

ProLion DataOptimizer

Manual



ProLion GmbH

A publication of:
ProLion GmbH
Maierhöfen 8/3
2813 Lichtenegg, Österreich

Mail: office@prolion.at
Web: www.prolion.at

POSSESSION, USE, DUPLICATION OR DISSEMINATION OF THIS DOCUMENTATION AS WELL AS THE SOFTWARE DESCRIBED IN THIS DOCUMENTATION IS AUTHORISED ONLY PURSUANT TO A VALID WRITTEN LICENSE FROM ProLion GmbH OR AN AUTHORISED SUB-LICENSOR. ProLion GmbH BELIEVES THE INFORMATION INCLUDED IN THIS PUBLICATION IS ACCURATE AS OF THE DATE OF PUBLICATION, IT IS SUBJECT TO CHANGE WITHOUT NOTICE. ProLion GmbH IS NOT RESPONSIBLE FOR ANY INADVERTENT ERRORS. ProLion GmbH MAKES NO REPRESENTATIONS THAT THE USE OF ITS PRODUCTS IN THE MANNER DESCRIBED IN THIS DOCUMENT WILL NOT INFRINGE ON EXISTING OR FUTURE PATENT RIGHTS. THE DESCRIPTIONS CONTAINED IN THIS DOCUMENT DO NOT IMPLY THE GRANTING OF LICENSES TO MAKE, USE, OR SELL EQUIPMENT OR SOFTWARE IN ACCORDANCE WITH THE DESCRIPTION.

Attention

A functioning storage workflow requires a well configured system environment with devices working free of faults and, if applicable, flawless storage media. Therefore it is of the essence that the user does backup all data by functions offered by ProLion software and/or (if required) by supplementary software products at adequate intervals (i.e. in accordance with the scope and frequency of changes), and thereby to facilitate the reinstatement of these data even in exceptional situations (i.e. in case of hardware malfunction).

© PoINT Software & Systems GmbH, 2017. All rights reserved.
Document Version: 6.0/10.02.2017/DG
Program Version: 6.0

Contents

1	Overview	5
1.1	Terms.....	5
1.2	Storage Vault	5
1.2.1	Storage Tiers.....	5
1.2.2	Policies.....	5
1.2.3	Job Cycles.....	5
1.3	Modules	5
1.3.1	ProLion DataOptimizer Service	6
1.3.2	Storage Agent.....	6
1.3.3	NetApp FAS Agent.....	6
1.3.4	ProLion DataOptimizer GUI	6
1.4	Performance Tier	6
1.5	Capacity Tier	7
2	Getting Started.....	8
2.1	Installation.....	8
2.1.1	Server Installation.....	8
2.1.2	Remote Administration	8
2.1.3	Uninstall ProLion DataOptimizer.....	8
2.2	Licensing	9
2.2.1	Storage Volume License	9
2.3	Configure ProLion DataOptimizer	9
2.3.1	Setup Directories.....	9
2.3.2	Storage Agent Account.....	9
2.3.3	Configure Storage Vaults.....	10
3	NetApp FAS Data Source.....	11
3.1	FPolicy Server for 7-Mode.....	11
3.1.1	Requirements.....	11
3.1.2	Functionality	12
3.1.3	Restrictions.....	13
3.1.4	Removing FPolicy Configuration	14
3.2	FPolicy Server for Cluster Mode.....	14
3.2.1	Requirements.....	14
3.2.2	Functionality	15
3.2.3	Restrictions.....	17
4	Using ProLion DataOptimizer.....	18
4.1	ProLion DataOptimizer GUI	18
4.1.1	Side Bar.....	18
4.1.2	Storage Vault Detail View	18
4.2	Setup User Accounts	19

4.3	Configure a Storage Vault	19
4.3.1	Performance Tier Settings.....	20
4.3.2	Migration Policies	20
4.3.3	Capacity Tier	21
4.4	Working with Policies	22
4.4.1	Overview	22
4.4.2	Templates	22
4.4.3	Conditions.....	23
4.4.4	Actions	24
4.5	Performing Job Cycles.....	25
4.6	Data Browser.....	26
4.6.1	Copy Data to another Location.....	26
4.6.2	Export List of Files.....	27
4.7	Reconfigure Storage Vault Data Source	27
5	Error Reporting.....	28
5.1	Log Files	28
5.2	Server Notifications.....	28
5.3	E-Mail Notifications.....	28
6	Backup and Recovery.....	29
6.1	Recovery of ProLion DataOptimizer Server.....	29
6.1.1	Recovery by Importing Database Files	29
6.1.2	Recovery by 3 rd Party Backup Programs.....	30
6.2	Recovery of Performance or Capacity Tier	31
6.2.1	Recover File Stubs on Performance Tier.....	31
6.2.2	Delete Orphaned File Stubs on Performance Tier.....	31
6.2.3	Recover Capacity Tier with Replication.....	32
7	Index.....	33

1 Overview

This chapter provides an overview of terms and concepts which help using ProLion DataOptimizer.

1.1 Terms

The following section describes the terms and definitions which are used throughout the manual.

1.2 Storage Vault

The core element in the concept of ProLion DataOptimizer is a **Storage Vault**. This object includes all parameters and defines a complete storage process. A Storage Vault is mainly specified by its Storage Tiers, Migration Policies and additional configuration parameters.

1.2.1 Storage Tiers

During their life cycle, data can be located on different storage tiers. Initially, data will be created on or copied to the Tier 1 (**Performance Tier**). This tier is also used to access data which have been migrated to another tier. Data which are less frequently used will be moved to Tier 2, the **Capacity Tier**. Tier 3, the **Archive Tier**, will be used for data which are not actively used anymore or which shall be stored in a compliant manner. Note that the Archive Tier is not supported by ProLion DataOptimizer.

1.2.2 Policies

Data can be copied or moved between different storage tiers using defined policies. Policies consist of a set of conditions and a set of actions. When applying a policy to a file, the specified conditions will be checked and if all conditions are matched, the actions will be performed on the given file.

Conditions can be used to filter files by their names, their attributes, their current location and other options.

Actions can be used to copy files to a specific tier or to remove them from a tier.

ProLion DataOptimizer uses different sets of policies which will be evaluated regularly based on a defined schedule or if specific events occur. Chapter 4.4, Working with Policies, describes the usage of policies in detail.

1.2.3 Job Cycles

The term job cycle describes the process of executing a set of policies on a set of files. Job cycles can be executed regularly based on a defined schedule and they can be triggered manually. Refer to chapter 4.5, Performing Job Cycles, for more information about job cycles.

1.3 Modules

This chapter describes the main modules of ProLion DataOptimizer and how they interact with each other.

1.3.1 ProLion DataOptimizer Service

The ProLion DataOptimizer Service is the main component of ProLion DataOptimizer. This service controls the connection between all other components and the user interface.

1.3.2 Storage Agent

The Storage Agent Service is a windows service which controls the execution of job cycles. It has a connection to the other Data Source dependent agents which are described below.

This service runs on the same computer where the ProLion DataOptimizer Service is installed and running.

1.3.3 NetApp FAS Agent

The NetApp FAS Agent controls data sources of type NetApp FAS. This agent is implemented as a part of the Storage Agent Service and therefore does not exist as separate executable.

1.3.4 ProLion DataOptimizer GUI

The ProLion DataOptimizer GUI is the user interface which is required to configure ProLion DataOptimizer. The user interface can be running on the computer where the ProLion DataOptimizer Service is running or on another computer in the network. It can be installed on client computers by performing a client or remote administration installation as described in chapter 2.1, Installation.

After installing it on a client computer it is necessary to select the computer where the ProLion DataOptimizer Service is running because the user interface cannot directly connect to the other agents.

