

# Symantec NetBackup PureDisk™ Best Practices Guide

Windows, Linux, and UNIX

Release 6.6.5

Revision 1



The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Documentation version: 6.6.5, revision 1

## Legal Notice

Copyright © 2013 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo and are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation  
350 Ellis Street  
Mountain View, CA 94043

<http://www.symantec.com>

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

# Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website at the following URL:

[www.symantec.com/business/support/](http://www.symantec.com/business/support/)

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

## Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

[www.symantec.com/business/support/](http://www.symantec.com/business/support/)

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
  - Error messages and log files
  - Troubleshooting that was performed before contacting Symantec
  - Recent software configuration changes and network changes

## Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

[www.symantec.com/business/support/](http://www.symantec.com/business/support/)

## Customer service

Customer service information is available at the following URL:

[www.symantec.com/business/support/](http://www.symantec.com/business/support/)

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs, DVDs, or manuals

## Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan [customercare\\_apac@symantec.com](mailto:customercare_apac@symantec.com)

Europe, Middle-East, and Africa [semea@symantec.com](mailto:semea@symantec.com)

North America and Latin America [supportsolutions@symantec.com](mailto:supportsolutions@symantec.com)

# Contents

Technical Support .....	4	
Chapter 1	PureDisk best practices overview .....	9
	About this guide .....	9
Chapter 2	Planning deployment, data selections, and backups .....	11
	Deploying PureDisk with the best results .....	11
	About performing initial backups manually or with the wizard .....	12
	About policy tuning .....	12
	About specific data selections when you use snapshot support .....	13
	Simulating a traditional backup plan .....	13
	Creating and copying data selection templates for daily, weekly, and monthly backups .....	14
	Applying data selection templates to the clients .....	15
	Creating policies for daily, weekly, and monthly backups .....	16
	Creating data removal policies for daily, weekly, and monthly backups .....	18
	About creating policy escalation actions .....	22
	About configuring users to run under the Backup Operator .....	22
Chapter 3	Creating data removal policies for specific files .....	23
	Creating data removal policies for specific types of files .....	23
Chapter 4	Comprehensive policy scheduling .....	25
	Backup, replication, and maintenance policy scheduling .....	25
Chapter 5	Tuning PureDisk .....	29
	Configuration parameters .....	29
	Changing TCP/IP settings to improve replication job performance .....	32
	About deduplication results for Oracle backups with PDDO .....	33
	Running the Proxy Copy with RMAN .....	34

	Using the Image copy .....	34
Appendix A	Third-party legal notices .....	37
	Third-party legal notices for the Symantec NetBackup PureDisk product family .....	37
	Third-party trademarks for the Symantec NetBackup PureDisk product family .....	37
	Glossary .....	39
	Index .....	49



# PureDisk best practices overview

This chapter includes the following topics:

- [About this guide](#)

## About this guide

The *PureDisk Best Practices Guide* includes information about how to employ PureDisk's features so that they interoperate effectively. This guide expands on the information that is included in the other PureDisk guides.

The complete PureDisk documentation set is as follows:

- *PureDisk Administrator's Guide*
- *PureDisk Backup Operator Guide*
- *PureDisk Best Practices Guide*
- *PureDisk Client Installation Guide*
- *PureDisk Getting Started Guide*
- *PureDisk Deduplication Option Guide*
- *PureDisk Storage Pool Installation Guide*
- *PureDisk Virtual Appliance Getting Started Guide*



# Planning deployment, data selections, and backups

This chapter includes the following topics:

- [Deploying PureDisk with the best results](#)
- [About performing initial backups manually or with the wizard](#)
- [About policy tuning](#)
- [About specific data selections when you use snapshot support](#)
- [Simulating a traditional backup plan](#)
- [About configuring users to run under the Backup Operator](#)

## Deploying PureDisk with the best results

The following procedure describes the general steps to deploying PureDisk with the best results

### To deploy PureDisk with the best results

- 1 Design your storage pool capacity strategy carefully.  
Use the storage pool calculator to determine an initial storage pool size that can accommodate the data you plan to back up and store.  
Be conservative in your estimates so that you can allow for growth. Use a larger storage pool capacity when in doubt.
- 2 Gradually deploy PureDisk on a small group of servers. A gradual, measured approach helps you to avoid filling up storage and lets you see how different data types affect capacity and deduplication rates.

- Start with a few servers that represent a variety of data types in your environment. For example, you can see the best deduplication rates from file servers, so deploy PureDisk first on a large file server.
  - Next, add another large file server and observe the deduplication rate.
  - Add additional servers that contain other types of data, such as a database server, a mail server, or a VMware server. Again, observe the effect of these data types on storage capacity and deduplication rates.
- 3 Based on your observations, adjust your storage pool strategy if necessary. You may decide to add more space to the storage pools if it appears that the initial storage estimates are too low. Or you may accelerate the deployment if your storage pool strategy seems to accommodate the additional data.
  - 4 Continue to deploy PureDisk to additional servers in a gradual manner.

## About performing initial backups manually or with the wizard

For an initial backup, Symantec recommends that you run the job manually or from the backup wizard.

Information about backups, timeouts, and how to enable timeout notifications is available.

See the *PureDisk Administrator's Guide*.

PureDisk terminates a backup or restore job if it does not complete in a reasonable amount of time, as follows:

- For a manual job, PureDisk generates an error message if the job does not complete within seven days. PureDisk terminates the job if it does not complete in 14 days.
- For a job that is run from a policy, PureDisk generates an error message if the job does not complete within six hours. PureDisk terminates the job if it does not complete in five days.

## About policy tuning

If you have more than 200 clients, make sure that each PureDisk backup policy includes no more than 200 data selections.

# About specific data selections when you use snapshot support

Snapshot support enables PureDisk to back up those files that are open at the time the backup is performed. If you use snapshot support, make sure that your data selections are specific. Preferably, specify your data selections on a per-drive basis.

More information about how PureDisk uses snapshot support is available.

See the *PureDisk Backup Operator Guide*.

## Simulating a traditional backup plan

Typical backup plans usually include daily, weekly, and monthly backups. Such plans usually include a policy for each backup interval. That policy specifies the data to back up, the system affected, and how long to retain the data. Because these traditional plans write both the file content and the metadata to backup media, they usually consume large amounts of tape or disk media.

If you have several similar clients, you can use PureDisk to create a traditional backup plan for these clients. Unlike the typical backup methods that consume large amounts of backup media, PureDisk uses its data deduplication technology to store file content only one time. When you run PureDisk backups at daily, weekly, and monthly intervals, you ensure that you have a record of the metadata at these intervals. In addition, PureDisk also backs up any file content that is new or changed since a previous backup was performed.

To implement this plan, create data selections for each of these intervals, apply them to clients, and create backup policies. To ensure that these backups are removed from the system after a suitable retention period, also create three data removal policies.

The following table lists the steps to simulating a traditional backup plan.

