistat -session R-8 -monitor`

SEE ALSO

omniabort(1)

omniupload(1)

omniupload — uploads information about a backup device from an ASCII file to the Data Protector Internal Database (IDB)
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omniupload -version | -help

omniupload -create_device *FileName*

omniupload -modify_device *BackupDevice* [-file *FileName*]

omniupload -remove_device *BackupDevice*

omniupload -create_library *FileName*

omniupload -modify_library *Library* [-file *FileName*]

omniupload -remove_library *Library*

DESCRIPTION

Uploads a backup device file to the Data Protector Internal Database (IDB).

Information on Data Protector backup devices is stored in the IDB. To configure a backup device, information on this device must be downloaded into a file. This is done using the `omnidownload` command. The file is then modified and uploaded back to the IDB.

OPTIONS

`-version`

Displays the version of the `omniupload` command.

`-help`

Displays the usage synopsis for the `omniupload` command.

`-create_device` *FileName*

Specifies the ASCII file containing the information about the backup device. This option is used to create a new backup device. If `-` is specified as *FileName* then data is read from stdin.

`-modify_device` *BackupDevice*

Uses the information in the uploaded file to modify an existing backup device in the IDB. If no filename is specified using the `-file` option the command searches the current directory for a file with the same name as the *BackupDevice*. Note that the media class may not be changed.

`-file` *FileName*

Specifies the ASCII file that will be parsed for information about the backup device (library). This option is used to modify an existing backup device (library). If - is specified as *FileName* then data is read from stdin.

`-remove_device` *BackupDevice*

Removes information about the *BackupDevice* from the IDB.

`-create_library` *FileName*

Specifies the ASCII file containing the information about the library. This option is used to create a new library. If - is specified as *FileName* then data is read from stdin.

`-modify_library` *Library*

Uses the information in the uploaded file to modify an existing library in the IDB. If no filename is specified using, the `-file` option the command searches the current directory for a file with the same name as the *Library*. Note that the media class may not be changed.

`-remove_library` *Library*

Removes information about the *Library* from the IDB.

EXAMPLES

The following examples illustrate how the `omniupload` command works.

1. To create a backup device using the information in the file "/tmp/Device", execute:

   `omniupload -create_device /tmp/Device`

2. To modify library "Exabyte1" using the information in the file "/tmp/EXA", execute:

   `omniupload -modify_library Exabyte1 -file /tmp/EXA`

3. To remove backup device "Stacker", execute:

   `omniupload -remove_device Stacker`

4. To create a virtual tape library named "VTL16" using the information in the file "lib16.txt", execute:

   `omniupload -create_library lib16.txt`

5. To modify the library capacity of a virtual tape library named "VTL" in an ASCII file named "libVTL.txt" in the directory "C:\Temp" to 50 TB, set the VTLCAPACITY parameter to 50:

   `VTLCAPACITY 50`

   and execute:

   `omniupload -modify_library VTL -file C:\Temp\libVTL.txt`

   Note that the VTLCAPACITY value in terabytes (TB) must be an integer.

SEE ALSO

omniamo(1), omnib2dinfo(1M), omnidownload(1), omnimcopy(1), omniminit(1), omnimlist(1), omnimm (1), omnimnt(1), omnimver(1), sanconf(1M), uma(1M)

omniusb(1)

omniusb — writes the DR ISO image to a USB drive, and makes the drive bootable
(this command is available on systems with the Data Protector `Automatic Disaster Recovery` component installed)

SYNOPSIS

omniusb --version|--help

omniusb --iso *Path* {--drive *VolumePath*|--disk *DiskNumber*} [--silent]

DESCRIPTION

The `omniusb` writes the disaster recovery OS, converted from the DR ISO image – which was created using the GUI or the `omniiso` command – , to a USB drive, and makes the drive bootable. You can then use the bootable USB drive to start your recovery process.

You can use this command to automate your backup and disaster recovery preparation.

Alternatively, you can create a bootable USB drive can using the EADR Wizard from the Data Protector GUI.

OPTIONS

`--version`

Displays the version of the `omniiso` command.

`--help`

Displays the usage synopsis for the `omniiso` command.

`--iso` *Path*

Specifies the location where the disaster recovery ISO image file is located.

`--drive` *MountPath*

Specifies the mount path to which the USB drive is mounted, for example `E:\.`

`--disk` *DiskNumber*

Specifies the USB drive by its disk number as reported by the Windows Disk Management Extension.

`--silent`

Suppresses any user interaction. This option is applicable if the command is used in a pre--exec script.

NOTES

The `omniusb` command is available on Windows systems only.

EXAMPLES

The following examples illustrate how the `omniusb` command works.

1. To save the USB drive image created from a disaster recovery ISO file, located in "C:\iso\dr\omnidr.iso", to a USB drive, mounted under "G:", execute:

   ```
   omniusb --iso c:\iso\dr\omnidr.iso --drive G:
   ```

2. To save a disaster recovery ISO file, located in "C:\iso\dr\omnidr.iso", to a USB drive with the disk number "6", execute:

   ```
   omniusb --iso c:\iso\dr\omnidr.iso --disk 6
   ```

SEE ALSO

omniiso(1), omnidr(1M), omniofflr(1M), omnisrdupdate(1M)

omniusers(1)

omniusers — adds or removes Data Protector users to or from an existing Data Protector user group, or lists the configured Data Protector users.
(this command is available on non-Windows systems with the Data Protector User Interface component installed)

SYNOPSIS

omniusers -version|-help

omniusers -add -type {U|W} -usergroup *DPUserGroup* -name *UserName* -group *GroupOrDomainName* -client *ClientName* [-desc *Description*]

omniusers -remove -name *UserName* -group *GroupOrDomainName* -client *ClientName*

omniusers -list

DESCRIPTION

The command adds, removes, or lists the Data Protector users configured on the Cell Manager. It does not create or remove Data Protector user groups.

Use the command to remotely add a new Data Protector user account from a system where the Data Protector GUI is not supported. You can then use the account to start the Data Protector GUI on another system, and connect to the Cell Manager.

OPTIONS

-version

Displays the version of the omniusers command.

-help

Displays the usage synopsis for the omniusers command.

-add

Adds a user to the specified Data Protector user group.

-remove

Removes a user from its Data Protector user group.

-name *UserName*

Specifies username of the user to be added/removed. By specifying asterisk (*) as the username, all users from the specified group (on UNIX systems) or domain (on Windows systems) will be granted/revoked access from the specified clients to the Cell Manager.

* corresponds to `<Any>` in the Data Protector GUI. Note that in some shells, backslash and asterisk (*) must be used instead of *.

Note: UNIX usernames and usernames of the configured Data Protector users are case-sensitive.

Note:The usernames and domain names of Windows GUI clients that are used with an HP-UX Cell Manager must be in capital letters.

`-type {U|W}`

Specifies the user type: a UNIX user (`U`) or a Windows user (`W`).

`-group` *GroupOrDomainName*

A group (on UNIX systems) or a domain (on Windows systems) the specified user belongs to. By specifying asterisk (*) as the group or domain name, the specified user will be granted/revoked access from any group or domain from the specified clients.

* corresponds to `<Any>` in the Data Protector GUI. In some shells, `backslash and asterisk` (*) must be used instead of *.

Note that domain names of Windows GUI clients that are used with an HP-UX Cell Manager must be in capital letters.

`-client` *ClientName*

Specifies the name of the client system from where the specified user will have access to the Cell Manager. By specifying asterisk (*) as the client name, the specified user will be granted/revoked access to the Cell Manager from any Data Protector client system.

* corresponds to `<Any>` in the Data Protector GUI. Note that in some shells, `backslash and asterisk (\*)` must be used instead of *.

If this option is used with the `-remove` option, *ClientName* must contain the fully qualified domain name (FQDN) of the client system.

`-usergroup` *DPUserGroup*

Specifies the Data Protector user group the user(s) will be added to.

`-desc` *Description*

Specifies the description for the added user(s).

`-list`

Lists users in all configured Data Protector user groups in the cell. For each configured Data Protector user the username, UNIX group or Windows domain, fully qualified domain name (FQDN) of the client system from which the user has granted access, and the user description are displayed. Asterisk (*) corresponds to the `<Any>` string in the Data Protector GUI.

RETURN VALUES

The return values of the `omniusers` command are:

`0` - The command operation completed successfully.

`1` - A generic error occurred.

`2` - The operation for adding or removing a user failed.

4 - Error parsing options.

EXAMPLES

The following examples illustrate how the `omniusers` command works.

1. To add the Windows user "win_user" from the domain "domain1" to the Data Protector "admin"
    user group and allow access only from the client system "client.company.com", execute:

    ```
    omniusers -add -type W -name win_user -usergroup admin -group domain1 -client
    client.company.com
    ```

2. To add the UNIX user "root" from the "sys" group to the Data Protector "admin" user group and
    allow access only from the client system "client.company.com", execute:

    ```
    omniusers -add -type U -name root -usergroup admin -group sys -client
    client.company.com
    ```

3. To add the UNIX user "root" to the Data Protector "admin" user group and allow access from any
    UNIX group but only from the system "client.company.com", execute:

    ```
    omniusers -add -type U -name root -usergroup admin -group \* -client
    client.company.com
    ```

4. To display the Data Protector users in all configured Data Protector user groups, execute:

    ```
    omniusers -list
    ```

SEE ALSO

ob2install(1M), omnigui(5), omniintro(9), omnimigrate.pl(1M), omnisetup.sh(1M), upgrade_cm_from_
evaa(1M)

SharePoint_VSS_backup.ps1(1)

SharePoint_VSS_backup.ps1 — creates backup specifications and starts backup sessions for Microsoft SharePoint Server
(this command is available on systems with the Data Protector `MS Volume Shadow Copy Integration` component installed)

SYNOPSIS

SharePoint_VSS_backup.ps1 –help|-version

SharePoint_VSS_backup.ps1 –createonly *CreateOptions*

SharePoint_VSS_backup.ps1 –backuponly *BackupOptions*

SharePoint_VSS_backup.ps1 -preview [–resumefarm]|-resumecert

CreateOptions

 {-device *DevName*|-hardware {no_keep|keep|ir} [-device *DevName*]}

 [-overwrite]]

 [-prefix *PrefixName*]

 [-excludeindex]

BackupOptions

 [-outfile *PathToFile*]

 [-prefix *PrefixName*]

 [-preview]

 [-snapshot {diskonly|disktape|tapeonly}]

 [-reduce]

 [-mode {full|incremental|incremental1…|incremental9}]

 [-timeout *Timeout*]

DESCRIPTION

The `SharePoint_VSS_backup.ps1` command creates backup specifications and start backup sessions for Microsoft SharePoint Server, using the Data Protector Volume Shadow Copy Service integration.

When you execute the command, Data Protector first queries for information about the Microsoft SharePoint Server environment. Then it creates backup specifications.

The newly created backup specifications are named `SharePoint_VSS_backup_ClientName` and have the same backup device specified for use (the one that you specified at command runtime). Once the

backup specifications are created, the command starts backup sessions (one session for each backup specification).

You can also only create the backup specifications first, modify them in the Data Protector GUI if necessary and then start the backup sessions.

OPTIONS

`-help`

Displays the `SharePoint_VSS_backup.ps1` command usage.

`-version`

Displays the `SharePoint_VSS_backup.ps1` version.

`–createonly`

If this option is specified, Data Protector only creates backup specifications. Backup is not started.

`–backuponly`

If this option is specified, Data Protector only starts backup sessions using the existing backup specifications. The `-device` option is not required.

`-device` *DevName*

Specifies which Data Protector device to use for backup. You can specify only one device.

`-hardware` {no_keep|keep|ir}

Specifies that the hardware provider should be used (instead of the software provider with `-device` option specified) and, consequently, ZDB options set. The default values for ZDB options are as follows:

- Keep the replica for instant recovery: selected if `ir` is specified.
- Keep the replica after the backup: selected if `ir` or `keep` is specified.
- Configuration check mode: Strict
- Replica type: Mirror/Clone (Plex)
- Numbers of replica rotated: 3

The default ZDB backup types are as follows (provided a device is also specified):

- `no_keep`: ZDB-to-tape
- `keep`: ZDB-to-disk+tape
- `ir`: ZDB-to-disk+tape

`–overwrite`

By default, Data Protector does not create backup specifications if they already exist. If this option is specified, Data Protector overwrites the existing backup specifications with the newly-created ones. Not applicable if `–backuponly` is specified.

`-prefix` *PrefixName*

With this option specified, the backup specifications are created under a different name: `SharePoint_VSS_backup_`*PrefixName*`_`*ClientName*.

In case of backup, this option specifies which backup specifications to use: those which name contains *PrefixName*.

Non-ASCII characters in *PrefixName* are not supported.

`-outfile` *PathToFile*

If this option is specified, backup specification names, errors, sessions outputs, and `omnir` restore commands are written to the specified file.

`-preview`

If this option is specified, Data Protector displays information about the Microsoft SharePoint Server environment and describes the related actions without actually performing them.

`-snapshot` {`diskonly`|`disktape`|`tapeonly`}

Applicable when starting ZDB backup sessions (that is, sessions that use backup specifications in which a hardware provider is specified for use). Performs a ZDB-to-disk (`diskonly`), ZDB-to-tape (`tapeonly`) or ZDB-to-disk+tape (`disktape`) session.

`-reduce`

Microsoft SharePoint Server 2010: If this option is specified, the command excludes mirrored query components from backup to reduce the backup size.

Microsoft SharePoint Server 2013: If this option is selected, the command selects primary index replicas of each index partition to reduce the backup size.

`-excludeindex`

Applicable only to a Data Protector standard filesystem backup of the FAST Search index files (Microsoft SharePoint Server 2010). If this option is specified, Data Protector excludes `data_index` folder contained in the `FASTSearch` home folder from backup specification. This way, the backup is faster, but the restore is more time consuming. The option enables balancing between a backup size and a time to recovery.

`-mode` {`full` | `incremental` | `incremental1...` | `incremental9`}

Applicable only to a Data Protector standard filesystem backup of the FAST Search index files (Microsoft SharePoint Server 2010). With this option specified, either a `Full` or `Incremental` or leveled incremental backup can be started. By default, the `Full` backup is performed.

When the `incremental` option is specified and the `Full` backup does not exist, the option is ignored and the `Full` filesystem backup of the FAST Search index files is started.

`-resumecert`

Applicable only to Microsoft FAST Search Server 2010. If this option is specified, the FAST Search certificates for the content and the query connectors are reinstalled.

`-resumefarm`

To be used after restore. This option returns the farm to a working state by resuming all background activities and crawling, unlocking sites, and starting Microsoft SharePoint Server services.

`-timeout` *Timeout*

This option sets the timeout in minutes after which the crawl of the FAST Search index files is aborted and the farm is resumed. If not specified, the default timeout is 15 minutes.

NOTES

The `SharePoint_VSS_backup.ps1` command is available on Windows systems only.

EXAMPLES

Creating backup specifications:

1. To create backup specifications in which the backup device "filelib_writer1" is specified for use, execute:

   `SharePoint_VSS_backup.ps1 –createonly –device filelib_writer1`

2. To create backup specifications with the label "weekly" in their names and in which the backup device "dev1" is specified for use, execute:

   `SharePoint_VSS_backup.ps1 –createonly –device dev1 –prefix weekly`

3. To create ZDB backup specifications in which the backup device "dev1" and the hardware provider (ZDB disk array) are specified for use, and in which the ZDB option "Keep the replica for instant recovery" is enabled, execute:

   `SharePoint_VSS_backup.ps1 –createonly -hardware ir –device dev1`

4. Applicable only to a Data Protector standard filesystem backup of the FAST Search index files (Microsoft SharePoint Server 2010).

   To create filesystem backup specifications in which the backup device "dev1" is specified for use and with the "data_index" folder, contained in the "FASTSearch" home folder, excluded from the backup of the FAST Search index files, execute:

   `SharePoint_VSS_backup.ps1 –createonly –device dev1 -excludeindex`

Starting backup sessions:

1. To preview the actions that are performed when a backup session is started, execute:
   `SharePoint_VSS_backup.ps1 –backuponly –prefix dev –preview`

2. To start backup sessions using the existing backup specifications that have no prefix in their names, execute:

   `SharePoint_VSS_backup.ps1 –backuponly`

3. To start backup sessions using the existing backup specifications that have the prefix `weekly` in their names, execute:

   `SharePoint_VSS_backup.ps1 –backuponly –prefix weekly`

4. To start backup sessions using the existing backup specifications that have no prefix in their names and to save the output of the sessions and the associated restore commands to the file "c:\logs\shp.log", execute:

   `SharePoint_VSS_backup.ps1 –backuponly –outfile C:\logs\shp.log`

5. To start ZDB-to-disk backup sessions using the existing ZDB backup specifications that have no prefix in their names, execute:

   `SharePoint_VSS_backup.ps1 –backuponly -snapshot diskonly`

6. To start incremental filesystem backup sessions of the FAST Search index files (Microsoft

SharePoint Server 2010), execute:

```
SharePoint_VSS_backup.ps1 -backuponly -mode incremental
```

SEE ALSO

omnib(1)

syb_tool(1)

syb_tool — a utility used to get ISQL command needed to restore a Sybase database that was backed up by Data Protector
(this command is available on systems with the Data Protector `Sybase Integration` component installed)

SYNOPSIS

syb_tool *dbname servername*

 -date *YYYY/MM/DD.hh:mm:ss*

[-new_db *dbname*]

[-new_server *servername*]

[-file *filename*]

[-media]

DESCRIPTION

The `syb_tool` is used to get the data needed for restore of Sybase databases.

OPTIONS

dbname

   The name of Sybase database.

servername

   The name of Sybase database server on which the backup was performed.

-date *YYYY/MM/DD.hh:mm:ss*

   The date until which your database will be restored. `syb_tool` will find the first backup done after this date.

-new_db *dbname*

   The name of the database that you want to restore to.

-new_server *servername*

   The name of the server that you want to restore to.

-file *filename*

   The name of the file where the ISQL statement needed for restore of desired database will be written to. The ISQL command can be started with the option -i, followed by the name of the file.

   See also the section "Notes".

```
-media
```

   This option returns the list of all media needed for restore.

NOTES

If the names of backed up objects contain characters that cannot be displayed using the current language group (on Windows) or code page (on UNIX):

1. Set the encoding used on the terminal to UTF-8.

2. **_Windows systems:_** Set the environment variable `OB2_CLI_UTF8` to 1.

3. Redirect the output of the `syb_tool` command to a text file using the `-i` option.

   If you need to edit the file containing the load command, use a UTF-8 aware editor that does not set the first byte ("BOM"), since such a file is not supported by `isql`. Note that the Windows Notepad editor cannot be used.

4. When restoring the objects, add the `-i` *file_name* `-J utf8` options to the `isql` command, where *file_name* is the file with the load command.

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

EXAMPLES

1. To get the ISQL statement needed for the restore of the last backup of the database named "database1" on the Sybase Adaptive Server named "server", execute:

   ```
   syb_tool database1 server -date
   ```

2. To get the ISQL statement needed for the restore of the database named "database1" on the Sybase Adaptive Server named "server", using the first backup performed after midday of May 07 2013, execute:

   ```
   syb_tool database1 server -date 2013/05/07.12:00:00
   ```

3. To get the ISQL statement needed for the restore of the database named "database1" on the Sybase Adaptive Server named "server", using the first backup performed after midday of May 07 2013 and restoring it as "database_one" on the Sybase server called "server_one", execute:

   ```
   syb_tool database1 server -date 2013/05/07.12:00:00 -new_db database_one -new_
   server server_one
   ```

4. To get the ISQL statement needed for the restore of the last backup performed for database named "database1" on the Sybase Adaptive Server named "server", saving the ISQL statement to file "/tmp/stat.isql", and getting the list of media IDs needed for restore, execute:

   ```
   syb_tool database1 server -date -file /tmp/stat.isql -media
   ```

   To start the restore, start the ISQL command, specifying the input file "/tmp/stat.isql" in the following way:

   ```
   isql -Usa -P -Sserver -i /tmp/stat.isql
   ```

Section 1M: Administrative commands

cjutil(1M)

cjutil — starts, stops, and queries the Windows Change Journal
(this command is available on systems with the Data Protector `Disk Agent` component installed)

SYNOPSIS

cjutil -volume *vol* {-start [-maxsize *max* -delta *del*]|-stop [-wait]|-query}

DESCRIPTION

The `cjutil` command is used to control and administer the Change Journal.

OPTIONS

-volume *vol*

Defines the volume name in the form `/C` or `/C:\`*mounted_folder*.

-start [-maxsize *max* -delta *del*]

Starts the Change Journal on the specified volume.

The `-maxsize` *max* option sets the maximum size of the Change Journal in bytes. The highest possible value is 4 GB (4 294 967 296 bytes). Any specified value greater than 4 GB will be rounded down to 4 GB. Note that a reasonable size for a 100 GB drive is an 85 MB Change Journal.

The `-delta` *del* option specifies the size in bytes to be purged from the Change Journal when it reaches its maximum size. HPE recommends the value be approximately one-eighth to one-quarter the value of the maximum size but not greater than one quarter the size of the maximum size. This value may be automatically adjusted to better correspond to the volume cluster size.

-stop [-wait]

Stops the Change Journal asynchronously.

The `-wait` specifies that the Change Journal will be stopped synchronously. The call returns only after the Change Journal has been deleted.

-query

Queries the status of the Change Journal.

NOTES

If the `-start` option is specified and the Change Journal is already active, the Change Journal is adjusted to the value of the maximum size and delta. Note that these values can only be adjusted to increase.

When starting the Change Journal, if you not specify `-maxsize` and `-delta`, or specify `0` for these parameters, the system chooses a default value based on the volume size.

As an alternative to the Data Protector `cjutil` command, you can also use the Windows `fsutil` command for administering the Change Journal.

EXAMPLES

1. To start the Change Journal with the maximum size of 8 MB (in bytes) and specify the size of 1 MB (in bytes) to be purged from the Change Journal when it reaches the specified maximum size, execute:

   ```
   cjutil -start -maxsize 8388608 -delta 1048576
   ```

SEE ALSO

omnicjutil(1M)

ob2install(1M)

ob2install — runs installation, removal, upgrade, or installation check of the specified Data Protector components to/from/on a remote UNIX system using the specified Installation Server
(this command is available on the Data Protector Installation Server)

SYNOPSIS

ob2install -version|-help

ob2install -server *InstallationServer* -input *Filename*

DESCRIPTION

The `ob2install` command can be used to remotely install, remove, upgrade, or check the installation of Data Protector components to/from/on a remote UNIX system. To run the desired operation, you need to specify a UNIX Installation Server appropriate for the platform the remote system is using.

OPTIONS

-version

> Displays the version of the `ob2install` command.

-help

> Displays the usage synopsis for the `ob2install` command.

-server *InstallationServer*

> Specifies the Installation Server used for the installation session. The Installation Server must belong to local cell.

> **Note**: If the Cell Manager and the Installation Server are two different systems in the cell, the Cell Manager hostname must be listed on the Installation Server in the file `/etc/opt/omni/client/cell_server` (UNIX systems), *Data_Protector_program_ data*`\Config\client\cell_server` (Windows systems).

> **Note**: If the Cell Manager and the Installation Server are two different systems in the cell, the Cell Manager hostname must be listed on the Installation Server in the file `/etc/opt/omni/client/cell_server`.

-input *Filename*

> Specifies the input file (plain text file) containing the data for the client installation. Each client is described in the input file with a newline-separated ASCII string, using the format described below.

INPUT FILE FORMAT SYNOPSIS

-host *Hostname* -*Component Version* [-*Component Version* ...][-encryption*EncryptionFlag*] - push_inst *RemoteInstallationParameters*

INPUT FILE OPTIONS

`-host` *Hostname*

> Specifies the system to which remote installation will be performed. The *Hostname* must be enclosed in double quotes.

-Component Version

> Specifies the components for the installation. The *Version* argument specifies the version of the product. Specify only the components that are supported on the target Data Protector system. The available components are:

> `cc` – User Interface

> `da` – Disk Agent

> `ndmp` – NDMP Media Agent

> `ma` – General Media Agent

> `sap` – SAP R/3 Integration

> `sapdb` – SAP MaxDB Integration

> `saphana` – SAP HANA Appliance Integration

> `emc` – EMC Symmetrix Agent

> `oracle8` – Oracle Integration

> `sybase` – Sybase Integration

> `ssea` – HPE P9000 XP Agent

> `informix` – Informix Integration

> `lotus` – Lotus Integration

> `db2` – DB2 Integration

> `smisa` – HPE P6000 / HPE 3PAR SMI-S Agent

> `netapp` – NetApp Storage Provider

> `vmwaregre_agent` – VMware Granular Recovery Extension Agent

> `vepa` – Virtual Environment Integration

> `StoreOnceSoftware` – StoreOnce software deduplication

> `autodr` – Automatic Disaster Recovery

> `docs` – English Documentation (Guides, Help)

> `jpn_ls` – Japanese Documentation (Guides, Help)

> `fra_ls` – French Documentation (Guides, Help)

> `chs_ls` – Simplified Chinese Documentation (Guides, Help)

`-encryption`

This option is not mandatory. If not specified, the default behavior is to enable encrypted control communication on the client with re-use of certificates, if they are already present. (This is possible, if

encrypted control communication is enabled on the Cell Manager). Also, this flag is applicable only for new installations.

The flag can have the following values:

16 - Does not enable encrypted communication on the client.

40 - Enables encrypted communication on the client using newly-generated certificates.

-push_inst *RemoteInstallationParameters*

This option specifies all parameters that are crucial for a successful remote client installation. The option must be used with all its parameters.

Note: All arguments except *GeneralInstallationType* and *InstallationType* must be enclosed in double quotes (" ").

SITE SPECIFIC PATCH TEST MODULE (SSPTM) OPTIONS

-testmodule Version -tm_ssp_num SSPTMNAME

Specifies the components for SSPTM package installation. The *Version* argument specifies the version of the product. The *SSPTMNAME* argument specifies the SSPTM package name without any extension.

> **Note:** The SSPTM specific options takes precedence over the "-Component Version" options. Therefore, the "-Component Version" options are ignored when the SSPTM specific options are provided.

RemoteInstallationParameters

InstallPath

Specifies the main installation path for remote installation to Windows systems—the *Data_Protector_home* directory. The path must end with a backslash (\). For remote installation to UNIX systems, for which this argument is ignored, you can use a placeholder ("-").

This argument is currently ignored. You can use a placeholder ("-").

InstallDataPath

Specifies the additional installation path for remote installation to specific Windows systems—the *Data_Protector_program_data* directory. The path must end with a backslash (\). For remote installation to UNIX systems or Windows systems other than Windows Vista, Windows 7, and Windows Server 2008, for which this argument is ignored, you can use a placeholder ("-").

UserName

Specifies the user name that is used by the Installation Server for remote installation. If not specified, a default value is used: root for UNIX systems and Administrator for Windows systems. If not specified, the default value root is used. If you perform remote installation using secure shell, use a placeholder ("-").

Password

Specifies the password that is used by the Installation Server for remote installation. If not specified, the ob2install command prompts for it during the installation process. If you want ob2install to prompt for the password interactively or you perform remote installation using

secure shell, use a placeholder ("-").

CellManagerName

Specifies the name of the Cell Manager to whose cell the remote system will be added. To only install components on the remote system without adding it to a cell, use a placeholder ("-").

GeneralInstallationType

Specifies the general installation type:

1 – currently unused value (reserved for future extensions)

2 – client installation

InstallationType

Specifies the installation type:

1 – new installation

2 – update

3 – delete

4 – check installation

NOTES

The ob2install command is available on UNIX systems only.

EXAMPLES

The following examples illustrate how the ob2install command works.