1.4 Performance Tier

The Performance Tier is the Tier 1 of a Storage Vault. This tier is always present because it defines where the data is located which shall be managed by the Storage Vault. The main parameter of a Performance Tier is its Data Source. Depending on the type of the Data Source there may be additional options which can be specified.

The NetApp FAS Data Source defines a file system ("NetApp FAS Volume") of a NetApp Storage System which is controlled by NetApp FAS Agent. NetApp FAS Agent supports both, NetApp FAS running in 7-Mode and Cluster Mode.

The NetApp FAS Data Source allows archiving of files stored on a NetApp FAS system to configured target devices. Archived files can be replaced by so called "stubs" in the NetApp FAS file system. This means that archived files are still transparently accessible through the original path name on the NetApp FAS file system.

The following list summarizes the general features of the NetApp FAS support:

- Archiving to Capacity Tier.
- Based on the specified parameters stored files are automatically substituted by a link (“stub”).
- Access to purged files is supported by a “pass through” mechanism which means the requested blocks of the file are directly delivery from the Storage Device to the application without first re-storing the complete file back into the NetApp FAS file system.
- Archived file versions can also be accessed by the Data Browser function of ProLion DataOptimizer.

Multiple Storage Vaults with configured NetApp FAS Data Sources can be created for different NetApp FAS Volumes.

1.5 Capacity Tier

The Capacity Tier can be configured as Tier 2 for a Storage Vault. Data can be copied or moved from the Performance Tier to the Capacity Tier controlled by policies.

Directories which are located on a locally attached hard disk or which are accessible using an UNC path can be used as Capacity Tier. The hard disk must be formatted with NTFS or a compatible file system.

It is also possible to specify a replication path for the Capacity Tier. All data from the primary Capacity Tier will then automatically be replicated to the specified replication path.

2 Getting Started

This chapter provides hints related to installation and first configuration steps for ProLion DataOptimizer.

2.1 Installation

2.1.1 Server Installation

The ProLion DataOptimizer server installation has to be performed on the system that configures the storage devices by executing the following steps:

1. Log in as a member of the local administrator's group.
2. Start the software installation by executing **SETUP.EXE** from the distribution package. During the setup you will be prompted to enter a Configuration Key. This key is handed to you together with the installation. It is used to personalize your copy of ProLion DataOptimizer. After providing the Configuration Key, the setup will prompt to select the components you want to install. The setup program will guide you through the rest of the installation process and will copy program related files/directories to hard disk and installs and configures the required services and device drivers.

Please note that specific options of ProLion DataOptimizer need to be enabled by the corresponding License Key. See chapter 2.2, Licensing, for more information.

Note: During installation, ProLion DataOptimizer automatically generates the System Key which is required for several operations of ProLion DataOptimizer. This key is essential for recovery purposes. Therefore it is highly recommended to save the generated System Key in a file on a separate storage medium.

2.1.2 Remote Administration

If you want to administer the ProLion DataOptimizer Server from a remote computer you can install the ProLion DataOptimizer Client on the remote computer.

Start the file SETUP.EXE to install the ProLion DataOptimizer. When the setup program prompts for the components to be installed, select only the **User Interface Components**. The setup program will guide you through the remaining installation process.

When you start the ProLion DataOptimizer for the first time it prompts you to select the computer which you intend to use as ProLion DataOptimizer Server.

2.1.3 Uninstall ProLion DataOptimizer

Before starting the un-installation of ProLion DataOptimizer you should stop and deactivate all Storage Vaults. Now open the Add/Remove Software Control Panel applet, find the entry for ProLion DataOptimizer and click **Change/Remove**. In the following dialog select **Remove All** to uninstall ProLion DataOptimizer from your system.

During the software removal you will be prompted to confirm remove of some directories. You can confirm this, if these directories have been solely used by ProLion DataOptimizer.

2.2 Licensing

ProLion DataOptimizer requires a License Key for operation. It is based on the ProLion DataOptimizer Serial Number and the machine-dependent Product Key.

To obtain the License Key, contact your dealer and refer to the Serial Number, Product Key and the connected storage systems.

The License Key can be entered during installation or at a later time in the ProLion DataOptimizer setup dialog after selecting **Tasks → Setup ProLion DataOptimizer → General → License**.

Note: The ProLion DataOptimizer Product Key is machine-dependent and may change whenever the (hardware) configuration is modified.

2.2.1 Storage Volume License

The ProLion DataOptimizer license key also contains information about the storage volume which can be managed by ProLion DataOptimizer. Once the present data volume exceeds the licensed data volume, ProLion DataOptimizer will not execute further migration or archive job cycles.

The licensed and the used data volume will be displayed in the main window after selecting **Tasks** in the menu bar on the left side.

2.3 Configure ProLion DataOptimizer

2.3.1 Setup Directories

The working directories of ProLion DataOptimizer can be configured by means of the installation wizard or at a later time in the setup dialog after selecting **Tasks → Setup ProLion DataOptimizer → General → Directories**.

The **Databases** directory stores the Storage Vault databases. The size of the databases depends on the number of managed files. Please refer to the ReadMe for detailed information about space requirements.

The directory for **Temporary Image Volumes and Volume Cache** is used to store temporary image files which are created during an archiving job cycle. These temporary image files serve as temporary storage for the migration process to the target devices. If you change this directory, new images will be stored there, while old ones remain in the previous directory and will be removed after they have been stored on the target devices.

The volume cache is used to cache data which have been read from tape or object store devices. The size of this cache is configurable using the **Advanced Settings** dialog of ProLion DataOptimizer. For optimal performance it is recommended to locate this directory on another physical hard disk than the database directory.

Configuration, Journal and Log Files are per default located in the installation directory of ProLion DataOptimizer. Because journal and log files may become large, it is recommended to move these directories to a partition with at least 100 GByte free space.

2.3.2 Storage Agent Account

The Storage Agent Account is the account which will be used by the Storage Agent Service.

This account must be a member of the local Administrators group. Additionally, this account must have permission to access the data on the configured Data Sources and it must have access to network shares or disks which will be used as Capacity Tiers.

You may need to create a new local user account or a domain account and grant access to the file and sharing permissions of the required directories.

Note: If your computers are not part of a Windows Domain, you are required to create this user account on all computers containing Storage Vault Data Sources. The account name and the password must be the same on all machines.

2.3.3 Configure Storage Vaults

The last step is to configure one or more Storage Vaults. Storage Vaults define the storage tiers and the policies used for data migration. Chapter 4.3, *Configure a Storage Vault*, describes how to create a new Storage Vault and how to configure its parameters.

