

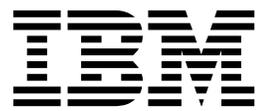
IBM Spectrum Protect Snapshot for Oracle
Version 8.1.0

Installation and User's Guide
UNIX and Linux

IBM

IBM Spectrum Protect Snapshot for Oracle
Version 8.1.0

Installation and User's Guide
UNIX and Linux



Note:

Before you use this information and the product it supports, read the information in “Notices” on page 203.

This edition applies to version 8, release 1, modification 0 of IBM Spectrum Protect Snapshot for UNIX and Linux (product numbers 5725-X22, and 5608-AB8) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2001, 2016.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures v

Tables vii

About this publication ix

Who should read this guide ix

Publications x

**Updates for IBM Spectrum Protect
Snapshot for Oracle V8.1.0 xi**

Chapter 1. Overview 1

Backup and restore methods with FlashCopy and
snapshots 2

Database cloning 3

Software components 4

Chapter 2. Planning. 7

IBM Spectrum Protect Snapshot Prerequisite Checker 7

Capacity planning 8

Required communication ports 9

Storage solutions 10

IBM XIV Storage System 10

SAN Volume Controller and Storwize V7000

storage systems 12

DS8000 storage system 17

Reconciliation of backups. 19

Remote mirror integration 20

Remote mirroring and consistency groups 22

Logical Volume Manager support (AIX only) 23

Preparing applications that run on VMware or KVM
virtual machines. 25

Checking the KVM setup. 26

Chapter 3. Preparation for installation 27

Required Operating System users and environment
variables 27

Prerequisite checker for Oracle 28

Installing the Prerequisite Checker. 28

Running the Prerequisite Checker 29

Interpreting the Prerequisite Checker output 29

Uninstalling the Prerequisite Checker. 30

Preparing Oracle systems (non-SAP, non-ASM) 31

Preparing Oracle Automatic Storage Management 32

ASM disk group layout 33

ASM in a DS8000 environment 33

ASM failure groups. 34

ASM cloning parameters and IBM Spectrum

Protect Snapshot. 34

Preparing Oracle RAC. 35

**Chapter 4. Preparing backup and
cloning servers 37**

Preparing backup and cloning servers 37

Installation prerequisites for backup and clone
servers 38

Preparing backup and cloning servers for VMware
or KVM 40

Chapter 5. Installing and upgrading . . 41

Installing on the production server 42

Adding or upgrading a new instance ID after
installation 43

Installing separately on backup or clone servers 44

Preparing a database or database instance for
configuration 45

Configuring a database or database instance 45

Installing in silent mode 46

Upgrading. 46

Uninstalling the software. 47

Migrating existing snapshot data 47

Editing **USE_CONSISTENCY_GROUP** before you
upgrade from IBM Spectrum Protect Snapshot

version 3.1, or earlier 47

**Chapter 6. Configuring IBM Spectrum
Protect Snapshot 49**

Running the setup script for Oracle 49

Configuring IBM Spectrum Protect Snapshot for
Oracle 50

Configuring Oracle RAC for daemon high
availability 53

Activating a database 54

Configuring storage environments. 55

Configuring Storwize V7000 and SAN Volume
Controller dynamic target allocation (SVCDTA) 55

Configuring the CIM adapter for SP 800-131A
compliant encryption 59

Defining Logical Unit Numbers on DS8000
storage subsystems 59

Defining virtual disks on SAN Volume Controller
and Storwize V7000 60

Select the **FLASHCOPY_TYPE** 61

Target set definitions 63

ASM failure group environment 65

LVM mirroring environments 66

Backup and clone server assignment 66

Managing backups and clones with the
DEVICE_CLASS parameter 68

Configuring for remote mirroring 70

Configure XIV remote mirroring 71

Setting up the daemons on the production and
backup systems 72

Postinstallation and configuration 72

Chapter 7. Backing up data 75

Backing up Oracle databases 75
 Backing up an Oracle database to IBM Spectrum
 Protect server. 76
Oracle Data Guard standby database backup and
restore 77
Support of Oracle Automatic Storage Management
failure groups 77
FlashCopy backup of individual mirrors. 78
 Backing up data with remote mirroring 79
Usability states of snapshot backup operations . . . 80

Chapter 8. Restoring Oracle databases 87

Restoring from a snapshot backup. 87
Restoring Oracle Real Application Clusters database
snapshots 88
Restoring Oracle databases from IBM Spectrum
Protect 90
 Restoring the entire database 90
 Restoring files 91
Restore Oracle Data Guard standby or primary
database 92
Restoring data with remote mirroring 92
 Restoring Oracle data with remote mirroring . . 93

Chapter 9. Cloning databases 95

Cloning databases with IBM Spectrum Protect
Snapshot 96
Cloning Oracle RAC databases 96
 Oracle instance names on the clone server for
 Oracle RAC 97
Cloning an Oracle Data Guard standby database . . 97
Database cloning preprocessing and postprocessing 98
Configuration files used for cloning 98
 Cloning processing example 100

Chapter 10. Troubleshooting 101

General troubleshooting procedure 101
Logging and tracing files 102
 Log files and trace files 102
 Storage system log and trace files 106
 CIM log and trace files 106
 IBM Spectrum Protect for ERP log and trace
 files 106
Troubleshooting mirroring relationships 106
Troubleshooting storage solutions 108
Troubleshooting connectivity problems 108
Troubleshooting tips for IBM Spectrum Protect
Snapshot for Oracle 109
 Guidelines for Oracle variables 110
 IBM Spectrum Protect Snapshot for Oracle
 miscellaneous errors 110
Internet Protocol Version 6 (IPv6) support 111

Appendix A. Configuration files . . . 113

Profile configuration file. 113
 GLOBAL. 116
 ACSD 117

CLIENT 119
ORACLE. 124
CLONING 126
DEVICE_CLASS *device*. 132
OFFLOAD 149
Changing profile parameters 152
Interdependency of LVM_FREEZE_THAW and
TARGET_DATABASE_SUSPEND 152
Target set and target volumes 153
 Manage target volumes files for your storage
 system 154
 DS8000 target volume parameter settings . . . 155
 SAN Volume Controller and Storwize V7000
 target volume parameter settings. 156
 Target set handling for cloning 158
IBM Spectrum Protect Snapshot password file . . 159
IBM Global Security Kit configuration 160
 Uninstall GSKit. 162
Option files used by Data Protection for Oracle . . 162

Appendix B. Commands and scripts 165

Backup, restore, cloning commands, and utilities 165
 The acsora user interface for Oracle 165
 acsutil Snapshot Object Manager for Oracle . . 167
Cloning commands 169
 FlashCopy cloning function-clauses 169
Deleting snapshot backups 171
 Deleting a target volume or target set 172
Snapshot backup status in the repository 172
Administrative commands 173
 Configuration commands 173
 Background daemons. 181
 Mounting and unmounting snapshots on a
 secondary system 185
 Integration with IBM Spectrum Protect. 189

Appendix C. Examples 193

Oracle overall disk layout example 193
Oracle profile example 194
Oracle ASM profile example 195
RMAN backup script example. 196
Oracle in an SAP environment (disk only) profile
example 196
DS8000 target volumes file example 197
SAN Volume Controller and Storwize V7000 target
volumes file example. 198

Appendix D. Accessibility features for the IBM Spectrum Protect product family. 201

Notices 203

Glossary 207

Index 209

Figures

1. IBM Spectrum Protect Snapshot backup and restore environment	3
2. IBM Spectrum Protect Snapshot and database cloning	4
3. IBM Spectrum Protect Snapshot system components	5
4. Remote mirroring using Metro Mirror and Global Mirror sources	22
5. IBM Spectrum Protect Snapshot in an LVM environment	24
6. IBM Spectrum Protect Snapshot host assignments. This example shows a DB2 configuration.	67
7. Cross-site mirrored SAP database protected with IBM Spectrum Protect Snapshot and IBM Spectrum Protect.	78
8. Usability States during snapshot backup	82
9. Usability states during snapshot restore	83
10. Usability states during snapshot delete	83
11. Usability states during snapshot mount	84
12. Usability states during snapshot offload	85
13. Example overall disk layout for a Oracle environment	193

Tables

1. Space requirements for a global product installation of IBM Spectrum Protect Snapshot	8	13. IBM Spectrum Protect Snapshot return codes	103
2. IBM Spectrum Protect Snapshot for UNIX and Linux port numbers	9	14. IBM Spectrum Protect Snapshot installer exit codes	104
3. Dynamic target volumes and predefined target volumes feature comparison.	13	15. DB2 vendor reason codes	104
4. Selecting the FLASHCOPY_TYPE for DS8000, SAN Volume Controller, and Storwize V7000	62	16. Actions taken depending on values of LVM_FREEZE_THAW and TARGET_DATABASE_SUSPEND	153
5. Supported storage subsystems and FlashCopy types.	63	17. Managing target volume LUNs by storage system.	154
6. Summary of Backup Commands for Oracle	75	18. Parameters of the 'VOLUMES_SET_x' Topic (DS8000)	156
7. Files used during a manual backup	76	19. Parameters of the 'VOLUMES_SET_x' topic (SAN Volume Controller and Storwize V7000)	157
8. Usability states	80	20. Parameters for Oracle databases	165
9. Summary of Restore Commands for Native Oracle	87	21. Options for starting the management agent, acsd, as a daemon process	182
10. Message prefixes used in the summary log file	101	22. Options for starting the generic device agent, acsgen	184
11. IBM Spectrum Protect Snapshot log files	102		
12. IBM Spectrum Protect Snapshot trace files	103		

About this publication

This publication provides information about installing, configuring, administering, and using IBM Spectrum Protect™ Snapshot for UNIX and Linux.

IBM Spectrum Protect Snapshot for UNIX and Linux is provided as a single installation package that supports the following database applications, storage systems, and operating systems:

- One of these applications:
 - DB2®, or DB2 in an SAP environment
 - Oracle or Oracle in an SAP environment
 - Custom applications such as file systems or other than DB2 or Oracle databases
- One of these storage systems or file systems that are used for the application:
 - IBM® System Storage® DS8000®
 - IBM System Storage SAN Volume Controller
 - IBM XIV® Storage System
 - IBM Storwize® family and IBM Storwize V7000 Unified
 - IBM General Parallel File System (GPFS™) in combination with Custom Applications on any storage system
- One of these operating systems:
 - AIX®
 - Linux
 - Oracle Solaris
 - HP-UX

IBM Spectrum Protect Snapshot runs online or offline backups of DB2, Oracle databases, or other applications that are on snapshot-oriented storage systems. Optionally, it backs up to IBM Spectrum Protect storage by using IBM Spectrum Protect for Enterprise Resource Planning, IBM Spectrum Protect for Databases, or IBM Spectrum Protect backup-archive client, as appropriate.

IBM Spectrum Protect is a client/server licensed product that provides storage management services in a multi-platform computer environment. It is required only if the offload backup function of IBM Spectrum Protect Snapshot is needed.

Who should read this guide

This guide is intended for system programmers and administrators who are responsible for implementing a backup and cloning solution in one of the supported environments.

The following list identifies hardware and software solutions and tasks that can be used with IBM Spectrum Protect Snapshot. The information that is presented in this publication assumes that you have an understanding of the following solutions and topics, as applicable.

- Storage systems or file systems that are used for the database or custom application:

- IBM System Storage DS8000
- IBM System Storage SAN Volume Controller or IBM Storwize family
- IBM XIV Storage System
- IBM System Storage N series
- NetApp systems
- IBM General Parallel File System (GPFS)
- Oracle or DB2 database administration
- IBM Spectrum Protect

Publications

The IBM Spectrum Protect product family includes IBM Spectrum Protect Snapshot, IBM Spectrum Protect for Space Management, IBM Spectrum Protect for Databases, and several other storage management products from IBM.

To view IBM product documentation, see IBM Knowledge Center.

Updates for IBM Spectrum Protect Snapshot for Oracle V8.1.0

Learn about new features and enhancements in IBM Spectrum Protect Snapshot for UNIX and Linux .

New and changed information is indicated by a vertical bar to the left of the change.

Remote mirroring with SAN Volume Controller dynamic target allocation adapter

Configure IBM Spectrum Protect Snapshot with SAN Volume Controller dynamic target allocation adapter for remote mirroring. The remote mirroring relationship between the primary and secondary SAN Volume Controller devices must be set up by using the SAN Volume Controller GUI or command-line interface. After you configure remote mirroring, you can run backup, restore, and cloning operations from the remote mirrored site.

For more information about remote mirroring, see “Configuring for remote mirroring” on page 70.

Support for partitioned ASM disks

Back up, restore, and cloning operations on databases that are set up on disk groups with partitioned disks are supported. For more information about ASM, see “ASM disk group layout” on page 33. Disks that are used for ASM can now contain a single partition or multiple partitions, or a combination of both types.

Deprecated storage systems

- NetApp storage system
- N Series storage system

New and modified parameters or functions

The following parameters are new:

SVC_REMOTE_SSHKEY_FULLPATH

This parameter for the SVCDTA storage adapter specifies a second SSH key file to be used for authentication on the remote site storage device.

For more information about this and other parameters, see “DEVICE_CLASS parameters for dynamic target allocation” on page 141.

Chapter 1. Overview

IBM Spectrum Protect Snapshot provides a method to back up and restore data by using the advanced snapshot technologies of storage systems.

IBM Spectrum Protect Snapshot can back up DB2 databases, Oracle databases, or other applications that are on snapshot-oriented storage systems or file systems.

Backup operations are based on volume-level copy operations that are provided by the storage system. For GPFS in combination with Custom Applications, the backup operations are based on GPFS file sets. In this scenario, any storage solution that is supported by the GPFS file system can be used. IBM Spectrum Protect Snapshot takes snapshots at a volume group or GPFS file set level for granular control.

When you use IBM Spectrum Protect Snapshot with other IBM Spectrum Protect products, snapshots can be sent to the server. Depending on the application, snapshots can be transferred by using IBM Spectrum Protect for Enterprise Resource Planning, IBM Spectrum Protect for Databases, or IBM Spectrum Protect backup-archive client. IBM Spectrum Protect Snapshot for Oracle uses Oracle RMAN Media Management API. Using RMAN maximizes the protection of Oracle data, and provides a comprehensive storage management solution. To send snapshot backups to IBM Spectrum Protect, you must configure a backup server or cluster.

The following list identifies the applications that can be protected and cloned with IBM Spectrum Protect Snapshot.

- Protect the following database applications with IBM Spectrum Protect Snapshot:
 - DB2, DB2 in an SAP environment, DB2 in a partitioned database environment. You can back up and restore data from single-partition databases, and logically or physically partitioned DB2 databases.
 - Oracle, Oracle with Automatic Storage Management (ASM), and Oracle in an SAP environment.
 - Any database application other than those database applications listed in the preceding list.
 - Any other applications that are on file systems that are supported by IBM Spectrum Protect Snapshot.
- Clone the following database applications with IBM Spectrum Protect Snapshot:
 - DB2, DB2 in an SAP environment, DB2 with the Database Partitioning Feature included.
 - Oracle and Oracle in an SAP environment that is on a file system that is supported by IBM Spectrum Protect Snapshot.

IBM Spectrum Protect Snapshot supports specific operating systems. All servers cooperating in an IBM Spectrum Protect Snapshot environment must be at the same operating system release level. Certain high availability (HA) environments are supported.

The following list identifies the storage solutions or file systems that you can use with IBM Spectrum Protect Snapshot software:

- IBM XIV Storage System

- IBM Storwize family
- IBM System Storage SAN Volume Controller
- IBM System Storage DS8000
- GPFS file system in combination with Custom Applications

Backup and restore methods with FlashCopy and snapshots

A snapshot or FlashCopy[®] is an instant, point-in-time copy of a logical unit (LUN) or a set of LUNs.

FlashCopy and snapshots

The term FlashCopy is used for IBM System Storage DS8000, IBM System Storage SAN Volume Controller, and IBM Storwize family storage devices. A FlashCopy creates a point-in-time copy in which the target volume represents an exact copy of the data on a source volume at the time the FlashCopy starts. Data that exists on the target volume is replaced by the copied data. When you create a FlashCopy of a source volume, the target volume must be the same size as the source volume. In addition, the target volume and source volume must have the same logical track format, and must be on the same storage system.