1. To start a remote installation to the UNIX system "unixsys.company.com" using the Installation Server "issys.company.com" and import the client into the cell of the Cell Manager "cmsys.company.com", use the default remote installation user name, make ob2install prompt for the password interactively, where the input file is named "infile.txt" and the specified components are User Interface, Disk Agent, and General Media Agent, execute the following command:

   ```
   ob2install -server issys.company.com -input infile.txt
   ```

   The input file "infile.txt" must contain the following line:

   ```
   -host "unixsys.company.com" -cc A.09.00 -da A.09.00 -ma A.09.00 -push_inst "-"
   "-" "-" "cmsys.company.com" 2 1
   ```

   ```
   -host "unixsys.company.com" -cc A.06.20 -da A.06.20 -ma A.06.20 -push_inst "-"
   "-" "-" "-" "cmsys.company.com" 2 1
   ```

2. To start a remote installation to the Windows Server 2003 system "winsys.company.com" using the Installation Server "issys2.company.com" and import the client into the cell of the Cell Manager "cmsys.company.com", use the user name "Administrator" and the password "q1w2e3r4", where the input file is named "infile.txt" and the specified components are HPE P6000 / HPE 3PAR SMI-S Agent, Automatic Disaster Recovery, and French Documentation

(Guides, Help), execute the following command:

```
ob2install -server issys2.company.com -input infile.txt
```

The input file "infile.txt" must contain the following line:

```
-host "winsys.company.com" -smisa A.06.20 -autodr A.06.20 -fra_ls A.06.20 -
push_inst "-" "Administrator" "q1w2e3r4" "cmsys.company.com" 2 1
```

```
-host "winsys.company.com" -smisa A.06.20 -autodr A.06.20 -fra_ls A.06.20 -
push_inst "-" "-" "Administrator" "q1w2e3r4" "cmsys.company.com" 2 1
```

3. To start a remote installation to the Unix system "unixsys.company.com" using the Installation Server "issys.company.com" and import the client into the cell of the Cell Manager "cmsys.company.com", use the default remote installation user name, password as specified in the NOTE, where the input file is named "infile.txt" and the specified SSPTM package is "QXCMUX0001", execute the following command:

```
ob2install -server issys.company.com -input infile.txt
```

The input file "infile.txt" must contain the following line:

```
-host "unixsys.company.com" -testmodule A.09.00 -tm_ssp_num QXCMUX0001 -push_
inst "-" "-" "-" "-" "cmsys.company.com" 2 1
```

> **Note:** You need not specify the password while installing SSPTM modules as Data Protector is already installed on your system. However, in case of connectivity issues to the remote host, the ob2install prompt interactively requests for a new password (if required).

SEE ALSO

omnigui(5), omniintro(9), omnimigrate.pl(1M), omnisetup.sh(1M), omniusers(1), upgrade_cm_from_evaa(1M)

omnib2dinfo(1M)

omnib2dinfo — displays information about CatalystStore and StoreOnce Software stores
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

omnib2dinfo -help|-h

omnib2dinfo-version|-v

omnib2dinfo -store_info -b2ddevice *libraryname*

omnib2dinfo -list_stores -b2ddevice *libraryname*

omnib2dinfo -list_objects -b2ddevice *libraryname* [-metadata] [-tags] [-modified_
since="dd.mm.yyyy hh:mm:ss" in UTC]

omnib2dinfo -get_server_properties -b2ddevice *libraryname*

DESCRIPTION

This command displays information about an ObjectStore or StoreOnce Software store – store details,
list of stores and associated team members (for teamed stores), objects within the store, and details
about the store host.

OPTIONS

-version

   Displays the version of the `omnib2dinfo` command.

-help

   Displays the usage synopsis for the `omnib2dinfo` command.

-store_info

   Displays detailed information about the store specified in the device configuration, such as the store
   name, description, and status, as well as size of the stored data, the actual size of the store, the
   deduplication ratio, and backup and store size quota (if set). Unlike `-list_stores`, this option lists
   only the store associated with the device, not other stores residing on the same system.

-list_stores

   Lists all stores that reside on the same system as the store to which the specified device points to.
   For example, if store `str1` is configured for device `dev1` and store `str2` for device `dev2` and both
   reside on the same system, both stores are listed regardless if you specify device `dev1` or `dev2`.
   Store `str3` which is configured for `dev3` but resides on a different system is not listed.

Additional details are displayed for each store such as the store name, description, status (online, offline), whether encrypted (Yes, No), whether teamed (Yes, No), list of team members (for teamed stores) and their status (ON, OFF), the user data stored, the store data size, and the deduplication ratio.

`-list_objects`

Lists details about objects in the store specified in the device configuration, such as the object key, creation date, last modified date, and the size of the object on the disk.

In addition, it displays details such as:

- `metadata`: Includes common section, format version, the Catalyst library version, name of the backup product, version of the backup product, OS information (Gateway OS), product-specific section, host name of Cell Manager, gateway name, and the deduplication mode (server/target).

- `tags`: Includes specification list, session type, backup type, and session ID.

- `modified since`: Retrieves objects modified on or after the specified date, which is defined in UTC.

`-get_server_properties`

Displays the server properties such as the hostname, B2D version, the B2D serial number, disk size, and the free space on the disk.

`-b2ddevice` *DeviceName*

Specifies the B2D device for which the information is displayed.

EXAMPLES

1. To list all StoreOnce Software stores that reside on the same system as the store to which the device "StoreDev8" points to, execute:

   `omnib2dinfo -list_stores -b2ddevice StoreDev8`

2. To list only the StoreOnce Software stores for the device "StoreDev8", execute:

   `omnib2dinfo -store_info -b2ddevice StoreDev8`

3. To list details about all objects in the store for the device "StoreDev45", execute:

   `omnib2dinfo -list_objects -b2device StoreDev45`

SEE ALSO

omniamo(1), omnidownload(1), omnimcopy(1), omniminit(1), omnimlist(1), omnimm(1), omnimnt(1), omnimver(1), omniupload(1), sanconf(1M), uma(1M)

omnicheck(1M)

omnicheck — performs a DNS connections check within a Data Protector cell and lists Data Protector patches installed on Data Protector clients
(this command is available on systems with any Data Protector component installed)

SYNOPSIS

omnicheck -version | -help

omnicheck -dns [-host *Client* | -full] [-verbose]

omnicheck -patches -host *Client*

omnicheck -ssphf -host *Client*

DESCRIPTION

The following tasks can be performed using the omnicheck command:

CHECKING DNS CONNECTIONS WITHIN A Data Protector CELL

To check DNS connections within a Data Protector cell, use the -dns option with the omnicheck command.

The omnicheck command does not verify DNS connections in general. It verifies that DNS information matches over all communications relevant for Data Protector among Data Protector cell members. The command reports only failed checks and the total number of failed checks unless the -verbose option is specified.

It is possible to verify the following DNS connections in the Data Protector cell, using the omnicheck command:

- To check that the Cell Manager and every Media Agent resolve DNS connections to every Data Protector client in the same cell properly and the other way round, use the -dns option.
- To check that a particular Data Protector client resolves DNS connections to every Data Protector client in the same cell properly and the other way round, use the -host option.
- To check all possible DNS connections in the cell, when every client resolves DNS connections to all other clients in the same cell, use the -full option.

LISTING PATCHES INSTALLED ON Data Protector CLIENTS

The omnicheck command can be used to list Data Protector patches installed on a particular client. The omnicheck option used to list Data Protector patches installed on a particular client is -patches.

LISTING SITE SPECIFIC PATCHES OR HOT FIXES INSTALLED ON Data Protector CLIENTS

The omnicheck command can be used to list Data Protector Site Specific Patches (SSPs) or Hot Fixes (HFs) installed on a particular client. The omnicheck option used to list Data Protector SSPs or HFs installed on a particular client is -ssphf.

OPTIONS

`-version`

Displays the version of the `omnicheck` command.

`-help`

Displays the usage synopsis for the `omnicheck` command.

`-dns`

Checks that the Cell Manager and every Media Agent resolve DNS connections to every Data Protector client in the same cell properly and the other way round. This option performs the same as running the `omnicheck -dns -host` *CellManager* and `omnicheck -dns-host` *MediaAgent1...* `omnicheck -dns -host` *MediaAgentN* commands.

`-dns -host` *Client*

Checks that a Data Protector client specified by the `-host` option resolves DNS connections to every Data Protector client in the same cell properly and the other way round.

`-dns -full`

Checks all possible DNS connections in the cell. Every client in the cell tries to resolve all other clients in the same cell.

`-verbose`

Returns all the messages when using the `-dns` option. If this option is not set (default), only the messages that are the result of failed checks are returned.

`-patches -host` *Client*

Returns Data Protector patches (patch level, patch description and number of all patches installed) installed on a Data Protector client specified by the `-host` option. To use this option, you need the `Client configuration` user right (by default only users in the `admin` user group).

`-ssphf -host` *Client*

Returns Data Protector Site Specific Patches or Hot Fixes (Site Specific Patch (SSP) or Hot Fixes (HF) name, status, and number of SSPs or HFs installed) installed on a Data Protector client specified by the `-host` option. To use this option, you need the `Client configuration` user right (by default only users in the `admin` user group).

RETURN VALUES

See the man page `omniintro` for return values.

Additional return values of the `omnicheck` command used to check the DNS connections are:

client_1 `cannot connect to` *client_2*

client_1 `connects to` *client_2,* `but connected system presents itself as` *client_3*

client_1 `failed to connect to` *client_2*

`checking connection between` *client_1* `and` *client_2*

all checks completed successfully.

number_of_failed_checks checks failed.

client is not a member of the cell.

client contacted, but is apparently an older version. Hostname is not checked.

Additional return values of the `omnicheck` command used to list the Data Protector patches are:

List of patches found on host *client*

Patch level Patch description

Number of patches found: *number_of_patches*

List of patches on host *client* is not available.

Host *client* is not a member of this cell.

Host *client* is unreachable.

Additional return values of the `omnicheck` command used to list the Data Protector SSPs or HFs are:

List of SSPs/HFs found on host *client*

Site Specific Patch or Hot Fix status

Number of SSPs and HFs found: *number_of_patches*

List of SSPs/HFs on host *client* is not available.

Host *client* is not a member of this cell.

Host *client* is unreachable.

NOTES

The `omnicheck` command can be used only within one Data Protector cell.

EXAMPLES

1. To check DNS connections needed for normal Data Protector operating (the Cell Manager and every Media Agent in the cell resolve DNS connections to every Data Protector client in the cell properly and the other way round), execute:

   `omnicheck -dns`

2. To check if the client with the hostname `backup.system.com` resolves DNS connections to every Data Protector client in the same cell properly and the other way round, and to get all relevant messages, execute:

   `omnicheck -dns -host backup.system.com -verbose`

3. To list the patches installed on client with the hostname `backup.system.com`, execute:

   `omnicheck -patches -host backup.system.com`

4. To list the SSPs/HFs installed on client with the hostname `backup1.system.com`, execute:

   `omnicheck -ssphf -host backup1.system.com`

SEE ALSO

omnicc(1), omnicellinfo(1), omnidlc(1M), omnisv(1M)

omnicjutil(1M)

omnicjutil — starts, stops, and queries the Windows Change Journal on Windows clients
(this command is available on the Data Protector Cell Manager)

SYNOPSIS

omnicjutil -help

omnicjutil -version

omnicjutil -file *filename*

omnicjutil -host *hostname* -volume *vol* {-start [-maxsize *max* -delta *del*]|-stop [-wait]|-query}

DESCRIPTION

The omnicjutil command is used to remotely control and administer the Change Journal on Windows clients.

OPTIONS

-help

   Displays the usage synopsis of the omnicjutil command.

-version

   Displays the version for the omnicjutil command

-file *filename*

   Defines the file containing multiple single line entries of this command. Each line must conform to the usage of the omnicjutil command. Note that no tabs are allowed. If a syntax error is found, none of the commands is executed.

-host *hostname*

   Defines the name of the system hosting the Change Journal.

-volume *vol*

   Defines the volume name in the form /C or /C:*mounted_folder*.

-start [-maxsize *max* -delta *del*]

   Starts the Change Journal on the specified volume.

   The -maxsize *max* option sets the maximum size of the Change Journal in bytes. The highest possible value is 4 GB (4 294 967 296 bytes). Any specified value greater than 4 GB will be rounded down to 4 GB. Note that a reasonable size for a 100 GB drive is an 85 MB Change Journal.

The `-delta` *del* option specifies the size in bytes to be purged from the Change Journal when it reaches its maximum size. HPE recommends the value be approximately one-eighth to one-quarter the value of the maximum size but not greater than one quarter the size of the maximum size. This value may be automatically adjusted to better correspond to the volume cluster size.

`-stop`

Stops the Change Journal asynchronously.

The `-wait` specifies that the Change Journal will be stopped synchronously. The call returns only after the Change Journal has been deleted.

`-query`

Queries the status of the Change Journal.

NOTES

If the `-start` option is specified and the Change Journal is already active, the Change Journal is adjusted to the value of the maximum size and delta. Note that these values can only be adjusted to increase.

When starting the Change Journal, if you not specify `-maxsize` and `-delta`, or specify `0` for these parameters, the system chooses a default value based on the volume size.

The command line tool gets the input either directly from the command line or from a file. Using input directly from the command line allows only one operation at a time. To perform more than one operation, create a file using the `-file` *filename* option and use it as an input. Note that the commands in the file are executed from top to bottom.

As an alternative to the Data Protector `omnicjutil` command, you can also use the Windows `fsutil` command for administering the Change Journal.

EXAMPLES

To start the Change Journal with the maximum size of 8 MB (in bytes) and specify the size of 1 MB (in bytes) to be purged from the Change Journal when it reaches the specified maximum size, execute:

```
cjutil -start –maxsize 8388608 –delta 1048576
```

SEE ALSO

cjutil(1M)

omnidbcheck(1M)

omnidbcheck — checks the consistency of the Data Protector Internal Database (IDB)
(this command is available on the Data Protector Cell Manager)

SYNOPSIS

omnidbcheck -version|-help

omnidbcheck [-quick|-extended]

omnidbcheck -sibf [-detail|-dumpmedia] [-summary]

omnidbcheck -smbf [-detail|-dumpmessages] [-summary]

omnidbcheck -keystore [-summary]

omnidbcheck -verify_db_files [-detail]

omnidbcheck -connection [-detail]

omnidbcheck -media_consistency [-detail]

omnidbcheck -schema_consistency [-detail]

omnidbcheck -database_consistency [-detail]

omnidbcheck -bf [-summary]

omnidbcheck -dc [-quick] [-media <list>] [-detail [-detail]] [-summary]

DESCRIPTION

The Data Protector Internal Database (IDB) consists of Media Management Database (MMDB), Catalog Database (CDB), Detail Catalog Binary Files (DCBF), and Session Messages Binary Files (SMBF). The MMDB and CDB objects, object versions and media positions form the core part of the IDB. The DCBF and SMBF form the detail part of the IDB.

The omnidbcheck command checks the status of the IDB or parts of the IDB and sends a report to the standard output.

Note that errors found during the database connection check and encryption keystore check are Critical, errors found during the schema check, database check, media check, and datafiles check are Major, errors found during the Detail Catalog check or binary files check are Minor, and errors found during the SMBF check are Warning.

Data Protector creates a log file for each part of the check on the Cell Manager in the default server log files directory:

Check_smbf.txt

Check_bf.txt

Check_dc.txt

There is a timestamp at the beginning of each log file stating when the check was performed.

OPTIONS

`-version`

Displays the version of the `omnidbcheck` command.

`-help`

Displays the usage synopsis for the `omnidbcheck` command.

`-quick`

Checks the database connection, schema consistency, data files consistency, presence and size of the DCBF part of the IDB, and displays the summary of the check.

`-extended`

Checks the entire IDB with the exception of the SMBF and displays the summary of the check.

`-sibf`

This option relates to the functionality that is no longer supported in the installed HPE Data Protector version.

If the `-detail` option is specified, it lists all SIBF and their status (`OK` or `corrupted/missing`). If the `-detail` option is not specified (default), only the corrupted SIBF and their status (corrupted or missing) are listed.

If the `-dumpmedia` option is specified with the `-sibf` option, it sends the SIBF filenames, object versions information, offset of the data in the SIBF file belonging to an object version as well as size of the data in the SIBF file belonging to an object version to the standard output.

If the `-summary` option is specified, the command sums up the data and displays the status of the SIBF.

`-summary`

Displays only the summary of the check (OK or failed/missing). The option does not impact the thoroughness of the check except for `omnidbcheck -dc`, where `-summary` implies a `-dc -quick` check.

`-smbf`

Checks the presence of the SMBF.

Note that if you have removed a SMBF in any way (for example, using Data Protector GUI or CLI or deleted the file manually), then this option reports the removed session message as missing. This does not mean that IDB is corrupted—it only indicates that a session has been removed.

If the `-detail` option is specified, it lists all SMBF and their status (OK or corrupted/missing). If the `-detail` option is not specified (default), only the SMBF and their status (corrupted or missing) are listed.

If the `-dumpmessages` option is specified with the `-smbf` option, it sends the session messages in the SMBF to the standard output.

If the `-summary` option is specified, the command sums up the data and displays the SMBF.

-keystore

>Performs a consistency check of the Data Protector's `keymap` index file and the encryption keys in the keystore. The following information is listed for each encryption key in the cell: key ID, store ID, KeyStore name, KeyFile name, and a result of the check (OK or corrupted).

>If the `-summary` option is specified, the command sums up the data and displays the status of the keystore.

-verify_db_files

>Checks for the existence of database datafiles.

>If the `-detail` option is specified, the command lists missing datafiles in case of a datafiles consistency failure.

-connection

>Checks the status of Data Protector database connectivity.

>If the `-detail` option is specified, the command lists errors in case of a connection failure.

-media_consistency

>Checks the consistency of media.

>If the `-detail` option is specified, the command lists inconsistent media names in case of a media consistency failure.

-schema_consistency

>Checks the consistency of schema and detects changes in the schema since its first creation during the Data Protector installation.

>If the `-detail` option is specified, the command lists the difference between the original and current schema in case of a schema consistency failure.

-database_consistency

>Checks the database consistency.

>If the `-force` option is specified, it overrides the default safety check. With this option, there is no confirmation request for the check of database consistency. If the `-force` option is not specified, the command displays a confirmation request for the check of database consistency.

>If the `-detail` option is specified, the command lists errors in case of a database consistency failure.

-bf

>Performs a presence and size check of the DCBF.

>If the `-summary` option is specified, the command sums up the data and displays the binary files.

-dc

>Checks consistency between the core and DC part of the database.

>If `-detail` option is specified once, the check lists all encountered errors. Without `-detail` option, it only displays the summarized error information. With `-detail` option specified twice, the check lists statuses of all checked items regardless of error (maximum verbosity).

If `-quick` option is specified, the check only checks the most recently completed segments for each DC binary file.

If `-media <list>` is specified, the check only checks DCBFs corresponding to the media matching the specified IDs in the list.

If `-summary` mode is specified, only summary is written to the console in `-extended` format, while the full output (depending on the specified options) is still written to the `Check_dc.txt`. In `-summary` mode, `-quick` option is implied.

EXAMPLES

1. To perform an extended check of the IDB, execute:

   `omnidbcheck -extended`

2. To perform a consistency check of the Data Protector's keymap index file and the encryption keys in the keystore, execute:

   `omnidbcheck -keystore`

3. To list all missing datafiles in case of a consistency failure, execute:

   `omnidbcheck -verify_db_files -detail`

4. To list all media with inconsistent names in case of a media consistency failure, execute:

   `omnidbcheck -media_consistency -detail`

SEE ALSO

omnidb(1), omnidbinit(1M), omnidbp4000(1), omnidbsmis(1), omnidbutil(1M), omnidbvss(1), omnidbxp(1), omnidbzdb(1), omniofflr(1M)

omnidbinit(1M)

omnidbinit — initializes the Data Protector Internal Database (IDB)
(this command is available on the Data Protector Cell Manager)

SYNOPSIS

`omnidbinit -version|-help`

`omnidbinit [-force]`

DESCRIPTION

The `omnidbinit` command initializes the Data Protector Internal Database (IDB). All information about sessions, media, and objects is lost after the initialization. The command does not delete the IDB archived log files but creates a gap in the sequence of them; when a rollforward operation is performed using the `omniofflr` command, the operation applies only the archived log files logs created before the initialization of the IDB.

In order to initialize the IDB successfully, the underlying structure of the embedded database has to exist at the Data Protector Internal Database location in the `pg` directory. Note that the actual location of the `pg` directory may not be the default one, if you have restored the IDB to some other location and registered the restored instance as the new IDB.

Make sure that the Data Protector Internal Database Service (`hpdp-idb`) is running and the connection to the current IDB is available. Verify it with the `omnisv -status` command. Any error found during `omnidbinit` execution is reported.

OPTIONS

`-version`

Displays the version of the `omnidbinit` command

`-help`

Displays the usage synopsis for the `omnidbinit` command

`-force`

Overrides the default safety check for the initialization. By default, the command displays a confirmation request. With this option, there is no confirmation request.

SEE ALSO

omnidb(1), omnidbcheck(1M), omnidbp4000(1), omnidbsmis(1), omnidbutil(1M), omnidbvss(1), omnidbxp(1), omnidbzdb(1), omniofflr(1M)

omnidbutil(1M)

omnidbutil — handles various Data Protector Internal Database (IDB) maintenance tasks
(this command is available on the Data Protector Cell Manager)

SYNOPSIS

omnidbutil -version | -help

omnidbutil -readdb *Directory*

omnidbutil -writedb *Directory*

omnidbutil -show_locked_devs [-all]

omnidbutil -free_locked_devs [-all | *DevName* | *MediumID* | *CartName PhyLocation* | *Serial_LDEV* | *WWW_LUN*]

omnidbutil -changebdev *FromDev ToDev* [-session *SessionID*]

omnidbutil -purge {-sessions [*NumberOfDays*] | -days [*NumberOfDays*] | -messages [*NumberOfDays*] | -daily | -dcbf}

omnidbutil -purge_failed_copies

omnidbutil -clear

omnidbutil -change_cell_name [*OldHost*]

omnidbutil -show_cell_name

omnidbutil -set_session_counter *NewSessionID*

omnidbutil -show_db_files

omnidbutil -free_pool_update

omnidbutil -list_large_mpos *MinNumberOfMpos* [-top *NumOfTopMedia*] [-detail] [-csv *CSVFile*]

omnidbutil -free_cell_resources

omnidbutil -mergemmdb *CellManagerHostname*

omnidbutil -cdbsync *CellManagerHostname*

omnidbutil -info

omnidbutil -autovacuum {-set -table *TableName* [-to_default] [-on_n_rows *NRows*] [-on_percentage *Percent*] [-freeze_max_age *FreezeMax*] | -get [-table *TableName* | -enabled | -disabled | -all]}

omnidbutil -list_dcdirs

omnidbutil -add_dcdir *PathName* [-maxsize *MaxSizeInMB*] [-maxfiles *NumberOfFiles*] [-spacelow *SpaceLowInMB*] [-seq *SeqNumber*]

omnidbutil -modify_dcdir *PathName* [-maxsize *MaxSizeInMB*] [-maxfiles *NumberOfFiles*] [-spacelow *SpaceLowInMB*] [-seq *SeqNumber*]

```
omnidbutil -remove_dcdir PathName

omnidbutil -remap_dcdir

omnidbutil -fixmpos

omnidbutil -cp {-set ParamName ParamValue|-get [-param ParamName]}

omnidbutil -set_passwd UserName

omnidbutil -set_passwd java -pass PreferredPassword

omnidbutil -sync_srv
```

DESCRIPTION

The `omnidbutil` command is used for Data Protector Internal Database (IDB) maintenance tasks. These tasks involve:

OPERATIONS ON DETAIL CATALOG BINARY FILES (DCBF)

The Detail Catalog part of the IDB is composed of two parts: 1) The Detail Catalog (DC) binary files, which stores pathnames of the backed up files and directories, together with client system names, and version information (size, modification time, attributes/protection, exact position on a medium (segment and block offset within a segment) of a backed up file or directory, and so on). 2) DC directories: registered directories that contain DC binary files. A DC directory is allocated when creating a new DC binary file using one of three possible allocation algorithms, specified by the `DCDirAllocation` global option.

Operations on DCBF include: 1) Registering, removing, and updating DC directories. 2) Locating DCBF across DC directories if they had been manually moved. 3) Removing invalid references to DC binary files. Invalid references may occur after an IDB recovery during which the replay of the archived logs is executed. In that case, CDB is newer than DCBF.

The `omnidbutil` options used for operations on DC are: `-list_dcdirs`, `-add_dcdir`, `-modify_dcdir`, `-remove_dcdir`, `-remap_dcdir`, and `-fixmpos`.

EXPORTING AND RE-CREATING THE CONTENTS OF THE MEDIA MANAGEMENT DATABASE (MMDB) AND CATALOG DATABASE (CDB)

The contents of MMDB and CDB can be exported and imported back. Data Protector uses the PostgreSQL `pg_dump` command to create files in the UTF-8 format for this purpose.

The `omnidbutil` options used for exporting and recreating the contents of MMDB and CDB are: `-readdb` and `-writedb`.

LISTING AND UNLOCKING BACKUP DEVICES, TARGET VOLUMES, MEDIA, AND LIBRARY SLOTS

Backup devices, target volumes, backup media, and library slots in use are locked during Data Protector sessions. In certain situations (backup or restore sessions end abnormally) devices remain locked, even though the Data Protector Media Agent or a Data Protector disk array integration agent is no longer running. By default, such devices are automatically unlocked after 60 minutes.

The `omnidbutil` options used for listing and unlocking backup devices, target volumes, backup media, and library slots are: `-show_locked_devs` and `-free_locked_devs`.

MERGING LOCAL MMDBS INTO A CENTRALIZED MMDB (CMMDB)

In large multicell environments with high-end backup devices, you may want to share these devices and media among several cells. This can be achieved by having one centralized MMDB (CMMDB) for all the cells and keeping an individual CDB for each cell. This allows backup media and backup device sharing while preserving the security capabilities of the multicell structure. To achieve this, merge the local MMDBs into the CMMDB.

The `omnidbutil` option used for merging MMDB into CMMDB is `-mergemmdb`.

SYNCHRONIZING CDB AND MMDB

In certain situations, the CDB and MMDB may be out of sync (the CDB and MMDB were imported from files generated in separate export sessions, the CMMDB was restored while leaving local CDB intact, and so on). In such cases, synchronize both databases.

The `omnidbutil` option used for synchronizing CDB and MMDB is `-cdbsync`.

MISCELLANEOUS TASKS

These tasks involve operations such as displaying the information about the IDB and the IDB upgrade process, updating backup device references in object versions, changing ownership of the CDB to the current Cell Manager, displaying the actual CDB ownership, reclaiming free disk space, setting parameters for the connection pool configuration file, changing the password for a configured user account, changing the password for Data Protector Web reporting, and more.

The `omnidbutil` options used for this group of tasks are: `-changebdev`, `-purge`, `-clear`, `-change_cell_name`, `-show_cell_name`, `-set_session_counter`, `-show_db_files`, `-free_pool_update`, `-list_large_mpos`, `-top`, `-csv`, `-free_cell_resources`, `-info`, `-autovacuum`, `-cp`, `-set_pasw`, `-set_passwd java -pass`, and `-sync_srv`.

The `-purge_failed_copies` option needs exclusive access to the IDB. Before you use it, make sure that no backup, restore, or media management sessions are in progress and that no graphical user interfaces are launched in the cell.

OPTIONS

`-version`

Displays the version of the `omnidbutil` command.

`-help`

Displays the usage synopsis for the `omnidbutil` command.

`-readdb` *Directory*

Reads and restores the IDB data from the `Directory` previously written using the `omnidbutil -writedb` command. Note that the `-readdb` command does not restore DCBFs or SMBFs but leaves the old files in place. You may need to backup and restore these manually using the paths listed by the `-writedb` command.

`-writedb` *Directory*

Writes the IDB data (without the DCBF and SMBF directories) to the specified `Directory`. The command lists all the DCBF and SMBF directory paths that need to be manually backed up, if needed for a restore later on. Note that the commands `-writedb` and `-readdb` are not a substitute for the IDB backup.

-show_locked_devs [-all]

Lists all locked devices, target volumes, media, and slots in the Data Protector cell.

The -all option applies only when you execute the command on the MoM system, in which case locked devices, target volumes, media, and slots from all cells are listed.

-free_locked_devs [-all | *DevName* | *MediumID* | *CartName PhyLocation* | *Serial_LDEV* | *WWN_LUN*]

Unlocks a specified device, target volume, medium, or slot, where *DevName* is the device, *MediumID* is the medium, *CartName* is the library name, *PhyLocation* is the number of the slot to be unlocked, *Serial_LDEV* is the target volume where *Serial* is the serial number of a disk array of the HPE P9000 XP Disk Array Family and *LDEV* is the HPE P9000 XP Disk Array Family volume number, *WWN_LUN* is the target volume where *WWN* is the world-wide-name of a disk array of the HPE P6000 EVA Disk Array Family and *LUN* is the logical unit number (LUN). If none of the above is specified, all devices, target volumes, media, and slots in the Data Protector cell are unlocked.

The -all option applies only when you execute the command on the MoM system, in which case all devices, target volumes, media, and slots from all cells are unlocked.

-changebdev *FromDev ToDev* [-session *SessionID*]

Changes all references in object versions from device *FromDev* to device *ToDev*. You can change the device name only for a single session by using the -session option.

-purge {-sessions [*NumberOfDays*] | -days [*NumberOfDays*] | -messages [*NumberOfDays*] | -daily | -dcbf}

This option allows you to remove obsolete backup, restore, and media management sessions, session messages, and obsolete DC binary files from the IDB.

The -sessions option removes media management sessions, restore sessions, and obsolete backup sessions (backup sessions without backed up data) older than *NumberOfDays*.