3 NetApp FAS Data Source

3.1 FPolicy Server for 7-Mode

3.1.1 Requirements

The following requirements must be fulfilled for operating the FPolicy Server for 7-Mode of ProLion DataOptimizer:

- ProLion DataOptimizer must have been installed and licensed. The license key must enable the FPolicy Server for 7-Mode.
- Tested NetApp ONTAP Versions: 7.3.2, 7.3.5, 7.3.6, 8.0.1, 8.1, 8.1.3 and 8.2. The 8.x versions are supported in 7-Mode.
- Both, the NetApp FAS and the Windows server (ProLion DataOptimizer) systems, have to be members of the same or a trusted computer domain. Otherwise, the following security policy of the Windows server system has to be enabled by the Administrator:
 - Network access: Let Everyone permissions apply to anonymous users
 - Security policies can be modified by means of the Security Policy Editor (secpol.msc). A reboot is required after changing this option. If this option has been enabled by the Administrator, the FPolicy Server automatically adds the names of the communication pipes to the following security option:
 - Network access: Named pipes that can be accessed anonymously
 - When using a cluster configuration make sure that these settings have been modified on all physical cluster nodes.
 - If the systems are not members of trusted domains and the security settings have not been adjusted, the NetApp FAS will not be able to send FPolicy requests.
- NetApp FPolicy Server for 7-Mode also creates the following registry value, if it does not exist:
 - HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\LanmanWorkstation\Parameters
 - Value Name: DirectoryCacheLifetime
 - Content: 0 (DWORD)
 - If this setting does not exist, FPolicy Server for 7-Mode creates it and requires a server reboot before it allows activating Storage Vaults. This value will not be changed if it already exists.
- The Storage Agent Account requires administrative privileges on the system where ProLion DataOptimizer is installed. The same account must exist on the NetApp FAS system with the same password and it must be in the group of Backup Operators on the NetApp FAS system. You can also use a domain account if both systems are members of trusted domains. On the NetApp FAS the account additionally requires the following capabilities:
 - api-system-get-version,api-fpolicy-*,api-cifs-share-list-*,api-nfs-exportfs-list-*,api-nfs-exportfs-storage-path
- File system access via CIFS to the Data Source on the NetApp FAS is required, even if clients are only using NFS.
- If the Data Source is also exported using NFS, the Storage Agent account requires read access to this export.

- The local firewall of the Windows server system must allow CIFS/SMB connections from and to the NetApp FAS system.
- NetBIOS must be enabled on the network interface used by the Windows server to connect to the NetApp FAS.
- The length of the name of the computer where ProLion DataOptimizer is running must not exceed 15 characters.
- The number of inodes for NetApp FAS volumes containing ProLion DataOptimizer Storage Vaults must be increased to at least three times the maximal number of files which will be stored on that volume. The maximal and current number of inodes can be viewed using the ONTAP command “maxfiles”. This command can also be used to increase the maximal number of inodes. No more files can be stored on the volume and ProLion DataOptimizer job cycles will fail with an error indicating insufficient disc space if the maximal number of inodes has been reached. Therefore it is recommended to regularly use the command “maxfiles” to make sure that sufficient inodes are available.

NFS Client Software

If the Data Source is also exported using NFS, the FPolicy Server for 7-Mode requires read access to this export. FPolicy Server for 7-Mode automatically selects the best matching NFS export, i.e. the export which is as close as possible to the selected Data Source. To mount the export the software detects if one of the supported NFS clients is present on the system and executes the appropriate command to perform the mount operation. If the mount operation requires additional parameters, it is possible to specify these parameters while configuring the Storage Vault. These parameters will be passed to the mount command line.

Microsoft NFS Client

The Microsoft NFS Client can be used if no other NFS software is installed. To use this client, the Windows server roles “Services for NFS” and “Client for NFS” must be enabled on the computer where ProLion DataOptimizer is running.

FPolicy Server for 7-Mode does not pass additional parameters when using Microsoft NFS. Parameters which have been specified in the Storage Vault configuration will be ignored.

OpenText NFS Solo

If this third party software is installed, FPolicy Server for 7-Mode will use it to mount the NFS export. It is possible to specify additional parameters in the Storage Vault configuration, for example to select a specific NFS protocol version or a character set. Please refer to the documentation of the OpenText NFS Solo command line tool “nfs”. The specified NFS parameters must not contain a drive letter or NFS export because this information will be inserted by FPolicy Server for 7-Mode.

This version of FPolicy Server for 7-Mode has been tested with OpenText NFS Solo 14.

3.1.2 Functionality

FPolicy Server for 7-Mode utilizes the NetApp FPolicy feature of Data ONTAP to allow transparent read access to files which have been copied to Capacity Tier. Whenever a copied file is read the NetApp FAS requests the data from the FPolicy Server for 7-Mode which then reads the data from the Capacity Tier device.

When a Storage Vault is configured for a NetApp FAS system, the FPolicy Server for 7-Mode creates “FPolicy policies” on the NetApp FAS, adds the related volume to these policies and registers itself

as an FPolicy Server. If the FPolicy feature is not enabled on the NetApp FAS system, FPolicy Server for 7-Mode automatically enables this feature.

FPolicy Configuration

FPolicy Server for 7-Mode configures the following policies on the NetApp FAS which has been specified as Performance Tier for the Storage Vault:

PSM_REV01_IDxxxxxxxxxxxxxxxx

The last part of the name of this policy differs on each installation. It is used to ensure that one NetApp FAS system will not be used by different ProLion DataOptimizer installations. This policy does not monitor any file system operations on the NetApp FAS.

PSM_P1

This policy registers for read and write requests for files which are marked as 'offline', i.e. which have been replaced by a stub. All NetApp FAS volumes which contain ProLion DataOptimizer Storage Vaults will automatically be added to the list of volumes monitored by this policy. Therefore, all file system operations for 'offline'-files on these volumes will be sent to ProLion DataOptimizer. Access to such files will fail, if ProLion DataOptimizer is not running or the Storage Vault is not active.

3.1.3 Restrictions

The following restrictions apply to the FPolicy Server for 7-Mode and must be considered:

- There should not be other FPolicy Servers configured which provide stubbing functionality using the offline-attribute mechanism.
- Only one ProLion DataOptimizer installation can access a NetApp FAS. The NetApp FAS must be manually reconfigured to allow access from another ProLion DataOptimizer installation.
- Up to four different NetApp FAS or vFilers can be accessed by one ProLion DataOptimizer installation.
- Files on the NetApp FAS must not be accessed from the system where ProLion DataOptimizer is installed. When accessing purged files from this system, the NetApp FAS would return files filled with 0-bytes. Additionally, if changes are performed on the NetApp FAS, these changes will not be reported to ProLion DataOptimizer. This may lead to inconsistencies and inaccessible data.
- The functionality to purge files, i.e. replace archived files by a stub file, cannot be used on SnapLock volumes. Trying to purge such files will result in error messages during job execution.
- Purging of files which are 4 GByte or larger is not supported by the tested NetApp ONTAP versions. An error message will be logged when trying to purge such a file. To avoid this error message add a condition to the policy so that only files smaller than 4 GByte will be purged.
- Because the Windows 'offline'-attribute is used to identify purged files, it is recommended that no files have this attribute set by other applications. Otherwise ProLion DataOptimizer may remove this attribute for files which have not been purged by ProLion DataOptimizer.
- FPolicy Server can only purge files with valid Windows file names. Therefore if the volume on the NetApp FAS will also be accessed by NFS clients, some restrictions apply: Files or directories which contain one of the characters :?*<>|\\ will not be processed. Also, if the names of two or more file or directory entries in the same directory differ only in capitalization, only one of

these entries will be processed. If such a file name has been detected, a warning message will be logged to the Migration job log file of the Storage Vault. Additionally, ProLion DataOptimizer prevents renames of files or directories via NFS if the resulting name would be an illegal name. In this case a warning message will be logged to the Migration job log file of the Storage Vault.

3.1.4 Removing FPolicy Configuration

Before removing the FPolicy configuration for a Storage Vault make sure that all archived data have been restored to the NetApp FAS system. Otherwise the NetApp FAS system would return incorrect data (files filled with zeros) when reading these files. After deleting the Storage Vault configuration, it is necessary to adjust the FPolicy policies on the NetApp FAS system.

If other Storage Vaults are still using the FPolicy Server for 7-Mode with other volumes on the NetApp FAS system it is necessary to change the configuration so that the volume, for which the Storage Vault has been deleted, will not be affected by the existing FPolicy policies. To do this, open a command line on the NetApp FAS system and execute these commands:

```
fpolicy vol inc remove PSM_P1 <vol_spec>
```

<vol_spec> must be replaced by the volume specification (e.g. vol1) of the NetApp FAS volume for which the Storage Vault has been deleted. Please note that the FPolicy PSM_P2 is optional and need not exist on your system.