For IBM XIV Storage System, and file systems such as GPFS, the term *snapshot* is used. A snapshot represents a point-in-time copy of a volume or set of volumes without having to define a specific target volume. The source volumes and snapshots are on the same storage system. Similarly, a file system snapshot represents a point-in-time copy of a file system or file set within a file system. The space that is required for the snapshot is allocated automatically within the same storage system or file system, and can increase over time.

Using a FlashCopy or snapshot you can back up data from source volumes to target volumes. Similarly, you can back up file systems or file sets within a file system. When data is restored, backup copies are retrieved and the data is copied to the source volume, or copied to the original location in the file system or file set.

Types of snapshot backups

Snapshot backups can be either full copy snapshots or space-efficient snapshots. The type of snapshot backups depends on the storage environment. During a full copy snapshot, all blocks of data on the source volume are copied to the target volume. During a space efficient snapshot, only blocks of data overwritten on the source volume are copied.

Transferring snapshots to IBM Spectrum Protect

When you use IBM Spectrum Protect Snapshot with IBM Spectrum Protect products, you can transfer snapshots to the IBM Spectrum Protect server. To send these snapshot backups to the IBM Spectrum Protect server, you must configure a backup server or cluster.

The following figure shows the relationship among the components in a production environment when you run a backup or restore snapshot.

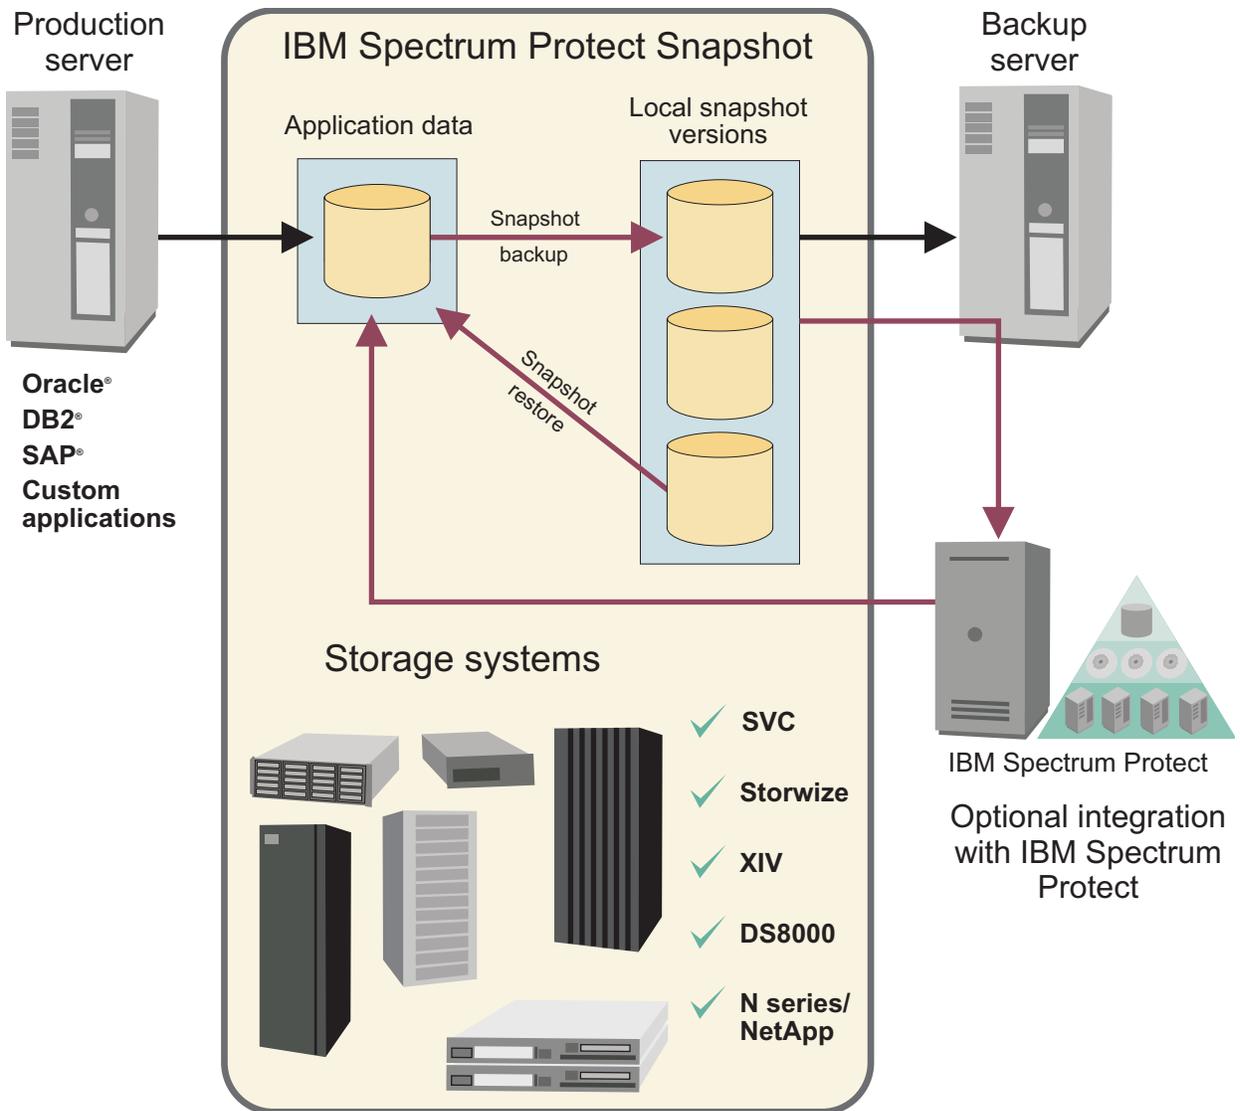


Figure 1. IBM Spectrum Protect Snapshot backup and restore environment

Database cloning

The database cloning process creates an exact copy of a database to provide near-production data.

IBM Spectrum Protect Snapshot uses the FlashCopy or snapshot function for database cloning. Choose to clone a database to create one of the following scenarios:

- To create a test system before you introduce a new product release or new functions into a production environment.
- To create an education system from a master training system. You can reset the cloned database before you start a new course.
- To create a dedicated reporting system to offload the workload away from the production environment.

Traditionally, the database cloning process redirected a restore operation to create the clone. This method has disadvantages, including system downtime and

degraded system performance. IBM Spectrum Protect Snapshot clones a database by using the storage system FlashCopy or snapshot capabilities to minimize the impact on the production database. A *clone server* or *clone system* is required by IBM Spectrum Protect Snapshot to mount a cloned database.

The following figure shows how IBM Spectrum Protect Snapshot creates and stores a cloned database on a clone server.

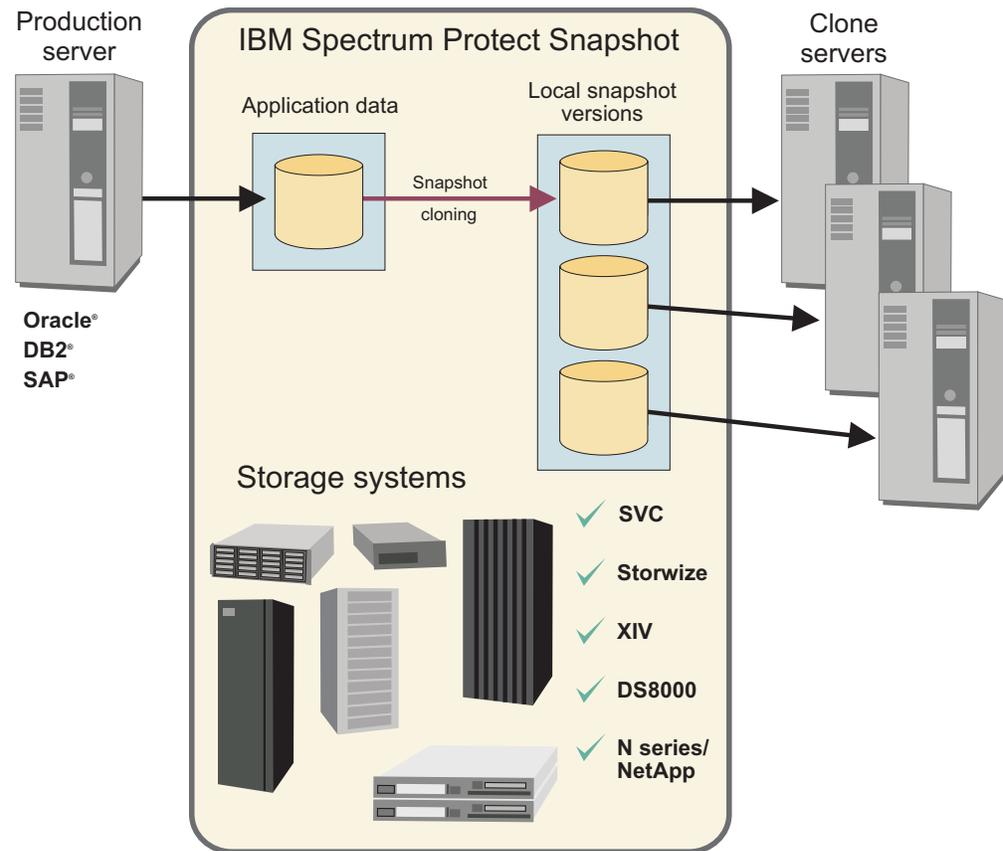


Figure 2. IBM Spectrum Protect Snapshot and database cloning

Software components

IBM Spectrum Protect Snapshot is composed of several software components.

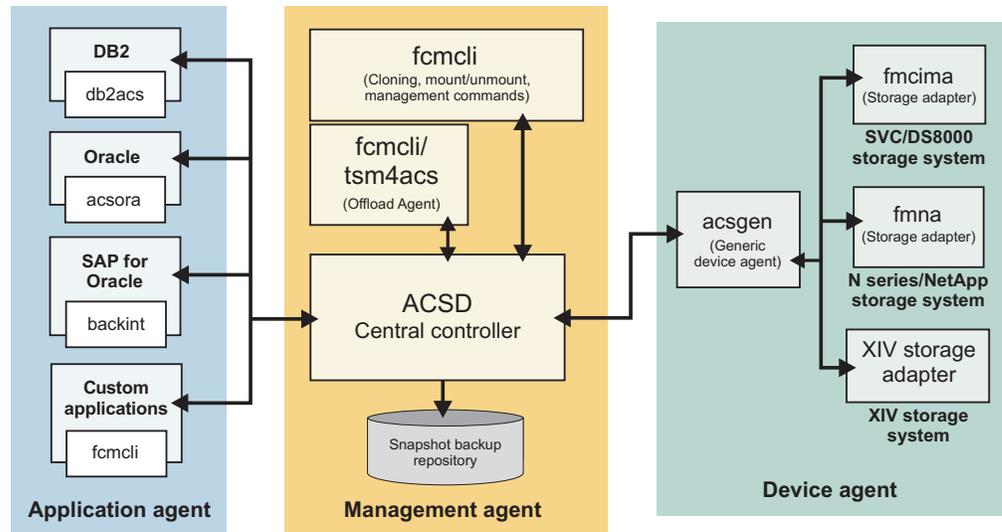


Figure 3. IBM Spectrum Protect Snapshot system components

Application agent

The application agent provides the necessary support to implement snapshot-based backup and restore operations. This agent interacts with the applications and tracks when an IBM Spectrum Protect Snapshot backup is created for a specific application.

Management agent

The management agent `acsd` coordinates all the components that are involved in backup, restore, and cloning operations. The agent controls the flow of information among the application and device agents, and other daemons. The agent provides access to the snapshot backup repository. This repository contains information about the snapshot backups and their relationships to snapshot-capable storage devices.

Device agent

The `acsgen` device agent is a generic agent that interacts with storage device-specific adapters and the central controller agent. This agent is also used to send and request updates of the progress and usability information that is stored in the local snapshot backup repository.

The following lists the storage device-specific agents that communicate with the `acsgen` agent:

- The CIM adapter `fmcima` is used with the generic device agent `acsgen`. This adapter sends commands to the supported storage device by using the CIM interface. Examples of supported storage include DS8000, Storwize V7000, and SAN Volume Controller.

Note: For Storwize V7000, and SAN Volume Controller storage systems, this communication using the CIM interface applies only in the case of static target allocation (device type SVC); the SVC adapter with dynamic target allocation (device type SVCDTA) uses the CLI interface via Secure Shell (SSH) rather than the CIM interface.

- The XIV system storage adapter is used with the generic device agent `acsgen`. This adapter communicates with the `acsgen` agent and issues commands to the XIV system Storage System by using the command-line interface XCLI.

Offload agent

The offload agent `tsm4acs` is used to send an existing snapshot to IBM Spectrum Protect. This agent also calls the generic device agent for mount and unmount operations on a backup system. From the command-line interface `fcml i`, you can manually start an offload backup to IBM Spectrum Protect.

IBM Spectrum Protect Snapshot command-line interface

The command-line interface `fcml i`, is used to issue various commands.

Chapter 2. Planning

Before you install IBM Spectrum Protect Snapshot for UNIX and Linux, review the system, application, and storage requirements.

Review the *Pre-installation Checklist* that is attached to the technote for the hardware and software requirements for IBM Spectrum Protect Snapshot. The detailed hardware and software requirements are published as a part of the *Hardware and Software Requirements* technote which can be found at this link: <http://www-01.ibm.com/support/docview.wss?uid=swg21427692>. From this technote, select the required software version and then select the required component link. The hardware and software requirements page contains the Pre-installation Checklist and an installation planning worksheet.

Note: The pre-installation checklist contains the most current requirement information, use this list to validate your environment.

The following conditions are the minimum environment requirements:

- A suitable disk layout of the application on the production server
- Correctly defined storage definitions on the storage system
- Connectivity from the production server to the storage system

The installation planning sheet helps you to determine the correct type of installation that is required for your environment. The following areas are covered in the planning sheet:

- How to determine the configuration mode for your environment.
- How to decide the parameters and settings for the specific application that you want to protect. The required parameters for each specific software application and custom application are outlined in the planning sheet.
- How to determine the parameters and settings for the specific storage system that you use in your environment.
- What passwords are required during the installation.

IBM Spectrum Protect Snapshot Prerequisite Checker

Run the checker tool to check the compatibility of the operating system, and available software that is to be used by IBM Spectrum Protect Snapshot in an AIX, or Linux environment. The Prerequisite Checker does not change the database or the system. Run the tool to retrieve information from the operating system and database in preparation for installing IBM Spectrum Protect Snapshot for DB2, and IBM Spectrum Protect Snapshot for Oracle and Oracle in an SAP environment.

The Prerequisite Checker is a tool that automatically checks your environment with a number of the checks that are documented in the IBM Spectrum Protect Snapshot *Pre-installation Checklist*. The *Pre-installation Checklist* is published as part of a release and is attached to the IBM Spectrum Protect Snapshot Hardware and Software Requirements technote.

The hardware and software requirements for IBM Spectrum Protect Snapshot for UNIX and Linux are published in the following technote: <http://www.ibm.com/support/docview.wss?uid=swg21427692>. Follow the link to the requirements technote for your specific release or update level. From there you will find the

Pre-installation Checklist and the Installation Planning Worksheet for the most recent version of the product.

Capacity planning

Ensure that there is sufficient storage space before you install and use IBM Spectrum Protect Snapshot.

The storage space that is required for IBM Spectrum Protect Snapshot can be divided into the following categories:

- Space that is required for the global product installation on the system.
- Space that is required to enable each individual database instance or custom application instance with IBM Spectrum Protect Snapshot. This enablement is referred to as activation.
- Space that is required on the storage system or in the GPFS file system to store the actual snapshot backups or clones.

Space requirement for global product installation

The space that is required for the product installation of IBM Spectrum Protect Snapshot varies depending on the underlying operating system. The following table shows the default installation paths and the average space requirements.

Table 1. Space requirements for a global product installation of IBM Spectrum Protect Snapshot

Operating system	Default installation path	Space required (MB)
AIX	<code>/usr/tivoli/tsfcm/acs_version_number</code>	1100
Solaris	<code>/opt/tivoli/tsfcm/acs_version_number</code>	700
Linux	<code>/opt/tivoli/tsfcm/acs_version_number</code>	500
HP-UX	<code>/opt/tivoli/tsfcm/acs_version_number</code>	1900

Space requirement for database instance or custom application installation

IBM Spectrum Protect Snapshot must also be installed on each database and custom application instance that is enabled for snapshot-based data protection or cloning. This process is called activation and must be started after the installation. During this process, all necessary files are copied from the installation path to a database instance-specific or custom application-specific installation directory. The space that is required for each IBM Spectrum Protect Snapshot enabled application is equal to the amount of space that is required for the global product installation.