The -days option removes media management sessions, restore sessions, obsolete backup sessions (backup sessions without backed up data), and session messages for all sessions older than *NumberOfDays*.

The -messages option removes session messages for all sessions older than *NumberOfDays*.

The -daily option starts the same purge session as started every day at 12:00 (depending on the Data Protector global option setting) and is a part of Data Protector daily maintenance tasks. This purge session deletes DCBF based on the catalog protection and removes obsolete sessions and their messages, by running the `omnidbutil -purge -sessions` *KeepObsoleteSessions* `-messages` *KeepMessages* `-dcbf` command, where *KeepObsoleteSessions* and *KeepMessages* are specified in the Data Protector global options. Default values for these two parameters are `30` and `0`, respectively.

The scheduled time at which the -daily option is started every day is defined by the `DailyMaintenanceTime` global option.

The -dcbf option removes the DC binary file of each media with expired catalog protection.

Specify at least one of these options. You can change or disable the global option `DailyMaintenanceTime` for the -daily option.

-purge_failed_copies

In certain circumstances the Data Protector IDB may hold multiple copies of objects made during a

backup. Use this option to remove all unrequired copies that may overload an IDB. This option requires exclusive access to the database.

-clear

Sets the status of all sessions that are actually not running but are marked *In Progress/Failed*, to *Failed*. It requires exclusive database access to ensure that no session is running.

-change_cell_name [*OldHost*]

This option changes the owner of the CDB to the current Cell Manager. It also changes all references in the CMMDB from *OldHost* to the current Cell Manager. It modifies all media entries within the MMDB or CMMDB associated with the original Cell Manager (old host).

If the *OldHost* parameter is not specified, omnidbutil determines the previous owner of the CDB (old host) from the database itself.

If you want to associate all media in a CMMDB with the current Cell Manager, it is necessary to execute the command once for each Cell Manager that has media associated with it, using the *OldHost* parameter.

Specify the *OldHost* parameter exactly the same as the owner of the media. If the system's Fully Qualified Domain Name (FQDN) is associated with the media, also use the FQDN with this command. If the *OldHost* parameter is not specified correctly, the operation will not be performed.

This command is used after moving databases from one Cell Manager to another or after using -readdb on files that were created on another Cell Manager.

-show_cell_name

Queries the CDB for its owner. If there is no information available, use the -change_cell_name option to update the information.

-set_session_counter *NewSessionID*

Sets a new value for the counter that is used for generating the sessionID. This option is used after the restore and recovery of the IDB to enable the import of tapes that were created on the same day. Suggested value is 100.

-show_db_files

Lists all directories and files that are backed up during an IDB backup, such as the IDB datafiles, IDB write-ahead logs, DCBF and SMBF . In effect they contain all components of IDB.

-free_pool_update

Finds any free (unprotected) media in pools with the free pool and move free media to free pool options set and by default deallocates the found free media to a free pool every day at 00:00.

-list_large_mpos *MinNumberOfMpos* [-top *NumOfTopMedia*] [-detail] [-csv *CSVFile*]

Lists top *NumOfTopMedia* media that has more than *MinNumberOfMpos* media positions. By default, positions used and medium are displayed. With the –detail option, additional fields are displayed: the total object versions, the data protected object versions, the catalog protected object versions, and the last-write time for medium. Every report is logged to the list_large_media.log file. Optionally, the report can be written to a comma separated values (CSV) file specified with the –csv option.

-free_cell_resources

Frees all resources that were allocated during backup and restore sessions. The option is used if a

session ends abnormally or a process is terminated unexpectedly.

-mergemmdb *CellManagerHostname*

Merges the local MMDB from the remote Cell Manager *CellManagerHostname* to the CMMDB. A MoM cell and a remote cell with a local MMDBF must exist for this action. All duplicated items (stores, media pools, devices) will have "_N" appended to their name, where N represents the number of the duplicate (starting with 1). Once the database is merged you cannot revert the operation. The merge operation preserves the local MMDB, which is no longer in use but must remain stored on the local system for the local IDB backup sessions to succeed.

-cdbsync *CellManagerHostname*

Synchronizes the centralized MMDB (CMMDB) and local CDB on the specified Cell Manager. In a MoM environment, the MMDB and CDB may be out of sync as a result of the centralized IDB restore.

Execute the command on the system where the CMMDB is installed.

If the CMMDB was changed, execute the command for each Cell Manager in this MoM cell that you want to use the central media managements by specifying each Cell Manager in the cell as the *CellManagerHostname* argument.

-info

Displays information about the IDB, such as MMDB, CDB, archived log files, datafiles, diskspace, DCBF, SMBF, and SIMBF usage.

-autovacuum {-set -table *TableName* [-to_default] [-on_n_rows *NRows*] [-on_percentage *Percent*] [-freeze_max_age *FreezeMax*] | -get [-table *TableName* | -enabled | -disabled | -all]}

Recovers or reuses disk space occupied by updated or deleted rows. The vacuum operation (periodic maintenance) is done automatically. By default, the option is enabled.

The -table option specifies the table to set or displays its autovacuum properties. You can provide the exact table name or use the asterisk (*) at the end to define a group of tables (for example, -table *"dp_catalog_*"*). Make sure that you quote the string if you use an asterisk. If a table was not yet customized with the -autovacuum option, Default is displayed as the value for all table properties.

If the -to_default option is specified together with -set and -table, the command resets autovacuum parameters for the selected table to default values.

If the -on_n_rows option is specified together with -set and -table, the command initiates the vacuum operation when the specified number of table rows is updated or deleted. By default, the value is set to 50.

If the -on_percentage option is specified together with -set and -table, the command initiates the vacuum operation when the specified percentage of table rows is updated or deleted. By default, the value is set to 20.

If the -freeze_max_age option is specified together with -set and -table, the command specifies the maximum age (in number of transactions) before autovacuum operation is forcibly invoked on the table. This happens even if autovacuum operation is disabled. Valid values for the option are between 100 million and 2 billion, by default it is set to 200000000 transactions.

If the -get option is specified, the command lists all table names:

-enabled lists all tables that are included in the autovacuum operation

-disabled lists all tables that are excluded from the autovacuum operation

-all lists all tables and their properties

-list_dcdirs

Lists all registered DC directories.

-add_dcdir *PathName* [-maxsize *MaxSizeInMB*] [-maxfiles *NumberOfFiles*] [-spacelow *SpaceLowInMB*] [-seq *SeqNumber*]

Creates a new directory at the specified path and registers it in the IDB as a new DC directory.

The -maxsize option specifies the maximum amount of disk space that can be used for DC binary files in this directory. When the specified size is reached, Data Protector stops creating new DC binary files here, and starts using the next DC directory defined by the effective allocation policy.

The -maxfiles option specifies the maximum number of DC binary files that can reside in the directory. When the specified number is reached, Data Protector stops creating new DC binary files here, and starts using the next DC directory defined by the effective allocation policy.

The -spacelow option defines the conditions under which the DC directory is considered to be full. It actually defines the minimum allowed difference between the actual size and the configured maximum size of the DC directory. When this threshold is reached, Data Protector starts using the next DC directory defined by the effective allocation policy. Additionally, this option defines the minimum amount of free space needed on the volume where the DC directory resides. Data Protector requires this space to log names of the backed up files and directories to the IDB. When free space for the last writable DC directory drops under this amount (meaning all other DC directories are considered to be full already), Data Protector automatically switches to the logging level No Log. HPE recommends to use 10% to 15% of the currently configured maximum DC directory size as a suitable value for this option.

If argument to any of the options -maxsize, -maxfiles, and -spacelow is omitted, the default value is used for the respective amount. For default values, see the *HPE Data Protector Product Announcements, Software Notes, and References*, chapter *Limitations and recommendations*, section *Internal Database scalability*.

The -seq option specifies the consecutive number that defines the order in which Data Protector chooses this DC directory to write new data to, provided that the effective DC directory allocation policy is fill in sequence (the DCDirAllocation global option is set to 0). The first DC directory to be used has the lowest allocation sequence number. If argument to this option is omitted, the value 0 is used.

-modify_dcdir *PathName* [-maxsize *MaxSizeInMB*]

[-maxfiles *NumberOfFiles*] [-spacelow *SpaceLowInMB*] [-seq *SeqNumber*]

Modifies properties of a DC directory that is registered with the specified path. The path itself cannot be changed.

The -maxsize option modifies the maximum amount of disk space that can be used for DC binary files in this directory. When the specified size is reached, Data Protector stops creating new DC binary files here, and starts using the next DC directory defined by the effective allocation policy. When you increase the maximum size for a specific DC directory, you should also adjust its minimum free disk space property by using the -spacelow option.

The `-maxfiles` option modifies the maximum number of DC binary files that can reside in the directory. When the specified number is reached, Data Protector stops creating new DC binary files here, and starts using the next DC directory defined by the effective allocation policy.

The `-spacelow` option modifies the conditions under which the DC directory is considered to be full. It actually defines the minimum allowed difference between the actual size and the configured maximum size of the DC directory. When this threshold is reached, Data Protector starts using the next DC directory defined by the effective allocation policy. Additionally, this option defines the minimum amount of free space needed on the volume where the DC directory resides. Data Protector requires this space to log names of the backed up files and directories to the IDB. When free space for the last writable DC directory drops under this amount (meaning all other DC directories are considered to be full already), Data Protector automatically switches to the logging level `No Log`. HPE recommends to use 10% to 15% of the currently configured maximum DC directory size as a suitable value for this option.

If argument to any of the options `-maxsize`, `-maxfiles`, and `-spacelow` is omitted, the respective default value is used. For default values, see the *HPE Data Protector Product Announcements, Software Notes, and References*, chapter *Limitations and recommendations*, section *Internal Database scalability*.

The `-seq` option modifies the consecutive number that defines the order in which Data Protector chooses this DC directory to write new data to, provided that the effective DC directory allocation policy is fill in sequence (the `DCDirAllocation` global option is set to `0`). The first DC directory to be used should have the lowest allocation sequence number. If argument to this option is omitted, the value `0` is used.

-remove_dcdir *PathName*

Withdraws registration of the specified DC directory in the IDB without removing the directory itself. The directory must not contain DC binary files in order to become unregistered.

-remap_dcdir

Locates DCBF across all DC directories and updates DCBF locations in the IDB if they had been moved manually (using the `mv` command or similar) between DC directories. This makes the IDB aware of the locations of each DCBF. This option requires exclusive access to the database.

-fixmpos

Removes invalid references to DCBF. This option should be used in the case of IDB recovery (after the dbreplay phase of the IDB restore process or `-import_logs`) or after a DCBF has been manually removed. This option requires exclusive access to the database.

-cp {-set *ParamName* *ParamValue* | -get [-param *ParamName*]}

Lists and sets parameters for the connection pool configuration file. The `idbhpdp-idb-cp.cfg` file is located in the Data Protector server configuration directory.

Certain parameters are predefined during the Data Protector installation and cannot be changed, such as `hpdpidb`, `service_name`, `auth_type`, `auth_file`, `admin_users`, and `stats_users`.

-set_passwd *UserName*

Changes the password for the configured Internal Database Service and Application Server user account.

-set_passwd java -pass *PreferredPassword*

Change the password for configuring the Data Protector Web reporting.

-sync_srv

> Synchronizes the location of the IDB data files for all nodes in a cluster environment. Execute it on the active cluster node.

> Use this option only when restoring the IDB to a different location in an HPE Serviceguard environment on UNIX systems.

EXAMPLES

The following examples illustrate how the omnidbutil command works.

1. To create a new DC directory in the "/var/opt/test" directory with maximum size 1000 MB, execute:

   omnidbutil -add_dcdir /var/opt/test -maxsize 1000

2. To list all locked devices, target volumes, media, and slots, execute:

   omnidbutil -show_locked_devs

3. To unlock a device, a medium, or library slot, respectively, execute:

   omnidbutil -free_locked_devs machine

   omnidbutil -free_locked_devs 0a1106452:5a45add9:2548:0007

   omnidbutil -free_locked_devs libraryName phyLocation

4. To unlock the target volume whose volume number is "288" and which resides on the HPE P9000 XP Disk Array Family storage system with the serial number "30658", execute:

   omnidbutil -free_locked_devs 30658_288

5. To manually change the maximum size for DC directory "dcbf13" in the "C:\Program Files\OmniBack\db46" directory to 48 GB and modify the free disk space needed for a DC binary file, execute:

   omnidbutil -modify_dcdir C:\Program Files\OmniBack\db46\dcbf13 -maxsize 49152 -
   spacelow 7372

6. To manually remove expired sessions and session messages older than 30 days, obsoleted data from the DCBF part of the IDB, and all the object versions for the overwritten tapes if the daily maintenance is disabled, respectively, execute:

   omnidbutil -purge –sessions 30

   omnidbutil -purge –messages 30

   omnidbutil -purge –dcbf

7. To remove all unrequired copies of objects that were made during a backup and may overload the IDB, execute:

   omnidbutil –purge_failed_copies

8. To export the IDB schema and its data into the dpidb.dat file and store it in the directory named "C:\dump_location", execute:

   omnidbutil –writedb C:\dump_location

9. To initiate a cleanup (vacuuming) of the table named "dp_catalog13" when 30% of table rows are updated or deleted, execute:

```
omnidbutil –autovacuum -table dp_catalog13 -set -on-percentage 30
```

10. To display information about autovacuum properties for all catalog tables named "dp_catalog*", execute:

```
omnidbutil –autovacuum -get -table "dp_catalog*"
```

11. To specify the maximum age in number of transactions (for example, 5000000000) before autovacuum operation is forcibly invoked on the specified table, execute:

```
omnidbutil –autovacuum -set -table dp_catalog_object_type -freeze_max_age
5000000000
```

12. To get the connection pool configuration, execute:

```
omnidbutil –cp -get
```

13. To set the connection pool parameter "max_client_conn" to "200", execute:

```
omnidbutil –cp -set -max_client_conn 200
```

14. To set a new password for the user named "hpdpidb_app", execute:

```
omnidbutil –set_passwd hpdpidb_app
```

15. To set a password for the Web reporting named "Pa55word", execute:

```
omnidbutil -set_passwd java –pass Pa55word
```

SEE ALSO

omnidb(1), omnidbcheck(1M), omnidbinit(1M), omnidbp4000(1), omnidbsmis(1), omnidbvss(1), omnidbxp(1), omnidbzdb(1), omniofflr(1M)

omnidlc(1M)

omnidlc — gathers or deletes Data Protector debug, log, and getinfo files from the Data Protector cell or from a MoM environment
(this command is available on the Data Protector Cell Manager)

SYNOPSIS

omnidlc -version | -help

omnidlc {-session *sessionID* | -did *debugID* | -postfix *string* | -no_filter} [-hosts *list*] [-pack *filename* | -depot [*directory*] | -space | -delete_dbg | -telemetry_files] [-no_logs] [-no_getinfo] [-no_compress] [-no_config] [-no_debugs | -debug_loc *dir1* [*dir2*...]] [-no_verbose] [-add_info [-any | *host*] *path*]

omnidlc -localpack [*filename*]

omnidlc -unpack [*filename*]

omnidlc -uncompress *filename*

omnidlc [-hosts *list*] -del_ctracelog

omnidlc [-module] -module 1, module 2

DESCRIPTION

The omnidlc command collects Data Protector debug, log, and getinfo files from the Data Protector cell (by default, from every client).

The Data Protector debug files are created during a Data Protector debug session. By default, the command collects debug files from the Data Protector default debug files directory. To collect debugs also from other directories, use the -debug_loc option.

Using the command, it is possible to collect Data Protector debug, log and getinfo files from selected clients in the Data Protector cell. In a MoM environment, you can only collect data for each Data Protector cell separately by running the command from the respective Cell Manager. On OpenVMS systems, getinfo files are not collected because the get_info utility is not available.

You can specify multiple filters for gathering the debug logs, for instance, Session ID, Postfix and Debug ID. A sample command may be –session <SessID> –postifx <string>. You can also specify multiple modules, for which you need the debug files. A sample command may be omnidlc –session <SessID> -module BSM,VBDA,DBSM .

Additionally, the Data Protector debug files to be collected can be limited to debugs that were generated within the specified Data Protector session or to debugs identified by a debugID or by a debug filename (debug postfix).

Typically, the Session Manager debugs and other non-session related executables do not have a sessionID along with the debug name. The debug collector also collects debugs that do not have a sessionID. While investigating the defects the IDB, MMD and other non-session related debugs are

also required. If you require debugs that belong only to specific sessions, then ensure the target debug folder is empty, before you run the session.

By default, every collected debug, log and getinfo file is then compressed and sent over the network to the Cell Manager. The final extension `.gz` is added on the Cell Manager, where all collected files with the `.gz` extension are, by default (if the `-depot` option is not specified), packed and saved in the current directory as the `dlc.pck` file. The file includes a generated directory structure that includes the hostnames, paths and the (compressed) collected files of the clients involved. This directory structure is described further on in this man page.

Optionally, files can be sent over the network to the Cell Manager uncompressed (if the `-no_compress` option is specified). Besides that (if the `-depot` option is specified), the transferred files can be left unpacked in the specified directory on the Cell Manager, in which the directory structure that includes the hostnames, paths and the collected files of the clients involved is generated as follows:

UNIX systems:

./dlc/*system_1*/tmp/*debug_files*

./dlc/*system_1*/log/*log_files*

./dlc/*system_1*/getinfo/get_info.txt

./dlc/*system_2*/tmp/*debug_files*

./dlc/*system_2*/log/*log_files*

./dlc/*system_2*/getinfo/get_info.txt

...

Windows systems:

.\dlc*system_1*\tmp*debug_files*

.\dlc*system_1*\log*log_files*

.\dlc*system_1*\getinfo\get_info.txt

.\dlc*system_2*\tmp*debug_files*

.\dlc*system_2*\log*log_files*

.\dlc*system_2*\getinfo\get_info.txt

...

If the file to be sent over the network is larger than 2 GB, the file is split in 2 GB chunks before it is compressed (it can be left uncompressed) and sent to the Cell Manager. Every chunk retains the file name and is added the first extension ranging from `s001` to `s999`. The second extension (`.gz`) is not added if the files are not compressed. Additionally, on the Cell Manager side, if the size of all collected compressed or uncompressed files exceeds 2 GB, the collected files are packed in 2 GB sized (original size) packages and added an extension ranging from `s001` to `s999`.

The collected debug files can also be deleted (if the `-delete_dbg` option is specified), or the disk space required on the Cell Manager for the collected files can be displayed (if the `-space` option is specified).

In these two cases, the selected files are neither transferred from the clients to the Cell Manager nor packed on the Cell Manager.

When collecting or deleting files or when displaying the required disk space, additional criteria can be defined to limit the files selection. Thus, it is possible to exclude the getinfo file, the log files, the debug files or any combination of the three groups of files from the selection.

Using the command, the collected files can then be additionally packed to be sent to the support center. The command provides also a means of unpacking the packed collected files.

OPTIONS

-version

Displays the version of the omnidlc command.

-help

Displays the usage synopsis for the omnidlc command.

-session *sessionID*

Limits the collected debug files to those that were produced during the Data Protector session identified by the *sessionID*. Note that on OpenVMS, the omnidlc command run with the -session parameter does not collect the debug files produced during specified session, because session names are not part of the OpenVMS debug filename. Instead, all available logs are collected.

-did *debugID*

Limits the collected debug files to those identified by the *debugID*.

-postfix *string*

Limits the collected debug files to the specified debug postfix.

-no_filter

Does not limit (select) the collected debug files.

-module

Allows you to specify multiple modules, for which you need the debug files. You can use comma (,) to separate the modules specified.

-hosts *list*

Limits the files to be collected to the clients specified in the *list*. The hostnames must be separated by spaces. The debug files collected are still subject to -session, -did or -postfix options.

-pack *filename*

All collected files are, by default (if this option is not specified), packed and saved in the current directory as the dlc.pck file. If this option is specified, the collected files are packed and saved in the specified file in the current directory on the Cell Manager. If the full path name is specified, the files are packed and saved in the specified file in the specified directory.

To add files other than the collected files to the package, copy the files to one of the following directories before running the command: dlc/*client*/getinfo, dlc/*client*/log, or

dlc/*client*/tmp (on UNIX), or .\dlc*client*\getinfo, .\dlc*client*\log, or .\dlc*client*\tmp (on Windows). You cannot add directories, but only files. If the files are not copied to one of the specified directories, the package cannot be unpacked during the unpack phase.

-depot [*Directory*]

If the *Directory* is specified, the collected files are not packed and are saved to the dlc directory of the specified directory. If the *Directory* is not specified, the files are saved on the Cell Manager in the default debug files directory.

-space

Displays the disk space required on the Cell Manager for the collected files.

-delete_dbg

Deletes the selected files on clients. On OpenVMS, if run together with the -session parameter, the command does not delete any debugs from the debug files directory.

-telemetry_files

Includes or excludes the telemetry data by enabling or disabling the selection. You cannot create telemetry when using the -depot option.

-no_getinfo

Excludes the getinfo file from the selection. For OpenVMS, this parameter is not applicable as OpenVMS systems do not have the get_info utility.

-no_config

Excludes the configuration information from the selection.

-no_logs

Excludes the log files from the selection.

-no_debugs

Excludes the debug files from the selection.

-no_compress

Disables the compression of the collected files on clients. By default, the compression is enabled.

-debug_loc *dir1* [*dir2*]...

Includes debugs not only from the default debug files directory but also from other directories, *dir1*, *dir2*,.... Note that the subdirectories are excluded from the search. If a specified directory does not exist on a particular client, the directory is ignored.

This option is valid only if the -no_debugs option is not specified.

-no_verbose

Disables verbose output. By default, verbose output is enabled.

-add_info *path*

Includes the additional information (for example, screenshots, pictures and the like) from a directory on client identified by *path*.

The `-any` option is used when the directory path is the same for all clients. It is important to make sure the path is not host-specific before using this option.

`-localpack` [*filename*]

Packs the directory structure from the current directory (must be the directory containing the `dlc` directory generated by the `-depot` option) to the *filename*. If the *filename* is not specified, the `dlc.pck` file is created in the current directory.

This option is equivalent to the `-pack` option, but is to be used only if the data is collected using the `-depot` option.

To add files other than the collected files to the package, copy the files to one of the following directories before running the command: `dlc/`*client*`/getinfo`, `dlc/`*client*`/log`, or `dlc/`*client*`/tmp` (on UNIX), or `.\dlc\`*client*`\getinfo`, `.\dlc\`*client*`\log`, or `.\dlc\`*client*`\tmp` (on Windows). You cannot add directories, but only files. If the files are not copied to one of the specified directories, the package cannot be unpacked during the unpack phase.

`-unpack` [*filename*]

Creates the `dlc` directory in the current directory, and unpacks the contents of the *filename* to the `dlc` directory. If the *filename* is not specified, the `dlc.pck` file in the current directory is unpacked.

Use this option when the collected (compressed or uncompressed) data was packed on the Cell Manager either using the `-pack` option or the `-localpack` option.

`-uncompress` *filename*

Uncompresses the unpacked compressed single file in the current directory.

Use this option after the packed data is unpacked using the `-unpack` option.

[`-hosts` *list*] `-del_ctracelog`

Deletes `ctrace.log` files containing the information where (on which clients) debug logs are generated and which debug prefixes are used. If the `-hosts`*list* option is specified, the command deletes `ctrace.log` files on specified clients only. Otherwise, `ctrace.log` files on all clients in a cell are deleted.

NOTES

The `omnidlc` command cannot be used to collect the Data Protector installation execution traces.

The Data Protector GUI debug files for systems other than Cell Manager can only be gathered using the `-hosts` option.

To collect debug files in a cluster, the command must be run using the `-hosts` option; the cluster nodes hostnames must be specified as the argument for the option. In a cluster, if the `-hosts` option is not specified, the data is collected from the active node.

Paths specified in postfix are not allowed.

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

EXAMPLES

1. To collect and compress all debug, log and getinfo files from the cell, and pack them in the "dlc.pck" file in the current directory on Cell Manager, using the verbose output, execute:

   ```
   omnidlc -no_filter
   ```

2. To collect only the log and debug files (without the getinfo files) from the clients "client1.company.com" and "client2.company.com" to the directory "c:\depot" on the Cell Manager, without compressing and packing the files, execute:

   ```
   omnidlc -no_filter -hosts client1.company.com client2.company.com -depot
   c:\depot -no_getinfo -no_compress
   ```

3. To collect log, debug, and getinfo files from the client "client1.company.com", compress and pack them to the "c:\pack\pack.pck" file on the Cell Manager, execute:

   ```
   omnidlc -hosts client1.company.com -pack c:\pack\pack.pck
   ```

4. To collect log, debug, and getinfo files from the default location and debugs from the additional directories, "C:\tmp" and "/temp/debugs", from the clients "client1.company.com" and "client2.company.com", and to compress and pack the files on the Cell Manager, execute:

   ```
   omnidlc -hosts client1.company.com client2.company.com -debug_loc C:\tmp
   /tmp/debugs
   ```

5. To delete all debug log files for the session with the ID "2013/04/27-9", execute:

   ```
   omnidlc -session 2013/04/27-9 -delete_dbg
   ```

6. To display disk space needed on the Cell Manager for the uncompressed debug files with the debugID "2351" from the client "client.company.com", execute:

   ```
   omnidlc -did 2351 -hosts client.company.com -space -no_getinfo -no_logs -no_
   compress
   ```

7. To pack the additional file located in the "C:\debug" directory on the client client1.company.com together with debug log files for the session with the ID 2013/05/17-24, execute:

   ```
   omnidlc -session 2013/05/17-24 -add_info -host client1.company.com C:\debug
   ```

8. To pack the directory structure in the current directory (must be the directory containing the `dlc` directory generated by the `-depot` option) to the "dlc.pck" file in the same directory, execute:

   ```
   omnidlc -localpack
   ```

9. To collect and pack telemetry files in "C:\tmp\dlc.dlc" on the Cell Manager "cellmanager.company.com", execute:

   ```
   omnidlc -no_filter -hosts cellmanager.company.com -no_compress -no_logs -no_
   config -no_getinfo -no_verbose -telemetry_files -pack C:\tmp\dlc.dlc
   ```

10. To unpack the "dlc.pck" file to the "dlc" directory of the current directory, execute:

    ```
    omnidlc -unpack
    ```

11. To specify multiple modules, execute:

    ```
    omnidlc bsm,vbda,dbsm
    ```

SEE ALSO

omnicc(1), omnicellinfo(1), omnicheck(1M), omnisv(1M)

omnidr(1M)

omnidr — a general purpose Data Protector disaster recovery command. Based on its input, it decides on what type of restore to perform (online restore using `omnir` or offline restore using `omniofflr`), as well as how to perform the restore (whether or not to use live operating system features).
(this command is available on systems with any Data Protector component installed)

SYNOPSIS

omnidr -version|-help

omnidr [-srd *FileName*] [-temp[os]] [-drimini *P1S*] [-map *OrgMnt_1 TrgMnt_1* [-map *OrgMnt_2 TrgMnt_2* ...]] [-[no_]cleanup] [-msclusdb] [-omit_deleted_files] [*GeneralOptions*]

GeneralOptions

 -target *ClientName*

 -local

 -report *Level*

 -omit_deleted_files

DESCRIPTION

The `omnidr` command is a general purpose Data Protector disaster recovery command that can be used in all recovery scenarios. Based on its input, `omnidr` decides what type of restore is going to be performed: online restore using `omnir` or offline restore using `omniofflr`, as well as how the restore is going to be performed (using or avoiding live operating system features).

OPTIONS

-version

   Displays the version of the `omnidr` command.

-help

   Displays the usage synopsis for the `omnidr` command.

-srd *FileName*

   Specifies the path to System Recovery Data (SRD) file that contains all required backup and restore object information to perform the restore.

   Note that `omnidr` always requires a valid SRD file with updated object information. By default the command searches the working directory for `recovery.srd` file. If it is not found, an error is reported. The `-srd` option overrides the default name `recovery.srd`.

-temp[os]

Specifies a temporary operating system used for disaster recovery. This way, the `omnidr` command can determine how to restore `CONFIGURATION` data. If this option is not specified, the active operating system is used.

`-drimini` *P1S*

Specifies the path to Phase 1 Startup (P1S) file if you have interrupted the `drstart` command during the 30 second pause and selected the `install only` option when performing EADR. In this case, the `drstart` command only installs disaster recovery files and exits. You have to start the `omnidr` command manually and provide the path to the P1S file using the `-drimini` option. The default path is `C:\$DRIM$.OB2\OBRecovery.ini` (Windows) or `/opt/omni/bin/drim/drecovery.ini` (Linux).

`-map` *OrgMnt TrgMnt*

Specifies mapping of original volumes to current volumes.