If no further Storage Vaults exist for this NetApp FAS or you are uninstalling the FPolicy Server for 7-Mode, it is necessary to disable and remove all FPolicy policies which have been created by ProLion DataOptimizer. To disable and remove a policy, open a command line on the NetApp FAS and execute the following commands:

```
fpolicy disable <policy_name>
fpolicy destroy <policy_name>
```

<policy_name> must be replaced by the policy to process. ProLion DataOptimizer creates the policy PSM_P1, the optional policy PSM_P2 and one policy whose name is similar to PSM_REV01_ID66F2A3629E74745C. The actual name of this policy is different on each system. You can determine this name by means of the command "fpolicy" without additional parameters.

3.2 FPolicy Server for Cluster Mode

3.2.1 Requirements

The following requirements must be fulfilled for operating the FPolicy Server for Cluster Mode of ProLion DataOptimizer:

- ProLion DataOptimizer must have been installed and licensed. The license key must enable the FPolicy Server for Cluster Mode.
- NetApp ONTAP Version 8.3 or later running in Cluster Mode. Versions 8.3 and 9.0 have been tested.
- File system access via CIFS to the Data Source on the NetApp FAS is required, even if clients are only using NFS.
- File and Printer Sharing must be allowed in the Windows firewall settings.
- The Storage Agent Service (PntStorageAgent.exe) must be allowed to accept TCP connections on port 8632. You may either create a firewall exception for the process PntStorageAgent.exe

or the port number. The default port number can be changed using the following registry value on the ProLion DataOptimizer server:

HKEY_LOCAL_MACHINE\SOFTWARE\POINT\ProLion DataOptimizer\ArchiveAgent

Value Name: FPolicyLocalPort

Value Type: DWORD

Content: <Port Number> (default: 8632)

If Storage Vaults have already been created, it is also necessary, to adjust the existing FPolicy configuration.

- The firewall of the NetApp cluster must allow ONTAPI connections from ProLion DataOptimizer. These settings can be changed using the following commands: ::>system services firewall policy create -vserver <vserver> -policy <policy> -service https -allow-list <IP Address/Mask of DataOptimizer Server>
::>network interface modify -vserver <vserver> -lif <lif-name> -firewall-policy <policy>
- On the NetApp FAS the Data Network Interface of the Storage Virtual Machine must allow management access (Enable Management Access).
- The Storage Agent Account requires administrative privileges on the system where ProLion DataOptimizer is installed. The same account must exist on the NetApp FAS system with the same password and it must be in the group of Backup Operators on the NetApp FAS system. You can also use a domain account if both systems are members of the same domain and the account has the required privileges.
- The maximum number of inodes for NetApp FAS volumes containing ProLion DataOptimizer Storage Vaults must be increased to at least three times the maximal number of files which will be stored on that volume. The maximal and current number of inodes can be viewed using the ONTAP command “volume show -fields files, files-used”. The command “volume modify -files” can be used to increase the maximal number of inodes. No more files can be stored on the volume and ProLion DataOptimizer job cycles will fail with an error indicating insufficient disc space if the maximal number of inodes has been reached.

3.2.2 Functionality

FPolicy Server for Cluster Mode utilizes the NetApp FPolicy feature of Data ONTAP to allow transparent read access to files which have been archived to the configured archive devices. Whenever an archived file is read the NetApp FAS requests the data from the FPolicy Server for Cluster Mode which then reads the data from the archive device.

When a Storage Vault is configured for a NetApp FAS system, the FPolicy Server for Cluster Mode automatically creates appropriate “FPolicy policies” on the NetApp FAS and registers itself as an FPolicy Server.

ONTAPI Login

During configuration of a new Storage Vault you will be prompted for an ONTAPI Login which will be used by the FPolicy Server to connect to the Storage Virtual Machine to configure FPolicy. This login must have been created on the Storage Virtual Machine using a NetApp command shell or the NetApp OnCommand System Manager. It is not necessary to create a related Windows or domain account.

This login must be allowed to log in using ONTAPI with password authentication and it must be assigned a role which provides at least the following permissions:

version: readonly

volume: readonly

vserver: readonly

vserver fpolicy: all

Please note that all Storage Vaults on a Storage Virtual Machine must use the same login.

FPolicy Scope

When configuring a Storage Vault the Data Source path of the Storage Vault will be specified. FPolicy Server uses ONTAPI to map this CIFS path to the local path on the Storage Virtual Machine and determines the volume which contains this path and all volumes which are mounted below in the hierarchy. The list of volumes will be specified as 'volumes-to-include' for the FPolicy scope when configuring the FPolicy policy. The scope will automatically be updated if volume mount points have been changed.

Change the FPolicy Configuration

It may be necessary to change the configuration of the FPolicy policies. For example, if the IP address of the computer running the FPolicy Server has changed or another local port number shall be used by the FPolicy Server. When activating a Storage Vault the FPolicy Server verifies the FPolicy configuration and eventually reports that the configuration must be updated. It does not automatically change the configuration because this requires deactivating the policy, so that the storage system may return wrong data for purged files. Therefore, the administrator must first ensure that no clients access the storage system and then manually disable or delete the FPolicy policies created by ProLion DataOptimizer. When activating the Storage Vault again, FPolicy Server automatically creates the new FPolicy policies.

The names of the policies created by the FPolicy Server will be logged to the "Log File for Storage Agent".

Volume Junction Points

FPolicy Server detects and handles changes of volume mount points (junction paths) and it automatically adjusts the FPolicy scope if it detects that a volume has been moved or that a new volume has been mounted. However, it is recommended to prevent all file system accesses and to deactivate all Storage Vaults before changing volume mount points. Also, when moving a volume outside of the scope of a Storage Vault, you must ensure that the volume does not contain purged files because access to these files will not be possible after moving the volume outside of the Storage Vault scope.

Snapshots

FPolicy Server supports access to purged files (stubs) in snapshots as long as that file version exists in the Archive or Capacity Tier.

During job cycles, ProLion DataOptimizer ignores hidden directories with name "~snapshot", because these directories contain read-only copies of older file versions which cannot be archived.

Removing a Storage Vault

When deleting a Storage Vault in ProLion DataOptimizer, FPolicy Server adjusts the FPolicy scope accordingly or deletes the configured FPolicy policy when removing the last Storage Vault on a Storage Virtual Machine.

Before removing a Storage Vault make sure that all purged files have been recalled. It is also recommended to execute the command line tool TagRemover.exe to remove metadata which have been appended to archived files. Refer to the ReadMe file of ProLion DataOptimizer for a detailed description of this command line tool.

3.2.3 Restrictions

The following restrictions apply to the FPolicy Server for Cluster Mode and must be considered:

- There must not be other FPolicy Servers which provide stubbing functionality configured for the same volumes. However, other FPolicy Servers – including other ProLion DataOptimizer installations – may register policies for other volumes.
- Up to four different NetApp Storage Virtual Machines (VServers) can be accessed by one ProLion DataOptimizer installation.
- The FPolicy functionality is not supported for Infinite Volumes. Therefore FPolicy Server can only be used for Flex Volumes.
- The functionality to purge files, i.e. replace archived files by a stub file, cannot be used on SnapLock volumes. Trying to purge such files will result in error messages during job execution.

Because the Windows 'offline'-attribute is used to identify purged files, it is recommended that no files have this attribute set by other applications. Otherwise ProLion DataOptimizer may remove this attribute for files which have not been purged by ProLion DataOptimizer.

4 Using ProLion DataOptimizer

This chapter provides detailed information about the usage of ProLion DataOptimizer.