**Table 2-1** Simulating a traditional backup plan

Step	Description
Step 1	<p>Create a template for daily backups and copy that template for weekly and monthly use.</p> <p>See <a href="#">“Creating and copying data selection templates for daily, weekly, and monthly backups”</a> on page 14.</p>

**Table 2-1** Simulating a traditional backup plan (*continued*)

Step	Description
Step 2	Apply the daily, weekly, and monthly data selection templates to each client. See <a href="#">“Applying data selection templates to the clients”</a> on page 15.
Step 3	Create a policy for daily backups and copy that policy for weekly and monthly use. See <a href="#">“Creating policies for daily, weekly, and monthly backups”</a> on page 16.
Step 4	Create a removal policy to remove any unneeded data that was backed up in daily, weekly, and monthly backups. See <a href="#">“Creating data removal policies for daily, weekly, and monthly backups”</a> on page 18.
Step 5	Create a policy escalation action for each policy. See <a href="#">“About creating policy escalation actions”</a> on page 22.

## Creating and copying data selection templates for daily, weekly, and monthly backups

The following procedures explain how to create a template for daily backups and how to copy that template for weekly and monthly use.

### To create a template for daily backups

- 1 Select **Manage > Data Selection Templates**.
- 2 In the right pane, click **Add Template**.
- 3 In the **Data Selection Template** panel, specify the files you want to back up. For example, specify the following items:
  - The name of the template. For example, specify **Daily My Documents backup** or **Daily data backup**.
  - The operating system. Select the operating system to which the data selection applies from the drop-down list
  - The directories to include in the data selection and the directories to exclude from the data selection.

More information about how to create a data selection template is available.

See the *PureDisk Backup Operator’s Guide*.

- 4 Click **Add** to save the new template.

**To copy a daily-use template and modify it for weekly and monthly use**

- 1 Select **Manage > Data Selection Templates**.
- 2 In the left pane, expand the tree and select the daily template that you previously created.
- 3 In the right pane, click **Copy Template**.  
 A copy of the template that you chose appears in the list of data selection templates in the left pane.
- 4 Select the copy of the template.
- 5 In the **Data Selection Template** panel, change the name of the template.  
 For example, change the template name to **Weekly - My Documents backup** or **Weekly - data backup**.  
 Do not change either the inclusion rules or the exclusion rules.
- 6 Click **Save**.
- 7 Perform the preceding steps again to create a monthly template.  
 For example, change the template name to **Monthly - My Documents backup**.

## Applying data selection templates to the clients

---

**Note:** If you did not yet create data selection templates, see the following topic:  
 See [“Creating and copying data selection templates for daily, weekly, and monthly backups”](#) on page 14.

---

**To apply a data selection template to a client**

- 1 Select **Manage > Agent**.
- 2 Expand the tree to locate the client or the group of clients to which you want to apply the template. You can apply the template to all the clients in a department or to all the clients in a storage pool.

- 3 Select the client name, the department name, or the storage pool name.  
To apply the template to more than one client, select a department name or the storage pool name.

---

**Note:** You may find it faster to apply the template at the department level than on each client individually. You may have several clients in a department, and you may have applied this template to some of the clients in a previous session. PureDisk does not reapply the template on the older clients.

---

- 4 In the right pane, select **More Tasks > Apply Data Selection Template**.
- 5 Select the template that you want to apply.
- 6 Under **Created data selections should inherit from the selected template**, select **Yes** to use the data selections that are specified in the template. This value is the default value. When this value is **Yes**, PureDisk applies any changes to the template to all of the clients that use the template.
- 7 Click **Apply**.  
Wait until PureDisk applies the template before you continue to the next step.

## Creating policies for daily, weekly, and monthly backups

The following procedures explain how to create a policy for daily backups and how to copy that policy for weekly and monthly use.

### To create a policy for daily backups

- 1 Select **Manage > Policies**.
- 2 In the left pane, expand **Policies > Backup Policies** and select a policy type.
- 3 In the right pane, click **Create Policy**. The policy creation dialog box opens.
- 4 Complete the fields on the **General** tab.  
Perform the following steps:
  - Type a name for the policy. For example, **Daily backup**.
  - Select **Enabled**.
- 5 Complete the fields on the **Data Selections** tab.  
Perform the following steps:
  - Under **Policy applies to the following data selections**, expand the tree to locate the client or the group of clients to which you want to apply the policy.



- Select the client or the client group. For example, you can apply the policy to all the clients in a department or to all clients in a storage pool.
  - Select **Apply all inclusion rules below to dataselections selected above**, and then select the name of the daily backup data selection from the **Data selections based on template** drop-down list.  
 See [“Creating and copying data selection templates for daily, weekly, and monthly backups”](#) on page 14.
- 6 Complete the fields on the **Scheduling** tab.  
 Perform the following steps:
- Select **Weekly schedule**.
  - Specify a **Start time**.
  - Click in the boxes to select **Monday, Tuesday, Wednesday, Thursday, and Friday**.
- 7 On the **Parameters** tab, select the parameters that are appropriate to this data selection for your site.
- 8 Click **Add**.

**To copy a daily use backup policy and modify it for weekly use**

- 1 Select **Manage > Policies**.
- 2 In the left pane, expand **Policies > Backup Policies** and select the daily policy that you previously created.
- 3 In the right pane, click **Copy Policy**.  
 A copy of the policy that you chose appears in the list of policies in the left pane.
- 4 Select the copy of the policy.
- 5 In the **Policy** panel, on the **General** tab, change the policy name.  
 For example, change the policy name to **Weekly backup**.
- 6 On the **Data Selections** tab, select the name of the weekly backup data selection from the **Data selections based on template** drop-down list.  
 See [“Creating and copying data selection templates for daily, weekly, and monthly backups”](#) on page 14.
- 7 Complete the fields on the **Scheduling** tab.  
 Perform the following steps:
  - Select **Weekly schedule**.
  - Specify a **Start time**.

- Select **Saturday**.
- 8 On the **Parameters** tab, select the parameters that are appropriate to this data selection for your site.
  - 9 Click **Save**.
- To copy a weekly use backup policy and modify it for monthly use**
- 1 Select **Manage > Policies**.
  - 2 In the left pane, expand **Policies > Backup Policies** and select the weekly policy that you previously created.
  - 3 In the right pane, click **Copy Policy**.  
A copy of the policy that you chose appears in the list of policies in the left pane.
  - 4 Select the copy of the policy.
  - 5 In the **Policy** panel, on the **General** tab, change the policy name.  
For example, change the policy name to **Monthly backup**.
  - 6 On the **Data Selections** tab, select the name of the monthly backup data selection from the **Data selections based on template** drop-down list.  
See [“Creating and copying data selection templates for daily, weekly, and monthly backups”](#) on page 14.
  - 7 On the **Scheduling** tab, perform the following steps:
    - Select **Monthly schedule**.
    - Specify a **Start time**.
    - Specify a day of the month.
    - Select all months in the year.
  - 8 On the **Parameters** tab, select the parameters that are appropriate to this data selection for your site.
  - 9 Click **Save**.

## Creating data removal policies for daily, weekly, and monthly backups

A data removal policy removes data from PureDisk storage if the data is no longer needed.

The following procedures show how to create a removal policy to remove any unneeded data that was backed up in daily, weekly, and monthly backups:

- On the **Data Selections** tab, you can identify files to remove from a particular data selection.
- On the **Metadata** tab, you can specify to remove only specific files. PureDisk displays this tab only for **Files and Folders** or **UNC Path** backups.