`-[no_]cleanup`

When the `-cleanup` option (default) is specified during disaster recovery of an active operating system, the `omnidr` command prepares a cleanup script and stores it into the `%ALLUSERSPROFILE%\Start Menu\Programs\Startup` folder. At first logon after the boot, the Data Protector disaster recovery installation is removed.

When this option is specified during disaster recovery of a temporary operating system, a cleanup command is written into restored software hive in the registry at `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\RunOnce`. The cleanup command is executed at first logon after the boot and it removes the temporary operating system installation together with Data Protector disaster recovery installation.

The cleanup script/command is not generated in the following cases:

- If Data Protector installation was found on the system during the `omnidr` command initialization.

- If the `-no_cleanup` option has been specified.

- If Data Protector disaster recovery installation does not reside in the *%SystemRoot%* folder (in this case it was most likely not installed during Data Protector disaster recovery).

- If the `-debug` option has been specified, the cleanup is not performed, because you would loose the debug information at next logon.

- if the `Minimal Recovery` option has been selected during EADR or OBDR, meaning that only boot and system disks would be recovered.

When the `omnidr` command is used on a dual-boot machine, it is strongly recommended to use the `-no_cleanup` option.

`-msclusdb`

If this option is specified, the `omnidr` command restores the Microsoft Cluster Service database.

GeneralOptions

`-target` *ClientName*

Specifies the target client system name. All objects will be restored to a computer specified by the -target parameter. If this parameter is not specified, the data will be restored to the system specified in the SRD file.

This option is used in two cases:

- During Disk Delivery disaster recovery the disks being restored can be installed into a client with a different hostname as original, therefore the name of the client must be specified.

- During Manual Disaster Recovery, it is possible, that DHCP protocol is installed. In this case, the hostname can be generated automatically by the DHCP server and is different from the original system hostname.

`-local`

Forces offline recovery from a local device. The `devbra` command is used to automatically scan for and configure attached devices. A list of detected devices is displayed if more than one is found and you must select one of them. If this option is not specified, the device used for the restore is going to be the same as the device used during backup.

`-report` *Level*

Specifies the error reporting level. This is useful if you want to reduce the number of messages written during recovery. For example, since practically all operating system files are overwritten during the active operating system recovery, this means that innumerable warnings bringing no useful information will be displayed, thus slowing down the recovery. Messages are classified (in ascending order) as: `1` (warning), `2` (minor), `3` (major) and `4` (critical). For example, if `3` is selected, only major and critical messages are reported. By default, all messages are reported.

`-omit_deleted_files`

Specifies that the files that were deleted between incremental backups or between a full and incremental backup are not restored. Note that this may slow down the recovery process.

NOTES

The `omnidr` command is available on Windows and Linux systems only.

EXAMPLES

The following examples illustrate how the `omnidr` command works.

1. To use the SRD file stored on a floppy drive for the restore, execute:

   ```
   omnidr -srd "A:\recovery.srd"
   ```

2. To use the local backup device, execute:

   ```
   omnidr -local
   ```

SEE ALSO

omniiso(1), omniofflr(1M), omnisrdupdate(1M)

omnihealthcheck(1M)

omnihealthcheck — checks the status of Data Protector services, the consistency of the Data Protector Internal Database (IDB), and if at least one backup of the IDB exists
(this command is available on the Data Protector Cell Manager)

SYNOPSIS

omnihealthcheck -version|-help

omnihealthcheck [-config *ConfigFile*]

DESCRIPTION

The `omnihealthcheck` command reads the specified configuration file where each line of the file is treated as a separate command and is executed. The commands must be listed with full paths except if they reside in the Data Protector default commands directory. With the Windows Cell Manager, the configuration file must be in the Unicode format. If the configuration file is not specified, the default `HealthCheckConfig` file located in the server configuration directory on the Cell Manager is used.

If the default file is used, `omnihealthcheck` checks if Data Protector services (`CRS`, `MMD`, `hpdp-idb`, `hpdp-idb-cp`, `hpdp-as`, `KMS`, `omnitrig`, and `omniinet`) are active, if the Data Protector MMDB is consistent, and if at least one backup of the Data Protector Internal Database (IDB) exists.

Exit codes of individual commands are inspected at the end.

There are 3 different exit codes for the `omnihealthcheck` command:

`0` - All listed commands and their exit codes have been executed

`1` - At least one of the commands in the configuration file could not be executed or has completed with an exit code other than `0`.

`2` - The configuration file could not be read.

The final health check exit code is `0` (OK) only if all executed commands from the configuration file completed successfully (exit codes of all executed individual commands from the configuration file are `0`).

Output of the `omnihealthcheck` command is saved on the Cell Manager in the `HealthCheck.log` file located in the default server log files directory.

If a timeout occurs, `omnihealthcheck` fails.

`omnihealthcheck` is by default scheduled to run daily at 12:00 (noon) as a part of the Data Protector check mechanism. The default schedule value can be changed by changing the `DailyCheckTime` global option.

OPTIONS

`-version` - Displays the version of the `omnihealthcheck` command.

`-help` - Displays the usage synopsis for the `omnihealthcheck` command.

`-config`*`ConfigFile`* - Specifies an alternative configuration file for the `omnihealthcheck` command.
Note that you can define the commands to be executed in the health check.

SEE ALSO

omnirpt(1), omnitrig(1M)

omniinetpasswd(1M)

omniinetpasswd — manages the local Data Protector Inet configuration on Windows systems where the Inet process must be run under a specific user account, and sets a user account to be used by the Installation Server during remote installation
(this command is available on systems with any Data Protector component installed)

SYNOPSIS

omniinetpasswd -version | -help

omniinetpasswd -add {*User@Domain* | *Domain\User* ...} [*Password*]

omniinetpasswd -delete {*User@Domain* | *Domain\User* ...}

omniinetpasswd -modify {*User@Domain* | *Domain\User* ...} [*Password*]

omniinetpasswd -list [*Domain*]

omniinetpasswd -clean

omniinetpasswd -[no_]inst_srv_user {*User@Domain* | *Domain\User* ...}

DESCRIPTION

On specific Windows operating systems, the Data Protector Inet process must be run under a specific operating system user account rather than under the local user account SYSTEM. Additionally, on Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012 systems, the Data Protector Installation Server must use a specific operating system user account for remote installation. The omniinetpasswd command provides functionality for management of Inet configuration on the local system, and functionality for setting a user account that will be used by the Installation Server during remote installation. Use command options -add, —delete, -modify, -list, and -clean for local Inet configuration management, and options -inst_srv_user and -no_inst_srv_user for setting a user account to be used for remote installation.

Note that omniinetpasswd does not add, remove, or change user accounts in the operating system configuration.

OPTIONS

-version

Displays the version of the omniinetpasswd command.

-help

Displays the usage synopsis for the omniinetpasswd command.

-add {*User@Domain* | *Domain\User*} [*Password*]

Adds the specified user account from the local Inet configuration. omniinetpasswd prompts for the

password if not specified in the command line.

-delete {*User@Domain* | *Domain\User*}

Removes the specified user account from the local Inet configuration. omniinetpasswd prompts for the password if not specified in the command line.

-list *Domain*

Lists user accounts from the local Inet configuration: either all or only the accounts belonging to the specified domain.

-modify {*User@Domain* | *Domain\User*} [*Password*]

Changes the password for a configured user account. omniinetpasswd prompts for the password if not specified in the command line.

-clean *Domain*

Removes all operating system user accounts from the local Inet configuration.

-inst_srv_user {*User@Domain* | *Domain\User*}

Sets the specified user in the local Inet configuration to be used by the Installation Server during remote installation.

This option can only be used on Windows Server 2008 and Windows Server 2012 systems.

-no_inst_srv_user {*User@Domain* | *Domain\User*}

Marks the specified user in the local Inet configuration not to be used by the Installation Server during remote installation.

This option can only be used on Windows Server 2008 and Windows Server 2012 systems.

NOTES

The omniinetpasswd command is available on Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012 systems only.

EXAMPLES

1. To remove the user "User1" from the Inet configuration, execute:

   omniinetpasswd -delete CompanyDomain\User1

2. To delete all operating system accounts from the local Inet configuration, execute:

   omniinetpasswd -clean

3. To set the user "User1" from the domain "CompanyDomain" to be used by Installation Server, execute:

   omniinetpasswd —inst_srv_user User1@CompanyDomain

omniintconfig.pl(1M)

omniintconfig.pl — configures and checks the configuration of one or multiple Oracle databases (this command is available on systems with the Data Protector `User Interface` component installed).

SYNOPSIS

omniintconfig.pl -version|-help

omniintconfig.pl [-encode] [-chkconf] [-force] {-passwordfile *FileName*|*Param*=*Value* [*Param*=*Value*...]}

Param

MoM

CellManager

Client

Instance

OSUSER

OSGROUP

TGTUser

TGTPasswd

TGTService

RCUser

RCPasswd

RCService

ORACLE_HOME

ClusterNodes

DESCRIPTION

Use the `omniintconfig.pl` command to configure and check the configuration of one or multiple Oracle databases at the same time. If configuration already exists then script does not perform the Oracle integration reconfiguration instead it changes only the specified configuration parameters directly in the existing configuration file.

On Windows systems, you must use the `perl` command to run `omniintconfig.pl`. An example of the command line is `perl omniintconfig.pl -help`.

OPTIONS

`-version`

Displays the version of the `omniintconfig.pl` command.

`-help`

Displays the usage synopsis for the `omniintconfig.pl` command.

`-encode`

Encodes passwords before they are saved to Data Protector Oracle database specific configuration files. Omit this option if the provided passwords are already encoded.

`-chkconf`

Performs a configuration check for specified Oracle databases. Provided parameter values are saved to corresponding Data Protector Oracle database configuration files, regardless of whether the check succeeds or not. By default, the session ends if a configuration check for a database fails. However, if you specify the `-force` option, Data Protector continues configuring other Oracle databases.

`-passwordfile` *FileName*

Specifies that configuration parameters should be read from a file. The file must be in XLS or CSV file format.

Alternatively, parameters can be specified at run time, however, only for one Oracle database at a time. See the parameters description below.

PARAMETERS

`MoM`

Manager of Managers (optional).

`CellManager`

Data Protector Cell Manager. Default: Cell Manager of the local client.

`Client`

Client with Oracle Server installed. In cluster environments, specify the virtual server or, in RAC, one of the cluster nodes. Default: local client.

`Instance`

Oracle database name (mandatory).

`OSUSER, OSGROUP`

(Applicable for UNIX clients.) The UNIX user account under which you want the configuration and browsing of Oracle databases to start. This user will be automatically added to the Data Protector `admin` user group for the client specified in `Client`.

`TGTUser, TGTPasswd`

Login information for the target database (username and password).

TGTService

Target database service(s). If there is more than one service, separate them with a semicolon (service1;service2...).

RCUser, RCPasswd

Login information for the recovery catalog database (username and password).

RCService

Recovery catalog database service.

ORACLE_HOME

Oracle Server home directory.

ClusterNodes

Cluster nodes (applicable in cluster environments). The user OSUSER, OSGROUP will be automatically added to the Data Protector admin user group for each cluster node listed here. Separate cluster nodes with a semicolon (node1;node2...).

If you do not specify this parameter, you need to add these users manually.

EXAMPLES

1. Suppose the file "C:\My_documents\Oracle_instances.csv" contains configuration parameters for the Oracle databases "IN1" and " IN2". The passwords in the file are encoded.

   To configure the Oracle databases "IN1" and "IN2" using the file "C:\My_documents\Oracle_instances.csv", log in to the Windows client on which the file is saved and execute:

   ```
   perl omniintconfig.pl -passwordfile C:\My_documents\Oracle_instances.csv
   ```

2. Suppose you are logged in to a UNIX client. To configure the Oracle database "IN2" by specifying parameters at run time, execute:

   ```
   omniintconfig.pl -encode CellManager=galaxy Client=star Instance=IN2 ORACLE_
   HOME=C:\oracle\product\10.2.0\db_1 TGTUser=system TGTPasswd=BlueMoon
   TGTService=IN2_1;IN2_2
   ```

   Note that the password "BlueMoon" is not encoded. Therefore, you must specify the "-encode" option.

3. Suppose you are logged in to a Windows client. To configure and check the configuration of all Oracle databases specified in "C:\My_documents\Oracle_instances.xls", execute:

   ```
   perl omniintconfig.pl -chkconf -force -passwordfile C:\My_documents\Oracle_
   instances.xls
   ```

4. Suppose you are logged in to a UNIX client. To check the configuration of the Oracle database "IN2", execute:

   ```
   omniintconfig.pl -chkconf CellManager=galaxy Client=star Instance=IN2
   ```

SEE ALSO

util_cmd(1M), util_oracle8.pl(1M), vepa_util.exe(1M)

omnikeytool(1M)

omnikeytool — manages keys used to encrypt backup data
(this command is available on the Data Protector Cell Manager)

SYNOPSIS

omnikeytool -version | -help

omnikeytool -create *EntityName* [-description *Description*]

omnikeytool -activate *EntityName* -keyid *KeyID StoreID*

omnikeytool -deactivate *EntityName*

omnikeytool -export *CSVFile ExportOptions*

omnikeytool -import *CSVFile*

omnikeytool -modify -keyid *KeyID StoreID* -description *Description*

omnikeytool -list [-active | -unused]

omnikeytool -delete -keyid *KeyID StoreID*

ExportOptions

 -keyid *KeyID StoreID*

 -active

 -entity *EntityName*

 -time *Day Hour Day Hour*

 -all

Date= *[YY]YY/MM/DD (1969 < [YY]YY < 2038)*

Hour= *HH:MM*

DESCRIPTION

The omnikeytool command manages keys used for encryption. You must create the key by using the omnikeytool command prior to performing an encrypted backup.

OPTIONS

-version

    Displays the version of the omnikeytool command.

-help

   Displays the usage synopsis for the omnikeytool command.

-create *EntityName* [-description *Description*]

   EntityName can be:

   - A *ClientName* value for the specified filesystem, disk image, or the IDB
   - An *AppType*:*DatabaseID* pair or an *AppType*:*ClientName*:*AppName* trinity for the specified application integration
   - A *MediumID* value, if you use drive-based encryption

   Ensure that the value of ClientName matches the name that was specified for the client system in the correspondent backup specification.

   If the -description option is specified, you can provide a description string for the new encryption key.

-activate *EntityName* -keyid *KeyID StoreID*

   Associates the specified encryption key with the specified entity name string and activates the key.

-deactivate *EntityName*

   Disassociates the specified entity name string from the current active backup encryption key.

   The password encryption is done by AES keys with entity name "Data Protector Passwords". Do not delete or deactivate this key, as Data Protector will not be able to decrypt encrypted strings. Deactivating the password key affects backup only; it does not affect restore, nor does it disable the key or act as a key revocation. For Cloud devices, deactivating or deleting the password key renders the Cloud device unusable.

-export *CSVFile ExportOptions*

   Exports encryption key records into the specified comma separated values (CSV) file. The file is exported to the Data Protector encryption keys directory. Exporting does not delete encryption keys from the keystore.

-import *CSVFile*

   Imports encryption key record matching the key number from the specified keystore file. The file is imported to the Data Protector encryption keys directory.

-modify [-description *Description*]

   Modifies the description for the specified encryption key.

-list [-active | -unused]

   Lists encryption keys related information from the cell.

   The command lists the following information for each encryption key in the keystore file: key status (active, inactive, migrated), key ID, date and time of creation, type of encryption, and the key description. For greater scrutiny, the above-mentioned information is listed for each client in the cell separately.

   If the -active option is specified, the command just lists currently active keys and the entity names associated with them.

If the `-unused` option is specified, the command lists all encryption keys which are present in the keystore file on the Cell Manager, but have never been used for encryption.

`-delete`

Deletes the record of an inactive encryption key identified by key ID.

Ensure that the key you intend to delete is not in use. If the encryption key is not available, restore of encrypted data is not possible.

ExportOptions

`-keyid` *KeyID StoreID*

Exports all encryption key records with the specified key ID.

`-active`

Exports all currently active encryption keys.

`-entity` *EntityName*

Exports only the active key record identified by the *EntityName* string.

`-time` *Day Hour Day Hour*

Exports all encryption key records in the specified time frame.

`-all`

Exports all encryption key records.

EXAMPLES

The following examples illustrate how the `omnikeytool` command works.

1. To activate the encryption key "10B536738F8831478408000000000000 5B9381955B9381955B9381955B938195" for the client system "proxima", execute:

   ```
   omnikeytool -activate proxima -keyid 10B536738F8831478408000000000000
   5B9381955B9381955B9381955B938195
   ```

2. To deactivate an encryption key for the client system "stella", execute:

   ```
   omnikeytool -deactivate stella
   ```

3. To modify your description of the encryption key "10B53673B8232747A806000001000000 5B9381955B 9381955B9381955B938987", execute:

   ```
   omnikeytool -modify -keyid 10B53673B8232747A806000001000000 5B9381955B
   9381955B9381955B938987 -description key_number_1
   ```

4. To export the active encryption key "10B53673B8232747A806000001000000 5B9381955B 9381955B9381955B938321" to a comma-separated values (CSV) file "a.csv", execute:

   ```
   omnikeytool -export a.csv -keyid 10B53673B8232747A806000001000000 5B9381955B
   9381955B9381955B938321
   ```

5. To list all encryption keys which are present in the keystore file on the Cell Manager, but have never been used for encryption, execute:

   ```
   omnikeytool -list -unused
   ```

SEE ALSO

omnib(1), omniobjconsolidate(1), omniobjcopy(1), omniobjverify(1), omnir(1)

omnimigrate.pl(1M)

omnimigrate.pl — migrates the Data Protector Internal Database (IDB) from the format used in earlier versions to the PostgreSQL relational database format used in Data Protector 8.00 and later.
(this command is available on the Data Protector Cell Manager)

SYNOPSIS

omnimigrate.pl -help -version

omnimigrate.pl -cleanup

omnimigrate.pl -export [-new_cm *NewCmHostName*] [-output_dir *OutputDirectoryPath*]

omnimigrate.pl -import [-input_dir *InputDirectoryPath*]

omnimigrate.pl -export_critical_part [-output_dir *OutputDirectoryPath*]

omnimigrate.pl -import_critical_part [-input_dir *InputDirectoryPath*]

omnimigrate.pl -start_catalog_migration

omnimigrate.pl -report_catalog_migration_progress

omnimigrate.pl -stop_catalog_migration

omnimigrate.pl -report_old_catalog [media|sessions|objects]

omnimigrate.pl -remove_old_catalog

DESCRIPTION

The omnimigrate.pl command helps you migrate the Data Protector Internal Database (IDB) from the format used in earlier versions to the PostgreSQL relational database format used in Data Protector 8.00 and later.

After the upgrade, all objects whose catalogs are still protected and were backed up prior to the migration, have their catalogs stored in the old format. When catalogs expire, the old DC binary files are automatically deleted (through daily maintenance tasks).

Old database files are however kept as long as there are filenames which are protected.

A warning is displayed in the Event Log which lists the protected media, object, and sessions that still need the old DC binary files and IDB files, and the amount of space they occupy.

Use this command after an upgrade from earlier versions of Data Protector to:

- list the protected media, objects, and sessions that still need the old DC binary files and IDB files, and the amount of space they occupy
- migrate the Detail Catalog Binary Files (DCBF) to the new format.

> **Tip:** HPE recommends that you wait until most of your old media expire and you trigger the migration only for permanently protected media. For the conversion duration and space requirement

> estimates, see the *HPE Data Protector Installation Guide*.

OPTIONS

`-help`

Displays the usage synopsis for the `omnimigrate.pl` command.

`-version`

Displays the version of the `omnimigrate.pl` command.

`-cleanup`

Removes all export and import files from the specified directory. Use this command if the IDB export fails.

`-export` [`-new_cm` *NewCmHostName*] [`-output_dir` *Directory*]

The command exports the database, prepares it for import, and converts the IDB backup filesystem specification to IDB backup application integration specification. If no output directory is specified, the database is exported to the default *Data_Protector_home*\tmp (Windows systems) or /var/opt/omni/tmp (UNIX systems) directory.

If the `-new_cm` option is not specified, the command assumes that the system is being upgraded.

Use this command to export the database of Data Protector versions prior to 8.00.

For Data Protector versions 8.00 and later, there is no need to use this command because the database format is already in the PostgreSQL relational database format.

`-import` [`-input_dir` *Directory*] [`-force`]

Imports the data exported by the `–export` option into the new database. If no directory is specified, data is searched in the default *Data_Protector_home*\tmp (Windows systems) or /var/opt/omni/tmp (UNIX systems) directory, in the `cdb` and `mmdb` subdirectories. To force an import after a failure during the upgrade process, specify `–force`.

`-export_critical_part` [`-output_dir` *OutputDirectoryPath*]

This option is used during the upgrade process. It exports only the critical part of the old IDB without the catalog. If `-output_dir` is not specified, the files will be by default exported to the *Data_Protector_home*\tmp (Windows systems) or /var/opt/omni/tmp (UNIX systems) directory.

`-import_critical_part` [`-input_dir` *InputDirectoryPath*]

Imports only the critical part of the old IDB without the catalog (exported with the `–export_critical_part` option) into the PostgreSQL relational database. If `-input_dir` is not specified, the command searches by default in the *Data_Protector_home*\tmp (Windows systems) or /var/opt/omni/tmp (UNIX systems) directory.

`-start_catalog_migration`

Starts the catalog migration.

`-report_catalog_migration`

Displays the progress of the catalog migration.

`-stop_catalog_migration`

> Stops the catalog migration process. The current DCBF upgrade is finished and logged. If you run the `omnimigrate` command with the `-start_catalog_migration` the migration continues from where it stopped.

`-report_old_catalog [media|sessions|objects] [-shared_dir]`

> Displays the usage and statistics of the old catalog. The `media` option displays the list of media whose catalog is still in the old format and their expiration date. The `sessions` option displays the sessions and their expiration date. The `objects` option displays the objects and their expiration date.

`-remove_old_catalog`

> Removes all of the old DC binary files and all of the old database data files.

RETURN VALUES

`0` - Successfully finished.

`(1-4)` - An error occurred.

ERRORS

`1` - A generic error occurred.

`2` - Migration of IDB catalogs failed.

`3` - Configuration error (Cell Manager configuration error or an error during the import of clients) occurred.

`4` - Error parsing options.

NOTES

The `omnimigrate.pl` command *cannot* be used to migrate the Cell Manager from obsolete platforms to supported ones. You need to migrate a Cell Manager to Data Protector 9.00 before you upgrade to Data Protector 9.08. For details, see the documentation of the respective Data Protector version.

EXAMPLES

1. To list all media, whose catalog is still in the old format and their expiration date, execute:

   `omnimigrate.pl -report_old_catalog -media`

2. To start the catalog migration on the Cell Manager after upgrade, execute the following command:

   `omnimigrate.pl -start_catalog_migration`

   > **Note:** Once the full catalog migration is done (after there are no old catalogs), change the global variable `SupportOldDCBF` to `0`.

SEE ALSO

ob2install(1M), omnigui(5), omniintro(9), omnisetup.sh(1M), omniusers(1), upgrade_cm_from_evaa (1M)

omniofflr(1M)

omniofflr — enables restore of any type of Data Protector backup objects in the absence of operable
Data Protector Internal Database (IDB), including the IDB itself
(this command is available on systems with any Data Protector component installed)

SYNOPSIS

omniofflr -version | -help

omniofflr *DeviceOptions MediaOptions1* [*MediaOptions2*...] *ObjectOptions1*
[*ObjectOptions2*...] [*GeneralOptions*]

omniofflr -idb -autorecover [*AutorecoverOptions*] [[-changedevhost *MAClientName*]]
[*GeneralOptions*]

omniofflr -idb -read *OptionFile* [*GeneralOptions*]

omniofflr -idb *DeviceOptions MediaOptions1* [*MediaOptions2*...] *ObjectOptions1*
[*ObjectOptions2*...] [*GeneralOptions*]

AutorecoverOptions

 [-force]

 [-session *SessionID*]

 [-save *OptionFile*]

 [-skiprestore]

 [-logview]

 [-optview]

DeviceOptions

 -name *DeviceName*

 -dev *PhysicalDevice1* [*PhysicalDevice2* ...]

 -mahost *DeviceHostName*

 -policy *LogicalDevicePolicy*

 -type *LogicalDeviceType*

 [-ioctl *RoboticsSCSIAddress*]

 [-description *DeviceDescription*]

 [-blksize *BlockSize*]

MediaOptions

  -maid *MediumID1* [*MediumID2* ...]

  [-slot *Slot1[:Flip1]* [*Slot2[:Flip2]*...]]

  [-position *Segment1:Offset1* [*Segment2:Offset2*...]]

ObjectOptions

FILESYSTEM RESTORE

  {-filesystem|-winfs} *Client:MountPoint Label*

  -daid *DAID*

  -tree *TreeName1* [*TreeOptions1*] [-tree *TreeName2* [*TreeOptions2*...]]

  [-merge]

  [-[no_]overwrite]

  [-move_busy]

  [-omit_deleted_files [-time *ObjectBackupStartTime*]]

  [-var *OptName OptValue*]

DISK IMAGE RESTORE (WINDOWS SYSTEMS)

  -rawdisk *Client Label*

  -section [*ToSection1=*]*Section1*

  [-section [[*ToSection2=*]]*Section*2...]

  -daid *DAID*

TreeOptions

  -exclude *TreeName1* [*TreeName2*...] {-as|-into} *NewTreeName*

GeneralOptions

  -verbose

  -preview

  -report

  -target *TargetHostName*

  -[no]ok[mediumlist]

DESCRIPTION

The `omniofflr` command is a standalone command. On Windows and Linux systems, it can also be used indirectly by the higher-level `omnidr` command, which automatically generates appropriate `omniofflr` command-line options, based on the information retrieved from the SRD file.

The `omniofflr` command enables you to restore any type of Data Protector backup objects in the absence of operable Data Protector Internal Database (IDB), including the IDB itself. The IDB may not be functioning as a result of a disaster, loss of connectivity with the Cell Manager, or other undesirable circumstances.

To execute the `omniofflr` command, you need to specify the details about the backup (restore) device and the backup media, including backup object positions on the media. You can obtain these details automatically from the SRD files location, or supply them manually in the `omniofflr` command line. Query the IDB using the `omnidb` command after the backup session, and write down the results. You can also prepare a script that queries the IDB and generates another script within which the `omniofflr` command is invoked with appropriate options.

OFFLINE RESTORE OF THE IDB

The IDB restore process, when invoked by `omniofflr`, consists of four phases:

1) Stopping the Data Protector services/daemons (with the exception of the Data Protector Inet service on Windows system).

2) Restoring the IDB.

3) Performing a rollforward operation on the IDB using transactions from the available archived log files. You can choose to skip this phase by responding to the `omniofflr` command prompt.

4) Starting the Data Protector services/daemons.

A new archived log file is created every time an IDB backup session is started, the IDB is initialized or checked for consistency, or an existing archived log file reaches its maximum size. Archived log files reside on the Cell Manager in the directory *Data_Protector_program_data*\server\db80\pg\pg_xlog_archive (Windows systems) or /var/opt/omni/server/db80/pg/pg_xlog_archive (UNIX systems).

For restoring the IDB, `omniofflr` can operate in three modes: autorecovery mode, read mode, and manual mode.

AUTORECOVERY MODE

In this mode, the `omniofflr` operation is fully automated. The command retrieves all required IDB restore parameters from the IDB recovery file named `obrindex.dat`, residing on the Cell Manager in the `rlog` directory on the IDB recovery files location. It is updated during each IDB backup session and contains all required IDB restore parameters, including filenames of the archived log files created during the IDB backup session. You can create and maintain a duplicate of this file by configuring the `RecoveryIndexDir` global option. `omniofflr` can use the duplicate if the original is missing or corrupted. HPE also recommends to place the IDB recovery file on a physical disk separate from the core part of the IDB.

READ MODE

You can use this mode to direct `omniofflr` to obtain the parameters from the file *OptionFile* that has been created either manually or using the `omniofflr` options `-idb -autorecover -save`. Use this mode when, for example, the restore devices differ from the backup devices (or they are attached to a different system). In such a case, you have to manually update the file *OptionFile* appropriately before invoking the restore session.

MANUAL MODE

You can use this mode when neither the `obrindex.dat` file nor the file *OptionFile* are available, and you need to manually specify all required parameters in the `omniofflr` command line.

OPTIONS

`-version`

Displays the version of the `omnioﬄr` command.

`-help`

Displays the usage synopsis for the `omnioﬄr` command.

`-idb`

Selects the Data Protector Internal Database (IDB) for restore. If no additional options as `-autorecover` or `-read` are specified, the `-idb` option starts the IDB restore in the manual mode.

`-autorecover`

This option can be only used in combination with the `-idb` option.

Starts the IDB restore in the autorecovery mode. To use this mode, the IDB recovery file `obdrindex.dat` (the original or its duplicate) should exist on the Cell Manager.