4.1 ProLion DataOptimizer GUI

The ProLion DataOptimizer GUI (graphical user interface) is used to configure ProLion DataOptimizer. It can be started by selecting the appropriate command from the Start Menu. Provided that the permissions have been granted, the user interface can be used to administer a ProLion DataOptimizer Service which runs on the local computer or on another computer in the network. To connect to another computer select **Tasks > Select Computer** from the menu.

After installation ProLion DataOptimizer allows only users with administrative privileges to connect to the service. Therefore on computers with enabled UAC (User Account Control) it is necessary to start the user interface with elevated rights (As Administrator). After connecting the first time, ProLion DataOptimizer remembers the account and allows connecting even if the account does not have administrative privileges. You can configure the list of users in the setup dialog (see 4.2, Setup User Accounts).

4.1.1 Side Bar

The Side Bar is on the left side of ProLion DataOptimizer user interface. It is used to group commands which belong together.

The topic **Tasks** provides some global tasks which can be performed. A click on a command will start a wizard or dialog which helps performing the task. After clicking on **Tasks** itself, the right pane displays the start screen of the user interface which provides global status information.

The topic **Storage Vaults** lists all configured Storage Vaults. After selecting one Storage Vault the right pane of the user interface opens the Storage Vault detail view which in turn provides status information and allows performing operations on the Storage Vault (see 4.1.2, Storage Vault Detail View)..

The topic **Devices** lists all configured Archive Devices. After selecting a device the right pane displays the device details view which provides status information about the selected device and allows performing operations on the device. Note, that Archive Tier is not supported.

The topic **Current Activity** can only be selected while an Archiving Job Cycle is running or a migration of Archive Volumes is in progress. After selecting the topic, the right pane shows information about the current activity.

After a click on the topic **Log Files**, the right pane displays all available log files. They are grouped into global log files and log files which belong to a Storage Vault.

4.1.2 Storage Vault Detail View

The Storage Vault detail view can be opened by selecting the Storage Vault in the topic Storage Vaults in the Side bar. It provides one section for each of the configured tiers of the Storage Vault as described below.

The detail view also provides buttons to activate or deactivate the Storage Vault and to view or modify its configuration. While the Storage Vault is active you can click on the button **Deactivate Storage Vault** to deactivate it. Note that data which has been migrated from Performance Tier to Capacity Tier will not be accessible while the Storage Vault is not active. Also, if you try to deactivate a Storage Vault while a Migration Job Cycle is currently running, that job cycle will be cancelled.

The button **Configure Storage Vault** allows viewing or changing settings of the Storage Vault.

Performance Tier Section

The Performance Tier section provides an overview of the selected Data Source of the Storage Vault. A progress bar visualizes the capacity usage of the Data Source.

Capacity Tier Section

This section provides information about the configured Capacity Tier.

The **Status** field displays whether a Migration Job Cycle is currently running and the time the next cycle has been scheduled. If no cycle is currently running, it is possible to trigger a Migration Cycle by clicking the button **Execute Migration Policies**. See also chapter 4.5, Performing Job Cycles, for more information about job cycles.

The **Log Files** button allows opening log files related to the Capacity Tier. Refer to chapter 5, Error Reporting, for more information about logging and log files.

The button **Actions** provides the command **Generate HTML View** to open a web browser to view the contents of the Capacity Tier.

4.2 Setup User Accounts

By default, ProLion DataOptimizer allows users with administrative privileges to administer the software. To allow other users the administration of ProLion DataOptimizer, open the setup dialog, select **General** and then click on **User Accounts**.

If this right is granted to a user, this user will be able to manage all configured Storage Vaults and change all configuration options of ProLion DataOptimizer.

4.3 Configure a Storage Vault

To start configuring a new Storage Vault select **Tasks** → **Create Storage Vault** from the **Side Bar** menu. A wizard will guide you through the mandatory settings for the new Storage Vault. You find a description of the wizard steps in the following chapters. Please note that the available options and settings vary depending on the selected Data Source Type.

You can directly jump to a topic using the left navigation menu or by selecting an item in the pictogram on the top of the dialog. However, the recommended method is to completely fill out the active view and then click on **Next** to step to the next view.

In the first step of the wizard you will be prompted to enter a name for new Storage Vault and to select the type of the Data Source that will be used. Please refer to chapter 1.4, Performance Tier, for a description of available Data Source Types.

The Storage Vault name can be any text which is a valid Windows file name. The name must not start or end with a white space and it must not contain any of the following characters: \ / : " ' * ? > < | ,

The length is limited to 50 characters.

The Storage Vault name will be used to identify the Storage Vault in the user interface. That's why you should select unique names for your Storage Vaults, even though this is not enforced by ProLion DataOptimizer. Internally ProLion DataOptimizer uses a unique Storage Vault ID which will be automatically created and assigned.

4.3.1 Performance Tier Settings

All Data Source Types require the specification of a Data Source directory. To add a Data Source directory click on **Add Directory**. To change a directory, first click on Remove and then add another directory.

For NetApp FAS (7-Mode) Data Source it is possible to specify additional parameters which will be passed to the NFS mount command. Please refer to the documentation of your NFS Client software and the documentation of the respective Agent for more information.

In case of NetApp FAS (Cluster Mode) you need to specify an ONTAPI login which allows configuring FPolicy on the Storage Virtual Machine. Please also refer to the documentation for the NetApp FAS Agent and the NetApp FPolicy Server for Cluster Mode.

4.3.2 Migration Policies

The specification of Migration Policies is required to automatically copy or migrate files between Performance Tier and Capacity Tier. These policies will only be evaluated if a Capacity Tier has been specified for the Storage Vault. Therefore you don't need to specify any policies here if you do not plan to use a Capacity Tier.

The following different sets of Migration Policies can be assigned to a Storage Vault:

Default Migration Policies

These policies will be regularly executed for all files based on the specified schedule. You can use these policies to perform operations on a daily basis, for example to copy all new files to the Capacity Tier.

Capacity Threshold Policies

The Capacity Threshold Policies will be executed if the used space on the Performance Tier reaches the specified threshold. Therefore, these policies can be used to purge specific files from Performance Tier to release disk space.

These policies will be executed for all files, unless specific conditions have been specified to filter some files.

Access to Migrated Files

The **On Access Policies** will be executed whenever a file is accessed which has been purged from the Performance Tier. The policies will only be executed for the files which have been accessed.

Use these policies to initiate retrieve to Performance Tier after a file access.

Please refer to chapter 4.4, Working with Policies, for more details related to policies.

4.3.3 Capacity Tier

In case you want to use a Capacity Tier for the Storage Vault, click on **Add Directory** to select the directory which will be used as Capacity Tier for the Storage Vault. Click on **Remove** to clear the current Capacity Tier setting.

The selected directory must not be specified as Data Source or as Capacity Tier for another Storage Vault. Also, the directory must not be used for other purposes or by other applications.

The directory shall be empty during first configuration of a Storage Vault. If you are reconfiguring the Storage Vault after a recovery, you should create and run migration policies which ensure consistency between the data on the Performance Tier and the files which already exist in Capacity Tier. Please refer to the chapter 6, Backup and Recovery, for more information about recovery situations. When reconfiguring a Storage Vault make sure to specify the storage layout which has been used when initially creating the Storage Vault.

Also, you should not specify the root directory of a volume as Capacity Tier because there are usually hidden files like "System Volume Information" which are required by the operating system.

Capacity Tier with Replication

If two paths have been specified the first path will be used as primary Capacity Tier path, i.e. files are copied there during Migration Job Cycles and purged files are read from this path. The second path will be used as replication target and all files from the primary path will be replicated there during Migration Job Cycles. Replication is supported for Capacity Tiers using Object Based storage layout or Original File Structure with WORM Compatibility.