IBM Spectrum Protect Snapshot must also be installed on application instances that are running on a backup server.

Extra space is required for IBM Spectrum Protect Snapshot log files. Log files are written continuously by IBM Spectrum Protect Snapshot without automatically deleting the older ones. You must monitor periodically the amount of space that is used by these log files and manually delete them if required.

Space requirement for snapshot copies

The snapshot copies of your application data or databases require the most space. The space that is required depends on the following factors:

- The total size of all storage system source volumes that are part of the volume group on the storage system. The volume groups contain the application data.
- The type of snapshot whether it is a full copy or a space-efficient snapshot.
- The number of backup copies.
- The number of changes that occur on the source volumes after a snapshot is taken. This factor applies to space-efficient snapshots only.
- For IBM Spectrum Protect Snapshot for Custom Applications, when snapshots are stored in the GPFS file system, that file system must have sufficient space to store all the snapshots. The size of a snapshot depends on the number of changes to the GPFS file system content that occur after the snapshot was taken. As a consequence, space requirements for a single snapshot can increase over time.

For remote mirroring with an XIV system, each backup copy uses space on the remote site storage and on the local site until it is deleted.

Use the **MAX_VERSIONS** parameter in the IBM Spectrum Protect Snapshot profile file to limit the number of snapshots that are stored on a storage system or in a GPFS file system.

On SAN Volume Controller, IBM Storwize family, and IBM System Storage DS8000, full snapshot copies require the same amount of space as the corresponding source volumes. If there is not enough storage space available, you can increase the capacity on the requested storage pool, or free up some items that are using existing capacity.

Required communication ports

IBM Spectrum Protect Snapshot for UNIX and Linux uses ports for communication between its daemon processes on backup or cloning systems, and the production system, and the storage systems. Port numbers are defined during the installation of IBM Spectrum Protect Snapshot for UNIX and Linux.

To determine the port number for the ports that are used for IBM Spectrum Protect Snapshot for UNIX and Linux see the following table:

Table 2. IBM Spectrum Protect Snapshot for UNIX and Linux port numbers.

TCP Port	Initiator: Out-Bound (From Host)	Target: In-Bound (To Host)
57328	Production server and backup/cloning server	ACSD port on production system
5989 (HTTPS port) ^[1] 5988 (HTTP port) ^[1] Note: Not applicable if you are using the new SVC storage adapter, in which case port 22 must be accessible on SAN Volume Controller storage for SSH access.	Production server and backup/cloning server	SAN Volume Controller Storwize family cluster CIM agent

Table 2. IBM Spectrum Protect Snapshot for UNIX and Linux port numbers (continued).

TCP Port	Initiator: Out-Bound (From Host)	Target: In-Bound (To Host)
6989 (HTTPS port) ^[1]	Production server and backup/cloning server	DS8000
6988 (HTTP port) ^[1]		DS8000 CIM Agent
7778	Production server and backup/cloning server	XIV system XIV system CLI
[1] Where COPYSERVICES_COMMPROTOCOL is the corresponding parameter name in the profile.		

Storage solutions

Before you install and configure IBM Spectrum Protect Snapshot software, review the storage solution setup. When the data to be protected is in a GPFS filesystem, IBM Spectrum Protect Snapshot is independent of the underlying storage that is used by the GPFS file system.

IBM XIV Storage System

When IBM Spectrum Protect Snapshot creates a backup on an IBM XIV Storage System, a snapshot of all source volumes that belong to the protected application is created on the storage system. This snapshot is a space-efficient read-only copy of the application.

The storage device and its storage volumes must be accessible from either a storage area network (SAN) zone, network or both. During the IBM Spectrum Protect Snapshot configuration process, if you set the **USE_WRITABLE_SNAPSHOTS** parameter to NO, the snapshots are not mounted directly on another host. Instead, IBM Spectrum Protect Snapshot creates duplicates from the snapshots as part of the mount procedure, these duplicates are removed when the backup is unmounted. A duplicate is a space-efficient logical copy of the snapshot and this copy is writable.

The **USE_WRITABLE_SNAPSHOTS** parameter specifies whether writable snapshots can be used for mount or restore operations. If writable snapshots are used, no duplicates are created during mount operations and all changes that are applied to the snapshot are preserved. Writable snapshots are only required in LVM mirroring environments. A typical IBM XIV Storage System profile entry is provided here:

```
>>>
DEVICE_CLASS                XIV01
COPYSERVICES_HARDWARE_TYPE  XIV
PATH_TO_XCLI                 path where XCLI is installed
COPYSERVICES_SERVERNAME     xiv_hostname
COPYSERVICES_USERNAME       admin
COPYSERVICES_REMOTE         YES
COPYSERVICES_PRIMARY_SERVERNAME xiv_hostname
COPYSERVICES_REMOTE_SERVERNAME xiv_remote_hostname
COPYSERVICES_REMOTE_USERNAME admin
USE_WRITABLE_SNAPSHOTS      AUTO
BACKUP_HOST_NAME            backup_host
<<<
```

To offload backups to IBM Spectrum Protect, IBM Spectrum Protect Snapshot must be installed on a backup server. You must also configure the **TSM_BACKUP** profile

parameter to YES and set the **BACKUP_HOST_NAME** profile parameter to the name of the hostname or cluster name as defined on the storage system.

For remote mirroring with an XIV system, each backup copy uses space on the remote site storage and on the local site until it is deleted.

Dependent software packages

IBM Spectrum Protect Snapshot requires the IBM XIV Storage System command-line interface (XCLI) to be installed on all hosts, production, backup, or clone servers where IBM Spectrum Protect Snapshot is installed.

Support for LVM mirroring (AIX only) and ASM failure groups

If AIX Logical Volume Manager (LVM) mirroring is used in the environment, IBM Spectrum Protect Snapshot can create separate snapshots of either mirror. In an Oracle ASM environment, a snapshot of selected failure groups is created. However, there must be enough remaining failure groups to mount the corresponding disk group for this image to be created. Each mirror or failure group must be located on a different XIV Storage System.

In LVM mirroring environments, the use of writable snapshots is required. IBM Spectrum Protect Snapshot uses IBM XIV Storage System capabilities to restore writable snapshots. For writable snapshots, a mount operation directly mounts the original snapshot to another host. All changes to the snapshot are preserved, and a subsequent mount or backup operation contains all changes that occurred to the snapshot while mounted. For more information about using writable snapshots, see information about the **USE_WRITABLE_SNAPSHOTS** parameter in **DEVICE_CLASS** section.

(AIX only) Support for virtual I/O

IBM XIV Storage System and IBM Spectrum Protect Snapshot support virtual I/O with n-port ID virtualization. On the production server, IBM Spectrum Protect Snapshot supports virtual I/O with N_Port ID Virtualization (NPIV) and Virtual I/O Server (VIOS). There is a one to one relationship between the virtual I/O logical volume and the storage LUN. On the backup server, IBM Spectrum Protect Snapshot supports virtual I/O with NPIV only.

Remote access to snapshot backups

Mounting a backup image onto another host with IBM Spectrum Protect Snapshot. IBM Spectrum Protect Snapshot creates a duplicate from the snapshot, which is then mounted on the host. As the duplicate is effectively another image, changes to the duplicate are not reflected in the snapshot. As a result, the mounted image can be altered without affecting the backup image and any subsequent restore of that backup. IBM Spectrum Protect Snapshot removes the duplicate during the unmount operation. All changes that were made on the duplicate are undone. A subsequent mount operation, presents the image as created when the snapshot occurred.

Best practices for IBM Spectrum Protect Snapshot with IBM XIV 11.6 Real-time Compression™

You can use IBM XIV 11.6 Real-time Compression with IBM Spectrum Protect Snapshot. The usage of IBM Spectrum Protect Snapshot with compressed volumes

does not change. However, when you transform volumes managed by IBM Spectrum Protect Snapshot from the uncompressed state to the compressed state (or if you transform from compressed to uncompressed), use the following list of behaviors as a guide:

1. When source volume transformation is in progress (from uncompressed to compressed, or compressed to uncompressed), most IBM Spectrum Protect Snapshot operations (for example, back up, restore, and mount) fail. The XIV adapter returns the **FMM18137E** message. Perform the volume transformation at a time that does not overlap with scheduled backups or other IBM Spectrum Protect Snapshot actions running on the volume that is being transformed.
2. With the XIV system, you can transform a volume from uncompressed to compressed state (or compressed to uncompressed state) using one of the following options:
 - With the `delete_source=yes` option, delete all volume backups. If you do not delete the volume backups, the transform is unsuccessful. You can use the IBM Spectrum Protect Snapshot GUI or CLI to manually delete the backups before the transform operation runs.
 - With the `delete_source=no` option, the volume backups are retained. After the transform completes, the original (source) volume is hidden from the host system. The original volume is replaced by the transformed volume. Any instant restore operation completed with the backups made before the transformation are restored to the hidden volume on the storage device. The restore is not made to the volume seen by the host. Note that the restore to the volume seen by the host appears to be successful, but the source volume visible to the host system is unchanged.

When using IBM Spectrum Protect Snapshot to protect volumes to be transformed, delete the existing snapshot backups, regardless of the `delete_source` option setting.

SAN Volume Controller and Storwize V7000 storage systems

IBM Spectrum Protect Snapshot restores point-in-time copies from backups on SAN Volume Controller, and Storwize V7000 storage systems. You can also mount images on a remote server and back up the images to IBM Spectrum Protect.

SAN Volume Controller storage adapter device types

IBM Spectrum Protect Snapshot for UNIX and Linux offers two backup solutions with Storwize V7000 and SAN Volume Controller storage systems.

When you configure IBM Spectrum Protect Snapshot, you can select one of the following device types (**COPYSERVICES_HARDWARE_TYPE**):

SVCDTA

Storwize V7000 and SAN Volume Controller: dynamic target allocation. During the backup process, target volumes are created dynamically and allocated on demand.

SVC Storwize V7000 and SAN Volume Controller: static target allocation. You must manually create target volumes on the storage system before the backup process.

The device type (**COPYSERVICES_HARDWARE_TYPE**) that you select is added to the device class section of the profile. The **COPYSERVICES_SERVERNAME** parameter stores the TCP/IP host name of the physical disk storage system.

For more information about configuring IBM Spectrum Protect Snapshot, see 'Configuration tasks > Running the setup script' for the application that is being protected.

Restriction: Both SVC and SVCDTA values are considered as different hardware types so limitations apply when they are used on the same storage system. For more information about the restrictions, see "Migrating from SVC with static target allocation to SVC with dynamic target allocation (SVCDTA)" on page 57

For a predefined target solution, before you start a backup operation you must ensure that the following tasks are completed:

- Target volumes are created on the storage system
 - Target sets for the volumes on the storage system are created
- A *target set* represents the mapping from the production host to the target volume on the storage system. You must specify a new target set for each backup generation to be retained on the storage system.

The following table provides a feature comparison between dynamic target volumes and predefined target volumes.

Table 3. Dynamic target volumes and predefined target volumes feature comparison.

Feature	Dynamic target volumes	Static target volumes
Command line interface	Storwize V7000 or SAN Volume Controller command-line interface (CLI)	Common Information Model (CIM) interface
Number of FlashCopy snapshot images retained	Specify an upper limit with MAX_VERSIONS	Limited by the number of target sets defined
Selectively restore a single FlashCopy snapshot image	Yes	Yes, however any FlashCopy image in the target set that is newer than the FlashCopy restored is deleted

Support for LVM mirroring (AIX only) and ASM failure groups

If AIX Logical Volume Manager (LVM) mirroring is used in the environment, IBM Spectrum Protect Snapshot can create separate FlashCopy images of either mirror. In an Oracle Automatic Storage Management (ASM) environment, a FlashCopy image of selected failure groups is created. However, there must be enough remaining failure groups to mount the corresponding disk group for this image to be created. Each mirror or failure group must be located in a different storage system.

Support for virtual I/O (AIX only)

DS8000, SAN Volume Controller, and Storwize V7000 logical unit numbers (LUNs) can be attached to a host directly or by using Virtual I/O (VIO). Both setups are supported, when there is a 1-1 relation between VIO logical volumes and storage LUNs on the storage subsystem.

A VIO is a logical partition (LPAR) on a pSeries system that is controlled by the IBM Hardware Management Console (HMC) or IBM Integrated Virtualization Manager (IVM). It owns the hardware adapters and allows access for other logical partitions. This feature allows the device to be shared. The LPAR associated with

the resources is the VIO Server and the logical partitions that use it are VIO Clients. For example, they can share one disk on the VIO Server instead of rebooting each logical partition from a Small Computer System Interface (SCSI) adapter and SCSI disk. This function eliminates the number of required adapters, adapter slots, and disks.

IBM Spectrum Protect Snapshot uses virtual SCSI adapters to map disks from a VIO to a client LPAR. Physical volumes are required to be mapped from the VIO to the client. However, mapping logical volumes or storage pools is not supported. On the production server, IBM Spectrum Protect Snapshot supports virtual I/O with N_Port ID Virtualization (NPIV) and Virtual I/O Server (VIOS). There is a one to one relationship between the virtual I/O logical volume and the storage LUN. On the backup server, IBM Spectrum Protect Snapshot supports virtual I/O with NPIV. In addition, VIOS is supported when you configure the **BACKUP_HOST_NAME** parameter to use the **PREASSIGNED_VOLUMES** in the IBM Spectrum Protect Snapshot profile file.

More details about supported combinations of operating system and storage subsystem levels, are available in the Pre-installation Checklist that is available at this URL <https://www.ibm.com/support/docview.wss?uid=swg21427692>. From this technote, select the required software version and then select the required component link. The hardware and software requirement page contains the Pre-installation Checklist and an installation planning worksheet.

Remote access to FlashCopy images

For static target allocation, IBM Spectrum Protect Snapshot allows mounting a FlashCopy backup image to another host. This image is writable and any changes that are made on that image are reflected in the backup and are included in the subsequent restore.

For dynamic target allocation, a writable duplicate is mounted which is dismissed on unmount. As a consequence, the original backup is not altered. For cloning operations, the backup is directly mounted in the same way as for static target allocation.

Related information:

 <https://www.ibm.com/support/docview.wss?uid=swg21427692>

Dynamic target allocation

This solution creates dynamic target volumes on the storage system during a backup operation.

During the backup process, target volumes are created dynamically and allocated on demand. IBM Spectrum Protect Snapshot uses the Storwize V7000 or SAN Volume Controller command line interface (CLI) to communicate with the storage system. You do not need to install a Common Information Model (CIM) server.

Tip: Ensure that OpenSSH is installed on the Production and Backup servers. During the configuration process, you are prompted for the location of the OpenSSH binary.

Important: You must set a specific number of backup generations to retain because of the space calculations for dynamic target allocation. The configuration wizard prevents you from using the **ADAPTIVE** option if at least one **DEVICE_CLASS** is **SVCDTA**.

In SAN Volume Controller environments where the source volumes of a backup are mirrored internally and the copies are in two different SAN Volume Controller storage pools, the storage pool for the target volumes is not automatically determined. You must specify the target storage pool with the **SVC_POOLNAME** parameter in the **DEVICE_CLASS** section of the IBM Spectrum Protect Snapshot profile when the **COPYSERVICES_REMOTE** is *YES*.

Space-efficient multi-target FlashCopy on SAN Volume Controller and Storwize V7000

Space-efficient targets that are part of a multi-target FlashCopy cascade might be deleted by SAN Volume Controller and Storwize V7000 if other targets of the same cascade are restored or overwritten by a new snapshot.

In a SAN Volume Controller or a Storwize V7000 environment, the following situations might cause space-efficient targets to be deleted:

Backup operations and cloning operations

An IBM Spectrum Protect Snapshot backup operation uses the oldest target set that is available for the specified **DEVICE_CLASS**. However, that target set might not be the oldest target set that is associated with the source volumes. This scenario is possible when more than one **DEVICE_CLASS** is specified in the IBM Spectrum Protect Snapshot profile. When the FlashCopy backup that is available on the target set is not the oldest backup, then the older backups are deleted during the backup operation. The oldest target set is the set that is used for the oldest FlashCopy backup in a multiple target set configuration. This situation can also happen when a new FlashCopy cloning operation is started with the force option (-F).