**To create a data removal policy for the daily backups**

- 1 Select **Manage > Policies**.
- 2 In the left pane, expand **Policies > Data Management Polices** and then select **Data Removal**.
- 3 In the right pane, click **Create Policy**.
- 4 In the **Policy** panel , on the **General** tab, perform the following steps:
  - Specify a name for the policy. For example, **Daily backup removal policy**.
  - Select **Enabled**.
- 5 On the **Data Selections** tab, perform the following steps:
  - Under **Policy applies to the following data selections**, expand the tree to locate the client or the group of clients to which you want to apply the policy.
  - Select the client or the client group. For example, you can apply the policy to all the clients in a department or to all clients in a storage pool.
  - Select **Apply all inclusion rules below to dataselections selected above**, and then select the name of the daily backup data selection from the **Data selections based on template** drop-down list.  
 See [“Creating and copying data selection templates for daily, weekly, and monthly backups”](#) on page 14.
- 6 On the **Scheduling** tab, perform the following steps:
  - Select **Weekly schedule**.
  - Specify a **Start time**.
  - Click in the boxes to select **Monday, Tuesday, Wednesday, Thursday, and Friday**.
- 7 On the **Parameters** tab, select the parameters that are appropriate to this data selection for your site. In **Remove versions backed up**, select **Older than (in days)** and specify 14 days.
- 8 (Conditional) On the **Metadata** tab, specify only specific files to remove.
- 9 Click **Add** to save the new policy.

### To copy a daily data removal policy and modify it for weekly use

- 1 Select **Manage > Policies**.
- 2 In the left pane, expand **Policies > Data Management Polices > Data Removal** and then select the daily data removal policy that you previously created.
- 3 In the left pane, click **Copy Policy**.  
A copy of the policy that you chose appears in the list of policies in the left pane.
- 4 Select the copy of the policy.
- 5 In the **Policy** panel , on the **General** tab, perform the following steps:
  - Specify a name for the policy. For example, change the policy name to **Weekly backup removal policy**.
  - Select **Enabled**.
- 6 On the **Data Selections** tab, perform the following steps:
  - Under **Policy applies to the following data selections**, expand the tree to locate the client or the group of clients to which you want to apply the policy.
  - Select the client or the client group. For example, you can apply the policy to all the clients in a department or to all clients in a storage pool.
  - Select **Apply all inclusion rules below to dataselections selected above**, and then select the name of the weekly backup data selection from the **Data selections based on template** drop-down list.  
See “[Creating and copying data selection templates for daily, weekly, and monthly backups](#)” on page 14.
- 7 On the **Scheduling** tab, perform the following steps:
  - Select **Weekly schedule**.
  - Specify a **Start time**.
  - Select **Saturday**.
- 8 On the **Parameters** tab, select the parameters that are appropriate to this data selection for your site.  
Under **Remove versions backed up**, select **Older than (in days)** and specify 35 days.
- 9 (Conditional) On the **Metadata** tab, specify only specific files to remove.
- 10 Click **Save**.

**To copy a weekly data removal policy and modify it for monthly use**

- 1 Select **Manage > Policies**.
- 2 In the left pane, expand **Policies > Data Management Policies > Data Removal** and then select and select the weekly data removal policy that you previously created.
- 3 In the left pane, click **Copy Policy**.  
 A copy of the policy that you chose appears in the list of policies in the left pane.
- 4 Select the copy of the policy.
- 5 In the **Policy** panel , on the **General** tab, perform the following steps:
  - Specify a name for the policy. For example, change the policy name to **Monthly backup removal policy**.
  - Select **Enabled**.
- 6 On the **Data Selections** tab, perform the following steps:
  - Under **Policy applies to the following data selections**, expand the tree to locate the client or the group of clients to which you want to apply the policy.
  - Select the client or the client group. For example, you can apply the policy to all the clients in a department or to all clients in a storage pool.
  - Select **Apply all inclusion rules below to dataselections selected above**, and then select the name of the weekly backup data selection from the **Data selections based on template** drop-down list.  
 See [“Creating and copying data selection templates for daily, weekly, and monthly backups”](#) on page 14.
- 7 On the **Scheduling** tab, perform the following steps:
  - Select **Monthly schedule**.
  - Specify a **Start time**.
  - Specify a day of the month.
  - Select all months in the year.
- 8 On the **Parameters** tab, select the parameters that are appropriate to this data selection for your site. In **Remove versions backed up**, select **Older than (in days)** and specify 365 days.
- 9 (Conditional) On the **Metadata** tab, specify only specific files to remove.
- 10 Click **Save**.

## About creating policy escalation actions

You can request that PureDisk notify you when a policy fails to run by creating a policy escalation action for each policy.

For example, you can create policy escalation actions for the three backup policies and the three data removal policies that you created.

More information about how to create policy escalation actions is available:

See the *PureDisk Backup Operator's Guide*.

## About configuring users to run under the Backup Operator

On Windows platforms, the PureDisk agent runs under `LocalSystem` by default. Configuring users to run under `Backup Operator` allows read permissions and permits PureDisk to read all of the ACL settings. The change also limits the user permissions that PureDisk allows for security purposes.

You are required to configure users and the PureDisk agent to run under `Backup Operator` when you back up mapped drives. You also may want to configure users and the agent to run under `Backup Operator` as part of your normal site procedures.

More information about how to configure users to run under the Backup Operator is available:

See the *PureDisk Getting Started Guide*.

# Creating data removal policies for specific files

This chapter includes the following topics:

- [Creating data removal policies for specific types of files](#)

## Creating data removal policies for specific types of files

After you have defined data selections and run a backup job, you may determine that you want to exclude certain types of files from backups. You can create one or more data removal policies to remove these files from PureDisk storage.

For example, to remove all files with the `.mp3` file extension from your PureDisk storage, create a data removal policy in the following manner:

- On the **Data Selections** tab, identify files to remove from a particular data selection.
- On the **Metadata** tab, specify to remove only specific files. PureDisk displays this tab only for **Files and Folders** or **UNC Path** backups.

**To create a data removal policy for a specific type of file**

- 1 Select **Manage > Policies**.
- 2 In the left pane, expand **Policies > Data Management Policies** and then select **Data Removal**.
- 3 In the right pane, click **Create Policy**.
- 4 In the **Policy** panel, on the **General** tab, perform the following steps:
  - Specify a name for the policy. For example, **Removing MP3 files**.

- Select **Enabled**.
- 5 On the **Data Selections** tab, select **Include all data selections selected above**.
  - 6 On the **Scheduling** tab, specify a schedule.
  - 7 (Conditional) On the **Metadata** tab, click **Add**.  
The **Metadata inclusion rule** dialog box opens.
  - 8 Complete the following fields in the dialog box to describe the files you want to remove:
    - For **Rule name**, specify a name for this filter. For example, **MP3 Files**.
    - For **Folder name**, specify an asterisk (\*).
    - For **File name**, specify \*.mp3.
  - 9 Click **OK** on the **Metadata inclusion rule** dialog box.
  - 10 Click **Add** on the **Policy** panel to save the new policy.



# Comprehensive policy scheduling

This chapter includes the following topics:

- [Backup, replication, and maintenance policy scheduling](#)

## Backup, replication, and maintenance policy scheduling

You can create several PureDisk policies for backups, replication, and other activities. For example, you can create your own policies for backups, garbage collection, and other tasks.