`-changedevhost` *MAClientName*

This option can be only used in combination with the `-idb` and `-autorecover` options.

Specifies the hostname of the Data Protector Media Agent system to be used for the IDB restore instead of the Media Agent system specified in the IDB recovery file.

`-read` *OptionFile*

This option can be only used in combination with the `-idb` option.

Starts the IDB restore in the read mode using IDB restore parameters from the specified file. To use this mode, the IDB restore parameter file should be available on the local system. The file should reflect the current configuration of the Data Protector cell with regards to the Media Agent system, the backup (restore) device, and the backup media on which IDB backup image is stored.

AutorecoverOptions

`-force`

This option can be only used in combination with the `-idb` and `-autorecover` options.

Forces Data Protector to overwrite existing Internal Database files that reside at their original location. If you omit this option, only the missing files are recreated with data from the IDB backup images.

`-session` *SessionID*

This option can be only used in combination with the `-idb` and `-autorecover` options.

Omits selecting the last valid IDB backup session for the restore process, and instead selects the specified IDB backup session (full or incremental). Make sure the specified session is referenced in the IDB recovery file.

`-save` *OptionFile*

This option can be only used in combination with the `-idb` and `-autorecover` options.

Saves the IDB restore parameters retrieved in the autorecovery mode to the specified parameter file in order to enable starting the IDB restore with `omniofflr` in the read mode later. Note that unless you specify the option `-skiprestore`, the IDB restore is also performed.

`-skiprestore`

This option can be only used in combination with the `-idb` and `-autorecover` options.

Skips starting the actual IDB restore process. Usually, you may want to specify this option together with the `-save` option.

`-logview`

This option can be only used in combination with the `-idb` and `-autorecover` options.

Displays the contents of the IDB recovery file.

`-optview`

This option can be only used in combination with the `-idb` and `-autorecover` options.

Displays the parameters used for the IDB restore session.

DeviceOptions

`-name` *LogicalDeviceName*

Parameter that specifies the logical device name.

`-dev` *PhysicalDevice*

Specifies the pathname of the device file. For example, `c:\temp\dev1`, `scsi1:0:0:0`, or `/dev/tape0`.

> **Note:** The Drive index number must be specified when restoring the IDB without an IDB recovery file, if the drive index is not 1. For more details, see Example 5.

`-mahost` *DeviceHostName*

Specifies the name of the client, where the restore device is attached and a Media Agent started.

`-policy` *LogicalDevicePolicy*

Specifies the policy ID for the device specified by the `-dev` option. Policy can be defined as:

1 (Standalone)

3 (Stacker)

5 (6300 magneto-optical jukebox)

6 (Exchange through cmd execution)

8 (GRAU DAS exchanger library)

9 (Silo medium library)

10 (SCSI exchanger)

11 (RSM exchanger)

`-type` *LogicalDeviceType*

Specifies the media type for the media in the device specified by the `-device` option. Media type

numbers are defined as media class in the `scsitab` file. For location, see the help index "support of new devices".

-ioctl *RoboticsSCSIAddress*

Specifies the pathname of the robotics control device file for library device. For example, `c:\temp\roboticsdev`, `scsi1:2:0:0`, or `/dev/dlt_robotics`

-description *DeviceDescription*

This is an optional parameter that specifies the logical device description.

-blksize *BlockSize*

This is an optional parameter that specifies the block size the device is going to use when accessing media.

MediaOptions

-maid *MediumID*

Specifies the medium identification number of the medium that contains the object data; for example `8c04110a:3b0e118b:041c:0001`. If `unknown` is specified, each medium will be accepted as valid and restore will be attempted. Whole medium will be scanned for the requested object and it may take a very long time, if the object is not on the medium. Mount prompt in such case will request the next medium, without specifying the medium label.

-slot *Slot*[:*Flip*]

Specifies the slot identifier of the slot, where the required media is located, thus enabling Data Protector to automatically load media from the exchanger slots. Note that the sequence has to match the sequence in the list created using the `-maid` option.

-position *Segment*:*Offset*

Specifies the segment and offset position of the restore object data on the medium; for example `67:20`. If the position is not specified, the position `1:0` is assumed, thus prolonging the restore time. Note that the sequence has to match the sequence in the list created using the `-maid` option.

ObjectOptions

-filesystem *Client*:*MountPointLabel*

Selects the filesystem identified with *Client:MountPoint Label* for restore. Client determines the name of the system where the object was backed up. *MountPoint* specifies the mount point name of the volume to be restored (for example `/C`, `/tmp`, `/`, and so on). It must be in the same format as stored in the IDB. *Label* specifies the backup/restore objects description that uniquely defines an object (`-filesystem computer.domain.net:/mount label`)

-winfs *Client*:*MountPoint Label*

Selects the Windows filesystem identified with *Client:MountPoint Label* for restore. Client determines the name of the system where the object was backed up. *MountPoint* specifies the mount point name of the volume to be restored (for example `/C`, `/tmp`, `/`, and so on). It must be in the same format as stored in the IDB. Therefore, for example, on Windows systems `C:` translates into `/C`. Label specifies the backup/restore object's description that uniquely defines an object (`-winfs computer.domain.net:/C:`, and so on)

-rawdisk *Client Label* **-section** [*ToSection*=]*Section*

Selects the disk image identified by *Host* and *Label* for restore. Specifies the disk image section to be restored. To restore the section to a new section, include both the source and destination sections.

This option is available only for Windows Vista, Windows 7, Windows 8, Windows Server 2008, and Windows Server 2012.

`-daid` *DAID*

Specifies the identification number of the Disk Agent process that backed up an object. The identification number must be in the POSIX time format. You can obtain it by invoking the `omnidb -session` *SessionID* `-detail` command while the Data Protector Internal Database is available.

`-merge`

This option merges files from the backup medium to the target directory and replaces older versions that exist in the directory with newer (if they exist on the medium) files. Existing files are overwritten if the version on the medium is newer than the version on disk. No existing directory is deleted. If a directory or file doesn't exist on disk (but is on the backup medium) it is restored (created).

`-overwrite`

By default, or if the `-overwrite` option is specified, the already existent files on the disk are overwritten by the restored files.

`-no_overwrite`

If the `-no_overwrite` option is specified, only the files that do not exist on the disk are restored.

`-move_busy`

This option is used with the `-omit_deleted_files` or `-overwrite` option. A problem can occur if, for example, a file to be overwritten cannot be deleted because it is currently in use. If this option is specified, Data Protector moves busy file *filename* to *#filename* on UNIX systems (adding a hash-mark in front of the filename), or to *filename.001* on Windows system. On UNIX systems the original file can thus be deleted as the lock is transferred to the corresponding file starting with the #sign. For example, */tmp/DIR1/DIR2/FILE* would be moved to */tmp/DIR1/DIR2/#FILE*. On Windows system the application only uses the newly-restored file after the file is restored and the system is restarted.

`-omit_deleted_files`

This option can only be used in combination with the `-overwrite` and `-time` options. 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 specified, 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 not restored.

If this option is not specified, 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.

If you use this option in combination with the `-as` or `-into` option, carefully choose the restore location to prevent accidental removal of existing files.

`-time` *ObjectBackupStartTime*

This option should be used in combination with the `-omit_deleted_files` option.

Specifies the backup start time for the backup object you are restoring. *ObjectBackupStartTime* must be a value in the POSIX time format. You should obtain it by invoking the `omnidb -session` *SessionID* `-detail` command while the Data Protector Internal Database is available, where *SessionID* is implicitly defined by the `-daid` *DAID* option.

`-var` *OptName OptValue*

This option lets you specify a variable name and its value for proper operation of some platforms and integrations. Setting user definable variables (a variable name and its value) enables flexible operation on some platforms and integrations with Data Protector. The list of variables and their values that are configurable with Data Protector is dynamic and comes with Data Protector patches.

`-tree` *TreeName TreeOptions*

Specifies the starting root directory of data restore. Note that this starting directory is also restored.

TreeOptions

`-exclude`*TreeName*

Specifies trees excluded from the restore.

`-as` *NewTreeName*

This is an optional parameter that restores the selected fileset as the specified tree. This parameter is of vital importance for the Disk Delivery disaster recovery, since without it the restore to the original location would be performed.

`-into` *NewTreeName*

This is an optional parameter that restores the selected fileset into the given directory. This parameter is of vital importance for the Disk Delivery disaster recovery, since without it the restore to the original location would be performed.

GeneralOptions

`-verbose`

Specifies the verbose level of progress reporting.

`-preview`

Specifies that the preview mode of the restore is entered.

`-report`

Displays a report of the disaster recovery using the `omniofflr` command.

`-target`

Specifies the target system name which is different than the original.

`-[no]ok[mediumlist]`

By default the options are parsed and displayed so that the user can check them and confirm the start of restore. This means that the `omniofflr` command used from a script could not be executed because it would wait for the confirmation before starting the restore. This option has to be used to skip confirmation, thus enabling the execution of the `omniofflr` command from a script.

NOTES

The omniofflr command is available on Windows systems, HP-UX systems, and Linux systems.

The omniofflr command does not support robotic loaders of backup media. You need to ensure the appropriate media is loaded into the specified backup devices. For this purpose, you can use the uma command on the system to which robotics is attached.

In computer clusters, execute the omniofflr command on the active cluster node.

EXAMPLES

The following examples illustrate how the omniofflr command works.

1. To restore the "c:/temp" directory of the system "computer.company.com" without the "c:/temp/vnc" directory, which was backed up using an HPE Ultrium standalone backup device on a STK Ultrium backup medium, attached to the Cell Manager "cm.company.com", into the "c:/test/temp" directory, execute:

    ```
    omniofflr -verbose -name HP:Ultrium -dev scsi2:0:4:0C -mahost cm.company.com -
    policy 1 -type 13 -maid 9e03110a:3b5ee669:05ac:0001 -computer.company.com:/C C:
    -daid 996144004 -tree /temp -exclude /temp/vnc -into c:/test/temp
    ```

    To get the logical device name and its SCSI address, execute:

    ```
    devbra -dev
    ```

    The output of the command looks something like this:

    ```
    HP:Ultriumscsi2:0:4:0cLTO : HP LTO drive
    ```

    "HP:Ultrium" is the logical device name of the backup device while "scsi2:0:4:0c" specifies the SCSI address of the device.

    To obtain the medium ID (MAID), execute the omnidb command with the appropriate backup session ID:

    ```
    omnidb -session 2013/05/06-1 -media
    ```

    To obtain all backup session IDs for the winfs computer.domain.com:/C computer.domain.com [/C], execute:

    ```
    omnidb -winfs computer.domain.com:/C "computer.domain.com [/C]"
    ```

    To obtain the Disk Agent ID (DAID) and the object name, execute the omnimm command with the relative MAID:

    ```
    omnimm -catalog 9e03110a:3b5ee669:05ac:0001
    ```

2. To perform a disaster recovery, using the Enhanced Automated Disaster Recovery (EADR) method, of the system "computer.company.com" (including its disk image sections E: and F:) using the standalone file device "E:\Devices\file_FR_EADR2.fd" and the Media Agent residing on the system "device.company.com", execute:

    ```
    omniofflr -name "file_FR_EADR2" -dev "E:\Devices\file_FR_EADR2.fd" -policy 1 -
    type 7 -mahost rdevice.company.com -blksize 1024 -maid
    f178b09b:4d6a83ce:0dd8:0001 -position 9:0 -winfs computer.company.com:"/C" "C:"
    -daid 1302093850 -tree / -overwrite -move_busy -rawdisk computer.company.com "
    ```

```
[Disk Image E, F]: computer.company.com" -section \\.\D:=\\.\e: -section
\\.\E:=\\.\f: -daid 1302093851 -report 1 -debug 1-200 dr.txt
```

3. To restore the Data Protector Internal Database with `omniofflr` in the autorecovery mode, using the backup image created in the backup session "2013/04/12-44", using a different backup device attached to the Media Agent system "newmasys.company.com", to save the IDB restore parameters to the file "new_ma_option_file.dat", and to display the IDB recovery file contents, execute:

```
omniofflr -idb -autorecover -changedevhost newmasys.company.com -session
2013/04/12-44 -save new_ma_option_file.dat -logview
```

4. To restore the Data Protector Internal Database with `omniofflr` in the read mode to a different system "newcmsys.company.com", using the IDB restore parameters from the file "new_ma_option_file.dat", and to monitor restore progress details, execute:

```
omniofflr -idb -read new_ma_option_file.dat -verbose -target
newcmsys.company.com
```

5. Consider a device configuration with two drives: drive 1 with path `dev/rtape/tape37_BESTn` and drive 2 with path `dev/rtape/tape_74_BESTn`. If we use drive 1 to restore the IDB by following the "Restoring the IDB Without IDB Recovery File", everything works fine. But if we use drive 2, the restore fails with "Cannot unload exchanger medium" error. So, if you use drive 2, you need to use the following:

```
-dev /dev/rtape/tape74_BESTn 2
```

Here, `2` represents the `drive_index`. The `drive_index` can be 2 or more depending on the number of drives.

SEE ALSO

omnidr(1M), omniiso(1), omnisrdupdate(1M), omniusb(1)

omniresolve(1M)

omniresolve — resolves a filesystem object or a list of filesystem objects and writes the results to the standard output or to a Unicode file
(this command is available on systems with any Data Protector integration component installed)

SYNOPSIS

omniresolve -version|-help

omniresolve {-files *filename* [*filename2* ...]|-inputfile *datafile*} [-verbose] [-unicodefile *outfile*]

DESCRIPTION

The `omniresolve` command reads the filesystem structures locating the physical disks (on Windows) or volumes (on UNIX) on which a filesystem object resides. If the files reside on a logical volume which is a part of a volume group (disk group), all volumes in a volume group are displayed.

You can list the filesystem objects to be resolved either in the CLI (on UNIX and Windows systems) or using a Unicode file (on Windows systems only). The results are written to standard output (on UNIX and Windows systems) or to a Unicode file (on Windows systems only).

OPTIONS

-version

Displays the version of the `omniresolve` command.

-help

Displays the usage synopsis for the `omniresolve` command.

-files *filename* [*filename2*...]

Resolves a list of files separated by spaces and writes the results to the standard output.

-inputfile *datafile*

Resolves all objects listed in *datafile* in and writes the results to the standard output.

Note that on Windows systems, if *datafile* is in the Unicode format, the output is by default written to the file `uniout.dat`. You can redirect the output to a different file by using the -unicode option.

-verbose

Provides a more detailed report (displaying details such as WWNs, LUNs, or LDEVs) using SCSI inquiry on physical disks.

-unicodefile *outfile*

Defines the file to which the output is redirected if the input file is a Unicode file.

NOTES

The resolve process requires root permissions on UNIX systems to get access to the disk device files. Therefore, the SUID flag is set on for `omniresolve`.

EXAMPLES

1. To resolve a list of three files ("system01.dbf", "redo01.log", and "control01.ctl") located in "/opt/oracle10g/oradata/dbname", execute:

   ```
   omniresolve -f '/opt/oracle10g/oradata/dbname/system01.dbf'
   '/opt/oracle10g/oradata/dbname/redo01.log'
   '/opt/oracle10g/oradata/dbname/control01.ctl' -v
   ```

omnirsh(1M)

omnirsh — returns the hostnames of the physical and virtual nodes for the specified cluster hostname, or returns the cell information stored in the cell_info file on the specified cluster
(this command is available on the Data Protector Cell Manager)

SYNOPSIS

omnirsh -version|-help

omnirsh *cluster_hostname* {INFO_CLUS|INFO}

DESCRIPTION

The `omnirsh` command returns the hostnames of the physical and virtual nodes for the specified cluster hostname, together with the flag indicating whether a specific node is a physical node or virtual node. The command can also be used to list the contents of the cluster `cell_info` file.

OPTIONS

-version

    Displays the version of the `omnirsh` command.

-help

    Displays the usage synopsis for the `omnirsh` command.

cluster_hostname

    Sets the hostname of the cluster for this command.

INFO_CLUS

    Lists the hostnames of the physical and virtual nodes for the specified cluster hostname, together with the flag indicating whether a specific node is a physical node or virtual node. Flag value 1 indicates a physical node, whereas flag value 8 indicates a virtual node.

INFO

    Displays the contents of the `cell_info` file for the system specified by the *cluster_hostname* parameter. The file resides on the Cell Manager on server configuration files location in the `cell` directory.

SEE ALSO

omniclus(1M)

omnisetup.sh(1M)

omnisetup.sh — installs or upgrades a Data Protector UNIX Cell Managers, UNIX Installation Servers, UNIX and Mac OS X client systems locally; installs and removes patch bundles.
(this command is available on the Data Protector installation DVD-ROMs for UNIX systems or is provided together with a patch bundle)

SYNOPSIS

omnisetup.sh -version|-help

omnisetup.sh [-source *directory*] [-server *name*] [-install *Component_list*] [-CM] [-IS] [-autopass] [-bundleadd *BundleTag*|-bundlerem *BundleTag*]

Component_list

cc = User Interface

da = Disk Agent

ndmp = NDMP Media Agent

ma = General Media Agent

sap = SAP R/3 Integration

sapdb = SAP MaxDB Integration

emc = EMC Symmetrix Agent

oracle8 = Oracle Integration

sybase = Sybase Integration

ssea = HPE P9000 XP Agent

informix = Informix Integration

lotus = Lotus Integration

db2 = DB2 Integration

smisa = HPE P6000 / HPE 3PAR SMI-S Agent

netapp = NetApp Storage Provider

vepa = Virtual Environment Integration

StoreOnceSoftware = StoreOnce software deduplication

autodr = Automatic Disaster Recovery

docs = English Documentation (Guides, Help)

jpn_ls = Japanese Documentation (Guides, Help)

fra_ls = French Documentation (Guides, Help)

chs_ls = Simplified Chinese Documentation (Guides, Help)

DESCRIPTION

The command first checks if Data Protector is already installed on the system.

New Installation or Re-installation of the same version of Data Protector

If Data Protector is not installed, then the command, depending on the selected options, installs the Cell Manager, Installation Server, or every Data Protector software component specified with the `-install` option. If none of these options are specified, the command issues a prompt for every Data Protector software component supported on the current system OS. Using this prompt, software components supported on the current system OS can be confirmed or rejected for installation, or the execution of the command can be canceled. There is no such prompt if the `-install` option is specified.

Upgrade from an earlier version of Data Protector

To upgrade your cell from the earlier versions of Data Protector, proceed as follows:

- Upgrade the Cell Manager
- Upgrade the Installation Server
- Upgrade the clients

To upgrade all Data Protector components on the system, run `omnisetup.sh` without options. If the Installation Server is installed together with the Cell Manager, or if it is installed without client components, it is upgraded automatically during the Cell Manager upgrade.

If the Installation Server is installed with the client components, it is removed during the Cell Manager upgrade. In this case, a new Installation Server must be installed using the `-IS` option, after the upgrade finishes.

To add a client to the Cell Manager, specify the `-install` option. If the client not residing on the Cell Manager is to be upgraded, the `-install` option does not need to be specified. In this case, the setup selects the same components as were installed on the system before the upgrade without issuing a prompt.

In all cases (new installation, re-installation, or upgrade), the following applies when using this command:

- When using the `-install` option, the software components not supported on the current system OS and mistyped software components are skipped.
- After the client (re-)installation or upgrade is finished, the system is imported to a Data Protector cell if the `-server` option was set, or if the `/etc/opt/omni/client/cell_server` (HP-UX, Solaris, and Linux clients) or the `/usr/omni/config/cell/cell_server` (other UNIX clients and Mac OS X clients) file exists on the system.
- The first time any software component is selected for installation or re-installation, the `core` component is automatically installed (or re-installed). Similarly, the first time any integration software component is selected for installation or re-installation, the `core-integ` component is automatically installed (or re-installed).
- ***HPE AutoPass supported operating systems:*** When the installation or upgrade is started on Cell Manager, you are prompted to install the HPE AutoPass utility (unless the `-autopass` option is specified — if it is, the HPE AutoPass utility is installed or upgraded without issuing a prompt). If

AutoPass is already installed on the system, it is automatically upgraded, if the prompt is confirmed. When Data Protector is uninstalled from the system, the HPE AutoPass utility is neither unregistered nor uninstalled. It must be uninstalled using UNIX utilities, for example `sd`.

If the HPE AutoPass utility is installed in a cluster environment, it must be installed on every node in the cluster.

Installation and removal of Data Protector Patch Bundles

If Data Protector is already installed on your system, you can also install a Data Protector patch bundle (a set of Data Protector patches) on this system by using the `-bundleadd` option. It is not possible to install individual patches from the patch bundle.

You can install a Data Protector patch bundle only on the Installation Server and the Cell Manager. If the installation fails or you stopped it, you can continue with the installation and install the rest of the patches (on Linux systems only), roll installed patches back to the previous patch level, or exit the installation without completing it.

You can remove the Data Protector patch bundle using the `-bundlerem` option. After removing the patch bundle, the base Data Protector release version remains on the system. For details, see the instructions coming with the patch bundle.

OPTIONS

`-version`

Displays the version of the `omnisetup.sh` command.

`-help`

Displays the usage synopsis for the `omnisetup.sh` command.

`-source` *directory*

Sets the location of the Data Protector installation files (DVD-ROM mountpoint). If this option is not specified, the current directory is set as the location of Data Protector installation files.

`-server` *name*

Sets the hostname of the Cell Manager of the cell to which the installed or upgraded client is to be imported after the installation or upgrade.

If this option is not specified, and the `/etc/opt/omni/client/cell_server` (HP-UX, Solaris, and Linux clients) or the `/usr/omni/config/cell_server` (other UNIX clients and Mac OS X clients) file does not exist on the system, the installed or upgraded client is not imported to any cell and has to be imported manually.

`-install` *Component_list*

Sets Data Protector software components that you want to install or upgrade on the current system. If more than one software component is to be installed or upgraded, a listing of software components, delimited by comma (without spaces) must be entered as the argument. If this option is not specified (except for the case when the client not residing on the Cell Manager needs to be upgraded), the command issues a prompt for every Data Protector software component supported on the current system OS; prompting whether to install or upgrade certain Data Protector software component or not.

If the client is to be upgraded, this option does not need to be specified. In this case, the setup selects the same components as were installed on the system before the upgrade without issuing a prompt.

`-CM`

Installs/upgrades the Data Protector Cell Manager.

`-IS`

Installs/upgrades the Data Protector Installation Server with *all* remote installation packages.

Use `-IS` if you install from the DVD-ROM.

Use `-IS1` when installing from the *first* Installation Server CD-ROM and `-IS2` when installing from the *second* Installation Server installation CD-ROM.

If you have copied the `DP_DEPOT` directory from *both CD-ROMs to one directory* on your local disk, use `-IS1 -IS2`.

Note that the Installation Server can be upgraded only after the Cell Manager in the Data Protector cell is upgraded.

`-autopass`

If this option is specified, the HPE AutoPass utility is automatically installed. If HPE AutoPass is already installed on the system, it is automatically upgraded.

This option is to be used only on the Cell Manager operating systems where HPE AutoPass is supported.

`-bundleadd` *BundleTag*

Installs the Data Protector patch bundle (a set of Data Protector patches) on the Cell Manager and the Installation Server.

`-bundlerem` *BundleTag*

Removes the Data Protector patch bundle (a set of Data Protector patches) from the Cell Manager and the Installation Server.

After removing the patch bundle, the base Data Protector release version remains on the system.

NOTES

This command requires that the

- Data Protector UNIX installation DVD-ROM is mounted on the system.
- `DP_DEPOT` and `LOCAL_INSTALL` folders are copied to the disk.

Before running the command make sure that no Data Protector backups or restores are running on the system. The command must be executed using the default POSIX `ksh` or `pdksh` shell.

EXAMPLES

1. To upgrade a system, execute:

```
omnisetup.sh
```

2. To install or re-install the General Media Agent, Disk Agent, HPE P6000 / HPE 3PAR SMI-S Agent, and SAP R/3 Integration software components, execute:

```
omnisetup.sh -install ma,da,smisa,sap
```

3. To install the Cell Manager and Installation Server together with the HPE AutoPass utility, mount the DVD-ROM and execute the following command from the `LOCAL_INSTALL` directory:

```
omnisetup.sh -CM -IS -autopass
```

4. To install the Data Protector patch bundle b701 on the Cell Manager, execute:

```
omnisetup.sh -bundleadd b701
```

SEE ALSO

ob2install(1M), omnigui(5), omniintro(9), omnimigrate.pl(1M), omniusers(1), upgrade_cm_from_evaa (1M)

omnisrdupdate(1M)

omnisrdupdate — updates the System Recovery Data (SRD) file

SYNOPSIS

omnisrdupdate -version|-help

omnisrdupdate [-session *FSSessionID* [*IDBSessionID*]] [-cell *CMName*]

[-host *ClientName*] [-location *Path_1* [-location *Path_2* ...]]

[-asr] [-use_raw_object] [-anyobj]

DESCRIPTION

The omnisrdupdate command is used to update System Recovery Data (SRD) file. An SRD file, which is a text file in the Unicode (UTF-16) format, is generated during CONFIGURATION backup, and saved to the Cell Manager to the SRD files directory.

The SRD filename is identical to the name of the system where it was generated, for example computer.company.com. After the CONFIGURATION backup, the SRD contains only the system information required for system configuration and installation of the operating system needed for disaster recovery. To be able to perform a disaster recovery without a functioning Data Protector Internal Database (IDB), additional information about backup objects and corresponding media must be added to the SRD by running this command. The name of the updated SRD file is recovery.srd.

OPTIONS

-version

Displays the version of the omnisrdupdate command.

-help

Displays the usage synopsis for the omnisrdupdate command.

-session *FSSessionID* [*IDBSessionID*]

Specifies IDs of the backup sessions that serve as the basis for updating the SRD file. All object backed up in the specified sessions and included in the SRD file are used for the update. This option must be specified when the omnisrdupdate command is run interactively, and must be omitted when the omnisrdupdate command is run from a post-exec script. In the latter case, Data Protector automatically obtains the required information from the environment.

If you are updating the SRD file for a Data Protector client, specify the *FSSessionID* argument for the most recent full or incremental filesystem backup session that involves the entire client. If you are updating the ISO image file for the Data Protector Cell Manager, additionally specify the

IDBSessionID argument for an appropriate full or incremental Data Protector Internal Database backup session.

CAUTION: The specified Data Protector Internal Database backup session must be a session that was run after the specified filesystem backup session had completed. To ensure the highest consistency of the included data, the time frame between both sessions' start times should be minimal.

Updating the SRD file succeeds only when all critical backup objects (as specified in the SRD file) were actually backed up in the specified sessions. To view which objects are marked as critical for the SRD update, open the SRD file in a text editor. All critical objects are listed under the -section objects section. Note that the database is represented as "/".

-cell *CMName*

Specifies the Cell Manager to connect to in order to obtain the required information about backup objects and the corresponding media from the IDB.

If this option is omitted, Data Protector automatically obtains the required information from the current environment.

-host *ClientName*

Specifies the system for which the SRD file is to be updated.

If this option is omitted, Data Protector automatically obtains the required information from the current environment.

-location *Path*

Specifies the location where the updated SRD file is saved. A local directory or a network share can be specified. To create several copies of the updated SRD file on different locations, use multiple -location *Path* argument pairs. It is recommended that, in addition to the Cell Manager, the updated SRD file is copied to several safe storage locations as a part of disaster recovery preparation policy. For example, assuming that this storage location is considered safe, you can copy the updated file to the SRD files directory on the Cell Manager.

When the omnisrdupdate command is run from a pre-exec or post-exec script, this option can be omitted. In this case, the omnisrdupdate command updates System Recovery Data internally in the Data Protector session, but does not save it to any SRD file. System Recovery Data updated in such a way can only be used for subsequent processing within the same session.

If you are running the omnisrdupdate command in a pre-exec or post-exec script, do not add a backslash at the end of the path.

-asr

If specified, the ASR archive file (a collection of files required for proper reconfiguration of the replacement disk packed in a single archive file) is downloaded from the Cell Manager and ASR files are extracted and stored to all destinations, specified by the -location option. At least one -location option must be specified otherwise the -asr option is ignored. If the ASR archive file on the Cell Manager does not exist, the omnisrdupdate command fails and the SRD file is not updated.

-use_raw_object

If the specified backup session contains both filesystem and disk image backup objects for the same volume, this option specifies that a disk image backup object should be used. If this option is

not specified, filesystem backup objects have a priority. If only one backup object for the same volume is present in the specified backup session, this option is ignored.

`-anyobj`

Enables you to create a recovery image even if the specified backup session does not contain all client volumes. Note that all host critical volumes must be part of the specified backup session:

- the boot and system volumes
- the Data Protector installation volume
- the volume where the CONFIGURATION object is located
- the Active Directory database volume (in case of an Active Directory controller)
- the quorum volume (in case of a Microsoft Cluster)

NOTES

- The `omnisrdupdate` command is available on Windows and Linux systems only.
- If the BTRFS volume is detected, you get the following **Warning** message:

  Warning: BTRFS volume detected. Make sure that you have included all the BTRFS sub volumes in the specified version.