Refer to chapter 6.2, Recovery of Performance or Capacity Tier, for more information about recovery in case of a Capacity Tier failure.

Storage Layout is Object Based

ProLion DataOptimizer stores files on the Capacity Tier using an object based file structure. Each file copied to the Capacity Tier will be assigned a unique ID which will then be stored in the files metadata on the Performance Tier. This allows access to purged files even if the name of the file has changed outside of the control of ProLion DataOptimizer, for example after restoring a backup to a different location or by accessing files through a snapshot.

Update Database during Jobs

If this option is enabled, file names of archived files will be updated in the database when running migration job cycles. Additionally, this operation will update a flag in the database which indicates that the file was found on the Performance Tier during this job.

This operation may slightly decrease the overall performance of migration jobs and it may increase the CPU and hard disk usage of the computer running ProLion DataOptimizer.

Create Content View after Job

This option is available when using the Object Based Storage Layout. If enabled, Storage Agent automatically creates/updates the index files for browsing the Capacity Tier contents after each job cycle. Otherwise these files will be created on demand when selecting the command **Actions > Browse Content of Capacity Tier** in the Capacity Tier section of the Storage Vault detail view.

Depending on the number of files in the Capacity Tier, creating the browse index may take several minutes. Therefore this option should only be enabled if it is required to browse the Capacity Tier contents regularly.

Jobs Before Deleting Files

This option can be used to specify after how many jobs a file will be removed from the Capacity Tier, if it has been deleted from the Performance Tier. Specify 0 to prevent removing files from Capacity Tier.

If a value other than 0 has been specified, files will be removed from the Capacity Tier if they have not been enumerated on the Performance Tier for the specified number of jobs. Note that only those jobs will be counted which had the option **Update Database during Jobs** enabled and which were not limited to a specific directory.

4.4 Working with Policies

4.4.1 Overview

ProLion DataOptimizer uses user defined policies to control the migration of the data between different storage tiers. They can be used to copy files from one tier to another, to replace files by a stub or even to delete files from a tier.

Migration Policies are used to move data between the Performance Tier and Capacity Tier. Policies always consist of one or more conditions and one or more actions. The conditions will be evaluated for each file until one condition does not match. The actions will be performed on a file if all conditions did match.

While it is possible to create policies with more than one action, it is recommended to use only one action per policy and create multiple policies if more actions are required. The policies should then be assigned in the correct order. For example, a policy which purges files from the Performance Tier shall be inserted after the policy which copies the file from Performance to Capacity Tier. Also, policies should always check the status of a file before performing an action. For example, the policy which purges a file must ensure that a copy of the file exists on the Capacity Tier and that the file has not already been purged.

While executing actions for a file, ProLion DataOptimizer does not perform prerequisite operations which are eventually required to successfully perform an action. For example, the policy which tries to purge a file from Performance Tier would fail if the file has not been copied to another tier yet.

To test configured policies, it is also possible to simulate the execution of policies. This can be done by manually triggering a job cycle and checking the option to simulate the job execution. Refer to chapter 4.5, Performing Job Cycles, for more information.

4.4.2 Templates

ProLion DataOptimizer automatically stores all policies which are assigned to a Storage Vault as global templates. These templates can be loaded and modified when creating or editing a policy. Because templates are identified by their names, new or modified policies will replace existing policy templates with the same name.

Existing templates can be modified and deleted in the ProLion DataOptimizer setup dialog after selecting **Tasks** → **Setup ProLion DataOptimizer** → **General** → **Policies**.

Note: Changes of templates will not change policies which are already assigned to Storage Vaults.

4.4.3 Conditions

The following list shows all conditions which can be selected in Migration Policies.

File name matches pattern <param>

<param> is a list of patterns, e.g. '*.txt; *.doc; temp\inbox'. The condition is met, if the file name matches at least one pattern. If a given pattern contains a path separator (backslash), then the pattern will be compared against the relative file path based on the root directory of the Storage Vault. Otherwise the pattern will be compared against the file name without its path.

Examples ('C:\Vault' is the Storage Vault root):

'*.txt' will match all files with extension 'txt'

'temp\inbox.txt' matches only the file 'C:\Vault\temp\inbox.txt'.

'temp*' matches all files in the directory 'C:\Vault\temp' and its sub-directories.

Status on TIER is <param>

TIER is one of the two storage tiers.

<param> is a combination of 'File is present', 'File stub is present', 'Old version is present' or 'File is not present'.

The status 'File is present' indicates that the current file version exists on the specified tier.

Status 'File stub is present' indicates that the file has been replaced by a stub on this tier. This status can only occur on the Performance Tier.

'Old version is present' indicates that this tier has a copy of the file, but another version exists on Performance Tier. Consequently, this status cannot occur on Performance Tier.

Status is set to 'File is not present', if no copy and no stub exists on this tier.

ProLion DataOptimizer compares creation and modification time stamps to determine whether a file has been changed or not.

Free space on Performance Tier is less/higher than <param>

<param> specifies the used capacity in percent.

This condition can be used execute the actions only if the specified capacity threshold has been reached on the Performance Tier.

File is older than <param>

If the file has not been modified since the time period specified, the condition is true. By default, ProLion DataOptimizer checks both, the creation time (C-Time) and the modification time (M-Time) of the file. The condition becomes true if both timestamps are older than the specified value. By clicking on the appropriate field in the policy parameter list, it is possible to specify that only one of the timestamps shall be evaluated.

File has not been accessed for <param>

In this condition <param> specifies a time period. In case the file has not been accessed within the time period, the condition is met. ProLion DataOptimizer uses the most recent timestamp (creation, modification or access) of a file to evaluate this condition.

File has been accessed within <param>

<param> specifies a time interval. The condition evaluates to true, if one of the timestamps of the file is more recent than the specified time interval.

File attributes <param> are set

<param> is a combination of following file attributes: 'Archive bit', 'Read-only', 'System', 'Hidden' or 'Offline'. The condition is true, if all specified file attributes are set on the file on the Performance Tier.

File attributes <param> are not set

Analogue to the condition above, <param> is a combination of 'Archive bit', 'Read-only', 'System', 'Hidden' or 'Offline'. This condition is met, if none of the specified file attributes is set.

File is larger than <param>

This condition checks the file size. In case file is larger than the specified amount, the condition is true.

File is smaller than <param>

In case the file is smaller than the specified amount, the condition is met.

File belongs to <param> oldest files

<param> a percent value.

The condition is matched if the file belongs to the given percent number of oldest files based on the used capacity.

File had more than <param> accesses per <param1>

<param> value in the range 1 through 9

<param1> can be set to number hours, days or weeks

The condition becomes true, if the file has been read more often than the specified number of times within the given time. Note that read requests will be counted as independent requests if there was no request for at least 60 seconds. Also, read requests may not be counted if the client file system caches read data for a long time.

4.4.4 Actions

The following list shows all actions which can be selected in Migration Policies.

Purge the file

This action replaces the file on the Performance Tier by a stub.

Delete the file and parent

The file will be removed from the Performance Tier.

Additionally this action removes the parent directory chain up the root of the Storage Vault if these directories are empty after deleting the file.

Copy to <Tier>

This action can be used to copy the current version of the file to the specified tier. If the file is present on the Performance Tier, that copy will be treated as the current version.

This action can also be used to retrieve a file on Performance Tier which had been previously purged.

Delete on <Tier>

Use this action to remove a file from either the Performance or Capacity Tier. This function fails if you try to delete the last copy of a file which has been purged from Performance Tier. To permanently remove such a file it is necessary to delete it on Performance Tier first.