For best performance, do not exceed 200 agents in a single backup policy. If you have more than 200 clients, make sure that no more than 200 jobs start at the same time. More than one job can run at the same time, but schedule a 10-minute interval between the start times of each job.

PureDisk includes some system policies and some default policies, many of which are for maintenance.

[Table 4-1](#) explains how to use and run various PureDisk policies.

**Table 4-1** PureDisk recommended practices for policies and scripts

Policy	Default	Recommended frequency	Notes
Backup policies	No default	Run backup policies on work days when system activity is low.	<p>Schedule the backup policies to run when system activity is low. For example, schedule the backup policies at night between 8:00 P.M. and 8:00 A.M.</p> <p>For example, configure the backup window in the <b>Parameters</b> tab for 8:00 P.M. and 8:00 A.M..</p>
<b>Data Removal</b> policy	No default	Run a removal policy one, two, or three times a week and when few backup policies run.	<p>Configure this policy immediately after you deploy a storage pool. Do not wait for content routers to fill with unneeded backups.</p> <p>The more frequently you run a data removal policy, the less time each run takes because less data is removed in each run.</p> <p>If you cannot schedule frequent data removal policy runs, run the data removal policy once per week on the weekend. If you schedule this policy to run less frequently, you risk overloading the system when it runs.</p>
<b>Replication</b> policy	No default	Daily.	<p>Schedule replication policies to run immediately after your backups complete.</p> <p>You can implement replication as additional disaster recovery support.</p>
<b>Disaster Recovery Backup</b> policy	No default	Daily.	<p>This policy consumes a lot of system resources, and Web UI performance can be slower.</p> <p>Run the policy when no backups run and when you do not perform system maintenance tasks.</p>
<b>MB Garbage Collection</b> policy	Monthly	One to three times each week.	<p>If you check <b>Enable extensive cleanup</b>, be aware that this method consumes many system resources. Enable this option on an infrequent basis, for example once every three months. You can copy the default policy and enable this capability only in a policy that runs every three months.</p>

**Table 4-1** PureDisk recommended practices for policies and scripts (*continued*)

Policy	Default	Recommended frequency	Notes
<b>Content router queue processing</b> policy	No default	Two times each day. Schedule this policy to run outside your file backup and disaster recovery backup windows.	Performs content router queue maintenance. Content router queue processing  Do not disable this policy. Make sure that this policy runs regularly.
<b>Server Maintenance</b> policy	Daily	Daily.	Performs server agent maintenance on the PureDisk nodes. It removes unreferenced (unneded) data from the databases.  PureDisk has a default server maintenance policy, and you can add additional server maintenance policies. For example, you can create the following policies: <ul style="list-style-type: none"> <li>■ Create a server maintenance policy to clean up only the content router database. Run this policy once every two weeks.</li> <li>■ Create a second server maintenance policy to clean up only the metabase engine database. Run this policy every week.</li> <li>■ Create a third server maintenance policy to clean up the storage pool authority database. Run this policy once a month.</li> </ul> You can delete this policy, but Symantec does not recommend deletion.



# Tuning PureDisk

This chapter includes the following topics:

- [Configuration parameters](#)
- [Changing TCP/IP settings to improve replication job performance](#)
- [About deduplication results for Oracle backups with PDDO](#)

## Configuration parameters

Be sure to review the topic Reconfiguring your PureDisk environment in the *PureDisk Administrator's Guide* before you make changes to configuration parameters. It contains procedures for editing and pushing configuration changes.

---

**Note:** You edit configuration parameters through the **Settings > Configuration** menu of the Web UI. You can also edit configuration files with any text editor. However, Symantec recommends that you edit these files with a text editor only if instructed to do so by a Symantec Technical Support representative. Locations of configuration files, as well as procedures for changing them, are provided in the *PureDisk Administrator's Guide*.

---

**Table 5-1** PureDisk configuration parameters

Configuration parameter	Description
FingerprintType	This parameter determines the fingerprinting algorithm to be used. Type 0 is the old algorithm without fingerprint collision detection. Type 1 is for PureDisk release 6.5 and later with fingerprint collision detection. This value is set during agent installation. Do not change it. Symantec recommends that you do not change this parameter.

**Table 5-1** PureDisk configuration parameters (*continued*)

Configuration parameter	Description
Port	<p>This parameter determines the port on which the content routers listen for incoming connections. If you change this parameter, you must enter the same number in the configuration file of all content routers and all agents. Symantec recommends that you do not change this parameter.</p>
TCPKeepAlive	<p>This parameter determines whether or not TCP keep-alive probes are to be sent during connection idle time. It lets an agent detect if an existing connection with a content router is still valid.</p> <p>For more information, search Google for <code>SO_KEEPALIVE</code> and <code>TCP_KEEPALIVE</code> or visit the Web pages that are listed in the <code>TCPSendBufferSize</code> parameter.</p>
TCPSendBufferSize	<p>This parameter sets the maximum socket send buffer size in bytes. If you increase this value, the agent can transmit more data before the backup blocks. The backup blocks while it waits for more buffer space to become available after PureDisk sends the currently buffered data to the content router. The default value is dependent on the operating system. For UNIX, it usually depends on other kernel configuration settings. Search Google for <code>SO_SNDBUF</code> for more information on this subject.</p> <p>On high-latency lines (or high-bandwidth lines), it is definitely worthwhile to experiment with various values for the send buffer size. For optimal behavior, adjust the <code>TCPReceiveBufferSize</code> value as well.</p> <p>A general formula for establishing an initial value for this parameter in kilobytes is as follows:</p> $\text{size} = b * d * 10^3 / 8$ <p>The formula assumes a delay <i>d</i> in microseconds and a bandwidth <i>b</i> in kilobytes per second. You can measure the delay in microseconds by using the <code>ping</code> command.</p> <p>If you have a 45-megabit line and a 30-ms delay, the optimal send buffer size is 165 KB.</p> <p>If you specify a value of 0, it enables the operating system default.</p>
TCPReceiveBufferSize	<p>This parameter sets the maximum socket receive buffer size in bytes. If you increase this value, it may reduce the amount of disk writes that the agent must perform during a restore. Using a small value has a negative effect on performance. The default value is dependent on the operating system. For UNIX, it usually depends on other kernel configuration settings. Search Google for <code>SO_RCVBUF</code> for more information on this subject.</p> <p>If you specify a value of 0, it enables the operating system default.</p>