EXAMPLES

1. To update the SRD file for the Data Protector Cell Manager with the backup object information belonging to the sessions "2013/05/02-5" and "2013/05/02-6", execute:

   `omnisrdupdate -session 2013/05/02-5 2013/05/02-6`

   To obtain the sessions IDs, execute the `omnidb` command with the `-session` option. To obtain the latest session ID, execute:

   `omnidb -session -latest`

2. To update the SRD file for a Data Protector client with the backup object information which belongs to the session "2013/05/02-5" and save the updated SRD file on a diskette as well as to the network share "srdfiles" on the system with the hostname "computer", execute:

   `omnisrdupdate -session 2013/05/02-5 -location A: -location //computer/srdfiles`

3. To update the first diskette from the ASR set for a Data Protector client with the backup object information and ASR files which belong to the session "2013/05/02-5", ensure the first diskette is not write-protected, insert it into the floppy disk drive, and execute:

   `omnisrdupdate -session 2013/05/02-5 -location A: -asr`

SEE ALSO

omnidr(1M), omniiso(1), omniofflr(1M), omniusb(1)

omnistoreapputil(1M)

omnistoreapputil — acts as a user interface to Storage Appliances, such as IAP and VLS
(this command is available on the Data Protector Cell Manager)

SYNOPSIS

omnistoreapputil -version|-help

omnistoreapputil [-check_connection] -hostname *HostName* -port *PortNumber* -user *UserName* -
password *Password* -certificate_name *CertificateName* {-check_iap|-check_vls} -check_vls

omnistoreapputil [-download_certificate] -hostname *HostName* -port *PortNumber* -user
UserName -password *Password* -certificate_name *CertificateName*

omnistoreapputil [-get_iap_client_id] -hostname *HostName* -port *PortNumber* -user *UserName*
-password *Password* -certificate_name *CertificateName* -client *ClientHostNameorIPAddress*

DESCRIPTION

The omnistoreapputil command is used as a user interface for the Storage Appliances, such as IAP
and VLS. It is used to check the connection to the Storage Appliance, download IAP certificates, and
get the IAP Deduplication Agent ID.

The omnistoreapputil command is part of the Cell Manager installation package and is available on
the operating systems supported by the Data Protector Cell Manager.

OPTIONS

-version

   Displays the version of the omnistoreapputil command.

-help

   Displays the usage synopsis for the omnistoreapputil command.

-check_connection

   Checks the connection between Data Protector and the Storage Appliance.

-hostname *HostName*

   Specifies a name of an IAP/VLS client.

-port *PortNumber*

   Sets the TCP/IP port number for the Storage Appliance.

-user *UserName*

   Sets the username that is used by Data Protector to establish the connection to the Storage

Appliance.

`-password` *Password*

Sets the password for the above specified username.

`-certificate_name` *CertificateName*

Specifies the name of the certificate that will be used for connecting to the IAP Server.

`-check_iap`

Specifies that the connection to the IAP Server needs to be checked.

`-check_vls`

Specifies that the connection to the VLS Device needs to be checked.

`-download_certificate`

Downloads the certificate from the IAP that will be used for connecting to the IAP Server.

`-get_iap_client_id`

Gets the IAP Deduplication Agent ID for a specific client.

`-client` *ClientHostName* | *IPAddress*

Specifies the name or the IP address of the client imported into the Data Protector cell.

EXAMPLES

1. To check the connection to the VLS Device, execute:

   ```
   omnistoreapputil -check_connection -hostname client.company.com -port 5988 -
   user Admin -password *** -check_vls
   ```

2. To get the ID of the IAP Server "client.company.com", execute:

   ```
   omnistoreapputil -get_iap_client_id -hostname client.company.com -port 8081 -
   user Admin -password *** -certificate_name New1 -client client.company.com
   ```

SEE ALSO

omnicc(1), uma(1M)

omnisv(1M)

omnisv — starts or stops the Data Protector services or daemons, displays their status, or turns the maintenance mode on or off
 (this command is available on the Data Protector Cell Manager)

SYNOPSIS

omnisv -version|-help

omnisv -status

omnisv {-start|-stop|-start_mon}

omnisv -maintenance [*GracefulTime*|-mom|-stop|-mom_stop]

DESCRIPTION

The `omnisv` command enables you to start or stop the Data Protector services or daemons, display their status, or turn the maintenance mode on and off.

`omnisv` can start or stop the `CRS`, `MMD`, `KMS`, `hpdp-idb`, `hpdp-idb-cp`, `hpdp-as`, and `omniinet` services on the Cell Manager. Note that the `MMD` service can only be started or stopped locally on the Cell Manager with the MMDB.

On UNIX Cell Managers, the `omnisv` command also adds the `omnitrig` process to the cron table and schedules it (on Windows systems, the `omnitrig` command is started by the `CRS` service). You can modify the scheduler granularity by changing the `SchedulerGranularity` global option. By default, the granularity is 15 minutes, but it can be modified to 1 minute.

On Windows Cell Managers, `omnisv` also starts he `Inet` service (the Data Protector Inet program (`/opt/omni/lbin/inet`), which is on UNIX systems started by the system inet daemon when an application tries to connect to the Data Protector port; normally, these daemons are started automatically during the system startup).

The `omnisv` command can also initiate maintenance mode, which prepares your environment for maintenance tasks on Cell Manager that require preventing changes to the Internal Database. For more
 information about the maintenance mode, see the *HPE Data Protector Installation Guide*.

OPTIONS

-version

   Displays the version of the `omnisv` command.

-help

   Displays the usage synopsis for the `omnisv` command.

-status

Displays the status and PID of the services.

`-start`

Starts the Data Protector services or daemons. On UNIX systems, it also adds the `omnitrig` command to the cron table, thus configuring it as a cron job.

`-stop`

Stops the Data Protector services or daemons. On UNIX systems, it also removes the `omnitrig` command from the cron table.

When used with the `-maintenance` option, `-stop` exits the maintenance mode.

`-start_mon`

Waits in loop until the `CRS`, `MMD`, `KMS`, `hpdp-idb`, `hpdp-idb-cp`, and `hpdp-as` services are up and running. If any daemon or service stops, `omnisv` exits with an exit code `1`. Exit code `0` means that all relevant Data Protector daemons/services are up and running, whereas the exit code `1` means that at least one of the relevant Data Protector daemons or services is not running.

`-maintenance`

Initiates the maintenance mode. The optional *GracefulTime* parameter overrides the `MaintenanceModeGracefulTime` global option and specifies the seconds given to the Data Protector services to abort the running sessions.

`-mom`

Initiates the maintenance mode in the entire MoM environment.

`-mom_stop`

Exits the maintenance mode in the MoM environment.

NOTES

On Windows systems, only the users in the Data Protector `admin` group can execute this command. On UNIX systems, only the `root` user can execute this command. It is not possible to start or stop services in clusters using this command.

SEE ALSO

omnicc(1), omnicellinfo(1), omnicheck(1M), omnidlc(1M),

omnitrig(1M)

omnitrig — triggers Data Protector scheduled backups
(this command is available on the Data Protector Cell Manager)

SYNOPSIS

omnitrig -version|-help

omnitrig [-start] [-log]

omnitrig -stop

omnitrig -run_checks

DESCRIPTION

The omnitrig command checks and triggers scheduled backups.

OPTIONS

-version

   Displays the version of the omnitrig command

-help

   Displays the usage synopsis for the omnitrig command

-start

   Adds the omnitrig command to the cron table and schedules it. You can modify the scheduler
   granularity by changing the SchedulerGranularity global option. By default, the granularity is 15
   minutes, but it can be modified to 1 minute.

   On Windows, scheduled backups will be run.

-log

   If this option is specified, the omnitrig will save information about each start of the omnitrig
   command and the backups started by the omnitrig command into the Data Protector log files
   directory to the omnitrig.log file.

-stop

   Removes the omnitrig command from the cron table.

   On Windows systems, scheduled backups will not be run.

-run_checks

   Starts checks for the following Data Protector notifications: IDB Space Low, Not Enough Free
   Media, Unexpected Events, Health Check Failed, IDB Limits, IDB Backup Needed, License Will

Expire, License Warning, and User Check Failed (if configured) every day at 12:30 P.M.

Starts the check for the IDB Reorganization Needed notification every Monday at 12:30 P.M.

You can change the time of these checks or disable them by changing the `DailyCheckTime` global option.

SEE ALSO

omnihealthcheck(1M), omnirpt(1)

omniwl.pl(1M)

`omniwl.pl` - can be used for bulk modification of filesystem backup specifications (datalists). The input data for modifying specifications must be present in a CSV file, with semicolon (;) as the delimiter and double quotes (") as text separator.

This command is available only on the Data Protector Cell Manager.

SYNOPSIS

`omniwl.pl –file <filename> {-datadir <directory_name>| -replace}`

DESCRIPTION

The `omniwl.pl` command allows you to modify multiple backup specifications. The input for modifying specifications must be present in a CSV file, with semicolon (;) as the delimiter. It is recommended that you use a spreadsheet application to create the input document, and subsequently export the document as a CSV file with semicolon as the delimiter.

When you execute the `omniwl.pl` command the specifications that need to be modified and operations that need to be performed on them are read from the CSV input file. The modified specifications can be saved in the original location (`-replace`) or to alternative one (`-datadir`). In case of `-replace` option the original specifications are overwritten.

The input document must be accurate for the correct execution of the command. Ensure that you read and understand the Creating the input document section before you create the input document.

> **Note**: Perl is installed with Data Protector. You do not need to install it separately. Only Perl that is available with Data Protector is supported.

OPTIONS

`-file <file_path_name>`

This option points to the CSV file.

`-datadir <directory_name>`

The directory where the modified specifications are saved. If you use `-datadir` in a MoM environment while editing datalists from multiple Cell Managers in a single input document, then ensure that all of the datalists have unique names.

`-replace`

The original specifications are overwritten.

INPUT DOCUMENT

This section describes the format and rules for creating the input document.

- **Input document format**: This section describes the columns that must be present in the input document.

- **UTF-8 support**: This section describes the supported encoding formats for the CSV file.

- **Usage of wildcard characters and multiple values**: This section describes the columns that support the usage of wildcard characters and multiple values.

- **Rules for the usage of quotes**: This section describes the rules for the usage of quotes in the input document.

Input document format

The input document may contain multiple rows. Each row contains details of one client. Ensure that the column headings are capitalized. Here is a sample input document format from a spreadsheet editor:

CELL MANAGER	SPECIFICATION	CLIENT	MOUNTPOINT	**DESCRIPTION**	**OPTION**	**VALUE**
computer1.company.com	new_backup_1	MOD computer1.company.com	/		tree	(/data1/file1)
computer2.company.com	new_backup_2	MOD computer2.company.com	/	*	exclude	(/data1/file2)

View the CSV format of the above document in the CSV format section.

The input document column description is listed below:

- **Cell Manager**: This column should provide the Cell Manager host name. A single value can be specified. By default, Data Protector uses fully qualified domain names for Cell Manager and clients in the backup specifications. If you specify an IP address or an alias name, the `omniwl.pl` command will not resolve it, and the specified client may not be found.

  Example: computer.company.com

- **Specification**: This column should provide the filesystem backup specification name.

  Example: new_backup_1

- **Client**: This column should provide the operation, followed by the host name of the target client. If a valid operation is not present before the host name, an error is displayed. The operation and host names must be separated by single space. The supported operations are DEL, ADD, and MOD. Note that the operator names must be capitalized. The description of each operation is listed here:

  - If the operation is DEL, the complete client section (mountpoint or full host) is deleted. If the specification remains without a single object, a warning is displayed.

  - If the operation is ADD, then the client section will be created if it does not exist. If the client section is present, an error message is displayed.

  - If the operation is MOD, you can specify the Option that needs to be modified. Ensure that you specify the Value.

- Unless the default value is used for the description field (empty description field) MOD and DEL operations do not require the target client to be a part of the Cell Manager. The ADD operation always requires a valid Cell Manager client.

> **Example**: MOD `computer1.company.com`

- **Mountpoint**: This column should provide the mountpoint name. For full host backups, you must leave this column blank.

> **Examples**: `C:\<mountpoint>` and `/C:\<mountpoint>`

- **Description**: For full host backups, the default description is a full client name. If the field is empty, the default value is taken. For file system backups, the default value for Windows is *mountpoint, and the volume label enclosed in brackets*. The default value for Linux or Unix is the *mountpoint*.

> **Windows example**: `/I: [New Volume]`

> **Unix/Linux example**: `/tmp`

- **Option**: This column should provide the option to be modified. The option and value fields must be specified in JSON format. If DEL is specified at the beginning, the value is ignored, and the option is deleted from the specification. The following options are supported:

JSON format	Data Protector specification format
tree	-trees
exclude	-exclude

- **Value**: This column should provide the values that need to be modified. Multiple values can be specified by enclosing them in parentheses, and separating them using a comma (,). For overwriting the list value, specify the equal sign (=) at the beginning. Duplicate values are not allowed.

An example of the Option-Value pair is shown below:

OPTION	VALUE	OPTION	VALUE
tree	(/tmp,/home/file1,DEL/home/work)	exclude	=(/var/opt)

When the `omniwl.pl` command is executed with the above Option-Value pair, `/tmp` and `/home/file1` is added, and `/home/work` is deleted from the tree section. The exclude section is overwritten with the new value, `/var/opt`.

CSV format

The CSV format of the sample input document is shown here:

```
CELL MANAGER;SPECIFICATION;CLIENT;MOUNTPOINT;DESCRIPTION;OPTION;VALUE

computer1.company.com;new_backup_1;MOD computer1.company.com;/;;tree;(/data1/file1)

computer2.company.com;new_backup_2;MOD computer2.company.com;/;*;exclude;
(/data1/file2)
```

UTF-8 support

UTF-8 and plain ASCII encoding formats are supported for the input document. If Microsoft Excel is used for editing the input document, exporting the document to the CSV format with UTF-8 format may not work as expected. Therefore, Excel is not recommended. Plain text editors are also not recommended, since using and editing CSV files from text editors may result in errors. But, OpenOffice Calc can handle UTF-8 documents correctly. Therefore, it is recommended that you use this tool to convert the input document into the CSV format.

Wildcard and multiple value support

- The Specification, Client, Mountpoint, and Description fields support wildcard characters.
- The Value field supports multiple values.
- If the asterisk symbol (*) is specified, any sequence of characters is matched. For

> **Example**: If ***p*** is specified in the Client field, all the clients that have "p" in their names are considered for the operation.

- If the question mark symbol (?) is specified, any single character will be matched.

> **Example**: If **backup?** is specified in the Specification field, all the specification names that have "backup followed by any character" are considered for the operation.

- In the Specification field, wildcard patterns are matched against the specifications present on the Cell Manager.
- In the Client field,
  - If MOD or DEL is specified, wildcard patterns are matched against the client section in the specification, even if host names and mountpoints are not present in the Cell Manager. Exception to this rule is when the Description field is blank. Then, the default value is extracted from the existing host and mountpoint. Therefore, wildcard patterns in the Mountpoint and Client fields are matched only against existing clients on the Cell Manager and the respective mountpoints. For more details, see the example Deleting the client section using wildcards.

  - If ADD is specified, wildcard patterns are not allowed in the Description field. Wildcard patterns in the Client field are matched against the existing clients on the Cell Manager. Wildcard patterns in the Mountpoint field are matched against the existing mountpoints from clients that belong to the Cell Manager. For more details, see the example Adding the client section using wildcards.

Input document rules

- The comma (,) is used as the separator for list values.
- List values are supported only in the Value field. The list values must be specified in a parentheses. Even if there is only one list value, it must be enclosed in parentheses.
- The DEL command is supported in the Option and Client fields. It must be used before the list values.
- Option and Value fields must not be empty. The only exception is when ADD or DEL is specified before the client name or the option name, or when DEL is specified before the option name.

Rules for the usage of quotes in field values

- For the Specification, Client , Mountpoint, and Description fields, the asterisk and the question mark symbols are reserved symbols. Therefore, these symbols must be enclosed in quotes if they are part of the string.

- In the Value field, strings that contain comma (,) must be double quoted.

- The list values have to be quoted for each list element separately.

- In the Client field, ADD, MOD and DEL followed by a space at beginning are reserved keywords. But, they don't have to be quoted if they are part of the host name, because host names cannot have spaces.

  Examples: `ADD host1.company.com`, `MOD DEL*`, `DEL DEL.company.com`.

- In the Value field, DEL followed by a space at the beginning is a reserved keyword. Therefore, if DEL followed by a space is part of a string, the content has to be double quoted. The same rule applies for the usage of comma in the string.

  Example: If the file name is *DEL File*, the file name should be *"DEL File"* in the CSV file.

Rules for the usage of quotes in CSV

- If a column delimiter or a double quote character is included in the fields, the entire field content has to be double quoted. For instance, *abc;* must be changed to *"abc;"*.

- If any string has double quotes, the double quotes also have to be doubled. For instance, *abc"test"* must be changed to *"abc""test"""*. After doubling all double quotes, the resulting string must be enclosed in double quotes.

In the Value field, list elements have to be quoted according to the CSV rules with comma as the separator. List elements that have characters like comma have to be quoted first.

Example: If the list elements are:

`/tmp`

`/tmp,coma`

`/file"quotes"`

The list elements must be represented as `/tmp,"/tmp,coma","/file""quotes"""`.

EXAMPLES

This section has a sample specification, an input document to modify the specification, and a modified specification after the command is executed.

Executing the command with the input document

Consider that you have named the input document as *Modifications.csv*. To update backup specifications with the changes specified in this file, and to save the new specifications to the directory named *ModifiedSpecifications*, execute the following command:

For Windows

```
cd <Data_Protector_home>/bin

perl omniwl.pl –file C:\Modifications.csv –datadir C:\ModifiedSpecifications
```

For UNIX

```
cd /opt/omni/lbin

omniwl.pl –file /tmp/Modifications.csv –datadir /tmp/ModifiedSpecifications
```

Modifying the `tree` and `exclude` sections of the specification

In this example, the tree and exclude sections of the specification "new_backup_1" are modified.

Original specification

```
FILESYSTEM "/" comp.company.com:"/"

{

-trees

        "/e"

        "/var"

        "/home/work"

-exclude

        "/e/restore"

        "/e/target"

}
```

Input document

CELL MANAGER	SPECIFICATION	CLIENT	MOUNTPOINT	DESCRIPTION	OPTION	VALUE	OPTION	VALUE
comp.company.com	new_backup_1	MOD comp.company.com	/		tree	(/tmp,/home/file1,DEL/home/work)	exclude	=(/var/opt)

When the `omniwl.pl` command is executed with the above input document, `/tmp` and `/home/file1` is added, and `/home/work` is deleted from the tree section. Also, `/var/opt` is added in the exclude section, and the old values are deleted.

Modified specification

```
FILESYSTEM "/" c3po.hermes.si:"/"

{

-trees

        "/e"

        "/var"

        "/tmp"

        "/home/file1"
```

```
-exclude

        "/var/opt"

}
```

Deleting the client section using wildcards

In this example, wildcards are used in the input document to delete client sections in the specification.

Consider a Cell Manager `computer1.company.com` that has 2 clients: `computer2.company.com`, and `computer3.company.com`.

Original specification

This specification has one client section in which the host name does not belong to the Cell Manager.

```
FILESYSTEM "/" server1.company.com:"/"

{

-trees

        "/e"

        "/var"

        "/home/work"

-exclude

        "/e/restore"

        "/e/target"

}
```

Input document

CELL MANAGER	SPECIFICATION	CLIENT	MOUNTPOINT	DESCRIPTION	OPTION	VALUE
computer1.company.com	new_backup_1	DEL server*	/	/		

When the `omniwl.pl` command is executed with the above input document, client sections with host names starting with "server", and having "/" in the mountpoint and description sections, are deleted. The `omniwl.pl` code searches all the host names in the client sections in specified backup specifications, and matches and deletes the client section specified above.

The client needs to be member of the Cell Manager in order to add it to a backup specification. If it is not found in the *cell_info* file, then the following error message is displayed:

```
Can not add object: there is no client hostname matching <hostname from input
document> pattern
```

Adding the client section using wildcards

In this example, wildcards are used in the input document to add new client sections to the specification.

Consider a Cell Manager `computer1.company.com` that has 2 clients: `computer2.company.com`, and `computer3.company.com`.

Input document

CELL MANAGER	SPECIFICATION	CLIENT	MOUNTPOINT	**DESCRIPTION**	**OPTION**	**VALUE**
computer1.company.com	new_backup_1	ADD computer*				

When the `omniwl.pl` command is executed with the above input document, all host names from the Cell Manager clients list that have " computer" are matched, and full host backup client sections are created in the specification. Note that only the host names that belong to the Cell Manager are matched.

Modified specification

HOST " computer1.company.com " computer1.company.com

{

}

HOST " computer2.company.com " computer2.company.com

{

}

sanconf(1M)

sanconf — auto-configures a library, modifies an existing library or drive configuration, or removes drives from a library configuration within a SAN environment
(this command is available on systems with the Data Protector `User Interface` component installed)

SYNOPSIS

sanconf -version|-help

sanconf [-mom] -list[_devices] [*ListFileName*] [-hosts *host_1* [*host_2...*]|-hostsfile *HostsFileName*]

sanconf [-mom] -configure [*ListFileName*] -library *LibrarySerialNumber LibraryName* [*RoboticControlHostName*] [*DeviceTypeNumber*|"*DeviceTypeExtension*"] [-hosts *host_1* [*host_ 2...*]|-hostsfile *HostsFileName*] [-drive_template *DriveTemplateFileName*] [-library_ template *LibraryTemplateFileName*] [-[no_]multipath] [-sanstableaddressing]

sanconf [-mom] -remove_drives *LibraryName* [-hosts *host_1* [*host_2...*]|-hostsfile *HostsFileName*]

sanconf [-mom] -remove_hosts *host_1* [*host_2 host_3 ...*] -library *LibSerNo* [-[no_] multipath]

DESCRIPTION

The `sanconf` command is a utility that provides easier configuration of libraries in SAN environments. 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.

The `sanconf` command can be run on the Data Protector Cell Manager or on Data Protector clients.

You can perform the following tasks using the `sanconf` command:

- Scan the specified Data Protector clients, 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.

All `sanconf` sessions are logged to the `sanconf.log` file in the Data Protector log files directory.

OPTIONS

-version

   Displays the version of the `sanconf` command.

-help

Displays the usage synopsis for the sanconf command.

-mom

Switches sanconf to operate in the MoM mode. This allows listing all devices connected to a MoM environment (see -list) and to configure devices in cells utilizing CMMDB (see - configure, - remove_hosts, -remove).

[-mom] -list[_devices] [*ListFileName*]

This option scans Data Protector clients to gather information on SCSI addresses of drives and robotic controls connected to the clients in the SAN environment and lists the gathered information. The information is uploaded to the Media Management Database on the Cell Manager. When *ListFileName* parameter is specified, the information acquired during the scan of clients is saved to the configuration file, which will be then used for configuring the library.

It is recommended to scan all clients that you want to configure, those that can see the robotics and those that can see the drives.

Note: When the option -mom is specified, sanconf lists clients and devices of all Data Protector cells in the MoM environment, even if they do not use CMMDB.

-hosts *host_1* [*host_2...*]

Specify the -hosts option if you want to limit the sanconf actions only to specified clients. Other clients in the Data Protector cell are skipped.

-hostsfile *HostsFileName*

Specify the -hostsfile option if you want to limit the sanconf actions only to clients specified in the *HostsFileName*. Other clients in the Data Protector cell are skipped. The *HostsFileName* is comprised of an ASCII list of clients, one client per line. It is recommended that all clients are specified in the clients list before you save the scan information to the configuration file.

For multipath devices, the path order is determined by the order in the given list or file.

[-mom] -configure [*ListFileName*]

This option scans, lists, configures, or reconfigures the specified library. Only one library can be configured with each invocation of the command line. If the *ListFileName* option is not specified, the sanconf command will dynamically scan, list, and configure the library. If this option is specified, the scan and data information that was saved to a file during the scan of the specified clients is used to configure the library and scan is not performed. If a client is not scanned, the library will not be configured.

Important: [*RoboticControlHostName*] and -hosts or -hostfile information must be specified during configuration.

Note: When reconfiguring a library, it is recommended that configuration information is first stored in the configuration file in case of configuration failure. It is also recommended that a different filename is used so that the initial configuration can be restored without any complications. sanconf reuses the custom settings when reconfiguring a library.

-library *LibrarySerialNumber LibraryName*

[*RoboticControlHostName*]

[*DeviceTypeNumber* | "*DeviceTypeExtension*"]

Specify the `-library` parameter to configure or reconfigure the specified library. Only one library can be configured with each invocation of the command line. `sanconf` creates only one logical library per physical library in the system and all devices on all specified clients. If the *RoboticControlHostName* parameter is specified, the specified client, which is connected to the specified library, will control the robotics for the library being configured. If this parameter is not specified, the library will be created with robotics on all clients, which are connected to the specified library, within the Data Protector cell, the Cell Manager will be used as a control host. If no library is installed on the Cell Manager in a multipath library, another client will be used as a control host.

If the *ListFileName* parameter is used together with the *RoboticControlHost* but without the `-hosts` or `-hostsfile` option specified, the *RoboticControlHost* parameter will be ignored and a library will be created on all clients which are connected to the library.

When the *RoboticControlHostName* is used with the `-hosts` or `-hostsfile` parameter (option) it limits a library configuration on a specified client. Robotics will be configured on the host which is specified with the *RoboticControlHostName* and on the drives on the host specified with the `-hosts` or `-hostsfile` option. The configuration will be successful only in case that the *RoboticControlHostName* and the *Hosts* have the specified library installed.

In case that you try to configure a library with a robotic control host on a client which does not have a library installed, the configuration will not be successful (parameter `-hosts` is used).

In a MoM environment and with the `-mom` option specified, if the *RoboticControlHostName* parameter is specified without the `-hosts` or `-hostsfile` options, the `sanconf` command will configure a library on all hosts which are connected to it. For example, we have two hosts using the same CMMDB, but they can be on a different Cell Manager. If the hosts are both connected to the same library and only one of them is specified in the *RoboticControlHostName*, `sanconf` will configure two libraries with a robotic control on each host. The same happens in case of a host name which does not have the specified library installed.

If the *ListFileName* parameter is used together with the *RoboticControlHost* but without the `-hosts` or `-hostsfile` option, the *RoboticControlHost* parameter will be ignored and a library will be created on all clients which are listed in the file.

When the *RoboticControlHostName* is used with the `-hosts` or `-hostsfile` parameter it limits a library configuration on a specified client. Robotics will be configured on the host which is specified with the *RoboticControlHostName* and on the drives on the host specified with the `-hosts` or `-hostsfile` option. The configuration will be successful only in case that the *RoboticControlHostName* and the *Hosts* have the specified library installed.

In case that you try to configure a library with a robotic control host on a client which does not have the library installed, the configuration will not be successful.

Additionally, if you try to configure a library on a host which does not use the CMMDB, but its own (local) MMDB, the configuration will fail, whether you try to configure a library which is also installed on clients in the same CMMDB or not.

When the *DeviceTypeNumber* parameter is used, the drives of that type will be configured in the library. When the *DeviceTypeNumber* is not specified, the LTO drive types are used as the default. Only one type number may be specified per library. If you use the "`.DeviceTypeExtension`" parameter instead of the *DeviceTypeNumber* parameter, you can specify the device type extension of the tape device to be configured in the library.

In the following table, *DTN* stands for *DeviceTypeNumber*, and *DTE* stands for *DeviceTypeExtension*.

DTN DTE

1 DDT

2 QIC

3 EXA

4 AIT

5 3480

6 RDSK

7 REGFILE

8 9840

9 TAPE

10 DLT

11 D3

12 3590

13 LTO

14 SDLT

15 VXA

16 DTF

17 9940

18 SAIT

19 3592

When drives in the library are not of the same type as specified, an error is reported.

-drive_template *DriveTemplateFileName*

This option alters the default configuration of each tape device added to the library. You can alter the default configuration of the library only at the initial configuration. After the library is configured, you can no longer change the configuration of the library using the sanconf command.

The *DriveTemplateFileName* must be an ASCII file with one parameter specified per line.

Drive template supports the following parameters:

VERIFY

This parameter corresponds to the CRC Check option in the Data Protector GUI.

CLEANME

This parameter corresponds to the Detect dirty drive option in the Data Protector GUI.

RESCAN

This parameter corresponds to the Rescan option in the Data Protector GUI.