Note: Use this action with care because it can be used to irrecoverably destroy data!

Recover file

This action can be used to recover files on the Performance Tier. Depending on the status of the file on the Performance Tier, the action does one of the following: If a file stub exists, this action ensures that the stub points to the correct file in either the Capacity Tier. If the file does not exist, the action creates a file stub. This action does nothing if the file already exists on Performance Tier.

Log a message

When you add this type of action to a policy, you can specify a message text. When this action is executed, the message is logged into the Job Cycle's log file. The message is composed of the file name and the specified message text.

Skip further policies

All further actions and policies will be skipped for this file.

4.5 Performing Job Cycles

Job Cycles are used to execute defined policies on the files in the Storage Vault. They will usually be triggered using a defined schedule, but it is also possible to trigger them manually.

Migration Job Cycles are used to execute Migration Policies (see also 4.4, Working with Policies). These policies provide actions to copy or migrate files between Performance and Capacity Tier.

To manually trigger a Migration Job Cycle, click on the button **Execute Migration Policies** (green arrow) in the section **Capacity Tier** on the Storage Vault details view.

After selecting the command, a dialog appears which allows modifying the list of policies to be executed. Please note that changes will not be stored to the Storage Vault configuration. To permanently change the policies for the Storage Vault open the properties dialog and change the default Migration Policies there.

Besides modifying the policy list, this dialog allows specifying a file or a directory within the Storage Vault. If specified, the policies will only be executed for that file or for files in the specified directory, respectively. Please note that the function which helps to select a file or directory only shows files and directories on the Performance Tier. If you want to recover a file or directory, which has been removed from the Performance Tier, you need to enter the path name directly into the text field.

By default, recover policies will only process files which are not yet marked as 'deleted' in the database. This means, files which were not present during the last job, which had the option **Update Database during Jobs** enabled, will not be recovered. To recover even those files check the option **Restore deleted files** when executing the recover policy.

By checking the **simulate** option, it is possible to simulate the policy execution to test the configured parameters. If this option has been checked, all conditions will be evaluated, but actions will not be performed. The job summary statistics will then indicate how many files would have been affected by the configured policies.

After a job cycle has been triggered, either manually or by schedule, the progress and error information can be monitored using the **Migration job log file**. The file can be opened using the button **Log Files** in the section **Capacity Tier** or by means of the **Log Files view** (see 5.1, Log Files).

4.6 Data Browser

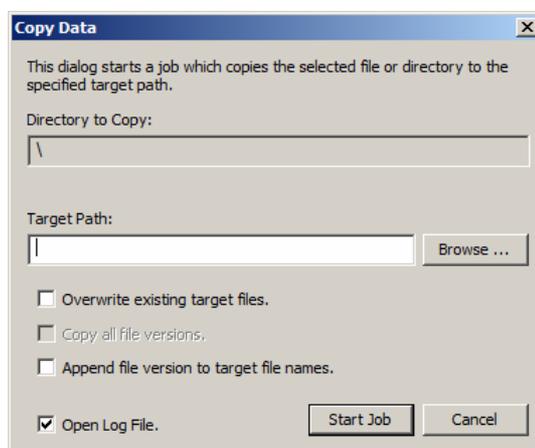
The Data Browser can be used to view all files which are stored in either the Capacity Tier. Additionally, it allows performing data authentication jobs, to copy data from the Capacity Tier to another location and it provides functions to select and to delete files from the Capacity Tier.

To open the Data Browser click on the following icon the Storage Vault detail view:



4.6.1 Copy Data to another Location

This function can be used to copy data from the Capacity Tier to another location. It is available in the context menu of a directory, to recursively copy the whole directory, and in the context menu of selected files to copy those files. After selecting the command, the following dialog appears:



The specified target path must be accessible for the Storage Agent account. Otherwise the operation will fail. Note that the files and directories will be created on the target location using their full path names. For example, if you copy a file "dir1\file1", ProLion DataOptimizer will create a folder "dir1" below the specified target path and put the file "file1" into that directory. This way it is possible to specify the root path of the Storage Vault if you want to copy the files to their original location.

Overwrite existing target files

If this option is checked, files with the same name in the target location will be replaced. Otherwise they will be skipped.

Copy all file versions

Check this option to copy all archived versions of the selected file to the target location. The time stamp of the file version will be appended to the names of the generated files.

Append file version to target file names

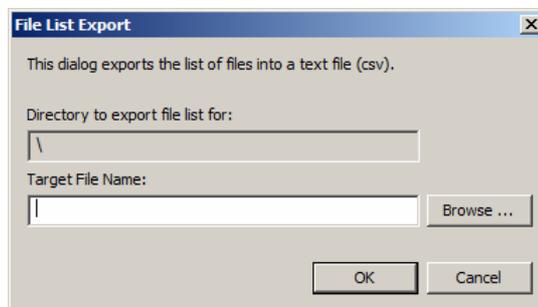
If checked, the target file names will include the time stamp of the file version.

Finally select **Start Job** to perform the operation. The job will internally be performed as a migration/recovery job. Progress and error information will be written to the Storage Vaults **Copy data job log**.

4.6.2 Export List of Files

This command allows exporting a list of files in the Storage Vault to a comma separated text file (csv). It is possible to export the list for selected directories only or for the whole Storage Vault by executing the command from the context menu of the root node in the directory tree.

After executing the command, the following dialog appears:



Select a file to copy the file list to and click on OK. Depending on the number of files, the operation may take several minutes. The Administration Tool will be blocked during that operation.

4.7 Reconfigure Storage Vault Data Source

This procedure describes how to reconfigure a Storage Vault so that archived data can be restored to a primary storage system which requires a different Storage Vault data source type. For example, it is possible to reconfigure a Storage Vault which has originally been configured for a NetApp FAS 7-Mode, so that it then allows restoring archived data to a NetApp cDot system.

To achieve this it is necessary to first delete the Storage Vault configuration using the command "Reconfigure Storage Vault" from the application menu. After removing the configuration, reconfigure the Storage Vault by means of the button **Configure Storage Vault** in the Storage Vault Detail View.

Note: Data on the old primary storage system will not be accessible after reconfiguring the Storage Vault, unless purged files have been recalled and metadata (reparse points) have been removed.

Execute recover policies to restore the data to the new primary storage system as described in chapter 4.4, Working with Policies.

5 Error Reporting

The following chapter provides information about logging and error notification mechanisms of ProLion DataOptimizer.

5.1 Log Files

ProLion DataOptimizer creates several log files which contain informational, warning or error messages. Severe error messages will also be written to the Windows Application Event Log.

The ProLion DataOptimizer log files can be viewed after selecting the topic **Log Files** in the Side Bar of the ProLion DataOptimizer user interface. The log files view on the right side now shows a list of all log files. You can double click on an entry to view the related log file. It is also possible to clear or permanently remove log files here.

The list begins with the **Global Log Files**, which contains log files of each of the services belonging to ProLion DataOptimizer.

Under that, for every Storage Vault one or more log files are displayed. Depending on the configuration of the Storage Vault, the list may contain a migration log file and one or more archiving job log files. Log files created by other operations like Data Retrieval or Data Authentication will also be listed here if you chose to keep the log files.

5.2 Server Notifications

While the ProLion DataOptimizer user interface is running, the ProLion DataOptimizer service reports conditions which require user interaction as so called server notifications. These notifications pop-up in a small window which is only visible while notifications are present. Double click on a message to receive more information about the condition. If possible ProLion DataOptimizer also suggests actions which help to resolve the condition.

5.3 E-Mail Notifications

You can instruct ProLion DataOptimizer to send E-Mail notifications about certain events to computers and/or users that you specify in this dialog.