**Table 5-1** PureDisk configuration parameters (*continued*)

Configuration parameter	Description
MaxTransferRate	<p>This parameter specifies in kilobytes per second the maximum throughput on a connection with a content router (with a variance of about 10%). A value of 0 means that no bandwidth limit is applied.</p> <p>You may want to specify a bandwidth limit in the following situations:</p> <ul style="list-style-type: none"> <li>■ If your connections are slow (T1 or so)</li> <li>■ If you have a content router that shares connections with other services</li> <li>■ If backups take place during the day</li> </ul> <p>If you specify a bandwidth limit, it increases the backup and restore duration.</p>
ReadBufferSize	<p>This parameter specifies, in bytes, the maximum amount of data that can be read from disk in a single read operation. The default value is 64 KB, which corresponds to the default <code>TCPSendBufferSize</code> on most operating systems.</p> <p>If a backup consists mainly of large files, a higher value can improve performance slightly. If you specify a lower value, it has a negative effect on performance because PureDisk requires more read operations to process a file.</p> <p>An optimal value is not easy to establish. It requires experimentation with different values. If you specify large values, it may defeat the read-ahead capabilities of the operating system or hardware, and therefore diminish performance.</p> <p>Initially, the following values may be used:</p> <ul style="list-style-type: none"> <li>■ For an average file size &lt; 128 KB, use the default setting</li> <li>■ For an average file size &lt; 64 MB, use 128 KB</li> <li>■ For an average file size &lt; 128 MB, use 256 KB</li> <li>■ For an average file size &gt; 128 MB, use 512 KB</li> </ul> <p>No direct link exists between this configuration parameter and the <code>TCPSendBufferSize</code> parameter. However, there is an indirect connection between the two parameters. The TCP buffer configuration directives let you specify optimal network throughput settings. The Read (and Write) buffer directives let you specify optimal disk throughput settings.</p> <p>There also exists an indirect link with the specified segment size: the agent needs to buffer data until it has a full segment available so that it can compute the segment fingerprint. To some degree, the rules that are given here use this link.</p> <p><b>Note:</b> If you specify a higher value, more memory is used during backup.</p>

**Table 5-1** PureDisk configuration parameters (*continued*)

Configuration parameter	Description
WriteBufferSize	<p>This parameter specifies, in bytes, the maximum amount of the data the agent may buffer during a restore before it writes to disk. The default value is 32 KB, which corresponds to the default <code>TCPReceiveBufferSize</code> setting on most operating systems.</p> <p>If you specify a higher value, it may increase the amount of disk writes. However, a higher value might increase the time the application is blocked while it waits for the I/O to complete). If you specify a low value, it has a negative effect on performance.</p> <p>An optimal value is not easy to establish. It requires experimentation with different values.</p> <p>The following information may help you determine this value:</p> <ul style="list-style-type: none"> <li>■ During restore, the content router streams the segment data to the agent. While the agent receives the data, it passes it through the decryption and decompression layer, which places the result in the write buffer. After this buffer is full, it is written to disk.</li> <li>■ The value for <code>WriteBufferSize</code> may be set equal to the <code>TCPReceiveBufferSize</code>, but it should never be set to a value larger than the segment size.</li> </ul>

## Changing TCP/IP settings to improve replication job performance

You might experience replication job performance degradation if you have a high-latency communication network between the two storage pools. You can possibly improve performance by changing some default TCP/IP settings.

The script `/opt/pdconfigure/scripts/support/tcp_tune.sh` lets you change several default TCP/IP values. The script also lets you view current values and to restore default values. Refer to comments in the script for usage notes, such as valid parameters and the actual settings and values that this script modifies.

---

**Note:** You should run this script only on kernel release 2.6.16.

---

After TCP/IP settings were changed, Symantec noticed performance improvements in replication jobs between two multinode storage pools on a network with latency. However, the degree of performance improvement depends on the characteristics of the network. You might not notice an increase in job performance after you modify the TCP/IP settings with the supplied script. To ensure that the new settings



improve performance, compare the elapsed time during a replication job before and after you change the settings. The timing information appears on the replication job report.

**To change TCP/IP settings to improve replication job performance**

- 1 Log on to the first node of the source storage pool authority service as root.
- 2 Run the following script:

```
# /opt/pdconfigure/scripts/support/tcp_tune.sh modify
```

- 3 Repeat step 1 and step 2 for each remaining node of the source storage pool authority service, including the passive node if the storage pool is clustered.
- 4 Log on to the first node of the destination storage pool authority service as root.
- 5 Run the following script:

```
# /opt/pdconfigure/scripts/support/tcp_tune.sh modify
```

- 6 Repeat step 4 and step 5 for each remaining node of the destination storage pool authority service, including the passive node if the storage pool is clustered.

## About deduplication results for Oracle backups with PDDO

The setup and the configuration of your environment influences the deduplication rates with Oracle database backups through PDDO. The initial deduplication rates can sometimes seem lower than expected.

To improve deduplication results with Oracle backups through PDDO, Symantec recommends that you choose one of the following solutions:

- Run the Proxy Copy with RMAN and execute a script that is provided in a subsequent topic.  
 For this solution you do not need an extra staging space, however, you do need the Enterprise Client to leverage the snapshot support.  
 More information about this solution is available:  
 See [“Running the Proxy Copy with RMAN”](#) on page 34.
- Use the Image copy.  
 For this solution you need extra staging space but no Enterprise Client. An image copy is an exact copy of a single data file, archived redo log, or control file. Image copies are not stored in RMAN-specific format. They are identical

to the results of copying a file with operating system commands. RMAN can use image copies during RMAN restore and recover operations, and you can also use image copies with non-RMAN restore and recovery techniques.

More information about this solution is available:

See [“Using the Image copy”](#) on page 34.

## Running the Proxy Copy with RMAN

This solution moves the data through a proxy gateway, which can be any supported snapshot provider. For this solution you do not need an extra staging space, however, you do need the Enterprise Client to leverage the snapshot support.

### To run the proxy copy with RMAN

- ◆ Use the Proxy Copy with RMAN and use the following script.

```
RUN {
RESYNC CATALOG; ALLOCATE CHANNEL ch00 TYPE 'SBT_TAPE';
CONFIGURE DEVICE TYPE sbt PARALLELISM 1;
BACKUP PROXY ( DATABASE SKIP INACCESSIBLE FILESPERSET 50 FORMAT 'bk_%s_%p_%t'
TAG Full_Oracle_Backup );
RELEASE CHANNEL ch00;
ALLOCATE CHANNEL ch00 TYPE 'SBT_TAPE';
CONFIGURE DEVICE TYPE sbt PARALLELISM 1;
BACKUP ( SPFILE FORMAT 'sp_%s_%p_%t' TAG Oracle_SPFILE INCLUDE CURRENT CONTROLFILE );
RELEASE CHANNEL ch00; }
```

Refer to the *Symantec NetBackup for Oracle Administrator's Guide* for more information about using RMAN scripts

## Using the Image copy

For this solution, you need extra staging space, but no Enterprise Client. An image copy is an exact copy of a single data file, archived redo log, or control file. Image copies are not stored in RMAN-specific format. They are identical to the results of copying a file with operating system commands. RMAN can use image copies during RMAN restore and recover operations, and you can also use image copies with non-RMAN restore and recovery techniques.

**To create image copies and have them recorded in the RMAN repository**

- 1** Run the `RMAN BACKUP AS COPY` command. Alternatively, configure the default backup type for disk as image copies using `CONFIGURE DEVICE TYPE DISK BACKUP TYPE TO COPY` before you perform a backup.

A database server session is used to create the copy.