SANSTABLEADDRESSING

This parameter corresponds to the `Automatically discover changed SCSI address` option in the Data Protector GUI.

-library_template *LibraryTemplateFileName*

This option alters the default configuration of the library. You can alter the default configuration of the library only at the initial configuration. After the library is configured, you can no longer change the configuration of the library using the `sanconf` command.

The *LibraryTemplateFileName* must be an ASCII file with one parameter specified per line.

Library template supports the following parameters:

BARCODEREADER

This parameter corresponds to the `Barcode reader support` option in the Data Protector GUI.

BUSYDRIVETOSLOT

This parameter corresponds to the `Busy drive handling: Eject medium` option in the Data Protector GUI.

BUSYDRIVETOMAILSLOT

This parameter corresponds to the `Busy drive handling: Eject medium to mail slot` option in the Data Protector GUI.

SANSTABLEADDRESSING

This parameter corresponds to the `Automatically discover changed SCSI address` option in the Data Protector GUI.

-[no_]multipath

By default or if the `-no_multipath` option is given, `sanconf` does *not* configure multipath devices – a separate logical device will be configured for *each* path.

When reconfiguring a multipath library as a non-multipath library, only one path is created. Multipath drives contained inside a multipath library are not changed, while new drives are created. Only non-multipath drives are modified.

If the `-multipath` option is used, `sanconf` configures all paths pointing to a single physical device as a *single* multipath device.

When reconfiguring a non-multipath library as a multipath library, the library control host is used as the first path. Non-multipath drives are not changed or removed. Instead, new multipath drives are created. Only multipath drives are modified.

-sanstableaddressing

Enables automatic discovery of changed SCSI addresses for the devices being configured.

[-mom] -remove_drives *LibraryName*

This option removes all tape devices in the specified library. If you want to remove drives on specific clients, you can use the `-hosts` *host_1* [*host_2...*] or the `-hostsfile` *HostsFileName* option. This command cannot be used together with the `-multipath` option. Drives that are configured as multipath drives are not removed.

Note: No rescanning is required for this operation.

[-mom] -remove_hosts

All paths containing the specified hosts are removed. However, if the specified hosts cover all paths of the library, no paths are not removed from this library, instead a warning is displayed.

To remove paths only from *multipath* devices, add the `-multipath` option.

To remove paths only from *non-multipath* devices, add the `-no_multipath` option.

To remove paths from *both*, multipath *and* non-multipath devices, execute the command *without* the `-no_multipath` and `-multipath` options.

NOTE: No rescanning is required for this operation.

NOTES

The `sanconf` command is available on Windows, HP-UX, Solaris, and Linux systems only.

All drives created with the `sanconf` command are named automatically. Drive names must not be changed manually because the reconfiguration will not work. You must follow the drive naming convention.

- For *non-multipath* devices:

  `libname_index_host`

  `libname_index_busindex_host`

  The `busindex` number is used only if there is more than one path for the drive.

- For *multipath* devices:

  `libname_index`

EXAMPLES

The following examples illustrate how the `sanconf` command works.

1. To scan host(s) for robotic control(s) and tape device(s) and create a file that will be used by `sanconf -configure`, execute:

   `sanconf -list device.list`

   This will display the serial number for any library discovered in the SUMMARY REPORT.

2. To scan and configure a library using the library serial number and the library name, on all clients on which the library is installed and which use CMMDB, execute:

   `sanconf -mom -configure -library US9LS01033 SAN_STORE`

   Clients on which the library is installed and which use a local MMDB are skipped.

3. To scan the specified clients and then create a logical library named "SAN_STORE" with robotics configured on client "host33" and drives for that library configured on clients "host01", "host02" and "host03", execute:

   `sanconf -configure -library MPC0100013 SAN_STORE host33 -hosts host01 host02 host03`

   A device type is .lto. An extension parameter does not need to be added.

4. To scan the SAN environment for the configuration information on the specified clients "host01", "host02", "host03", and "host33" which use CMMDB, and save this information is into the

mySAN.cfg file, execute:

`sanconf -mom -list_devices mySAN.cfg -hosts host01 host02 host03 host33`

5. To use information stored in the mySAN.cfg file and create a logical library named "SAN_STORE" with robotics configured on client host33 and drives for the library configured on clients "host01", "host02", and "host03", execute:

`sanconf -configure mySAN.cfg -library MPC0100013 SAN_STORE host33 -hosts host01 host02 host03`

6. To scan all clients in the cell and then create a logical library named "SAN_STORE" on client "host33" with the parameters specified in the files DriveTemplate.txt and LibraryTemplate.txt, execute:

`sanconf -configure -library MPC0100013 SAN_STORE host33 -hosts host33 -drive_ template DriveTemplate.txt -library_template LibraryTemplate.txt`

7. To configure a tape library with the default tape device and library settings using the "device.list" file created by the example above, execute:

`sanconf -configure device.list -library MPC0220423 myLib1`

8. To configure a library with a specific drive type, execute:

`sanconf -configure -library MPC0100013 SAN_STORE host33 ".9840" -hosts host01 host02`

This command creates a library named "SAN_STORE" with robotics configured on client "host33" and STK drives configured on clients "host01" and "host02". The drives are named as follows:

SAN_STORE_1_host01

SAN_STORE_1_host02

SAN_STORE_2_host01

SAN_STORE_2_host02

9. To configure three libraries using the configuration options contained in the library template "myway", execute:

`sanconf -configure -library US9LS02033 mylib5 -library_template myway`

`sanconf -configure -library US9LS02034 mylib6 -library_template myway`

`sanconf -configure -library US9LS02035 mylib7 -library_template myway`

10. To configure a multipath LTO library with the serial number "LLL1", named "Library1", and connected to client "host1", execute:

`sanconf -configure -library LLL1 Library1 host1 ".LTO" -multipath`

11. To configure a multipath LTO library with the serial number "LLL1", named "Library1", and connected to client "host1" and "host2", execute:

`sanconf -configure -library LLL1 Library1 host1 ".LTO" -hosts "host1" "host2" - multipath`

This will configure a library and drives with multipath option checked and configured paths on host1 and host2.

12. To update an already configured library with the configuration information for new hosts or tape devices, execute:

`sanconf -configure -library US9LS01023 mylib2`

13. To reconfigure an already configured library after adding a new host "myhost" to a Data Protector cell, execute:

    `sanconf -configure -library US9LS01033 mylib2 -hosts myhost`

    This will scan and configure only the new host.

14. In a MoM environment, to reconfigure an already configured library on "host02" after adding a new host "myhost" to a Data Protector cell, execute:

    `sanconf -mom -configure -library US9LS01033 mylib2 host02 -hosts myhost`

    This will add drives from the host "myhost" to the library "mylib2" which is configured on the host "host2".

15. To configure only LTO Ultrium tape drives and add them into the library "myLTOlib", execute:

    `sanconf -list device.list`

    `sanconf -configure device.list -library MPC0230031 myLTOlib "libraryhost" ".LTO"`

16. To reconfigure a non-multipath library named "SAN_STORE" with serial number "MPC0100013" to a multipath library using the `-hosts` option, when new clients "host04" and "host05" are added to the cell, execute:

    `sanconf -configure -library MPC0100013 SAN_STORE host33 -hosts host04 host05 -multipath`

17. To delete all tape drives configured in the library "mylib2" related to the clients "host04" and "host05", execute:

    `sanconf -remove_drives mylib2 -hosts host04 host05`

18. To delete all tape drives configured in the library "mylib2", execute:

    `sanconf -remove_drives mylib2`

19. To remove all paths in the multipath library named "SAN_STORE" with serial MPC0230031 that are configured on clients "host04" and "host05", execute:

    `sanconf -remove_hosts -hosts host04 host05 -library MPC0230031 -multipath`

SEE ALSO

omniamo(1), omnib2dinfo(1M), omnidownload(1), omnimcopy(1), omniminit(1), omnimlist(1), omnimm (1), omnimnt(1), omnimver(1), omniupload(1), uma(1M)

uma(1M)

uma — controls the robotics of SCSI compliant autochangers
(this command is available on systems with the Data Protector `General Media Agent` or `NDMP Media Agent` component installed)

SYNOPSIS

uma -version|-help

uma [-policy *LogicalDevicePolicy*] -ioctl *deviceFile* [-interface {0|1}] [-tty] [-barcode] [-device *deviceFile_1* [*deviceFile_n*] -type *DeviceType*] [-ddt *NDMP_server_name NDMP_port_number backup_type username password*]

uma command line-interface commands:

help

inq

init

addr

offl *driveID*

sense

pos *slot*

move *source_slot destination_slot* [0|1]

stat [{*slot*|*drive*|*transport_element*|*mail_slot*}]

modesense [*page*]

test

bye|exit|quit

doorlock [0|1]

enter *slot*

eject *slot*

DESCRIPTION

The uma program is a standalone utility program which can be used to control the robotics of most SCSI compliant autochangers, also those which are not directly supported by Data Protector. It implements a shell-like user command interface and can be used both interactively and in batch mode.

uma is packaged and installed as part of a Data Protector Media Agent fileset. If you have received uma as a standalone program or if you run it on a system where Data Protector has not been installed, the

uma command is fully functional and behave as documented, but it is probably not able to locate and use Data Protector NLS message catalog.

On HP-UX and Solaris systems, the NLS message catalog is located in the `/opt/omni/lib/nls/C/` directory.

On other UNIX systems, the NLS message catalog is located in the `/usr/omni/lib/nls/C/` directory.

On Windows systems, the NLS message catalog is located in the *Data_Protector_home*\bin directory.

uma can be started both interactively or in batch mode. The only obligatory option is the pathname of the device file (UNIX systems) or the SCSI address (Windows systems) that controls the robotics of the target autochanger (the `-ioclt` option). For backup devices with library robotics connected to an NDMP Server (to a supported NAS device), the `-interface` and the `-ddt` options must also be specified.

For your convenience, the uma command allows you to specify symbolic instead of physical element addresses (slot IDs). Whenever you need to refer to the 1st drive of the autochanger, you can specify either the physical address '128' or the more convenient, symbolic 'D1'. The output of the addr command reflects this addressing convention.

OPTIONS

`-version`

 Displays the version of the uma command.

`-help`

 Displays the usage synopsis for the uma command.

`-policy `*LogicalDevicePolicy*

 Specifies the backup device policy ID. Policy can be defined as 6 (external control), 8 (Grau DAS exchanger library), 9 (STK Silo medium library), 10 (SCSI Library), or 13 (VLS).

 The default value for the `-policy` option is 10.

`-ioctl `*deviceFile*

 Specifies the pathname of the device file (UNIX systems) or the robotics SCSI address (Windows systems) that controls the robotics of the target autochanger.

`-interface {0 | 1}`

 Sets the type of SCSI interface used to access library robotics. This option is to be used only with backup devices with library robotics connected to an NDMP Server. 0 sets the standard SCSI interface (the default value). 1 sets the NDMP protocol interface and must be specified for backup devices with library robotics connected to an NDMP Server. The default value is 0.

`-tty`

 Forces the uma command to enter the command line interface mode or to read from script. This option is obligatory on UNIX systems. On Windows systems, this option is not to be used; the command line interface mode is invoked automatically.

`-barcode`

If this variable is set, the uma command's stat command displays also the barcode information for each medium.

-device *deviceFile_1* [*deviceFile_n...*]

Specifies the device file (UNIX systems) or the SCSI address (Windows systems) of one or more autochanger drives. For a multi-drive autochanger, you must specify a list of device files/SCSI addresses which correspond to the autochanger's drives in ascending order. The drives have to be known to uma in order for the offl command to work. This option is only to be used together with -type option.

-type *DeviceType*

Specifies the media type for the media in the device specified by the -device option. Media type numbers are defined as media class in the scsitab file. For location, see the help index "support of new devices".

-ddt *NDMP_server_name NDMP_port_number backup_type username password*

This option is mandatory for backup devices with library robotics connected to an NDMP Server (to a supported NAS device). It specifies the NDMP Server name, port number used by Data Protector to connect to the NDMP Server and username and password used by Data Protector to connect to the NDMP server. The *backup_type* parameter has to be set to dump.

uma command-line interface commands:

help

Displays the usage synopsis for the uma command.

inq

Performs a SCSI Inquiry operation on the device file/SCSI address specified with the -ioctl option. It returns the device's type, vendor ID, product ID and firmware revision number.

init

Performs a SCSI 'initialize element status' operation, which (if applied to an autochanger robotic device) forces the autochanger to reset its internal state and perform an inventory of its repository. This command should not be used if another process is accessing the autochanger at the same time, as the effects are unpredictable.

addr

Queries and displays the autochanger's element assignment page. Each addressable item inside the autochanger mechanism (drive, repository slot, robotic arm, import/export slot) has a unique integer number (slot ID) which can be used to address this specific item.

As the element assignment differs among different autochangers, the software, which is to control the movement of media inside the autochanger, must find out and use these numbers to perform move, pos and stat operations.

offl *driveID*

This command can be used only if at least one drive was specified using the -device option. If a medium is loaded in the specified drive, it will eject the medium just as if an UNIX mt offl command was specified. The mandatory argument is a symbolic drive ID (that is, D3 for the 3rd drive == the 3rd device file specified with the -dev option). If the drive specified is not defined by the -device option, then the last drive defined by the -device option will be used.

The `offl` command can fail with a message: `"No such device or address"` if it is issued immediately after the `move` command since it takes a certain time after the `move` command for the drive to be online. For more information, see the `move` command.

sense

Read the device's sense data and dump them in a hex- dump format.

pos *slot*

Positions the autochanger transport mechanism in front of the specified slot. This operation is only meaningful if the specified slot refers to an import/export, data drive or repository element. The actual meaning of this operation may differ among different autochanger models. This command is generally not required, but is provided for testing purposes and convenience. Both physical as well as symbolic slot addressing may be used.

move *source_slot destination_slot* [0 | 1]

Moves a medium from a source slot into a destination slot. This command has two mandatory arguments, the source and destination slot IDs (address numbers, as reported by the `addr` command described above) and an optional numeric Boolean argument which can be used to instruct the robotics to flip the medium before inserting it into the destination slot. By default (if no flipping argument is specified), flipping is disabled.

Note that when `move` command is issued to move a tape into a drive, it takes a certain time (around 30 seconds) for the drive to become online, because tape load and calibration/selftest have to be performed. The command prompt however, returns immediately after the command is issued.

Note: Flipping is supported only for double-sided optical media. For tapes, the effect of the flip command is not defined.

NOTE: Most autochanger do not allow you to `move` a tape from a drive to a repository location if the tape has not been dismounted and ejected by the drive. You might want to use the `offl` command on the drive device file/SCSI address to put the drive offline before executing the `move` command.

stat [{*slot* | *drive* | *transport_element* | *mail_slot*}]

Queries the device for information about the state of each of its addressable elements. The output of this command is a table of physical and symbolic element IDs and their states, indicating which elements are free (Empty) and occupied (Full).

Additionally, if barcode support is available and enabled, the barcode for each medium is displayed.

The `uma` command recognizes one specific environment variable which can be used to enable barcode support for autochangers which are equipped with barcode reading hardware. By default, `uma` barcode support is disabled. It can be enabled by exporting/setting the `OB2BARCODE=1` environment variable before starting the command or by using the `-barcode` option.

The `stat` command can be used to query the status of a specific slot (that is, 'stat 290' or 'stat S35') or a related group of slots (that is, 'stat D' will query all drives, 'stat S' will query all repository slots, and so on).

If no additional arguments are specified, the stat command will query and print the status information for all slot IDs it can address.

modesense [*page*]

Reads the vendor specific data and unit settings from the unit and displays them. You can limit the display only to certain pages by using the page parameter. If the page parameter is not specified, all

pages are displayed.

test

Checks if the unit is ready. If the unit is not used by any process, then the unit is ready. If it is, however, used either by the robotics, backup or restore processes then it is not ready.

exit|bye|quit

Exits the command mode.

doorlock [0|1]

If the input parameter is 1, this command locks the library mail slot door; if it is 0, it unlocks it.

-enter slot

Enters media into a specified library slot.

-eject slot

Ejects media from a specified library slot.

NOTES

Do not use the uma utility while Data Protector backup or restore is running. On UNIX systems the -tty option is obligatory. On Windows systems it is not used.

EXAMPLES

1. uma can be started both interactively or in batch mode. The only option which needs to be specified (except for backup devices with library robotics connected to an NDMP Server) is the pathname of the device file which controls the robotics of the target autochthons:

   ```
   UMA -ioctl /dev/spt/sctl0

   *** PROGRAM: UMA VERSION: HPE Data Protector 9.08

   *** Copyright (C) 1999 Hewlett Packard Enterprise Company

   *** License is restricted for use with licensed

   *** HPE Data Protector products.

   /dev/spt/sctl0> exit
   ```

2. To start uma for a backup device with the library robotics connected to the NDMP Server with the robotics SCSI address "mc2", the NDMP Server hostname "ndmpserver", the port number used by Data Protector to connect to the NDMP Server "10000", and username and password of the user used by Data Protector to connect to the NDMP Server "user password", enter the following command:

   ```
   UMA -ioctl mc2 -interface 1 -ddt ndmpserver 10000 dump user password
   ```

3. To let uma execute a batch script of its own commands, simply redirect its stdin to a file containing a list of uma commands separated with newlines:

   ```
   cat >/tmp/cmdFile

   inq
   ```

```
addrstat
```

```
<ctrl-D>
```

```
uma -ioctl /dev/spt/sctl0 </tmp/cmdFile >/tmp/outFile
```

4. The following output is obtained by executing the `addr` command on the UNIX device file referring to an ACL 4/52 DLT autochanger:

```
/dev/spt/sctl0>addr Element Addresses (T=Transport, X=Im/Export, D=Drive,
S=Storage):
```

```
Transport: 1 .. 1 (T1 .. T1)
```

```
Im/Export: 64 .. 67 (X1 .. X4)
```

```
Data Drive(s): 128 .. 131 (D1 .. D4)
```

```
Repository: 256 .. 303 (S1 .. S48)
```

The numbers returned by the `addr` command are the physical element addresses of different elements within the autochanger - that is, element address "256" would correspond to the first repository slot, element address "65" would correspond to the location of the second data drive, and so on.

5. To start `uma` for the Grau DAS exchanger library with the robotics device file "grauamu", execute:

```
uma -pol 8 -ioctl grauamu
```

SEE ALSO

omniamo(1), omnib2dinfo(1M), omnidownload(1), omnimcopy(1), omniminit(1), omnimlist(1), omnimm (1), omnimnt(1), omnimver(1), omniupload(1), sanconf(1M)

upgrade_cm_from_evaa(1M)

upgrade_cm_from_evaa — upgrades the EVADB entries created by the HPE EVA Agent (legacy) to the SMISDB entries created by the HPE P6000 / HPE 3PAR SMI-S Agent
(this command is available on the Data Protector Cell Manager)

SYNOPSIS

upgrade_cm_from_evaa -version|-help

upgrade_cm_from_evaa [-preview]

DESCRIPTION

The upgrade_cm_from_evaa command needs to be executed on any Cell Manager after completing the Cell Manager upgrade from the EVA Agent (legacy) to the HPE P6000 / HPE 3PAR SMI-S Agent. It upgrades the following:

- EVADB login entries into SMISDB login entries
- EVADB disk group rules into SMISDB disk group rules
- EVAA backup specifications into SMISA backup specifications
- EVADB backup sessions into SMISA backup sessions

OPTIONS

-version

   Displays the version of the upgrade_cm_from_evaa command.

-help

   Displays the usage synopsis for the upgrade_cm_from_evaa command.

-preview

   Gives a preview of what happens when the command is run.

EXAMPLES

The following examples illustrate how to use the upgrade_cm_from_evaa command.

1. To display the version information, execute:

   upgrade_cm_from_evaa -version

2. To preview what happens when the upgrade from the EVA Agent (legacy) to the HPE P6000 / HPE 3PAR SMI-S Agent is run on the Cell Manager, execute:

   upgrade_cm_from_evaa -preview

This command displays a list of actions that will be taken when the upgrade is run but it does not update the EVADB entries.

SEE ALSO

ob2install(1M), omnidbsmis(1), omnigui(5), omniintro(9), omnimigrate.pl(1M), omnisetup.sh(1M), omniusers(1)

util_cmd(1M)

util_cmd — sets, retrieves, or lists the parameters stored in the Data Protector Oracle, MySQL, SAP R/3, SAP MaxDB, Microsoft Exchange Server 2010/2013, Informix, and Sybase configuration files. In addition, it encodes passwords.
(this command is available on systems with any Data Protector component installed)

SYNOPSIS

util_cmd -version|-help

util_cmd -getconf[ig] {Oracle8|SAP|SAPDB|Informix|Sybase} *instance* [-local *filename*]

util_cmd -getopt[ion] [{Oracle8|SAP|SAPDB|Informix|Sybase} *instance*] *option_name* [-sub [list] *sublist_name*] [-local *filename*]

util_cmd -putopt[ion] [{Oracle8|-integ MySQL|SAP|SAPDB|Informix|Sybase} *instance*] *option_name* [*option_value*] [-sub[list] *sublist_name*] [-local *filename*]

util_cmd -encode *Password*

util_cmd -setomnirc *hostname name* [*value*]

util_cmd -getomnirc *hostname name*

util_cmd -delomnirc *hostname name*

DESCRIPTION

The util_cmd command is used to set, retrieve, or list the parameters stored in the Data Protector Oracle, SAP R/3, SAP MaxDB, Microsoft Exchange Server 2010/2013, Informix, and Sybase configuration files. In addition, it can be used to encode passwords.

Data Protector stores the integration parameters on the Cell Manager in the directory *Data_Protector_program_data*\Config\Server\Integ\Config*integration_name* (Windows systems) or /etc/opt/omni/server/integ/config/*integration_name* (UNIX systems).

ORACLE

For each configured Oracle database, the following configuration files are created:

- Target database configuration file: *client_name*%[*DB_NAME* | *INSTANCE_NAME*]

  For Oracle Data Guard, *client_name* is *primary_hostname* or *secondary_hostname*

  The parameters stored in the target database configuration file are:

  — Oracle home directory

  — encoded connection strings to the target database, recovery catalog, and standby database

  — variables, which are exported when you start a session using the Data Protector GUI or CLI

  OB2_RMAN_COMMAND_TIMEOUT (environment variable)

This variable is applicable when Data Protector tries to connect to a target or catalog database. It specifies how long (in seconds) Data Protector waits for RMAN to respond that the connection succeeded. If RMAN does not respond within the specified time, Data Protector aborts the session. Default: `300` s.

OB2_SQLP_SCRIPT_TIMEOUT (environment variable)

This variable is applicable when Data Protector issues an SQL*Plus query. It specifies how long Data Protector waits for SQL*Plus to respond that the query completed successfully. If SQL*Plus does not respond within the specified time, Data Protector aborts the session. Default: `300` s.

SBT_LIBRARY

Specifies which Data Protector MML should be used by RMAN, in case you want to override the default Data Protector selection.

- Global database configuration file: `client_name%_OB2_GLOBAL`

The parameters stored in the global configuration file are:

— instance list (all Oracle instances on the Oracle server)

— variables that need to be exported prior to starting a backup and which affect every Oracle instance on the Oracle server.

- In case of zero downtime backup, backup method configuration file: `zdb_methodORACLE_DBID`

- In case of zero downtime backup, for backup set method, the file: `client_name%initDB_NAME_bckp.ora`

SAP R/3

The SAP R/3 parameters stored are:

- Oracle home directory

- encoded connection string to the target database

- BRTOOLS home directory

- variables, which are exported when you start a session using the Data Protector GUI or CLI:

ORA_NLS_CHARACTERSET

After upgrading a Data Protector __DP_V55__ SAP R/3 client to the latest version of Data Protector, set this variable to the encoding used by the Oracle database.

OB2_MIRROR_COMP

This variable is applicable for ZDB sessions that use the SPLITINT functionality (`-t {online_mirror | offline_mirror}`). Set this variable to 1 if you want BRBACKUP to be started on the backup system and not on the application system. By default, BRBACKUP is started on the application system.

SBT_LIBRARY

Specifies which Data Protector MML should be used by RMAN, in case you want to override the default Data Protector selection.

- concurrency number and balancing (for each backup specification) and number of channels for RMAN backup

- speed parameters (time needed for a specific file to back up - in seconds)

- manual balancing parameters

SAP MaxDB

The SAP MaxDB parameters stored are:

- Username of the SAP MaxDB database user
- Password of the SAP MaxDB database user
- SAP MaxDB version
- SAP MaxDB independent program path parameter that was specified during the installation of SAP MaxDB Server
- Data Protector SAP MaxDB integration related environment variables

INFORMIX SERVER

The Informix parameters stored are:

- Informix Server home directory
- pathname of the `sqlhosts` file
- name of the Informix instance `ONCONFIG` file

SYBASE

The Sybase parameters stored are:

- Sybase home directory
- pathname for the `isql` command
- Sybase backup operator username and password
- name of the Sybase *SYBASE_ASE* directory (Sybase 12.x only)
- name of the Sybase *SYBASE_OCS* directory (Sybase 12.x only)
- environment variables

RETURN VALUES

The `util_cmd` command displays a short status message after each operation (written to the standard error):

- `Configuration read/write operation successful.`

  This message is displayed when all the requested operations have been completed successfully.

- `Configuration option/file not found.`

  This message appears when either an option with the specified name does not exist in the configuration, or the file specified as the `-local` parameter does not exist.

- `Configuration read/write operation failed.`

  This message is displayed if any fatal errors occurred, for example: the Cell Manager is unavailable, the Data Protector configuration file for a specific integration is missing on the Cell Manager, and so on.

OPTIONS

`-version`

Displays the version of the `util_cmd` command.

`-help`

Displays the usage synopsis for the `util_cmd` command.

`-getconf[ig]` *integration instance*

Lists the Data Protector configuration files parameters for the specified integration and instance to the standard output, unless the `-local` option is specified.

`-getopt[ion]` [*integration instance*] *option_name*

Retrieves the parameter (specified by the option_name) and its value from one of Data Protector configuration files and writes it to the standard output, unless the `-local` option is specified.

`-putopt[ion]` [*integration instance*] *option_name* [*option_value]*

Sets the specified parameter (specified by the option_name) and (optionally) its value to the Data Protector configuration files, unless the `-local` option is used.

To remove a value of a parameter, specify the *option_name*, without the *option_value*. However, if the option is in a sublist, you must specify an empty ("") *option_value* to remove a value.

`-sublist` *SublistName*

Specifies the sublist in the configuration file in which a parameter is written to or taken from.

`-local` *FileName*

If the `-local` option is used with the `-getconf` option, the command output is written to the file with the filename specified by the `-local` option. If the `-local` option is used with the `-getopt` option, the parameter and its value is taken from the file with the filename specified by the `-local` option. If the `-local` option is used with the `-putopt` option, the parameter and its value is written to the file with the filename specified by the `-local` option.

`-encode` *Password*

Returns the encoded form of the specified password.

`-setomnirc` *hostname name* [*value*]

This option sets the environment variable name with a value in the `omnirc` file on the host specified as hostname. If the environment variable is present in the `omnirc` file, its values gets updated, otherwise it gets created.

The following parameters are passed in the `-setomnirc` command:

- `hostname` - Specifies the place where the `omnirc` file is located.
- `name` - Specifies the environment variable name.
- `value` - Specifies the value of the environment variable.

The access to remote `omnirc` may fail due to the following reasons:

- When strict security is enabled. In this case, you need to start the `util_cmd` on the cell server.
- When the target host is on another domain due to secure communication issues. This is an issue with the Encrypted Control Communication (ECC) and domain configurations.

Note: The `omnirc` file does not support multiple occurrences of the same variables as it leads to errors. All these limitations apply to `getomnirc` and `delomnirc` commands too.

-getomnirc *hostname name*

This option reads the value of the environment variable from the `omnirc` file on the hostname.

-delomnirc *hostname name*

This option deletes the environment variable from the `omnirc` file on the hostname.

EXAMPLES

The following examples illustrate how the `util_cmd` command works.