Important: This does not apply if you select SAN Volume Controller and Storwize V7000 dynamic target allocation.

Restore operation

An IBM Spectrum Protect Snapshot restore operation deletes any FlashCopy backups that are newer than the backup that is being restored. In addition, the backup that is restored with the current operation can also be deleted.

Important: This does not apply if you select SAN Volume Controller and Storwize V7000 dynamic target allocation.

Target volume storage space exceeded

When the available storage capacity of a space-efficient FlashCopy target volume is exceeded, the target volume is taken offline. The data on the target volume that is taken offline is deleted.

SAN Volume Controller static target allocation and Storwize V7000

When you use SAN Volume Controller and Storwize V7000, IBM Spectrum Protect Snapshot software can restore FlashCopy backups before completion of a background copy.

When you restore FlashCopy backups before completion of a background copy, space-efficient volumes can be enabled as backup targets. The background copy rate is set to zero to prevent the FlashCopy target from becoming fully allocated. When you use either SAN Volume Controller or Storwize V7000, and IBM Spectrum Protect Snapshot software in this scenario, use the following guidelines for the environment:

Physical capacity

The physically allocated capacity of a space-efficient target volume must be large enough to contain all changes that occur to your production environment. Specifically, all changes that occur between the current and the subsequent backup. If the capacity is insufficient, the target volume goes offline and the corresponding backup becomes invalid.

SAN Volume Controller and Storwize V7000 support the creation of automatically expanding target volumes. If you create target volumes that automatically expand, more storage is assigned to the target when storage capacity decreases. This additional storage ensures that sufficient capacity is available.

Tip: If you select SAN Volume Controller and Storwize V7000 dynamic target allocation, all target volumes that were created dynamically will be auto-expandable.

FlashCopy relationships

During a restore, IBM Spectrum Protect Snapshot software stops FlashCopy relationships. These relationships include relationships that are established at the time when the backup is created to any subsequent relationships that are created on the same source LUN. All backups to space-efficient targets that are newer than the backup used for restore, and the backup from which you are restoring, are deleted. If the background copy was not completed, the same restriction applies to full and incremental FlashCopy backups.

To check whether a backup is going to be deleted, query the usability state of IBM Spectrum Protect Snapshot backups. If the backup is going to be deleted, during the restore process, the `DESTRUCTIVELY_RESTORABLE` state is set. Otherwise, the state is set to `REPETITIVELY_RESTORABLE`.

Important: This does not apply if you select SAN Volume Controller and Storwize V7000 dynamic target allocation. With SVCDTA, no backups are deleted during a restore operation.

Target sets

IBM Spectrum Protect Snapshot cannot reuse a target set for a new FlashCopy backup unless it corresponds to the last FlashCopy mapping in a cascaded FlashCopy relationship. This scenario implies that when IBM Spectrum Protect Snapshot reuses a target set, all backups that are created before this point in time are deleted. In a non-mirrored environment, all backups that are created before this point in time are deleted when the following conditions are met:

- The same profile for the IBM Spectrum Protect Snapshot backups is used.
- This profile contains only one **DEVICE_CLASS** statement in the `CLIENT` section.

In an LVM mirrored environment, all backups that are created before this point in time are deleted when the `CLIENT` section of the profile contains one **DEVICE_CLASS** statement for each LVM mirror. If multiple device classes are specified within this statement, each device class must manage the same number of target sets.

Important: This does not apply if you select SAN Volume Controller and Storwize V7000 dynamic target allocation.

Recommendations for setting up the environment with static target volumes

When you set up the SAN Volume Controller and Storwize V7000 environments for use with IBM Spectrum Protect Snapshot software, the following list identifies guidelines for the environment:

- If space-efficient source volumes are used in combination with space-efficient target volumes, IBM Spectrum Protect Snapshot can be configured to use **FLASHCOPY_TYPE** COPY, INCR, or NOCOPY. If fully allocated source volumes are used in combination with space-efficient target volumes, then IBM Spectrum Protect Snapshot can be configured to use **FLASHCOPY_TYPE** NOCOPY only.
- Decide whether you want to use space-efficient or fully allocated backup targets. In mirrored environments, a different choice can be made for each mirror.
- For each mirror, use one **DEVICE_CLASS** statement for disk-only backups. In addition, use one **DEVICE_CLASS** statement for dual backups. A dual backup is a disk backup and tape backup. Make sure that the schedule is defined so that the target sets are reused cyclically across both device classes per mirror.

For example:

- Define three target sets in the **DISK_ONLY** device class. Schedule these disk only backups to occur at 6:00, 12:00, and 18:00.
- Define one target set in a **DUAL_BACKUP** device class. Set this schedule to create a disk and IBM Spectrum Protect backup at 00:15.

If you retain only one target set generation for dual backups, do not specify six target sets to retain disk only backups (created at 6:00, 12:00, and 18:00) for two days. The second dual backup operation attempts to reuse the target set of the previous dual backup. If the version policy specifies ADAPTIVE, this action results in a deletion of all disk-only backups that are taken before that point in time. Otherwise, the version policy causes the dual backup to fail if **retain** specifies seven versions.

- If a backup that is characterized as DESTRUCTIVELY_RESTORABLE is restored, the backup you are restoring and all backups that are taken after that point in time are deleted. The backup is not deleted when the backup is created with **FLASHCOPY_TYPE** FULL or INCR, and the background copy completed.

DS8000 storage system

For the DS8000 storage system, it is not possible to restore point-in-time copies when you set the **FLASHCOPY_TYPE** parameter to *NOCOPY* in the IBM Spectrum Protect Snapshot profile file.

You can mount images on a remote server and back up the images to IBM Spectrum Protect when you use DS8000 storage systems.

CIM server

Starting with DS8000 R4.1 the Common Information Model (CIM) server is embedded with the storage device. It is not necessary to install and configure the CIM server separately. For earlier releases of DS8000, a proxy CIM server is required and must be configured to manage the necessary storage clusters. For more information about configuring a proxy CIM server, see the DS8000 documentation.

IBM Spectrum Protect Snapshot requires that FlashCopy backup target volumes be created in advance on DS8000. To provide a target set definition to IBM Spectrum

Protect Snapshot, organize target volumes into target sets, where each target set represents one backup generation. IBM Spectrum Protect Snapshot automatically matches source volumes to suitable target volumes. However, each target set must contain at least one suitable target volume for each source volume to be backed up. Additional target volumes in a target set are allowed, but these target volumes are ignored.

Support for LVM mirroring (AIX only) and ASM failure groups

If AIX Logical Volume Manager (LVM) mirroring is used in the environment, IBM Spectrum Protect Snapshot can create separate FlashCopy images of either mirror. In an Oracle Automatic Storage Management (ASM) environment, a FlashCopy image of selected failure groups is created. However, there must be enough remaining failure groups to mount the corresponding disk group for this image to be created. Each mirror or failure group must be located in a different storage system.

DS8000 allows one incremental FlashCopy per source volume. When production volumes are mirrored by using Logical Volume Manager (LVM) mirroring or ASM failure groups, only one FlashCopy backup of this type per volume mirror is created. For incremental snapshots with DS8000 storage, only one target set can be specified in the target volumes file (.fct).

Support for virtual I/O (AIX only)

DS8000 logical unit numbers (LUNs) can be attached to a host directly or by using Virtual I/O (VIO). Both setups are supported, when there is a 1-1 relation between VIO logical volumes and storage LUNs on the storage subsystem.

A VIO is a logical partition (LPAR) on a pSeries system that is controlled by the IBM Hardware Management Console (HMC) or IBM Integrated Virtualization Manager (IVM). It owns the hardware adapters and allows access for other logical partitions. This feature allows the device to be shared. The LPAR associated with the resources is the VIO Server and the logical partitions that use it are VIO Clients. For example, they can share one disk on the VIO Server instead of rebooting each logical partition from a Small Computer System Interface (SCSI) adapter and SCSI disk. This function eliminates the number of required adapters, adapter slots, and disks.

IBM Spectrum Protect Snapshot uses virtual SCSI adapters to map disks from a VIO to a client LPAR. Physical volumes are required to be mapped from the VIO to the client. However, mapping logical volumes or storage pools is not supported. On the production server, IBM Spectrum Protect Snapshot supports virtual I/O with N_Port ID Virtualization (NPIV) and Virtual I/O Server (VIOS). There is a one to one relationship between the virtual I/O logical volume and the storage LUN. On the backup server, IBM Spectrum Protect Snapshot supports virtual I/O with NPIV. In addition, VIOS is supported when you configure the **BACKUP_HOST_NAME** parameter to use the **PREASSIGNED_VOLUMES** in the IBM Spectrum Protect Snapshot profile file.

More details about supported combinations of operating system and storage subsystem levels, are available in the Pre-installation Checklist that is available at this URL <https://www.ibm.com/support/docview.wss?uid=swg21427692>. From this technote, select the required software version and then select the required component link. The hardware and software requirement page contains the Pre-installation Checklist and an installation planning worksheet.

Remote access to FlashCopy images

IBM Spectrum Protect Snapshot allows mounting a FlashCopy backup image to another host. This image is writable and any changes that are made on that image are reflected in the backup and are included in the subsequent restore.

Related information:

 <https://www.ibm.com/support/docview.wss?uid=swg21427692>

Reconciliation of backups

Reconciliation is the process where IBM Spectrum Protect Snapshot periodically verifies that backups on the storage system are valid.

Depending on the storage system, FlashCopy or snapshot backups can be deleted, withdrawn, or stopped by certain operations on the storage system. When these events occur, it invalidates the FlashCopy or snapshot backup. During reconciliation FlashCopy or snapshots backups that are no longer present or are invalid on the storage system are removed from IBM Spectrum Protect Snapshot repository.

The reconciliation process removes IBM Spectrum Protect Snapshot backups when the following events take place on storage systems:

All storage systems

Manual intervention causes the following events to occur:

- The source volume or target volume relationship is withdrawn.
- The snapshot or FlashCopy is deleted.
- The FlashCopy mappings are stopped.

IBM XIV Storage System

When there is no available space for snapshot backups, the XIV system Storage System deletes old snapshots to free space for new snapshots.

IBM System Storage SAN Volume Controller and IBM Storwize family storage systems

When either of the following events occur:

- When a FlashCopy backup becomes invalid, because it was created after the creation of the original backup that was later restored. This issue applies to backups with space efficient target volumes or if the background copy process is not yet finished. In addition, the backup that is subject to restore can also be invalidated by the storage system.
- In this environment FlashCopy mappings of target volumes are used by the storage system for FlashCopy backups. When used in a specific FlashCopy backup, then previous FlashCopy backups can become invalid if they are dependent on the same mapping. This issue applies to backups with space efficient target volumes or if the background copy process is not finished.

Restriction: This does not apply for the SVC storage adapter with dynamic target allocation. Neither backup nor restore operations using the SVCDTA adapter will invalidate other backups.

IBM System Storage DS8000

When a source target relationship is withdrawn. This process cannot happen automatically in this environment.

Remote mirror integration

When you use storage solutions with mirror technologies and IBM Spectrum Protect Snapshot, certain criteria must be met by the environment to integrate backup, restore and cloning operations. For IBM System Storage SAN Volume Controller, mirror technologies are labeled Global Mirror and Metro Mirror. For IBM XIV Storage System, mirror technologies are labeled Synchronous Remote Mirroring and Asynchronous Remote Mirroring.

SAN Volume Controller

IBM Spectrum Protect Snapshot backs up application data consistently on SAN Volume Controller storage solutions with volumes that are simultaneously used as Metro Mirror or Global Mirror sources. You can configure either the sources or the targets of the Remote Mirror to be selected as the sources for the FlashCopy backup. In addition, do not use FlashCopy targets as Global Mirror or Metro Mirror sources.

IBM System Storage DS8000

IBM Spectrum Protect Snapshot can back up DS8000 storage solutions with volumes that are simultaneously used as Global Mirror or Metro Mirror sources. In contrast to SAN Volume Controller, you can configure only the sources of the Global Mirror or Metro Mirror to be selected as the sources of the FlashCopy backup. When you use IBM Spectrum Protect Snapshot in this environment, do not use FlashCopy targets as Global Mirror and Metro Mirror sources.

IBM XIV Storage System

IBM Spectrum Protect Snapshot can back up application data consistently on XIV system storage solutions with volumes that are simultaneously used as Synchronous Remote Mirroring or Asynchronous Remote Mirroring sources. You can configure either the sources or the targets of the Remote Mirror to be selected as the sources for the FlashCopy backup.

Storage solutions that use mirror technologies with IBM Spectrum Protect Snapshot must have the correct environment. The following list describes the criteria that must be met to ensure mirroring works correctly.

- The connectivity state must be online.
- The cluster partnership between the primary and secondary clusters must be configured before you use IBM Spectrum Protect Snapshot. The following list identifies what you must configure when you are setting up the cluster partnership:
 - IBM Spectrum Protect Snapshot is installed on the production and backup host on the local site (primary cluster).
 - IBM Spectrum Protect Snapshot is installed on all systems, including the takeover and standby servers, running at the remote site (secondary cluster).
 - The local site contains the primary storage cluster for the production hosts. The primary cluster has data that is replicated to a secondary cluster on the remote site or to the same cluster.
 - For intersystem copying, the remote site contains the mirror volumes in another storage cluster. In addition, the remote site also hosts the takeover and standby servers.
 - SAN Volume Controller supports both intrasystem and intersystem Metro and Global Mirror.

- For XIV system Synchronous Remote Mirroring and Asynchronous Remote Mirroring, configure either the source or the targets as a source for the snapshot backup.
- IBM Spectrum Protect Snapshot uses a consistency group on the SAN Volume Controller and XIV system storage solutions for the FlashCopy or snapshot. A consistency group is a group of volumes that are associated with a FlashCopy pair. A FlashCopy pair is a group of two corresponding instant copies of data, that is, point-in-time copies of a volume. For the FlashCopy pair, the logically related data must be kept consistent across the volumes. The FlashCopy consistency group can be used for a consistent point-in-time copy for an application or database that spans multiple volumes. The following list identifies more information about using consistency groups with IBM Spectrum Protect Snapshot:

SAN Volume Controller

- A consistency group contains a list of FlashCopy or Remote Copy relationships.
- The IBM Spectrum Protect Snapshot software creates a FlashCopy consistency group on the secondary site to build a consistency unit between the source and target of the FlashCopy.
- You must define the consistency group for the mirror relationships between the master and auxiliary virtual disks.
- For Metro and Global Mirror, the state of the consistency group must be consistently synchronized.

XIV system

- The operational state of mirror must be operational.
- A consistency group contains a list of volumes.
- A consistency group that contains all of the remote copy target volumes must exist before starting the snapshot on the remote system. Apply the storage commands to the consistency group to simplify management.
- The mirror relationship between the master and slave volumes must be defined in the consistency group.
The master is where source volumes are located for the remote replication. The slave is where target volumes are located.
- For XIV system synchronous mirroring, the state of the consistency group must be consistently synchronized.
- For XIV system asynchronous mirroring, the state of the consistency group must be RPO_OK.
- For Metro Mirror and Synchronous Remote Mirroring, the write operation is committed to the host after the data is written to both the source and target volumes.
- For Global Mirror and Asynchronous Remote Mirroring, the write operation is committed to the host immediately after the data is written to the source volume.
- In terms of master and slave sites, the master site is where source volumes are located for the remote replication. The slave site is where target volumes are located. When a disaster occurs or when maintenance is necessary, the roles of master site and slave site can be changed.

The following figure illustrates the hosts and volumes that are involved in remote mirroring that uses Metro and Global mirrors.