Using the target database control file instead of the recovery catalog, the RMAN configuration parameters are as follows:

```

RMAN> show all;
CONFIGURE RETENTION POLICY TO REDUNDANCY 1; # default
CONFIGURE BACKUP OPTIMIZATION OFF; # default
CONFIGURE DEFAULT DEVICE TYPE TO DISK; # default
CONFIGURE CONTROLFILE AUTOBACKUP ON;
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '%F'; # default
CONFIGURE DEVICE TYPE DISK BACKUP TYPE TO BACKUPSET PARALLELISM 2;
    < backup type before changing to copy >
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
CONFIGURE MAXSETSIZE TO UNLIMITED; # default
CONFIGURE ENCRYPTION FOR DATABASE OFF; # default
CONFIGURE ENCRYPTION ALGORITHM 'AES128'; # default
CONFIGURE ARCHIVELOG DELETION POLICY TO NONE; # default
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/oracle/10g/dbhome/dbs/snapcf_vrts.f'; # default
RMAN> configure device type disk backup type to copy;
    
```

- 2** Change the following RMAN configuration parameters:

- **Before:**

```

CONFIGURE DEVICE TYPE DISK BACKUP TYPE TO BACKUPSET PARALLELISM 2;
    
```

- **After:**

```

CONFIGURE DEVICE TYPE DISK BACKUP TYPE TO COPY PARALLELISM 2;
    
```

- 3 After the new RMAN configuration parameters are successfully stored, run the following command to display the new configuration parameters:

```
RMAN> show all;
```

The output is similar to the following example:

```
CONFIGURE RETENTION POLICY TO REDUNDANCY 1; # default
CONFIGURE BACKUP OPTIMIZATION OFF; # default
CONFIGURE DEFAULT DEVICE TYPE TO DISK; # default
CONFIGURE CONTROLFILE AUTOBACKUP ON;
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '%F'; # default
CONFIGURE DEVICE TYPE DISK BACKUP TYPE TO COPY PARALLELISM 2;
--> backup type changed to copy(image copy)
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
CONFIGURE MAXSETSIZE TO UNLIMITED; # default
CONFIGURE ENCRYPTION FOR DATABASE OFF; # default
CONFIGURE ENCRYPTION ALGORITHM 'AES128'; # default
CONFIGURE ARCHIVELOG DELETION POLICY TO NONE; # default
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/oracle/10g/dbhome/dbs/snapcf_vrts.f'; # default
RMAN>
```

- 4 Save a copy of the output in a text file so that you can refer to it later if necessary.

# Third-party legal notices

This appendix includes the following topics:

- [Third-party legal notices for the Symantec NetBackup PureDisk product family](#)
- [Third-party trademarks for the Symantec NetBackup PureDisk product family](#)

## Third-party legal notices for the Symantec NetBackup PureDisk product family

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in *NetBackup Product Family Third-party Legal Notices*.

## Third-party trademarks for the Symantec NetBackup PureDisk product family

Active Directory, Excel, Hyper-V, Internet Explorer, Microsoft, Windows, Windows NT, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and other countries.

AIX, IBM, PowerPC, and Tivoli are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc., in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

AMD is a trademark of Advanced Micro Devices, Inc.

Firefox and Mozilla are registered trademarks of the Mozilla Foundation.

Intel, Itanium, Pentium, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java, Sun, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States and other countries.

Mac OS is a trademark of Apple Inc., registered in the U.S. and other countries.

Nessus is a trademark of Tenable Network Security, Inc.

NetApp is a registered trademark of Network Appliance, Inc. in the U.S. and other countries.

Novell and SUSE are registered trademarks of Novell, Inc., in the United States and other countries.

OpenLDAP is a registered trademark of the OpenLDAP Foundation.

Oracle is a registered trademark of Oracle and/or its affiliates.

Red Hat and Enterprise Linux are registered trademarks of Red Hat, Inc., in the United States and other countries.

UNIX is a registered trademark of The Open Group.

VMware, vSphere, and the VMware "boxes" logo and design are trademarks or registered trademark of VMware, Inc., in the United States and other countries.

# Glossary

<b>absolute path</b>	The location of a given file or directory on a file system regardless of the current working directory.
<b>ACL (access control list)</b>	The information that identifies specific users or groups and their access privileges for a particular file or directory.
<b>active agent</b>	The PureDisk software that is enabled for use within a PureDisk environment. The PureDisk “agent” (the software) is installed on “clients” (the hardware).
<b>active node</b>	The nodes in a high availability cluster on which services are running.
<b>Agent Files and Folders data selection</b>	A collection of files, folders, or directories to be backed up. PureDisk creates this data selection automatically when you perform a full system backup.
<b>All-in-one storage pool</b>	A PureDisk storage pool with all PureDisk services installed on one node.
<b>alternate client</b>	A client other than the source client that receives restored files. If the data is not restored to the original client, the client that is designated to receive the data is the alternate client.
<b>ATOP (All Through One Port)</b>	See single-port communication.
<b>attended installation</b>	An installation that requires active interaction with a person.
<b>authentication broker</b>	The process that communicates with a root broker to verify user identities.
<b>backup</b>	A process where selected files on a computer drive are copied and stored on a reliable form of media.
<b>backup operator</b>	A user or a group of users with the rights to initiate client backups.
<b>backup streams</b>	A connection between the PureDisk client and the PureDisk content router through which data is sent. In case of multistreaming, the client establishes multiple connections to the same content router and distributes the total volume of backup data over all available connections. Multistreamed backups (vs. single streamed) increase the aggregated throughput from client to content router, thus allowing backups to finish sooner.
<b>backup window</b>	The timeframe in which backups are permitted.
<b>central reporting</b>	A PureDisk feature in which one or more storage pools send reporting data to another storage pool.

<b>CIFS (Common Internet File System)</b>	A protocol that defines a standard for remote file access. CIFS allows users with different platforms and computers can share files.
<b>cluster</b>	A set of hosts (each termed a node) that share a set of disks and are connected by a set of redundant heartbeat networks.
<b>common root</b>	A shared directory structure. Common root is a concept applicable to data restore. When the user restores the data, there is a “Do not restore common root” option. The user can enable or disable the feature.
<b>configuration files</b>	The files that define PureDisk’s methods and assumptions. A custom configuration file for each component is created automatically by the storage pool authority. It is pushed out to the component for which this file is created. If you need to tune any of the default PureDisk processes, you can edit the configuration files.
<b>content router</b>	A service that stores and retrieves file content. PureDisk breaks larger files into segments and distributes the segments across the available content routers.
<b>content router garbage collection</b>	The process of removing unneeded data objects and files from the content router. This workflow removes files and objects that cannot be removed during the normal data removal process.
<b>CSV (comma-separated variable) file</b>	A text file that uses commas as data delimiters.
<b>data lock password</b>	An option that allows the administrator to require users to enter a password before they perform certain operations. These operations include any operation that exposes directory names, directory content, file names, or file content.
<b>data mining</b>	The process of collecting information about all the files in a PureDisk storage pool.
<b>data removal</b>	The process of removing old and unneeded versions of the files that PureDisk previously backed up.
<b>data selection</b>	A list of files, directories, or other data objects that you want PureDisk to back up. They can be used to specify files and folders, databases, system information, and other types of data. After a data selection is created, it can be backed up automatically (through a policy) or on demand (initiated by a backup operator).
<b>data selection removal</b>	The process of removing the entire data selection (including file content data) from the content router. This process also removes the associated metadata information from the metabase engine.
<b>data selection template</b>	A pattern that is used for creating the list of files, directories, or other data objects for PureDisk to back up. Templates can be developed that include or exclude certain file types, or that back up a specific directory. PureDisk includes some default data selection templates.