1. To set the Data Protector "OB2OPTS" parameter for the Oracle instance "ICE", execute:

   ```
   util_cmd -putopt Oracle8 ICE OB2OPTS "-debug 1-200 INSTANCE.txt" -sublist
   Environment
   ```

2. To set the Data Protector "OB2OPTS" parameter for the SAP R/3 instance "ICE", execute the following command on the Data Protector SAP R/3 client:

   ```
   util_cmd -putopt SAP ICE OB2OPTS '-debug 1-200 INSTANCE.txt' -sublist
   Environment
   ```

3. To set the "BR_TRACE" parameter for the SAP R/3 instance "ICE" to value "10" in the "Environment" sublist, execute the following commands on the Data Protector SAP R/3 client:

   ```
   util_cmd -putopt SAP ICE BR_TRACE "'10'" -sublist Environment
   ```

4. To list the Data Protector configuration file parameters for the Oracle instance "ICE", execute:

   ```
   util_cmd -getconf Oracle8 ICE
   ```

5. To retrieve the value of the "OB2OPTS" parameter for the Oracle instance "ICE", execute:

   ```
   util_cmd -getopt Oracle8 ICE OB2OPTS -sublist Environment
   ```

6. To remove the value of the "OB2OPTS" parameter for the SAP R/3 instance "ICE", execute the following command on the Data Protector SAP R/3 client:

   ```
   util_cmd -putopt SAP ICE OB2PTS "" -sublist Environment
   ```

7. To get the encoded form of the password "BlueMoon", execute:

   ```
   util_cmd -encode BlueMoon
   ```

8. To set the environment variable "OB2_RMAN_COMMAND_TIMEOUT" to "100" seconds for the Oracle database "INST2", execute:

   ```
   util_cmd -putopt Oracle8 INST2 OB2_RMAN_COMMAND_TIMEOUT 100 -sublist
   Environment
   ```

9. To set the environment variable "TEST1" with value "10" to omnirc file on the hostname "Win9", execute:

   ```
   util_cmd -setomnirc Win9 TEST1 10
   ```

10. To get the value of the environment variable "TEST1" from omnirc file on hostname "Win9", execute:

    ```
    util_cmd -getomnirc Win9 TEST1
    ```

11. To delete the environment variable "TEST1" from omnirc file on hostname "Win9", execute:

    ```
    util_cmd -delomnirc Win9 TEST1
    ```

SEE ALSO

omnib(1), omniintconfig.pl(1M), util_oracle8.pl(1M), vepa_util.exe(1M)

util_oracle8.pl(1M)

util_oracle8.pl — configures an Oracle database and prepares the environment for backup, and checks the configuration of an Oracle database
(this command is available on systems with the Data Protector `Oracle Integration` component installed)

SYNOPSIS

util_oracle8.pl -version|-help

util_oracle8.pl -chkconf -dbname *DB_NAME* [-client *CLIENT_NAME*]

util_oracle8.pl -chkconf_smb -dbname *DB_NAME* [-bkphost *BACKUP_SYSTEM*] [-client *CLIENT_NAME*]

util_oracle8.pl -chkconf_ir -dbname *DB_NAME* [-client *CLIENT_NAME*]

util_oracle8.pl -config -dbname *DB_NAME* -orahome *ORACLE_HOME PRIMARY_DB_LOGIN* [*CATALOG_DB_LOGIN*] [*STANDBY_DB_LOGIN*] [*ZDB_OPTIONS*] [*ASM_OPTIONS*] [-client *CLIENT_NAME*]

PRIMARY_DB_LOGIN

 -prmuser *PRIMARY_USERNAME*

 -prmpasswd *PRIMARY_PASSWORD*

 -prmservice *PRIMARY_NET_SERVICE_NAME_1*[,*PRIMARY_NET_SERVICE_NAME_2* ...]

CATALOG_DB_LOGIN

 -rcuser *CATALOG_USERNAME*

 -rcpasswd *CATALOG_PASSWORD*

 -rcservice *CATALOG_NET_SERVICE_NAME*

STANDBY_DB_LOGIN

 -stbuser *STANDBY_USERNAME*

 -stbpasswd *STANDBY_PASSWORD*

 -stbservice *STANDBY_NET_SERVICE_NAME_1*[,*STANDBY_NET_SERVICE_NAME_2* ...]

ZDB_OPTIONS

 -zdb_method {*PROXY*|*BACKUP_SET*}

 [-ctlcp_location *BACKUP_CONTROL_FILE_COPY_LOCATION*]

 [-pfile *PARAMETER_FILE*]

 [-bkphost *BACKUP_SYSTEM*]

ASM_OPTIONS

 [-asmhome *ASM_HOME*]

```
[-asmuser ASM_USERNAME -asmpasswd ASM_PASSWORD -asmservice ASM_NET_SERVICE_NAME_1
[,ASM_NET_SERVICE_NAME_2 ...]]
```

DESCRIPTION

Use the `util_oracle8.pl` command to configure an Oracle database and prepare the environment for backup, and to check the configuration of the database.

To back up a standby database, you must provide the *STANDBY_DB_LOGIN* information. For standby database backup, a recovery catalog must be used. Therefore, you must also provide the *CATALOG_DB_LOGIN* information.

To configure an Oracle database for ZDB, you must provide the *ZDB_OPTIONS* information. If your ZDB method is backup set, you must also provide the *BACKUP_SYSTEM* information.

The *ASM_OPTIONS* options are needed for instant recovery in Oracle Server configurations that use Automatic Storage Management (ASM).

On Windows systems, you must use the `perl` command to execute `util_oracle8.pl`. An example of the command line is `perl util_oracle8.pl -help`.

On HPE OpenVMS systems, you must omit the command's file extension to execute the command. An example of the command line is `util_oracle8 -help`.

OPTIONS

`-version`

Displays the version of the `util_oracle8.pl` command.

`-help`

Displays the usage synopsis for the `util_oracle8.pl` command.

`-client` *CLIENT_NAME*

Name of the Oracle Server system with the database to be configured. It must be specified in a cluster environment or if the ZDB configuration is run on the backup system.

In an RAC environment: Name of the node or the virtual server of the Oracle resource group. The latter can only be used on HP-UX systems: Name of the database to be configured.

In an Oracle Data Guard environment: Name of either a primary system or secondary (standby) system.

`-dbname` *DB_NAME*

Name of the database to be configured.

`-orahome` *ORACLE_HOME*

Pathname of the Oracle Server home directory.

`-config`

Configures an Oracle database.

`-chkconf`

Checks the configuration of an Oracle database.

-chkconf_smb

Checks if an Oracle database is properly configured for ZDB.

-chkconf_ir

Checks if an Oracle configuration is suitable for instant recovery.

-bkphost *BACKUP_SYSTEM*

Name of the backup system. It must be specified for a ZDB backup set configuration.

-prmuser *PRIMARY_USERNAME*

Username for login to the target or primary database. Note that the user must have been granted the SYSDBA privilege.

-prmpasswd *PRIMARY_PASSWORD*

Password for login to the target or primary database. Note that the user must have been granted the SYSDBA privilege.

-prmservice *PRIMARY_NET_SERVICE_NAME_1*[,*PRIMARY_NET_SERVICE_NAME_2* ...]

Net services names for the primary database.

In an RAC environment: Each net service name must resolve into a specific database instance.

-rcuser *CATALOG_USERNAME*

Username for login to the recovery catalog. This is optional and is used only if you use the recovery catalog database catalog as an RMAN repository for backup history.

-rcpasswd *CATALOG_PASSWORD*

Password for login to the recovery catalog. This is optional and is used only if you use the recovery catalog database catalog as an RMAN repository for backup history.

-rcservice *CATALOG_NET_SERVICE_NAME*

Net services name for the recovery catalog.

-stbuser *STANDBY_USERNAME*

Used in the Oracle Data Guard environment for backing up a standby database. Username for login to the standby database.

-stbpasswd *STANDBY_PASSWORD*

Used in the Oracle Data Guard environment for backing up a standby database. Password for login to the standby database.

-stbservice *STANDBY_NET_SERVICE_NAME_1*[,*STANDBY_NET_SERVICE_NAME_2* ...]

Net services names for the standby database.

-zdb_method {PROXY | BACKUP_SET}

Configures the Oracle database for ZDB environment and sets the ZDB method to Oracle proxy-copy or Oracle backup set.

-ctlcp_location *BACKUP_CONTROL_FILE_COPY_LOCATION*

The location on the source volumes where a copy of the current control file is made during ZDB to disk. This is optional and if not specified, `ob2rman.pl` will copy the copy of the control file from the application system to the backup system when it is needed. Thus, you do not need to create an additional disk for this location if you do not need the control file copy on a replica.

If you use a raw logical volume as the *BACKUP_CONTROL_FILE_COPY_LOCATION*, the logical volume must reside on a volume group that will be replicated. If there is no such raw logical volume available, create a new shared disk (volume group) residing on the disk that will be replicated and configure a raw logical volume on it. If you use a raw logical volume, in case of an ZDB to disk, you need to ensure enough free space in the `/var/opt/omni/tmp` directory on the backup host to hold the copy of the raw logical volume.

`-pfile` *PARAMETER_FILE*

Full name of the PFILE residing on the application system. This is optional and used if backup method is backup set and the database instance uses PFILE (and not SPFILE).

`-asmhome` *ASM_HOME*

Specifies the home directory of the ASM instance in an Oracle ASM configuration. Specify this option if the value differs from the home directory of the Oracle database instance.

`-asmuser` *ASM_USERNAME*

This option can be used only in combination with the `-asmpasswd` and `-asmservice` options.

Specifies the user name used by the Data Protector Oracle integration agent to connect to the ASM database. Note that the user must have been granted the `SYSDBA` privilege.

`-asmpasswd` *ASM_PASSWORD*

This option can be used only in combination with the `-asmuser` and `-asmservice` options.

Specifies the password used by the Data Protector Oracle integration agent to connect to the ASM database.

`-asmservice` *ASM_NET_SERVICE_NAME_1*[,*ASM_NET_SERVICE_NAME_2* ...]

This option can be used only in combination with the `-asmuser` and `-asmpasswd` options.

Specifies the name of the net service to be used to access the ASM database. For Oracle environments involving multiple net services, multiple names can be specified.

NOTES

- On HPE OpenVMS, to invoke the Data Protector CLI, execute:

  `$@OMNI$ROOT: [BIN] OMNI$CLI_SETUP.COM`

- *BACKUP_CONTROL_FILE_COPY_LOCATION:*

  This parameter is optional and if not specified, `ob2rman.pl` will copy the copy of the control file from the application system to the backup system when it is needed. Thus, you do not need to create an additional disk for this location if you do not need the control file copy on a replica.

  If you use a raw logical volume as the *BACKUP_CONTROL_FILE_COPY_LOCATION*, the raw logical volume must reside on a volume group that will be replicated. If there is no such raw logical volume available, create a new shared disk (volume group) residing on the disk that will be replicated and configure a raw logical volume on it. If you use a raw logical volume, in case of an ZDB to disk, you

need to ensure enough free space in the `/var/opt/omni/tmp` directory on the backup host to hold the copy of the raw logical volume.

- *PARAMETER_FILE:*

  This parameter is optional and used if backup method is backup set and the database instance uses PFILE (and not SPFILE).

EXAMPLES

The following names are used in the examples below:

- database name: `oracl`
- Oracle Server home directory: `/app10g/oracle10g/product/10.1.0`
- primary user name: `system`
- primary password: `manager`
- primary net service name 1: `netservice1`
- primary net service name 2: `netservice2`
- recovery catalog user name: `rman`
- recovery catalog password: `manager`
- recovery catalog net service name: `catservice`
- standby user name (Oracle Dataguard only): `system`
- standby password (Oracle Dataguard only): `manager`
- standby net service name 1 (Oracle Dataguard only): `netservicesb1`
- standby net service name 2 (Oracle Dataguard only): `netservicesb2`
- parameter file: `/app10g/oracle10g/product/10.1.0/dbs/pfile.ora`
- backup system name: `bcksys`
- ASM home directory: `/oracle/crshome/crshome/crs/app/10.1.0/grid`
- ASM user name: `sys`
- ASM password: `oracle`
- ASM net service name: `ASMSRV`

1. The following example illustrates the configuration of an Oracle database and its recovery catalog in an Oracle Data Guard environment and using the Oracle backup set ZDB method:

   ```
   util_oracle8.pl -config -dbname oracl -orahome app10g/oracle10g/product/10.1.0
   -prmuser system -prmpasswd manager -prmservice netservice1,netservice2 -stbuser
   system -stbpasswd manager -stbservice netservicesb1,netservicesb2 -rcuser rman
   -rcpasswd manager -rcservice catservice -zdb_method BACKUP_SET -pfile
   /app10g/oracle10g/product/10.1.0/dbs/pfile.ora
   ```

2. The following example illustrates the configuration of an Oracle database and its recovery catalog in an Oracle backup set ZDB environment:

   ```
   util_oracle8.pl -config -dbname oracl -orahome app10g/oracle10g/product/10.1.0
   -prmuser system -prmpasswd manager -prmservice netservice1,netservice2 -rcuser
   ```

```
rman -rcpasswd manager -rcservice catservice -zdb_method BACKUP_SET -pfile
/app10g/oracle10g/product/10.1.0/dbs/pfile.ora -bkphost bcksys
```

3. The following example illustrates the configuration of an Oracle database and its recovery catalog in an Oracle backup set ZDB environment which uses Automatic Storage Management (ASM):

```
util_oracle8.pl -config -dbname oracl -orahome app10g/oracle10g/product/10.1.0
-prmuser system -prmpasswd manager -prmservice netservice1,netservice2 -rcuser
rman -rcpasswd manager -rcservice catservice -zdb_method BACKUP_SET -pfile
/app10g/oracle10g/product/10.1.0/dbs/pfile.ora -bkphost bcksys -asmhome
/oracle/crshome/crshome/crs/app/10.1.0/grid -asmuser sys -asmpasswd oracle -
asmservice ASMSRV
```

SEE ALSO

omniintconfig.pl(1M), util_cmd(1M), vepa_util.exe(1M)

vepa_util.exe(1M)

vepa_util.exe — configures a VMware ESX(i) Server system, VMware vCenter Server system, VMware vCloud Director, Microsoft Hyper-V system, checks the configuration, configures virtual machines, browses and lists VMware datacenters
(this command is available on Windows and Linux systems with the Data Protector `Virtual Environment Integration` component installed)

SYNOPSIS

vepa_util.exe --version|--help|--details {command_opt|query_opt|browse_opt}

vepa_util.exe {command *COMMAND_OPTIONS*|query *QUERY_OPTIONS*|browse *BROWSE_OPTIONS*} *ENVIRONMENT_OPTIONS*

vepa_util.exe command --upgrade-cell_info

COMMAND_OPTIONS

--add-standalone-host *ESX_CONFIG_OPTIONS* [--ssl-thumbprint *ThumbPrint*]

--remove-standalone-host --esx-server *EsxName* [*EsxName*...]

--check-config

--config *CONFIG_OPTIONS*

--configvm *VM_CONFIG_OPTIONS*

--unlock-vmotion {--vm *VmPath*|--uuid *VmUUID*}

--show-incremental-flag [--uuid *VmGUID*|--uuid "*VmGUID*, *VmGUID* [, *VmGUID*...]"]

--enable-incremental {--uuid *VmGUID*|--uuid "*VmGUID*, *VmGUID* [, *VmGUID*...]"}

--disable-incremental {--uuid *VmGUID*|--uuid "*VmGUID*, *VmGUID* [, *VmGUID*...]"}

QUERY_OPTIONS

--list-organizations

--list-datacenters

--list-datastores

--list-esx-servers [--cluster *ClusterName*]

--list-resource-pools [--hypervisor *ClusterOrESXName*]

--list-clusters [--instance *DatacenterName*]

    --list-vms

  BROWSE_OPTIONS
   --root-node *NodePath*

  ENVIRONMENT_OPTIONS
   --virtual-environment {VMWare|vCD|HyperV}
   --host *HostName*

  CONFIG_OPTIONS
   --port *PortNumber*
   --username *UserName*
   --password *Password*
   --encoded-password *EncodedPassword*
   --webroot *Webroot*
   --security-model {0|1}

  ESX_CONFIG_OPTIONS (VMWare only)
   --esx-username *EsxUsername*
   --esx-password *EsxPassword*
   --esx-server *EsxName* [*EsxName*...]
   --datacenter *DatacenterName*
   --ssl-thumbprint *SslThumbprint*

  VM_CONFIG_OPTIONS
   --instance *DatacenterName*
   --vm *VmPath*
   --transportation-mode {san|nbd|nbdssl|hotadd}
   --quiescence
   --quiescenceErrLvl {0|1}
   --uuid *VmUuid*
   --default

DESCRIPTION

Use the `vepa_util.exe` command to configure a VMware ESX(i) Server system, VMware vCenter Server system, VMware vCloud Director, and Microsoft Hyper-V system, check the configuration, configure virtual machines, browse and list VMware datacenters.

OPTIONS

`--version`

Displays the version of the `vepa_util.exe` command.

`--help`

Displays the usage synopsis for the `vepa_util.exe` command.

`--details {command_opt|query_opt|browse_opt}`

Displays short descriptions for the specified `vepa_util.exe` options.

`--upgrade-cell_info`

Upgrades the `cell_info` file after upgrading Data Protector 6.20 to the latest product version.

The `cell_info` file upgrade is mandatory.

COMMAND_OPTIONS

`--add-standalone-host` *ESX_CONFIG_OPTIONS* [`--ssl-thumbprint` *ThumbPrint*]

This is a VMware specific option.

Adds the specified standalone ESX Server system to the datacenter.

`--remove-standalone-host --esx-server` *EsxName* [*EsxName*]...

This is a VMware specific option.

Removes the specified ESX Server system from a datacenter.

`--check-config`

Checks whether the specified application client is configured right.

`--config` *CONFIG_OPTIONS*

Configures the specified application client.

`--configvm` *VM_CONFIG_OPTIONS*

This is a VMware specific option.

Configures the backup options for VMware virtual machines.

Note that this option does not check the environment. If you mistype a virtual machine name or a virtual machine UUID the configuration reports success but it is useless.

`--unlock-vmotion {--vm` *VmPath* `|--uuid` *VmUUID*`}`

This is a VMware specific option.

Unlocks vMotion for the specified VMware virtual machine.

--show-incremental-flag [--uuid *VmGUID* | --uuid "*VmGUID*, *VmGUID*[, *VmGUID*]..."]

This is a Microsoft Hyper-V specific option.

Displays states of the specified virtual machines (GUIDs) with regard to their readiness for incremental backup. If the option --uuid is not specified, the states of all virtual machines residing on the specified Hyper-V system or in the specified Hyper-V cluster are listed.

--enable-incremental {--uuid *VmGUID* | --uuid "*VmGUID*, *VmGUID*[, *VmGUID*]..."}

This is a Microsoft Hyper-V specific option.

Prepares specified virtual machines (GUIDs) for Data Protector incremental backup sessions by enabling them for incremental backup in the Hyper-V environment. To complete the preparation process, you need to further run a full backup session for them.

For information on an alternative way of preparing virtual machines for incremental backup sessions, see the *HPE Data Protector Integration Guide*.

--disable-incremental {--uuid *VmGUID* | --uuid "*VmGUID*, *VmGUID*[, *VmGUID*]..."}

This is a Microsoft Hyper-V specific option.

Makes specified virtual machines (GUIDs) incremental backup-disabled in the Hyper-V environment. You cannot perform incremental backup on these virtual machines until you prepare them for incremental backup sessions again. Executing the vepa_util.exe command with this option specified is the only way to prevent the specified virtual machines from being backed up incrementally.

QUERY_OPTIONS

--list-organizations

This is a VMware specific option.

Lists all organizations in the vCloud Director.

--list-datacenters

This is a VMware specific option.

Lists all datacenters.

--list-datastores

This is a VMware specific option.

Lists all datastores.

--list-esx-servers [--cluster *ClusterName*]

This is a VMware specific option.

Lists all ESX Server systems.

--list-resource-pools [--hypervisor *ClusterOrESXName*]

This is a VMware specific option.

Lists all resource pools (including vApps).

The `--hypervisor` is a cluster or an ESX Server system. If specified as a cluster it lists all resource pools on the specified cluster. If specified as an ESX Server system it lists all resource pools on the specified ESX Server system. If the `--hypervisor` is not specified the `--list-resource-pools` option lists all datastores of the specified client.

`--list-clusters` [`--instance` *DatacenterName*]

This is a VMware specific option.

Lists all clusters on the specified client.

If the `--instance` option is specified, it lists all clusters on the specified datacenter.

`--list-vms`

This is a Microsoft Hyper-V specific option.

Lists names and GUIDs of all virtual machines configured on the specified Hyper-V system or in the specified Hyper-V cluster.

BROWSE_OPTIONS

`--root-node`*NodePath*

This is a VMware specific option.

Specifies a root node to start the browsing.

ENVIRONMENT_OPTIONS

`--virtual-environment` {vmware | vCD | hyperv}

Specifies the virtual environment type.

`--host` *HostName*

Specifies the application host (for example, a vCenter Server system, ESX(i) Server system, vCloud Director, or Microsoft Hyper-V system).

CONFIG_OPTIONS

`--port` *PortNumber*

This is a VMware specific option.

Specifies the port to connect to (for example, 443).

`--username` *UserName*

Specifies an operating system user account for the connection.

`--password` *Password*

Specifies the user's password.

`--encoded-password` *EncodedPassword*

Specifies the user's encoded password.

`--webroot` *WebRoot*

This is a VMware specific option.

Specifies the web service entry point URI (for example, /sdk).

`--security-model {0 | 1}`

> This is a VMware specific option.
>
> Specifies the security model.
>
> If the `0` option is specified, you have to specify all login credentials manually (standard security).
>
> If the `1` option is specified, Data Protector connects to the VMware vCenter Server system with the user account under which the Data Protector Inet service on the backup host is running (integrated security). Ensure this user account has appropriate rights to connect to the VMware vCenter Server system.

ESX_CONFIG_OPTIONS

`--esx-username` *EsxUserName*

> This is a VMware specific option.
>
> Adds a username for the ESX Server system.

`--esx-password` *EsxPassword*

> This is a VMware specific option.
>
> Adds a password for the ESX Server system.

`--esx-server` *EsxName* [*EsxName*]…

> This is a VMware specific option.
>
> Specifies ESX Server system(s) on which to execute a command.

`--datacenter` *DatacenterName*

> This is a VMware specific option.
>
> Adds a datacenter to the backup client.

`--ssl-thumbprint` *SslThumbPrint*

> This is a VMware specific option.
>
> Specifies the thumbprint of a SSL certificate.

VM_CONFIG_OPTIONS

`--instance` *DatacenterName*

> This is a VMware specific option.
>
> Specifies the datacenter that a virtual machine belongs to.

`--vm` *VmPath*

> This is a VMware specific option.
>
> Specifies the virtual machine (for example, `/vm/myTestVM`).

`--transportation-mode {san|nbd|nbdssl|hotadd}`

> This is a VMware specific option.

Specifies the transportation mode to be used for backup. If this option is not specified, the fastest available transportation mode is used.

`--quiescence`

This is a VMware specific option.

Specifies whether to use Microsoft Volume Shadow Copy Service (VSS) functionality to quiesce all applications with VSS writers before performing the backup.

`--quiescenceErrLvl {0|1}`

This is a VMware specific option.

Specifies the level of error message to be generated if the quiescence snapshot fails: 0 (warning), 1 (fatal). Default: 0.

`--uuid VmUuid`

This is a VMware specific option.

Specifies the UUID of the virtual machine.

`--default`

This is a VMware specific option.

Uses default virtual machine settings for all virtual machines.

`--forceNonCBTFull`

This option is used to disable CBT backups and force only non-CBT backups.

`--allowNonCBTFull`

This option is used to prevent failed CBT backup and allow fallback to non-CBT.

EXAMPLES

The following examples illustrate how the `vepa_util.exe` command works.

1. To configure the vCenter Server system "vc.company.com", execute:

   ```
   vepa_util command --config --virtual-environment vmware --host vc.company.com -
   -security-model 0 --username Administrator --password XYZ --webroot /sdk --port
   443
   ```

2. To check the configuration of the vCenter Server system "vc.company.com", execute:

   ```
   vepa_util command --check-config --virtual-environment vmware --host
   vc.company.com
   ```

3. To list all datacenters registered in the vCenter Server system "vc.company.com", execute:

   ```
   vepa_util query --virtual-environment vmware --host vc.company.com --list-
   datacenters
   ```

4. To browse the datacenter "PRODUCTION" registered in the vCenter Server system "vc.company.com", execute:

   ```
   vepa_util browse --virtual-environment vmware --host vc.company.com --root-node
   "PRODUCTION"
   ```

5. To list all vCloud Director organizations, execute:

```
vepa_util query --virtual-environment vcd --host vcd.vepa.company.com --list-
organizations
```

6. To check whether the configuration of the vCloud Director client "vcd.vepa.company.com" was successful, execute:

```
vepa_util command --virtual-environment vcd --host vcd.vepa.company.com --
username admin --encoded-password xaF3r3af
```

7. To list names and GUIDs of the virtual machines configured on the Microsoft Hyper-V virtual server "hyperclus3.company.com", execute:

```
vepa_util query --list-vms --virtual-environment HyperV --host
hyperclus3.company.com
```

8. To prepare virtual machine with the GUID "741FF564-DA19-45E5-B273-D72FA2D91998" on the Microsoft Hyper-V system "hypersysB.company.com" for incremental backup, execute:

```
vepa_util command --enable-incremental --uuid 741FF564-DA19-45E5-B273-
D72FA2D91998 --virtual-environment HyperV --host hyperclus3.company.com
```

9. To configure the virtual machine to use the "forceNonCBTFull" option, execute:

```
vepa_util command –configvm –virtual-environment vmware –host peh –instance /DP
DEV SA –vm /DP DEV SA/klaster/VM1_test1 –uuid '503eff1b-9bc8-eea5-a417-
7678da09d529'–quiescence 1 quiescenceErrLvl 0 –transportation-mode fastest --
forceNonCBTFull
```

10. To configure the virtual machine to use the "allowNonCBTFull" option, execute:

```
vepa_util command –configvm –virtual-environment vmware –host peh –instance /DP
DEV SA –vm /DP DEV SA/klaster/VM1_test1 –uuid '503eff1b-9bc8-eea5-a417-
7678da09d529'–quiescence 1 quiescenceErrLvl 0 –transportation-mode fastest --
allowNonCBTFull
```

SEE ALSO

omniintconfig.pl(1M), util_cmd(1M), util_oracle8.pl(1M)

Section 5: Miscellaneous

omnigui(5)

omnigui — describes usage of the commands that launch the Data Protector GUI

SYNOPSIS

GUICommand [-help]

GUICommand

 manager [*ContextOptions*] [-server *HostName*]

 mom [*ContextOptions*] [-server *HostName*]

ContextOptions

 -admin

 -backup

 -clients

 -copy

 -db

 -instrec

 -monitor

 -report

 -restore

 -users

DESCRIPTION

These commands are used to launch the Data Protector GUI and activate all or any combination of the Data Protector GUI contexts.

To use the Data Protector GUI functionality with UNIX Cell Manager systems, on which the Data Protector GUI is not available, use the `omniusers` command to remotely add a new Data Protector user to a Cell Manager on which the Data Protector GUI is not installed. You can then use the user account of the newly added Data Protector user to launch the Data Protector GUI on another system with the Data Protector GUI installed, and connect to the Cell Manager. For details, see the `omniusers` reference page. For a list of operating systems supported by the Data Protector user interfaces, see the *HPE Data Protector Product Announcements, Software Notes, and References*.

For more information on local language support and the usage of non-ASCII characters in file names, see the *HPE Data Protector Help*.

COMMANDS

`manager`

Launches the Data Protector GUI with all Data Protector contexts activated, or, when additional options are specified, with the specified Data Protector contexts activated.

`mom`

Launches the Data Protector Manager-of-Managers GUI with all Data Protector contexts activated (with the exception of the Internal Database and Devices & Media contexts) or, when additional context options are specified, with the specified Data Protector contexts activated.

OPTIONS

`-help`

Displays the usage synopsis for the specified command.

`-server` *HostName*

Connects to the specified Cell Manager.

`-display` *HostName:0*

Redirects the output to the display on the specified system.

`-admin`

Launches the Data Protector GUI with the Devices & Media contexts activated.

`-backup`

Launches the Data Protector GUI with the Backup context activated.

`-clients`

Launches the Data Protector GUI with the Clients context activated.

`-copy`

Launches the Data Protector GUI with the Object Operations context activated.

`-db`

Launches the Data Protector GUI with the Internal Database context activated.

`-instrec`

Launches the Data Protector GUI with the Instant Recovery context activated.

`-monitor`

Launches the Data Protector GUI with the Monitor context activated.

`-report`

Launches the Data Protector GUI with the Reporting context activated.

`-restore`

Launches the Data Protector GUI with the Restore context activated.

`-users`

Launches the Data Protector GUI with the Users context activated.

EXAMPLES

1. `manager`

   This command launches the Data Protector GUI with all contexts activated.

2. `manager -admin -monitor -report -server host3`

   This command launches the Data Protector GUI with the Devices & Media, Monitor, and Reporting contexts activated and connects to the Cell Manager with the hostname "host3".

SEE ALSO

ob2install(1M), omniintro(9), omnimigrate.pl(1M), omnisetup.sh(1M), omniusers(1), upgrade_cm_from_ evaa(1M)

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 Command Line Interface Reference (Data Protector 9.08)

Just add your feedback to the email and click send.

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

We appreciate your feedback!