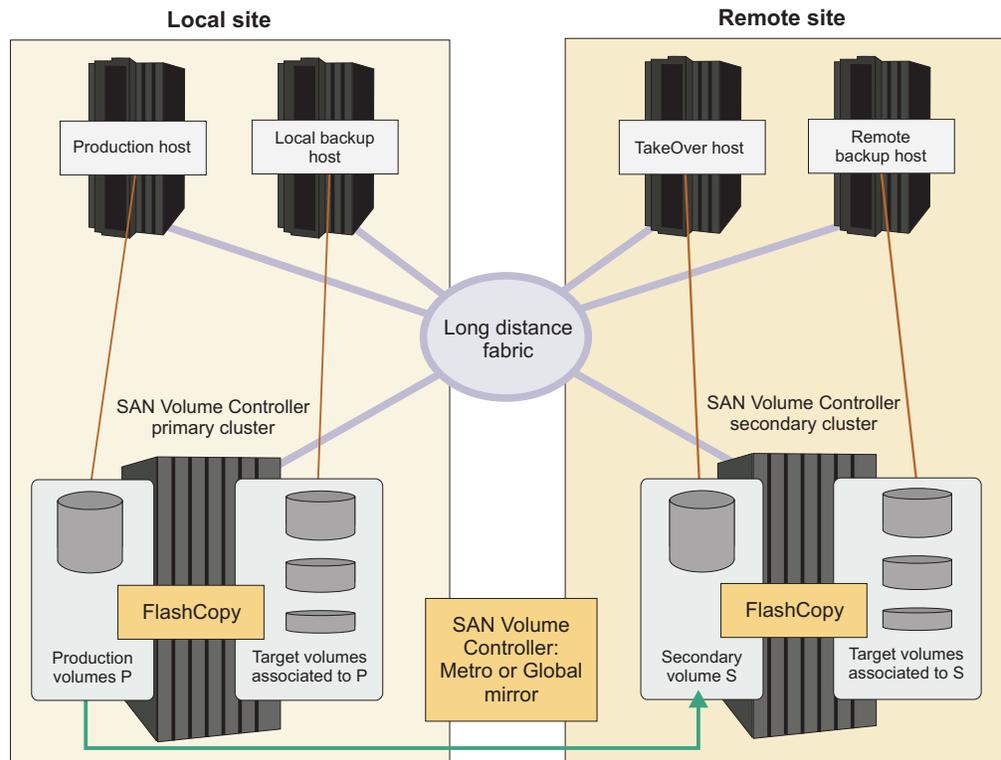


Figure 4. Remote mirroring using Metro Mirror and Global Mirror sources

Remote mirroring and consistency groups

You must verify the configuration of the consistency group on SAN Volume Controller and XIV systems that use mirroring functions before you run IBM Spectrum Protect Snapshot backup operations.

A *consistency group* is a group of copy relationships. You can group relationships into a consistency group that manages the consistency of dependent writes by creating a consistent point-in-time copy across multiple volumes or storage systems.

You must ensure that the connectivity state is online and configured for a SAN connection between the primary and secondary storage systems. The primary site contains the primary storage volumes for the production site. The volumes are then replicated to target volumes on the secondary site. IBM Spectrum Protect Snapshot requires the following configuration:

- For SAN Volume Controller, you must configure the consistency group:
 - For Metro Mirrors for static and dynamic target allocation, ensure that the state of the consistency group is consistently synchronized.
 - For Global Mirrors with dynamic target allocation, you must configure a *Global Mirror with Change Volumes* relationship:
 - Ensure that the consistency group for the relationship has cycling mode set to multiple by selecting the Global Mirror with Change Volumes option when you create the relationship between the volumes. Global Mirror with Change Volumes is the name for a point-in-time asynchronous volume

replication. You can create change volumes either when you create the Global Mirror relationships or you can add them to an existing relationship. Cycling mode and change volumes are not needed when you assign target allocation manually.

- The cycle period time set for the cycling mode and the number of I/O operations can influence the IBM Spectrum Protect Snapshot FlashCopy backup time. IBM Spectrum Protect Snapshot waits until the volumes at both sites are synchronized before a backup operation is completed. The cycle period is defined in seconds. The higher the cycle period the longer the time that is required for synchronization and to complete a FlashCopy backup. The factors that can influence the time are the number of I/O operations and the spread of the block-level changes across the storage system. The default value is 300 seconds.

Restriction: When you set the cycle period, the initial replication from the primary site change volume to the secondary change volume can take several hours before the volumes are synchronized. If you start an IBM Spectrum Protect Snapshot backup operation during this initial replication, the backup operation can fail due to the amount of time that is taken to complete the synchronization operation. Therefore, wait until the initial replication of change volumes is completed before you start a backup operation.

- For XIV systems, you must configure the consistency groups:
 - The consistency group must contain a list of mirrors.
 - The consistency group must contain a list of all of the remote copy target-volumes and this list must exist before you start the snapshot on the remote system.
 - The mirror relationship between the master (source) and slave (target) volumes must be defined in the consistency group. The master is on the source volume. The slave is on the target volume.
 - For synchronous mirroring, the state of the consistency group must be consistently synchronized.
 - For asynchronous mirroring, the state of the consistency group must be RPO_OK.

Logical Volume Manager support (AIX only)

You can use IBM Spectrum Protect Snapshot in environments where volume groups are mirrored between two storage clusters by using Logical Volume Manager (LVM) mirroring on AIX.

This support is provided on IBM System Storage DS8000, IBM System Storage SAN Volume Controller, IBM Storwize family, and IBM XIV Storage System. When LVM mirroring is used to mirror volume groups between two storage clusters, a FlashCopy backup is created such that only one mirror is being copied.

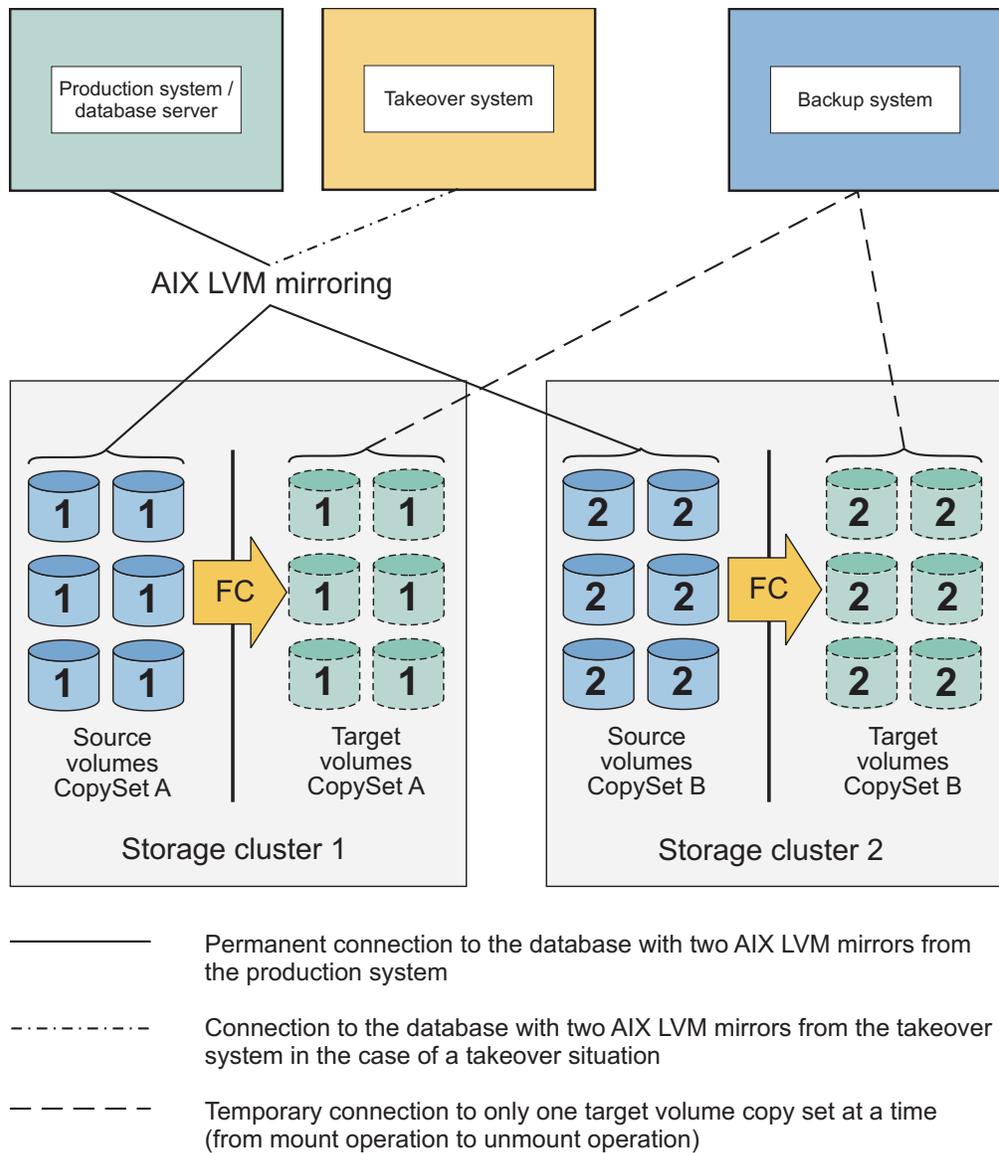


Figure 5. IBM Spectrum Protect Snapshot in an LVM environment

AIX LVM mirroring provides these advantages:

- Only one of the two LVM mirrors are used in the FlashCopy process. Using one mirror saves the number of needed target volumes and reduces the time that is needed for the FlashCopy process.
- Avoids unnecessary performance degradation within the storage system.
- All LVM mirrors on the production system remain synchronized during the FlashCopy backup process.
- Online or offline FlashCopy backups can be created in both LVM mirrored and non-LVM mirrored environments. There is no change in the backup and restore procedures as provided in the applicable documentation.
- The FlashCopy backup process at no time compromises the high-availability purpose for which the mirrors were set up. It is not necessary to resynchronize the logical volumes after the FlashCopy backup request.
- IBM Spectrum Protect Snapshot provides information about asymmetrical LVM mirror setups when encountered. This information can prevent the FlashCopy

backup from running in unfavorable situations but can also reveal a general deficiency of the high-availability setup as well.

IBM Spectrum Protect Snapshot requires that the LVM mirroring sets are in different storage subsystems. For example, different SAN Volume Controller clusters, Storwize V7000, DS8000, or XIV system. Complete mirrors are recommended to be stored on both storage clusters. If this setting is not possible, IBM Spectrum Protect Snapshot continues processing for those clusters where a complete image of the application can be found.

To configure IBM Spectrum Protect Snapshot for LVM mirroring, define both storage subsystems within the IBM Spectrum Protect Snapshot profile. Use the **DEVICE_CLASS** parameter to allow IBM Spectrum Protect Snapshot to select the storage subsystem. At least one backup server is required so that IBM Spectrum Protect Snapshot can mount a FlashCopy backup to verify the consistency of the backup and split the LVM mirrors.

During a restore operation, IBM Spectrum Protect Snapshot runs all the commands that are required to prepare the LVM environment again for the second mirror. The administrator is informed by message FMM0755I in the detailed restore log file that the volume groups are ready for synchronization. The administrator can run this operation at a more suitable time for instance after completion of the database recovery.

Note: The administrator must examine the log files for these messages. They do not display on the screen.

Preparing applications that run on VMware or KVM virtual machines

Before you install IBM Spectrum Protect Snapshot on VMware or KVM virtual machines that run Linux guest operating systems, you must verify the configuration of the application that you want to protect.

Before you begin

Different applications have specific IBM Spectrum Protect Snapshot configuration requirements. For more information about application-specific requirements, see Chapter 2, "Planning," on page 7.

Procedure

VMware

- Before you back up data or clone databases on VMware virtual machines, ensure that all source LUNs in the backup or clone operations are attached to the virtual machine with one of the following methods:
 - VMware physical mode raw device mapping (pRDM)
 - iSCSI
 - Network file system (NFS)
- Run an IBM Spectrum Protect Snapshot restore operation from a snapshot to an existing pRDM disk. The operation does not create a virtual machine or pRDM definition as part of the restore process.

KVM

- Before you back up data or clone databases on KVM virtual machines, ensure that all source LUNs in the backup or clone operations are attached to the virtual machine with one of the following methods:
 - Block device mapping (BDM)
 - iSCSI
 - Network file system (NFS)
 - PCI Passthrough
- Run an IBM Spectrum Protect Snapshot restore operation from a snapshot to an existing BDM disk. The restore operation does not create a virtual machine or BDM definition as part of the restore process.

Checking the KVM setup

Ensure that when the IBM Spectrum Protect Snapshot KVM setup uses Block Device Mapping, the LUNs are mapped to the KVM guest as multipath devices. The LUNs must be visible as multipath devices inside the KVM guest. Run the **multipath** command to check your setup for KVM.

Procedure

To verify your KVM setup, run the **multipath** command from within the KVM guest. The command output looks similar to the following example:

```
kvm-guest:~ # multipath -ll
mpathat (360050768018205de4000000000001949) dm-7 IBM ,2145
size=2.0G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  `-- 3:0:0:3 sdf 8:80 active ready running
```

In the example, *360050768018205de400000000001949* is the LUN identifier. It is a unique number that must not be overwritten by the KVM stack. The product storage identifier must be visible inside the KVM guest. In the example, this identifier is *IBM ,2145*.

Chapter 3. Preparation for installation

Before you install IBM Spectrum Protect Snapshot, review the hardware, software requirements, and application environment. You must complete the Pre-installation Checklist and Planning Worksheet before you install IBM Spectrum Protect Snapshot for UNIX and Linux.

The hardware and software requirements for IBM Spectrum Protect Snapshot for UNIX and Linux are published in the following technote: <http://www.ibm.com/support/docview.wss?uid=swg21427692>. Follow the link to the requirements technote for your specific release or update level. From there you will find the *Pre-installation Checklist* and the *Installation Planning Worksheet* for the most recent version of the product.

To help you to prepare your environment for IBM Spectrum Protect Snapshot for AIX and Linux, you can run the preinstallation checker tool. For more information about the preinstallation checker, see “IBM Spectrum Protect Snapshot Prerequisite Checker” on page 7.

Before you start the installation process, complete the following tasks:

- Review the requirements and ensure that all requirements are met.
- Complete the *Pre-installation Checklist*.
- Complete the *Installation Planning Worksheet*.

Important: You must complete the *Pre-installation Checklist* and *Installation Planning Worksheet* before you install the product.

Before you install IBM Spectrum Protect Snapshot, ensure that the volume and storage layout is correct for your application environment.

Related concepts:

“Prerequisite checker for Oracle” on page 28

Required Operating System users and environment variables

The *instance* is a term that is used for a set of background processes and shared memory. The term *database instance* refers to a set of data that is stored on disk. An Oracle instance can mount and open a single database, and a database can be mounted and opened by one or more instances. In Oracle RAC environments, multiple instances open and mount the same database.

Often the Operating System (OS) user *oracle* has administrator permissions for databases in Oracle installations. IBM Spectrum Protect Snapshot requires a dedicated OS user for each instance. The environment variables *ORACLE_HOME* and *ORACLE_SID* must be set permanently for that user, and refer to the instance that can be used to connect to the database that is to be protected by IBM Spectrum Protect Snapshot. The user requires administrator permissions.

If you want to protect one database on a host, and plan to use the backup server to back up only one database with one Oracle instance, you can reuse the *oracle* user.

Set the `ORACLE_SID` variable to the instance that is used to connect to that database permanently on both the production server and the backup server in the `oracle` user profile.

If you want to protect multiple databases on one production or backup server host, you need one OS user per database on both the production server and backup server. The `ORACLE_SID` that is set in the environment for each OS user, enables SQL*Plus connections to the specific database to be protected by IBM Spectrum Protect Snapshot. Export the Oracle-specific environment variables, such as `ORACLE_HOME` and `ORACLE_SID` so that they are accessible when you enter the `su - oracle_user -c` command.

Prerequisite checker for Oracle

Check your system by running the Prerequisite checker tool before you install IBM Spectrum Protect Snapshot for Oracle and Oracle in an SAP environment.

You must complete the *Pre-installation Checklist* checklist before you install IBM Spectrum Protect Snapshot. In AIX and Linux environments, running the Prerequisite Checker tool automatically runs some of the checks that are documented in the *Pre-installation Checklist*. Running the tool on your AIX or Linux system, automatically checks for compatible operating system, database instance, and volume group layout in preparation for installing the product.

The *Pre-installation Checklist* is published here: <http://www.ibm.com/support/docview.wss?uid=swg21427692>.

Installing the Prerequisite Checker

As part of your planning activities, install and run the Prerequisite Checker tool before you install or upgrade to a new version of IBM Spectrum Protect Snapshot for UNIX and Linux. Running the tool on your system automatically checks for compatible operating system, database instance, and volume group layout in preparation for installing IBM Spectrum Protect Snapshot for UNIX and Linux.