<b>deduplication</b>	The process of dividing a file into segments, comparing each segment with the previously stored file segments, and then storing only the unique segments. Deduplication significantly reduces the amount of data that is stored because redundant data is replaced with a pointer to the unique data copy.
<b>department</b>	A logical collection of client systems.
<b>disabled data selection</b>	A data selection that PureDisk ignores when it performs policy-based actions. If a policy includes deactivated data objects, the policy does not include them when the policy runs.
<b>disabled policy</b>	A policy that exists in the PureDisk environment but that is currently not activated. No jobs are created for this policy.
<b>disaster recovery</b>	The process of restoring information from a backup after the original data was lost (due to a disaster) or deleted.
<b>DMP (dynamic multipathing)</b>	An input/output (I/O) enhancement technique that balances I/O across many available paths from the computer to the storage device to improve performance and availability.
<b>enabled data selection</b>	An activated list of files, directories, or other data objects that PureDisk uses for back ups and other policy-based actions.
<b>enabled policy</b>	A policy that exists in the PureDisk environment that is currently activated and run according to a schedule. PureDisk automatically creates jobs to run this policy.
<b>escalation action</b>	A defined procedure that takes place when an event occurs.
<b>event</b>	A significant occurrence in a system or application that a program detects. Events typically trigger actions, such as sending a user notification or adding a log entry.
<b>event escalation action</b>	A defined procedure that takes place when a specific notable occurrence takes place.
<b>exclude files</b>	A listing of files or file patterns that are not included in a data selection.
<b>exclusion rules</b>	The means by which PureDisk determines the files or folders that should not be part of a given data selection. See also inclusion rules.
<b>expert installation method</b>	The process of loading the PDLinux software onto a computer with multiple hard drives where none of the options is predefined.
<b>external authentication</b>	A credential verification authority that resides on a computer that is not part of a PureDisk storage pool.
<b>failover</b>	The process of moving services from the active node in a cluster to a passive one.
<b>file change rate</b>	The frequency with which files on a client system are modified.
<b>file pattern</b>	A character sequence that includes wild cards and instructs PureDisk to select multiple files based on the character sequence.

<b>file system browsing</b>	The ability to search through a graphic representation of a computer system's file structure.
<b>Files and Folders data selection</b>	A data selection that is used to back up files, folders, and directories.
<b>fingerprint</b>	A unique sequence of digits identifying a file or a file segment. The fingerprint of a file or segment is computed from the file or segment's content and is unique for that file or segment.
<b>folder pattern</b>	An absolute path that may contain wild cards and instructs PureDisk to select folders based on the character sequence.
<b>FQDN (fully qualified domain name)</b>	An unambiguous domain name that specifies the exact location of a computer within the domain's hierarchy.
<b>garbage collection</b>	The process of removing stale data or records from PureDisk that cannot be removed during the normal data removal process.
<b>HCL (hardware compatibility list)</b>	A document that indicates the various components that are known to work with a given software product.
<b>heartbeat</b>	A signal sent at regular intervals to indicate that a host and its connections are operating normally.
<b>high availability</b>	A system or a resource that is continuously operational.
<b>host address</b>	The TCP/IP address of a computer.
<b>inactive agent</b>	A client computer that is registered to the storage pool authority, but which is not yet acknowledged as part of the PureDisk environment.
<b>include files</b>	A listing of files or file patterns that are included in a data selection.
<b>inclusion rules</b>	The means by which PureDisk determines the files to consider as part of a given backup. See also exclusion rules.
<b>inheritance</b>	The process of receiving attributes from a parent object, such as a template.
<b>job</b>	An operation that has been scheduled for processing. Jobs contain source or destination information, settings, and a schedule.
<b>LDAP (Lightweight Directory Access Protocol)</b>	A software protocol that enables anyone to locate organizations, individuals, and other resources such as files and devices in a network, whether on the Internet or on a corporate intranet. LDAP is a lightweight (smaller amount of code) version of Directory Access Protocol (DAP), which is part of X.500, a standard for directory services in a network.
<b>location</b>	A logical grouping of one or more departments. A location does not need to be a physical or a geographic location.

<b>mapped network drive</b>	A persistent connection within Microsoft Windows to a shared directory on a remote server that is assigned a drive letter. The drive mapping may or may not survive a restart, depending on how it is configured.
<b>mapping file</b>	A two-column file that lists both IP addresses and fully qualified domain names (FQDNs) for each node. Mapping files are used when redefining a PureDisk environment to use either FQDNs or IP addresses.
<b>metabase</b>	A distributed database that contains all the metadata about the files that are backed up by PureDisk.
<b>metabase engine</b>	The service that maintains and manages file metadata information in the metabase database. During backup, the PureDisk agent records all relevant metadata information (the file attributes) of each file it backs up. File attributes include the file name, its location, its size, its type, and so on. The metabase engine stores these metadata records in its database. The metabase engine manages an inventory of all files that are backed up.
<b>metabase garbage collection</b>	The process of removing the stale, unneeded information from the metabase that cannot be removed during the normal data removal process.
<b>metabase server</b>	The service that redirects metabase queries to the correct metabase engine(s). Each metabase engine in the storage pool is responsible for managing the metadata records from a group of agents. Queries for a file or files do not go directly to the metabase engine but instead are directed to the metabase server. The metabase server redirects the query to the metabase engine that manages the metadata records of the agent that made the query or to which the query is associated. In some cases the metabase server may have to redirect the incoming query to multiple metabase engines.
<b>metadata</b>	Structural data describing the attributes of files on a disk.
<b>Microsoft Exchange data selection</b>	A predefined data selection within PureDisk that backs up Microsoft Exchange server databases.
<b>Microsoft SQL data selection</b>	A predefined data selection within PureDisk that backs up Microsoft SQL server databases.
<b>multinode storage pool</b>	A storage pool that includes more than one PDLinux server node and can be clustered.
<b>multistreaming</b>	The process of establishing multiple connections to the content router for the purposes of backing up data faster.
<b>NetBackup export engine</b>	An optional service that allows backed up PureDisk Files and Folders data selections to be exported to NetBackup. The NetBackup export engine does not export other PureDisk data selection types.
<b>network drive</b>	A directory on a remote server that is designated as shared.