To use E-Mail notifications, an SMTP host and a sender address must be specified. The SMTP host should be configured to accept SMTP connections from the computer running ProLion DataOptimizer. Also, it must not require an SMTP-after-POP authentication. The sender address will be added to the generated e-mail so that the recipient knows where the e-mail originates from. It is not necessary to create a special e-mail account for ProLion DataOptimizer; you may also use your own account.

You can choose to be notified if one or both of the following events occur:

- A hardware failure occurs
- An input bin is empty, or the required media type is not available.

Use the **Add**, **Edit**, and **Delete** buttons to modify the list of users. If the list is empty, no notifications will be sent.

Use the **Test** button to ensure that messages can be sent to a specific user or computer.

6 Backup and Recovery

This chapter describes backup and recovery procedures and requirements for ProLion DataOptimizer server systems and for the configured storage devices.

6.1 Recovery of ProLion DataOptimizer Server

ProLion DataOptimizer deposits information it needs for proper operation in database and configuration files which are located on local hard disks of the server system. If the corresponding files are lost or damaged due to a system failure, there are three methods to restore parts or all of them.

Availability of the methods depends on the backup measures you have taken, and on whether ProLion DataOptimizer shall be recovered on the same or on a different (new) server system.

Beginning with an introduction, the following sections describe the three methods in detail:

- **System recovery by importing database files:**
This method recovers from server system failures and works on the same and on a different server system by re-importing a backup of database and configuration files which has been exported by ProLion DataOptimizer and saved earlier, e.g. in the course of a backup procedure. You additionally need to know the System Key to apply this method.
- **System recovery by third party backup program:**
This method recovers from failures of a server system by restoring a backup of database and configuration files which has been saved by a 3rd party backup program earlier, e.g. in the course of a backup procedure.

Note: It is absolutely necessary that the System Key is stored in a safe place in order avoid loss of archived data. Use the System Key Export function to save the System Key.

6.1.1 Recovery by Importing Database Files

Database Backup

You can save the database and configuration files of ProLion DataOptimizer by means of the **Database Backup** function, which you can find in the **System Recovery** dialog of **ProLion DataOptimizer Setup**.

It is recommended to backup the database at a time where no migration process is busy, because data which have not been migrated yet are not included in the exported database file and cannot be restored later. Also make sure that you store the System Key on a safe place, because it is required to successfully import the database later.

Database Import

If you had saved the database and configuration files of ProLion DataOptimizer by means of the **Database Backup** function, you can proceed as follows:

- Install the operating system and ProLion DataOptimizer.
- Import the System Key of the old installation during setup of ProLion DataOptimizer.
- Import the database and configuration files by means of the **Database Import** function.

After these steps, ProLion DataOptimizer is in the state as before the failure except the data that were not yet written to target devices at the time of the database export are lost. Also, any changes made to the settings of the Storage Vaults and device configurations export are lost. It is strongly recommended to review all settings before continuing to work with ProLion DataOptimizer.

Any data that were not written to target devices are lost.

Note: The Storage Agent is stopped during database import. Therefore, no Job Cycles are executed at this time. Database import and export can take a long time, up to several hours.

6.1.2 Recovery by 3rd Party Backup Programs

In contrast to the scenarios above, this method can only be used on the system where the backup was done.

Backup and restore should be done only if the ProLion DataOptimizer Service has been stopped. It can be stopped and started in the Services applet of the operating system or by the command line

“net stop PntArchiverService” or

“net start PntArchiverService”, respectively.

Note that the service should not be stopped when Job Cycles are being executed; they would be aborted.

Backup

All of the following files must be saved – and also restored later - at the same time. The files to include in the backup are:

- **conf_db.sys, database.sys**
located in the installation directory of ProLion DataOptimizer
- **vol_db.sys, dirCache.sys, *.db**
located in the Database Directory. Refer to the **Working Directories** dialog of the **Global Settings**.
- All files in the Image Directory. Refer to the **Working Directories** dialog of the **Global Settings**.
- All files in the directory **config**
which is located in the installation directory of ProLion DataOptimizer
- All files in the directory **logs**
which is located in the installation directory of ProLion DataOptimizer
- The system registry
- The program files of ProLion DataOptimizer

Restore

In order to reset the system to a specific state, all files in the Image Directory and directories “config” and “logs” must be deleted and then restored from the backup. If the total system is to recover, all the backed up files including the program files and the registry must be restored.

6.2 Recovery of Performance or Capacity Tier

ProLion DataOptimizer supports you in recovering from failures of Performance Tier by means of special recovery policies and based on the user data and metadata (e.g. files attributes and security information) it has stored on Capacity Tier at the time when it has copied the data in the course of a migration job cycle. Please note that this does neither cover new files and file versions nor metadata changes which have been created or applied on the Performance Tier since then. Consequently, you need to continue making regular backups of the Performance Tier by means of a 3rd party backup solution of your choice in order to be able to completely recover from failures of the Performance Tier. Depending on your requirements, the configuration of an existing backup solution may have to be adjusted in order to avoid backing up purged files (i.e. files carrying an offline file attribute) and in order to correctly include reparse tags in the backup.

Failures of a Capacity Tier device may be recovered from by using a replication device or, again, by backing up the Capacity Tier device by means of a 3rd party backup solution. In case that you use a backup solution, the backup of the Performance Tier should be performed before the backup of the Capacity Tier, and you should not execute migration job cycles by ProLion DataOptimizer in parallel to the backup processes. You may use the command line tools of ProLion DataOptimizer to synchronize backup processes and job cycles. Corresponding recovery processes should be performed in opposite order.

Please note: After restoring a Performance or Capacity Tier device from a backup, it is required to re-synchronize the tiers and, potentially, to repair damaged stub files on the Performance Tier. This can be achieved by manually performing Migration Jobs which execute appropriate policies as described in the following sections. Refer to section 4.4, Working with Policies, for more information on principles of creating and executing policies.

6.2.1 Recover File Stubs on Performance Tier

The policy action “Recover File” can be used to create a file stub on Performance Tier for files which exist on the Capacity Tier but not on Performance Tier. This action also repairs file stubs if they belong to another copy than the current file in Capacity Tier.

The predefined policy template “RECOVERY: Recover files on Performance Tier” can be used to execute this action for all files.

6.2.2 Delete Orphaned File Stubs on Performance Tier

If file stubs exist on the restored Performance Tier which have no counterpart on the Capacity Tier, e.g. because the file has been deleted after the backup has been taken from the Performance Tier, these stubs must be removed because they are not accessible anymore. This can be achieved using a policy which checks if the status on Performance Tier is “stub” and on the other tiers “not present”. If this is the case, the policy would use the action “Delete on Performance Tier” to remove the stub. The predefined template “RECOVERY: Delete orphaned file stubs” performs these steps.

6.2.3 Recover Capacity Tier with Replication

Failure of the primary path

If the primary Capacity Tier device failed, it is possible to switch to the specified replication path by simply removing the primary path from the configuration. However, depending on the time and the type of the failure of the primary device, it may be necessary to perform recovery policies as described above.

After configuring the replication path as primary path and running the recovery policies, it is possible to add another device as replication path.

Failure of the replication path

In case of a failure of the replication path, it is possible to remove the configured path from the configuration and specify another path. All data from the primary path will then be replicated to the replication path during the next Migration Job Cycle.

7 Index

B

Backup and Recovery 30

C

Capacity Tier 21

Capacity Tier Replication 21, 32

I

Installation 8

J

Job Cycles 25

L

Licensing 9

N

NetApp FAS Agent 6

P

Performance Tier 6, 20

Policies 20, 22

Product Key 9

S

Serial Number 9

System Key 8