Procedure

1. Download the IBM Spectrum Protect Prerequisite Checker file for your operating system from the download website, or extract the file from the product DVD. For information about downloading the Prerequisite Checker, see the Download Information.
2. Log on with the root user ID.
3. Start the installation wizard by running one of the following commands using the default swing console:

```
AIX: <VERSION>-FCM-PREREQ-AIX.bin [-i console | swing]
```

```
Linux: <VERSION>-FCM-PREREQ-Linux.bin [-i console | swing]
```

where `-i console` indicates that the Prerequisite Checker is installed with the console version of the installer. `-i swing` indicates that the Prerequisite Checker is installed with the GUI version of the installer, which is the default method.

4. Complete the steps of the installation wizard. Choose to install the Prerequisite Checker to an arbitrary `checker_path`.

Running the Prerequisite Checker

Run the Prerequisite Checker any number of times for any database instances on the production server, and review the `results.html` in your browser.

Before you begin

Log on to the production server that is to be supported by IBM Spectrum Protect Snapshot, with the root user ID. Check the following requirements:

- The database must be activated; Oracle databases must be mounted.
- The default environment of the database owner must contain all the environment settings necessary for interaction with the database. The default shell must be included. On AIX systems for example, set `BASH_ENV` to point to the user profile in the `/etc/environment` configuration file.
- The database owner must have the necessary access rights.

Procedure

1. Log on with the root user ID.
2. Change to the `checker_path` directory where the Prerequisite Checker was installed.
3. Run the `fcmprereqchecker.sh` script as follows:

Oracle environment:

```
fcmprereqchecker.sh -u dbusername -s storage_management_IP_address  
-p storage_management_port -r connection_string  
[-o output_path] [-d database_name]
```

Where,

`dbusername` is the name of the database owner.

`storage_management_IP_address` is the name or IP address of the storage subsystem that contains the database files.

`storage_management_port` is the management port of the storage subsystem that contains the database files.

`output_path` is used to specify a fully qualified directory where all output files and information are written. The default output path is `checker_path/logs`.

`database_name` is used to specify the name or alias of the database to be checked.

`connection_string` for Oracle databases is the Oracle RMAN connection information including user name, password, and SID of the catalog database. For example, `-r username/password@SID`.

Interpreting the Prerequisite Checker output

After you run the Prerequisite Checker, the results are stored to the `result.html` file that can be viewed in your default browser in the Prerequisite Checker Results page. In the case of passed checks, there is no corrective action required. For warnings and failures, you must modify your system before you proceed to install IBM Spectrum Protect Snapshot for UNIX and Linux.

About this task

The check results are stored in `result.html` and can be opened in your browser. The results are also available in the `result.txt` file in `output_path/`.

The Summary reports the overall result of the checks run; the status is either Failed or Passed. The machine name, Operating System, and serial number are listed. If your system was fully compliant and passed, the completed checks are listed followed by the next steps you must take.

If your system did not meet the prerequisites, the Summary status is failed. The Critical checks that failed are listed, followed by any warning checks that were unsuccessful. You must review all warnings and take appropriate action, such as running a check manually. All failed checks must pass successfully before you proceed to install the product.

Procedure

- Find the `result.html` file and open it in your browser.
The `result.html` file is stored in the installation path of the Prerequisite Checker, `<install_dir>/logs/`. There is also a text file version of the results stored there, `result.txt`. For information about fails and warnings, including more message information, see `<install_dir>/logs/logfile`.
If you specified a different output path with the `-o` option, the `result.html` and log files are stored there.
- If your system has Passed, you can proceed to work through the checks in the *Pre-installation Checklist* that were not covered by the Prerequisite Checker tool.
- If your system has Failed, you must fix the critical checks.
- Review each warning, and where possible fix the issues and rerun the checks for your system. In some cases, you must rerun a check manually. For more information about a check, go to the *Pre-installation Checklist* that is published at this link <http://www.ibm.com/support/docview.wss?uid=swg21427692>.
- Follow the Next Steps that are advised in the results page.

Uninstalling the Prerequisite Checker

You can uninstall the Prerequisite Checker tool independently of any action to the IBM Spectrum Protect Snapshot product.

Procedure

1. Log on with the root user ID.
2. Enter the following command:
`checker_path/uninstall/uninstall.bin [-i console | swing]`

where:

`checker_path` is the path where the Prerequisite Checker was installed.

`-i console` indicates that the Prerequisite Checker is uninstalled using the console version of the uninstaller.

`-i swing` indicates that the Prerequisite Checker is uninstalled using the GUI version of the uninstaller.

If option `-i` is not specified, the same method used for installing the Prerequisite Checker is used for uninstalling the tool.

Results

The Prerequisite Checker executable files are removed from your system.

Preparing Oracle systems (non-SAP, non-ASM)

IBM Spectrum Protect Snapshot processes database files at a volume level on Oracle systems and requires these database files to be on one or more dedicated volume groups.

The Oracle table space database files must be on file systems that are supported by IBM Spectrum Protect Snapshot on your system. However, Oracle active log files must not be in the same volume groups as the table space data files. If Oracle database control files are in the same volume groups as the table space data files, an advanced recovery procedure is required after a restore operation completes. This procedure is described here, *Step required for advanced recovery*.

If other data is stored in the table space data volume groups, it is processed by IBM Spectrum Protect Snapshot and included in the IBM Spectrum Protect Snapshot backup image. During a restore operation, this data is overwritten. Therefore, do not store objects such as database binary files, online redo logs or offline redo logs on the data volume groups.

Extra files that are not associated with the database backup operation but are stored within the same file system, can cause the backup to fail. Use the **NEGATIVE_LIST** parameter in the IBM Spectrum Protect Snapshot profile to control the processing of these files. For information about this parameter, see the profile section “CLIENT” on page 119 for details.

IBM Spectrum Protect Snapshot processes table spaces only. The volume group layout that is required for Oracle systems (non-SAP and non-ASM) is detailed in the *Volume group layout requirements* section of the *Pre-installation Checklist*.

The hardware and software requirements for IBM Spectrum Protect Snapshot for UNIX and Linux are published in the following technote: <http://www.ibm.com/support/docview.wss?uid=swg21427692>. Follow the link to the requirements technote for your specific release or update level. From there you will find the *Pre-installation Checklist* and the *Installation Planning Worksheet* for the most recent version of the product.

IBM Spectrum Protect Snapshot does not support a volume and storage layout where the database is spread across multiple storage devices. In an AIX logical volume manager mirroring environment, each mirror must be located within a separate storage cluster.

IBM Spectrum Protect Snapshot requires that an Oracle recovery catalog database is available when you back up data. This recovery catalog must be present regardless of the type of backup.

If a shell other than ksh or bash is used, Oracle-specific environment variables such as ORACLE_HOME and paths must be accessible when the `su - oracle_user -c` command is entered.

Preparing Oracle Automatic Storage Management

Ensure that you have your environment set up to protect your Oracle ASM data with IBM Spectrum Protect Snapshot.

Before you begin

Redundancy level is defined when the ASM disk group is created, and can be normal, high, or external. For more information about ASM disk group redundancies, see the Oracle product documentation. IBM Spectrum Protect Snapshot supports all three redundancy types.

Procedure

1. Define the storage clusters in the IBM Spectrum Protect Snapshot profile with the **DEVICE_CLASS** profile parameter. The definitions must specify that IBM Spectrum Protect Snapshot selects the cluster for the current operation.
For an ASM disk group with normal redundancy, the backup completes even if one of the failure groups is not on the storage device. For an ASM disk group with high redundancy, the backup completes even if two of the failure groups are not on the storage device.
2. Create the ASM disk groups. If you already created the ASM disk groups, verify that they meet the following criteria.
 - For an ASM backup, create a disk group on the backup system to host an Oracle control file. This disk group has the same name as the disk group where the control file is on the production system. During a backup operation, the Oracle control file is stored in the IBM Spectrum Protect Snapshot repository directory on the production server.
 - For ASM cloning, ensure that the ASM disk groups are on supported storage systems, otherwise the cloning operation fails.
 - Do not run any queries or operations on the Oracle database and ASM instance on the clone system while IBM Spectrum Protect Snapshot clone operations are running.
3. During the offload operation, the control file is transferred into a temporary directory within ACS_DIR on the backup server by IBM Spectrum Protect Snapshot. RMAN restores the control file into the locations that are specified within the database profile from this location as part of the offload operation.
4. IBM Spectrum Protect Snapshot must have an Oracle recovery catalog database available when you back up data. This recovery catalog must be present regardless of the type of backup. Typically one ASM instance can serve multiple databases.
IBM Spectrum Protect Snapshot supports this setup when dedicated disk groups are used for each database. For normal-redundancy disk groups and high-redundancy disk groups, IBM Spectrum Protect Snapshot handles failure groups by allowing backups of disk groups that are on supported storage systems.
5. If needed, you can change the redundancy level and add failure groups after a restore operation. IBM Spectrum Protect Snapshot does not re-create failure groups when a restore operation completes.
6. Export Oracle specific environment variables for the ASM instance and paths, such as *ORACLE_HOME*. Type in one of these commands to verify that the variables are available:

```
su - oracle_user -c
```

For ksh or bash, enter the following command with root user ID `su - oracle_user -c env | grep ORA`

7. Export Oracle Grid specific environment variables for the ASM instance and paths, such as `ORACLE_HOME` and `ORACLE_SID`. Type in one of these commands to verify that the variables are available:

```
su - grid_user -c
```

```
su - grid_user -c env | grep ORA
```

ASM disk group layout

An ASM disk group is a collection of disks that Oracle ASM manages. Because IBM Spectrum Protect Snapshot processes database files at the ASM disk group level, the database files are required to be on dedicated disk groups. Non-application data that is stored on these disk groups is also processed by IBM Spectrum Protect Snapshot and is included in the backup images. This non-application data is overwritten during a restore operation.

For ASM cloning, the control files must be stored in an ASM disk group. If any control file is not on an ASM disk group, the cloning operation fails with an error.

Extra files that are not associated with the database backup operation but are stored within the same file system, can cause the backup to fail. Use the **NEGATIVE_LIST** parameter in the IBM Spectrum Protect Snapshot profile to control the processing of these files. For information about this parameter, see the profile section “CLIENT” on page 119 for details.

IBM Spectrum Protect Snapshot processes table spaces only. The disk group layouts that are required when you are backing up or cloning in an ASM environment are detailed in the *Volume group layout requirements* section of the *Pre-installation Checklist*.

The hardware and software requirements for IBM Spectrum Protect Snapshot for UNIX and Linux are published in the following technote: <http://www.ibm.com/support/docview.wss?uid=swg21427692>. Follow the link to the requirements technote for your specific release or update level. From there you will find the *Pre-installation Checklist* and the *Installation Planning Worksheet* for the most recent version of the product.

ASM in a DS8000 environment

For IBM System Storage DS8000 environments, IBM Spectrum Protect Snapshot does not use FlashCopy consistency groups when a FlashCopy backup is created. All table space files must be on a single disk group. There is no consistency across multiple LUNs on this storage system.

Ensure that all table space files are on a single disk group on DS8000 system. In addition, verify that the disk group contains exactly one LUN on the DS8000. IBM Spectrum Protect Snapshot does not support databases that are distributed across multiple storage clusters. There is one exception to this statement, and that is in environments where failure groups are used in a manner where every storage cluster contains a complete image of the databases.

ASM failure groups

IBM Spectrum Protect Snapshot supports Oracle Automatic Storage Management (ASM) failure groups.

When the ASM database is configured with normal or high redundancy, the disk groups are composed of two or three failure groups. These groups can be on two or three respective storage clusters. IBM Spectrum Protect Snapshot creates a backup entirely within only one of the storage clusters.

If you use ASM failure groups, review these guidelines:

- In setups where all failure groups are on one storage cluster, IBM Spectrum Protect Snapshot backs up all failure groups.
- In setups where all failure groups are on different storage clusters, IBM Spectrum Protect Snapshot backs up all failure groups that are on the storage cluster that is configured for the current operation. For more information, see “*DEVICE_CLASS device*” on page 132.
- IBM Spectrum Protect Snapshot does not split the failure groups on a remote system.
- Failure group support does not require a backup server.
- IBM Spectrum Protect Snapshot does not re-create failure groups when a restore operation is completed. You can reset the redundancy level manually by adding the missing failure groups after the restore operation.

ASM cloning parameters and IBM Spectrum Protect Snapshot

The ASM parameters that point to ASM disk groups on a clone system affect the IBM Spectrum Protect Snapshot in certain scenarios. Review the list of ASM parameters when you are preparing IBM Spectrum Protect Snapshot for Oracle ASM.

IBM Spectrum Protect Snapshot parameters

ENHANCED_PARTITIONING

When **VOLUME_MGR** is set to *ASM*, the **ENHANCED_PARTITIONING** parameter is not evaluated by IBM Spectrum Protect Snapshot, and the default setting applies.

FLASH_DIR_LIST

For ASM cloning, ensure that this parameter is not configured. If this parameter is configured in the cloning section of the profile file when **VOLUME_MGR** is set to *ASM*, IBM Spectrum Protect Snapshot stops and sends an error message.

VOLUME_MGR

For the cloning section of the IBM Spectrum Protect Snapshot profile, this parameter must be set to *ASM* to support ASM cloning.

Oracle parameters

DB_CREATE_FILE_DEST

The parameter **DB_CREATE_FILE_DEST** setting on your production system is also used on the clone system when the parameter **OVERWRITE_DATABASE_PARAMETER_FILE** is set to *YES*. When you create new data files, the specified location must be available on the clone system.

The following situations are possible:

DB_CREATE_FILE_DEST points to the ASM disk group that is part of the snapshot. No action is required.

DB_CREATE_FILE_DEST points to the ASM disk group that is not part of the snapshot. The ASM disk groups must be created on the clone server before the clone operation.

DB_CREATE_FILE_DEST points to a directory on a file system. No action is required. The directory is created on the clone system.

DB_CREATE_ONLINE_LOG_DEST_n

The parameter **DB_CREATE_ONLINE_LOG_DEST_n** setting on your production system is also used on the clone system when the parameter **OVERWRITE_DATABASE_PARAMETER_FILE** is set to YES. When you create new log files, the specified location must be available on the clone system. The same situations as specified for parameter **DB_CREATE_FILE_DEST** are possible.

LOG_ARCHIVE_DEST

The log archive destination that is specified with **LOG_ARCHIVE_DEST** and **LOG_ARCHIVE_DEST_n** can point to regular file systems or ASM disk groups. IBM Spectrum Protect Snapshot turns off archive logging on the clone database as part of the clone operation. Archive logging can be activated manually after the clone operation finishes, or through postprocessing scripts.

CONTROL_FILES

All database control files for ASM cloning must be on ASM disk groups. The LUNs of these disk groups are added to the snapshot backup.

Preparing Oracle RAC

Before you install IBM Spectrum Protect Snapshot in an Oracle RAC environment, you must set up a shared file system. With a shared file system, you can share the IBM Spectrum Protect Snapshot repository, configuration, and binary files between all RAC nodes. IBM Spectrum Protect Snapshot can be used to work on clones that were created on any of the RAC nodes. A clone that is created on one RAC node can be inquired, refreshed, and deleted from every other node within the RAC cluster in one operation.

Before you begin

Before you proceed to install IBM Spectrum Protect Snapshot in an Oracle RAC environment, ensure that all the Prerequisite checks for ASM cloning pass successfully. ASM disk layout checks, and Operation System user and group requirements must pass before you proceed.

The hardware and software requirements for IBM Spectrum Protect Snapshot for UNIX and Linux are published in the following technote: <http://www.ibm.com/support/docview.wss?uid=swg21427692>. Follow the link to the requirements technote for your specific release or update level. From there you will find the *Pre-installation Checklist* and the *Installation Planning Worksheet* for the most recent version of the product.

Procedure

Create a shared file system for the IBM Spectrum Protect Snapshot repository, configuration, and binary files such as ACFS or NFS.

Related tasks:

“Cloning Oracle RAC databases” on page 96

Chapter 4. Preparing backup and cloning servers

Backup servers and clone servers are auxiliary hosts where IBM Spectrum Protect Snapshot can mount backups and clones.