<b>node</b>	A computer in a storage pool that hosts PDLinux and other PureDisk services.
<b>OpenLDAP (Open Lightweight Directory Access Protocol)</b>	A network protocol that is designed to work on TCP/IP stacks. It extracts information from a hierarchical directory such as X.500. This software gives users a single tool to find a particular piece of information. For example, a user can find a user name, an email address, security certificate, or other contact information.
<b>Oracle data selection</b>	A predefined data selection within PureDisk that backs up Oracle databases.
<b>parallel rerouting</b>	The process of redistributing data among the content routers in a PureDisk storage pool. All content routers are actively involved in the redistribution at the same time.
<b>passive node</b>	Any node in a highly available, clustered environment that is not running PureDisk services.
<b>path</b>	The directory location of a given file or directory on a file system. Paths can be either relative or absolute.
<b>PDDO (PureDisk deduplication option)</b>	A plug-in that uses the NetBackup OpenStorage API to enable NetBackup to write backups to a PureDisk storage pool. All NetBackup data that is written to a PureDisk storage pool is deduplicated.
<b>PDLinux (PureDisk Linux)</b>	The operating system that hosts the PureDisk application. Symantec developed PDLinux based on SUSE Linux Enterprise Server.
<b>policy</b>	A method for managing backup jobs and strategies. Policies contain settings for jobs.
<b>policy escalation action</b>	A procedure that is defined to take place when a specific policy event (such as a backup failure) takes place.
<b>private network</b>	A computer network that is accessible only by other servers, not by the general user base.
<b>private NIC</b>	A network interface card that communicates to a private network.
<b>public network</b>	A computer network that is accessible to the general user base.
<b>public NIC</b>	A network interface card that communicates to a public network.
<b>registration</b>	The process of enrolling with the storage pool authority.
<b>relative path</b>	The directory location of a given file or directory on a file system that depends on the current working directory.
<b>replication</b>	The process of copying backed up data selections from one storage pool to another storage pool.
<b>rerouting</b>	The process of redistributing data over all available content routers. When the process finishes, each content router stores a volume of data proportional to its relative capacity. Rerouting is necessary when a new content router is activated, or an existing content router is deactivated.

<b>resource name</b>	The unique identifier for a service on a PureDisk node.
<b>root broker</b>	The authentication authority in the network. A root broker is local when it resides on the same physical computer as the PureDisk storage pool authority. A root broker is remote when it resides on a different PureDisk computer in the network. A root broker is external when it does not reside on any PureDisk computer within the storage pool.
<b>segment</b>	A piece of a file.
<b>segmentation</b>	The process of breaking a file down into smaller pieces for backup.
<b>segmentation threshold</b>	The maximum allowable size for a file fragment.
<b>serial rerouting</b>	The process of redistributing data among the content routers in a PureDisk storage pool. Only one content router redistributes data at a time.
<b>service</b>	A PureDisk software component. The possible services are as follows: content router, metabase engine, metabase server, storage pool authority, and NetBackup export engine.
<b>service address</b>	The TCP/IP address associated with a service group.
<b>service group</b>	A collection of PureDisk services.
<b>shared disk</b>	A physical hard drive on a computer that can be remotely accessed from another computer. In a highly available cluster, the shared disk is normally a drive that does not physically reside in any of the cluster nodes. Any resources that can failover among cluster nodes must reside on a shared disk.
<b>shared folder</b>	A network directory, to which multiple users have read and write access, used to exchange files with other users.
<b>silent installation</b>	An installation in which the user sees no indication that the installation is occurring. The user is not prompted to enter any information and the user does not see status messages. See also attended installation.
<b>single port communication</b>	A PureDisk feature that directs all network communication through one port. Storage pools that implement single-port communication require fewer firewall ports to be open between PureDisk service agents and clients.
<b>snapshot</b>	A consistent point-in-time view of a volume that is used as the reference point for the backup operation. After a snapshot is created, the primary data can continue being modified without affecting the backup operation.
<b>SPA (storage pool authority)</b>	The service that manages a storage pool.
<b>SPAR (storage pool authority replication)</b>	The replication of storage pool authority configuration information from an all-in-one local storage pool to a main storage pool.

<b>storage pool</b>	The main data repository in PureDisk. PureDisk writes backup copies of content and metadata to the disk storage that is associated with a storage pool. A storage pool consists of one or more PureDisk nodes.
<b>stream</b>	A sequence of digital data.
<b>SUSE</b>	A distribution of Linux software. The name is an acronym for the German phrase “Software-und System-Entwicklung” (Software and system development).
<b>System State and Services data selection</b>	A data selection that is used to back up Microsoft system data on Windows platforms.
<b>template inheritance</b>	The process of receiving attributes from a parent template.
<b>TLS (transport security layer)</b>	An encrypted protocol that provides secure communications in the PureDisk environment.
<b>topology</b>	The types of PureDisk services that a storage pool includes. All storage pool topologies include one or more of the following services: storage pool authority, content router, metabase engine, metabase server. A controller is installed on a metabase engine. Optionally, a storage pool can also include a NetBackup export engine.
<b>unattended install</b>	An installation that does not require human interaction.
<b>UNC path data selection</b>	A data selection that backs up data on a CIFS network drive on a Windows client. Also use this data selection to indicate the path for a NetApp Filer.
<b>user</b>	An individual with rights to access your protected network resources. Users are defined by creating a user account that consists of a unique user name and authentication method.
<b>user group</b>	A collection of users with identical permissions. These users can perform common functions within a PureDisk environment.
<b>vacuuming</b>	The process of cleaning up and optimizing a database. Vacuuming removes the records that are no longer needed and results in better database performance.
<b>VCS (Veritas cluster server)</b>	High-availability cluster software developed by Symantec for UNIX, Linux, and Windows platforms.
<b>VEA (Veritas Enterprise Administrator)</b>	A separate middleware server used by the SAN Access Layer and other processes to provide client-server communication. The VEA infrastructure enables software components to share information about objects, manage those objects, and effect change on those objects.
<b>CommandCentral Console</b>	A graphical user interface that displays reports and other information for users of CommandCentral Service through a standard Web browser. The Console provides a central point to manage cost analysis and chargeback for services, managing workflow, displaying and managing reports, and other tasks.

<b>VSP (Volume Snapshot Provider)</b>	Symantec software that backs up open files. PureDisk uses VSP on Windows 2000 clients.
<b>VSS (Volume Shadow Copy Service)</b>	A set of application programming interfaces (APIs) that creates a framework. Within this framework, volume backups and application write can occur at the same time. PureDisk uses Microsoft's VSS technology to back up open files on Windows 2003 and Windows XP clients.
<b>VxVM (Veritas Volume Manager)</b>	A Symantec product installed on storage clients that enables management of physical disks as logical devices. It enhances data storage management by controlling space allocation, performance, data availability, device installation, and system monitoring of private and shared systems.
<b>workflow</b>	A collection of steps that the software completes to accomplish a task.
<b>XFS (Extended File System)</b>	A journaling file system that you can configure on a PureDisk node.
<b>YaST (Yet another Setup Tool)</b>	The operating system installation tool for SUSE Linux.





# Index

## B

- backup
  - implementing a traditional backup plan 13
  - initial 12
  - policy recommendations 26
  - PureDisk timeouts 12
  - snapshot support 13

## C

- configuration parameters 29
- content router queue processing
  - recommendations 27

## D

- data deduplication 13
- data removal policy
  - for specific file types 23
  - recommendations 26
- deduplication results
  - Oracle backups with PDDO 33
- deployment 11
- disaster recovery
  - policy recommendations 26

## F

- Fingerprint parameter 29

## G

- garbage collection
  - policy recommendations 26

## M

- MaxTransferRate parameter 31

## P

- Port parameter 30

## R

- ReadBufferSize parameter 31
- replication
  - job performance 32
  - policy recommendations 26

## T

- TCP/IP settings 32
- TCPKeepAlive parameter 30
- TCPReceiveBufferSize parameter 30
- TCPSendBufferSize parameter 30

## W

- WriteBufferSize parameter 32