A backup server or clone server is used to offload the workload of sending data to the server from the production server where the protected application is running. You must configure a backup or clone server when you want to offload snapshots to IBM Spectrum Protect. A clone server creates a clone of the productive database from a snapshot backup. You can share one backup or clone server among multiple applications or you can have multiple backup or clone servers. A backup server can also serve as a clone server. However, IBM Spectrum Protect Snapshot does not allow backup images to be mounted directly on the production server. A backup or clone server must be set up as a separate host.

A backup server is used to offload the workload of sending data to the server from the production server where the protected application is running. With the exception of environments on GPFS file systems, you must configure a backup server when you want to offload snapshots to IBM Spectrum Protect. A clone server creates a clone of the productive database from a snapshot backup. You can share one backup or clone server among multiple applications or you can have multiple backup or clone servers. A backup server can also serve as a clone server. However, IBM Spectrum Protect Snapshot does not allow backup images to be mounted directly on the production server. A backup or clone server must be set up as a separate host. In combination with GPFS filesystems, all IBM Spectrum Protect Snapshot actions take place in the productive GPFS cluster; a backup server is not required.

In combination with GPFS filesystems, all IBM Spectrum Protect Snapshot actions take place in the productive GPFS cluster; a backup server is not required.

Preparing backup and cloning servers

Backup servers and clone servers are auxiliary hosts where IBM Spectrum Protect Snapshot can mount backups and clones.

A backup server or clone server is used to offload the workload of sending data to the server from the production server where the protected application is running. You must configure a backup or clone server when you want to offload snapshots to IBM Spectrum Protect. A clone server creates a clone of the productive database from a snapshot backup. You can share one backup or clone server among multiple applications or you can have multiple backup or clone servers. A backup server can also serve as a clone server. However, IBM Spectrum Protect Snapshot does not allow backup images to be mounted directly on the production server. A backup or clone server must be set up as a separate host.

A backup server is used to offload the workload of sending data to the server from the production server where the protected application is running. With the exception of environments on GPFS file systems, you must configure a backup server when you want to offload snapshots to IBM Spectrum Protect. A clone server creates a clone of the productive database from a snapshot backup. You can share one backup or clone server among multiple applications or you can have multiple backup or clone servers. A backup server can also serve as a clone server.

However, IBM Spectrum Protect Snapshot does not allow backup images to be mounted directly on the production server. A backup or clone server must be set up as a separate host. In combination with GPFS filesystems, all IBM Spectrum Protect Snapshot actions take place in the productive GPFS cluster; a backup server is not required.

In combination with GPFS filesystems, all IBM Spectrum Protect Snapshot actions take place in the productive GPFS cluster; a backup server is not required.

Installation prerequisites for backup and clone servers

For hosts that are used as a backup server or clone server, the operating system version and maintenance level must be the same as the production server.

Backup server requirements

To run the software, the following settings are required on the backup server:

- The user name and group name of the database instance owner on the production server must be available on the backup server. The same user ID (UID) and group ID (GID) must be used.
- A database instance with the same version as the database instance on the production server must be installed on the backup server.

When IBM Spectrum Protect Snapshot is used in an environment with IBM Spectrum Protect, a backup server is required. This backup server is used to offload the backup workload from the production server to the backup server and sends the application critical backups to IBM Spectrum Protect.

The IBM Spectrum Protect for Databases agent for Oracle client is used by IBM Spectrum Protect Snapshot to initiate a subsequent backup to IBM Spectrum Protect, and must be installed and configured on the backup server.

IBM Spectrum Protect backup-archive client is used by IBM Spectrum Protect Snapshot to initiate a subsequent backup to IBM Spectrum Protect and must be installed and configured on the backup server.

Update the IBM Spectrum Protect Data Protection client node password on the production server and all backup servers whenever it changes. When IBM Spectrum Protect is configured to use the **PASSWORDACCESS GENERATE** parameter, the password can change without notification. If the IBM Spectrum Protect Data Protection client is configured to use the **PASSWORDACCESS GENERATE** parameter, use the IBM Spectrum Protect proxy-node capability to avoid authentication errors when the password is reset. Create one data node on the IBM Spectrum Protect where all Data Protection clients from all backup and production servers are sending and retrieving data. Create one authentication node for each production server and backup server that is configured as proxy node to this data node.

Offload backups to IBM Spectrum Protect with Oracle RMAN

To offload incremental backups from an Oracle environment, IBM Spectrum Protect Snapshot requires an Oracle database instance to be installed and configured on the backup server. The Oracle specific environment variables, for example **ORACLE_HOME**, and paths must be exported so that they are accessible if the `su - oracle_user -c` command is entered. This accessibility can be verified by running `su - oracle_user -c env | grep ORACLE` as root user. The Oracle recovery catalog

database must exist and must be accessible from the production and backup server for the Oracle user ID. In this database, Oracle RMAN records all offloaded backups. For details on the setup of a recovery catalog database, see the Oracle manuals. To verify this setup, run the following command as the Oracle user on the production host:

```
rman target / catalog catalog_user/catalog_password@
<catalog_connect_string>
```

To verify the setup of the backup system, run the following command as the root user on the backup host:

```
su -oracle user -c "rman target / catalog catalog_user/
catalog_password@catalog_connect_string"
```

The command must be able to connect to both the target and the recovery catalog databases, and show the RMAN prompt. This operation can be finished with the quit command. For example:

```
$ rman target / catalog rman/rman@catdb
,
Recovery Manager: Release 11.2.0.3.0 - Production on Mon Aug 5 13:59:53 2013

Copyright (c) 1982, 2011, Oracle and/or its affiliates. All rights reserved.

connected to target database: P01 (DBID=1213110920, not open)
connected to recovery catalog database

RMAN> quit
```

```
Recovery Manager complete.
$
```

Clone server requirements

For FlashCopy cloning, the database instances must be installed and configured on the clone server.

IBM Spectrum Protect Snapshot requires the following settings for cloning:

Oracle and Oracle in an SAP environment

A clone database instance with the same version and name as the database instance on the production server must be installed on the clone server.

The user name and group name of the clone database instance owner on the clone system must be available on the clone server. The user ID (UID) and group ID (GID) can be different from the UID and GID of the production server database instance.

For cloning in Oracle environments, the database SID cannot contain lowercase letters. If the SID contains lowercase letters, the clone database is created, but the lowercase letters in the SID of the clone system are converted to capitalized letters.

Preparing backup and cloning servers for VMware or KVM

If a backup or clone server you are using is a VMware or KVM virtual machine, the storage device must be attached to the virtual machine with either iSCSI or Network file system.

Before you begin

If physical hosts are used as backup or clone servers, see “Installation prerequisites for backup and clone servers” on page 38. These requirements are also required for backup servers on virtual machines.

Procedure

- Verify that all target LUNs in backup or clone operations are attached to the virtual machine with one of the following attachment methods:
 - iSCSI
 - Network file system (NFS)
- Verify that all target LUNs in backup operations are attached to the virtual machine with one of the following attachment methods:
 - iSCSI
 - Network file system (NFS)

Chapter 5. Installing and upgrading

To install IBM Spectrum Protect Snapshot you must follow the installation steps, run the setup script for your component, activate the applications you want to protect, and configure the product. The first step is to install IBM Spectrum Protect Snapshot on the production server. Depending on your environment, a separate installation of IBM Spectrum Protect Snapshot can be required on a backup or clone server. If you choose to, you can upgrade your system from a previous version of IBM Spectrum Protect Snapshot to version 8.1.0.

About this task

When you are installing IBM Spectrum Protect Snapshot software, the installation process varies, depending on the environment.

The following set of tasks are required to complete the installation process.

Procedure

- Install IBM Spectrum Protect Snapshot on the production server.
The production server is where IBM Spectrum Protect Snapshot protects critical business applications by providing a method to back up and restore these applications.
- Activate the applications that you want to protect with IBM Spectrum Protect Snapshot.
During the activation, all the necessary files are copied from the installation directory `FCM_INSTALL_DIR`, to the application-specific installation directory `INSTANCE_DIR`. The installation directory is referred to as the `FCM_INSTALL_DIR` directory, and the application-specific installation directory is referred to as `INSTANCE_DIR` directory.
- Configure IBM Spectrum Protect Snapshot.
The following files and directories are created during the configuration process:
 - An `ACS_DIR` configuration directory, if the `ACS_DIR` directory is not identical to the `INSTANCE_DIR` directory. The path for the `ACS_DIR` directory is specified in the IBM Spectrum Protect Snapshot profile file.
 - A profile file within the `ACS_DIR` configuration directory.
 - A symbolic link is created from the `INSTANCE_DIR/profile` file that points to the `ACS_DIR/profile` file when the two directories are not identical.
 - A password file within `ACS_DIR/shared` directory.
 - An entry `/etc/inittab` for daemon processes if requested.
For Red Hat Enterprise Linux 6, the daemon processes are started automatically by using the `upstart` program when requested.
- Install IBM Spectrum Protect Snapshot on a backup or clone server, if not automatically installed and configured. Backup servers or clone servers are auxiliary hosts that are required by IBM Spectrum Protect Snapshot to mount backup images and clone databases. A backup or clone server also is required to offload backups to IBM Spectrum Protect.

If Open Secure Shell (OpenSSH) is configured between the production and the backup or clone servers, IBM Spectrum Protect Snapshot is installed and configured automatically. Otherwise, a separate installation on a backup or clone server is required.

Installing on the production server

To install IBM Spectrum Protect Snapshot on the production server, you can use the graphical installation wizard, the console wizard, or the console in silent mode.

Before you begin

For the current requirements, review the *Hardware and Software Requirements* technote that is associated with the IBM Spectrum Protect Snapshot release. This technote is available in the *IBM Spectrum Protect Snapshot - All Requirement Documents* website at: <https://www.ibm.com/support/docview.wss?uid=swg21427692>. Follow the link to the requirements technote for your specific release or update level and review the pre-installation checklist and planning worksheet.

IBM Spectrum Protect Snapshot installation packages are delivered as individual files. They are provided on an installation DVD or from an image that is downloaded from IBM Passport Advantage®.

The files for *OS-platform* AIX and Linux are named:

```
<VERSION>-TIV-TSFCM-OS-platform.bin
```

The files for *OS-platform* Solaris and HP-UX are named:

```
<VERSION>-TIV-TSFCM-OS-platform.bin
```

Before you install IBM Spectrum Protect Snapshot on AIX or Linux, run the Prerequisite Checker to ensure that you have the prerequisites to proceed with the installation process.

Procedure

To install IBM Spectrum Protect Snapshot on the production server, complete the following steps.

1. Log on to the production server and use the root user ID. Change to the directory where you downloaded the package file or insert the DVD into the DVD drive. Use one of the following methods to start the installation:

Graphical user interface with the installation wizard

The installation wizard requires a graphical X Window System installation. Make sure the environment variable *DISPLAY* specifies *host:display*, where *host* identifies the host name of the X Server to be contacted and *display* is the display number. To use the graphical installation wizard, enter this command for AIX and Linux:

```
./<VersionAIXLinux>-TIV-TSFCM-OS-platform.bin
```

Enter this command for Solaris and HP-UX:

```
./<VersionSolarisHP>-TIV-TSFCM-OS-platform.bin
```

If the graphical X Window System is not present, the installation continues in console mode.

Console mode

To install in console mode, enter the following command for AIX or Linux:

```
./<VersionAIXLinux>-TIV-TSFCM-OS-platform.bin -i console
```

Enter this command for Solaris and HPUX:

```
./<VersionSolarisHP>-TIV-TSFCM-OS-platform.bin
```

2. Follow the prompts to install IBM Spectrum Protect Snapshot.
Activate the database or database instances with the setup script after the installation completes.
3. On the Summary page, review your installation settings. If an error occurs during the installation process, correct the errors and restart the installation procedure. You can find an `installation.log` file in the `FCM_INSTALL_DIR` directory to help you to troubleshoot any installation errors.

What to do next

After the installation, you must activate the database instance from the install directory of the new installed version, and configure the database instances to complete the installation.

Related concepts:

“Configuring storage environments” on page 55

Related tasks:

“Running the setup script for Oracle” on page 49

“Configuring IBM Spectrum Protect Snapshot for Oracle” on page 50

Adding or upgrading a new instance ID after installation

If you want to add or upgrade an instance ID after the global installation, you must configure the database instances to complete the action. During the installation with the installer, you input the instances to be activated and the directory for those instances. The installer does not activate the instance, you must run the activation command for the database instance from the install directory of the new version. The correct access rights for the directories are assigned

Procedure

1. Log in to the production server and use the root user ID. Change to the `FCM_INSTALL_DIR` directory.
2. Run the following command to activate the database:
 - Non-RAC Oracle and Oracle in an SAP environment:

```
./setup_ora.sh -a install -d Oracle_instance_owner_$HOME_directory
```
 - RAC Oracle:
To enable an instance for Oracle RAC with a shared file system, run the following command,

```
./setup_ora.sh -a install -d Oracle_instance_owner_$HOME_directory -t directory_in_shared_file_system
```

The `-t` parameter for Oracle RAC specifies a target directory within a shared file system, so that the IBM Spectrum Protect Snapshot binary files are copied to the shared file system. These files are then available on all nodes of the Oracle RAC cluster.

If the home directory of the database instance owner is not identical to the database instance directory, install the product in the database instance

directory. For example, `$ORACLE_HOME`. For installations where `$ORACLE_HOME` is shared between multiple database instances, any other directory that is unique to this instance can be used.

Installing separately on backup or clone servers

If IBM Spectrum Protect Snapshot is not installed remotely on the backup or clone server by using OpenSSH, use the following instructions to install the product on a backup or clone server.

Before you begin

For information about how to prepare a backup or clone server, see “Installation prerequisites for backup and clone servers” on page 38. If a setup using OpenSSH is not possible, install IBM Spectrum Protect Snapshot software separately on the backup or clone server. Information about when a backup server or clone server is needed is available at Chapter 4, “Preparing backup and cloning servers,” on page 37.

Tip: If you are using self-signed certificates, you must manually copy the `.arm` file from the production server to the auxiliary server `INSTANCE_DIR`. You must also copy the contents in `../acs/shared` must be copied to the backup system.

Procedure

To install IBM Spectrum Protect Snapshot on the backup or clone server, complete the following steps:

1. Log on to the production server and use the root user ID. Change to the directory where you downloaded the package file or insert the DVD into the DVD drive. Use one of the following methods to start the installation:

Graphical user interface with the installation wizard

The installation wizard requires a graphical X Window System installation. Make sure the environment variable `DISPLAY` specifies `host:display`, where `host` identifies the host name of the X Server to be contacted and `display` is the display number. To use the graphical installation wizard, enter this command for AIX and Linux:

```
./<VersionAIXLinux>-TIV-TSFCM-OS-platform.bin
```

Enter this command for Solaris and HP-UX:

```
./<VersionSolarisHP>-TIV-TSFCM-OS-platform.bin
```

If the graphical X Window System is not present, the installation continues in console mode.

Console mode

To install in console mode, enter the following command for AIX or Linux:

```
./<VersionAIXLinux>-TIV-TSFCM-OS-platform.bin -i console
```

Enter this command for Solaris and HP-UX:

```
./<VersionSolarisHP>-TIV-TSFCM-OS-platform.bin
```

2. Follow the prompts to install IBM Spectrum Protect Snapshot.
Activate the database or database instances with the setup script after the installation completes.

3. On the Summary page, review your installation settings. If an error occurs during the installation process, correct the errors and restart the installation procedure. You can find an `installation.log` file in the `FCM_INSTALL_DIR` directory to help you to troubleshoot any installation errors.

What to do next

After the installation, you must activate the database instance from the install directory of the new installed version, and configure the database instances to complete the installation.

Preparing a database or database instance for configuration

Before you configure the IBM Spectrum Protect Snapshot instance on the backup or clone system, you must prepare the instance as the application instance owner of the backup or clone system.

Procedure

1. As the application instance owner of the backup or clone system, copy the `fcmselcert.arm` file from the production server to the backup or clone server `INSTANCE_DIR` directory.

```
cd Oracle_instance_owner_${HOME_directory}/acs
scp Oracle_instance_owner@production_system:$PWD/fcmselcert.arm
```
2. Copy the password file from the production system to the backup or clone system. Paste the file into the `$/HOME/acs` of the application instance owner. If this directory does not exist, create it with the following command.

```
mkdir -p $HOME/acs/shared
cd $HOME/acs/shared
```

3. Run the setup script as the Application instance owner from the `INSTANCE_DIR` directory. Running the script from this directory, configures the IBM Spectrum Protect Snapshot instance.

```
cd Oracle_instance_owner_${HOME_directory}/acs
./setup_ora.sh
```

What to do next

You must configure the instance on the clone or backup server, “Configuring a database or database instance.”

Configuring a database or database instance

Before you begin

Run the setup script as the owner in the `INSTANCE-DIR` directory,
`./setup_ora.sh`.

About this task

When you have installed the product on the backup or clone server, and prepared for the configuration by activating the database or database instance, you are ready to configure the database.

Procedure

1. Choose a configuration type.

- Onsite Production System configuration, with optional remote backup system configuration.
 - Onsite Backup System configuration, to configure an onsite backup system configuration. Provide configuration parameters as required
2. Specify the hostname of the production system, and the port that is configured on the production system for IBM Spectrum Protect Snapshot communication. If the default port 57328 is used, then you do not have to specify it and it can be left blank.


```
***** Profile parameters for section GLOBAL :
***** Hostname and port of machine running Management Agent {ACSD}
(<hostname> <port>) = [] utprod2 57328
```
 3. Choose to configure the passwords, and enter the device class names that are used for this backup or clone system. The configuration completes with the installation and starting of the daemons.

Installing in silent mode

To install IBM Spectrum Protect Snapshot in silent mode you require a response or properties file.

About this task

You can generate a properties file during installation in either graphic or console mode by starting the executable file as follows:

```
./8.1.0-TIV-TSFCM-platform.bin [-i console]
-DRECORDFILE=/tmp/installer.properties
```

Procedure

1. To install in silent mode, set the variable for the license file
`LICENSE_ACCEPTED=TRUE`
2. Invoke the executable file with the `-i silent` option and the `-f` option to specify the properties file:
`./version-TIV-TSFCM-OS-platform.bin -i silent -f properties_file`

The *properties_file* specification must contain a full path.

3. Activate the instance with the following command, `./setup_ora.sh -a install -d /oracle/ACB/./setup_orc.sh -a install -d /orc/ACB/sqllib`

Upgrading

To upgrade to a newer version of IBM Spectrum Protect Snapshot, you must follow three steps. These steps are to install the new version, activate your application instances with the new version, and run the setup script. You can then proceed to uninstall the old version.

Procedure

1. Install the new version of the product as described here: “Installing on the production server” on page 42
2. After the product is installed successfully, the application-specific instances must be activated with the new version. “Adding or upgrading a new instance ID after installation” on page 43
3. Run the setup script from within each activated application instance, and choose the option to modify the profile. Step through the parameters in the

wizard. Upgrade the product on your backup or clone system by selecting it and choosing the option to **update IBM Spectrum Protect Snapshot installation**. This updates the profile with new parameters and removes deprecated parameters, or renames them if required. Upgrade to the new version on your backup or clone system by selecting it from the wizard and choosing the option to **update IBM Spectrum Protect Snapshot installation**. Follow the instructions to run the setup script as described here, “Configuring IBM Spectrum Protect Snapshot for Oracle” on page 50

4. Uninstall the older version of the product. “Uninstalling the software”

Uninstalling the software

When you are upgrading the product, you can complete the process by uninstalling the older version of the product to finalize the upgrade steps.

Procedure

1. Determine the installation path of the version of the product you want to uninstall. The following paths provide the default location of the installation files:
 - For AIX operating systems, it is this path, `/usr/tivoli/tsfcm/acs_version`.
 - For Linux operating systems, it is this path, `/opt/tivoli/tsfcm/acs_version`.
 - For Solaris and HP-UX operating systems, it is this path, `/opt/tivoli/tsfcm/acs_version`.
2. Run the appropriate command for your operating system from the installation path:
 - For AIX operating systems, use this command `/usr/tivoli/tsfcm/acs_version_number/uninstall/uninstaller.bin`.
 - For Linux, Solaris, and HP-UX operating systems, use this command `/opt/tivoli/tsfcm/acs_version_number/uninstall/uninstaller.bin`.

Migrating existing snapshot data

You can upgrade to IBM Spectrum Protect Snapshot and migrate data from IBM Spectrum Protect for Advanced Copy Services.

Editing `USE_CONSISTENCY_GROUP` before you upgrade from IBM Spectrum Protect Snapshot version 3.1, or earlier

If you are upgrading from IBM Spectrum Protect Snapshot, Version 3.1 you must set the `USE_CONSISTENCY_GROUP` parameter to `N0`, for version 3.2 or later of IBM Spectrum Protect Snapshot to work. IBM Spectrum Protect Snapshot Version 3.2 and later software requires the use of consistency groups.

About this task

Log in to the production server with the database user ID and go to the `INSTANCE_DIR` directory.

Procedure

1. Start the setup script by entering the following command:
 - Oracle in an SAP environment
`./setup_ora.sh`
 - Oracle

```
./setup_ora.sh
```

2. Follow the setup script instructions that are displayed. For each IBM Spectrum Protect Snapshot, Version 3.1 profile configuration that has the **USE_CONSISTENCY_GROUP** parameter, repeat these steps to automatically remove the **USE_CONSISTENCY_GROUP** parameter.

Chapter 6. Configuring IBM Spectrum Protect Snapshot

After the installation and activation procedure is complete, configure IBM Spectrum Protect Snapshot. To configure IBM Spectrum Protect Snapshot, use the setup script for your environment. The information that you enter is used to create the profile configuration file.

Before you begin

Review the installation planning sheet that is associated with the *Hardware and Software Requirements* technote. This sheet contains the required parameters for each specific software application and custom application that are required during the configuration.

For the current requirements, review the *Hardware and Software Requirements* technote that is associated with the IBM Spectrum Protect Snapshot release. This technote is available in the *IBM Spectrum Protect Snapshot - All Requirement Documents* website at: <http://www.ibm.com/support/docview.wss?uid=swg21427692>. Follow the link to the requirements technote for your specific release or update level. Use the *Pre-installation checklist*, and *Installation Planning worksheet* before you install IBM Spectrum Protect Snapshot.

About this task

When you configure IBM Spectrum Protect Snapshot, you are prompted to enter parameter values that are specific to your environment. Syntax and value ranges are checked during the setup. Also, you must enter password information that is used to create a password file. A separate IBM Spectrum Protect Snapshot profile is created for each application.

Running the setup script for Oracle

Run the Oracle setup script to configure IBM Spectrum Protect Snapshot for Oracle, or Oracle in an SAP environment.

Before you begin

Review the completed IBM Spectrum Protect Snapshot installation sheet to ensure that the product installed correctly.

In most cases, configure IBM Spectrum Protect Snapshot in basic mode. To display help for the parameters, enter the ? character. The help is best viewed in a window that is set for at least 130 characters. If you choose to configure IBM Spectrum Protect Snapshot in advanced mode, -advanced option, you can configure all parameters even ones that have default values. For this reason, the advanced mode takes longer to process.

Procedure

1. From the production database instance, log on as the database instance owner.
2. Go to your Oracle or Oracle in an SAP environment database, instance-specific installation directory:

```
INSTANCE_DIR: <Oracle_instance_owner_$HOME>/acs/
```

3. Start the setup script by entering the following command:
`./setup_ora.sh`
4. Follow the setup script instructions. For information about the configuration steps, see “Configuring IBM Spectrum Protect Snapshot for Oracle”

Results

The setup script creates the following directories on the instance directories:

- The `$HOME/acs` directory contains the IBM Spectrum Protect Snapshot binary files.
- The `ACS_DIR` directory is the IBM Spectrum Protect Snapshot configuration directory. It contains the following files and directories:
 - The profile configuration file.
 - The IBM Spectrum Protect Snapshot repository.
 - The logs directory. All newly started daemons and active daemons processes are recorded in the summary log file.
 - The configuration wizard registers the IBM Spectrum Protect Snapshot management daemon `acsd` and generic device agent `acsgen` in the `/etc/inittab` or creates and starts upstart jobs on the production server. These processes are started automatically even after a system restart.

If you are installing IBM Spectrum Protect Snapshot in an Oracle RAC environment, the setup script starts the daemons on the local node where the setup script is running. To ensure that the daemons are available on all other RAC nodes, you must add them to `/etc/inittab`, or create the upstart jobs manually on all other RAC nodes. A locking mechanism using the shared file system ensures that only one daemon is active at a time.

Configuring IBM Spectrum Protect Snapshot for Oracle

After you run the setup script, the configuration wizard leads you through the configuration of the IBM Spectrum Protect Snapshot for Oracle.

Before you begin

To start the configuration process, run the setup script for Oracle with the following command: `./setup_ora.sh`

If you are using the setup script to configure an onsite backup server and you do not use standard CA-signed certificates for server authentication, you must copy `fcmselcert.arm` from `INSTALL_DIR` on the production server to `INSTALL_DIR` on your backup or clone server. For information about IBM Global Security Kit configuration, see “IBM Global Security Kit configuration” on page 160.

About this task

For some parameters, you can create multiple entries with different values. To create these multiple entries, when prompted **Do you want to add another instance of this parameter?**, enter `y`. To delete a parameter entry, when prompted for the parameter value, enter `!d`.

Procedure

1. Choose one of the following options:

- (1) On-site Production Server configuration with optional remote Backup Server configuration.

This selection guides you through the configuration of IBM Spectrum Protect Snapshot on the production server. It also provides the option to remotely activate and synchronize the configuration of one or more backup servers by using the OpenSSH protocol.

- (2) Do you want to configure IBM Spectrum Protect Snapshot for Oracle RAC? [y/n]

When you choose to configure for Oracle RAC, IBM Spectrum Protect Snapshot determines the Oracle RAC node names. Each node is made available to the mount agent that is running on the backup or clone server. The mount agent then attempts to connect to the management daemon for each node name.

- (3) On-site Backup Server configuration.

This selection guides you through the configuration of IBM Spectrum Protect Snapshot on the backup server as a separate installation.

2. Select one of these configurations for the database instance:

- (1) Backup only
- (2) Cloning only
- (3) Backup and cloning

Note: In an environment, where cloning is not supported by IBM Spectrum Protect Snapshot choose (1) Backup only.

3. Choose if you are going to run offload backups.

Are you going to perform offloaded backups to IBM Spectrum Protect?
[Y|N]

- Specify yes to configure support for offloaded tape backups.

Note: For Oracle in an SAP environment, you must manually update the IBM Spectrum Protect for Enterprise Resource Planning profile `initSID.utl` file after the configuration completes.

- Specify no to configure support for disk-based snapshot backups only.

4. Choose if you want to start offloaded tape backups after the snapshot.

Do you want to start offloaded tape backups after the snapshot? [Y/N]

- Choose yes to start the offload immediately after the FlashCopy backup completes.
- Choose no if you want to schedule the offload operation to run later by scheduling backups individually. The backup to IBM Spectrum Protect can be delayed until the necessary resources in IBM Spectrum Protect server are available. This answer requires the scheduled backup process to be started manually. For example, add a crontab entry. The default value is to run `tsm4acs` as a daemon process on the production server.

5. Choose one of the following options:

- (Oracle and Oracle in an SAP environment) Linux: Do you want IBM Spectrum Protect Snapshot to create and start the upstart jobs for you? [y|n]
- (Oracle and Oracle in an SAP environment) AIX, Solaris, and HP-UX: Do you want IBM Spectrum Protect Snapshot to create the inittab entries for you? [Y/N]

Specify NO for the executable files that include command line options not to be added to the `/etc/inittab` and not to create upstart jobs. You must

make sure that they are started by your HA start-up scripts, and that they are restarted whenever they are ended.

Specify YES to enter the daemon processes in the /etc/inittab or to create and start upstart jobs.

- (Oracle RAC nodes) AIX: To use IBM Spectrum Protect Snapshot on non-local Oracle RAC nodes, create a link from Oracle_instance_owner_\$HOME_directory/acs to the target directory in the shared file system you specified during the installation or configuration. At least two daemon processes must be running at each point in time. IBM Spectrum Protect Snapshot creates and starts /etc/inittab entries for the necessary daemon processes on the current host. You must manually create /etc/inittab entries for the necessary daemon processes on all other Oracle RAC nodes. Manually stop and start these entries before and after running the setup script. Review the chapter "Configuring Oracle RAC for daemon high availability" in the user's guide for a detailed explanation. Press [Enter] key to continue.
- (Oracle RAC) Linux: To use IBM Spectrum Protect Snapshot on non-local Oracle RAC nodes, create a link from Oracle_instance_owner_\$HOME_directory/acs to the target directory in the shared file system you specified during the installation or configuration. At least two daemon processes must be running at each point in time. IBM Spectrum Protect Snapshot creates and starts upstart jobs for the necessary daemon processes on the current host. You must manually create upstart jobs for the necessary daemon processes on all other Oracle RAC nodes. Manually stop and start these entries before and after running the setup script. Review the chapter "Configuring Oracle RAC for daemon high availability" in the user's guide for a detailed explanation. Press [Enter] key to continue.

Important: After this procedure completes, you are prompted whether to deploy the configuration to one or multiple backup or clone systems. This deployment associates the device classes that are specified in the profile with the backup or clone systems. The following section describes the configuration of a backup system. When you configure a clone system, similar options are displayed.

6. Select the backup system to update or delete:

- n) To configure a new backup system
- b) Return to the previous menu
- q) To quit the configuration

IBM Spectrum Protect Snapshot requires a backup server to be available when the following conditions exist:

- Offload backups to IBM Spectrum Protect are run.
- FlashCopy backup consistency must be verified during a forced mount operation.

Select n to configure and activate IBM Spectrum Protect Snapshot on a remote site by using OpenSSH. OpenSSH must already be available for remote connections from the production system to the backup system. You are prompted to specify the **DEVICE_CLASS** to be enabled on the backup system. Select one or more **DEVICE_CLASS** parameters from the list that is displayed on the console.

Enter q to quit the configuration of the backup system and exit the setup script if one of the following conditions exist:

- OpenSSH is not available.
- You want to configure the backup system in a separate step.

When a backup system is configured, it is possible to run several actions on this backup system. For example, update, stop, start, delete IBM Spectrum Protect Snapshot agents that are running on the backup system or you can set up SSH key authentication to the backup system.

The following example illustrates these actions.

Select the backup system to update or delete:

```
1) acsback1
2) acsback2
3) acsback5

n) to configure a new backup system
q) to quit configuration
1
selected backup system: acsback1
```

The backup system on acsback1 is configured with the device class DISK_ONLY3. Select the action you want to take on the backup system acsback1:

```
1) update IBM Spectrum
Protect Snapshot installation
2) start IBM Spectrum
Protect Snapshot services
3) stop IBM Spectrum
Protect Snapshot services
4) uninstall IBM Spectrum
Protect Snapshot
5) setup the SSH key authentication

b) return to the backup system selection
q) quit the configuration
```

Select one of the options.

The same set of functions is provided for the configuration of the clone instances with SSH.

Configuring Oracle RAC for daemon high availability

You can configure IBM Spectrum Protect Snapshot in an Oracle RAC environment so that the IBM Spectrum Protect Snapshot daemons are highly available. When you choose this configuration, IBM Spectrum Protect Snapshot remains operable even if the RAC node that is hosting the daemons fails. You must stop and restart the daemons as part of this procedure.

Before you begin

Ensure that a shared file system exists for the IBM Spectrum Protect Snapshot repository, configuration, and binary files. For more information about this prerequisite, see “Preparing Oracle RAC” on page 35.

About this task

When you are running the setup script for configuring IBM Spectrum Protect Snapshot for Oracle, you can follow these steps to make the daemons highly available.

Procedure

1. When you are configuring IBM Spectrum Protect Snapshot for Oracle in a RAC environment, choose yes to the question Do you want to configure IBM Tivoli Storage FlashCopy(R) Manager for Oracle RAC? [y|n] .

If you are installing IBM Spectrum Protect Snapshot in an Oracle RAC environment, the setup script starts the daemons on the local node where the setup script is running. To ensure that the daemons are available on all other RAC nodes, you must add them to `/etc/inittab`, or create the upstart jobs manually on all other RAC nodes. A locking mechanism using the shared file system ensures that only one daemon is active at a time.

The setup script creates a link from `Oracle_instance_owner_${HOME}_directory/acs` to the target directory in the shared file system on the local node where the setup script is running. You must manually create that link on all other RAC nodes to operate IBM Spectrum Protect Snapshot from other nodes.

2. To start or stop the daemons from the setup script, log on as the RAC database instance owner on each production database.

The setup script cannot stop and start the daemons automatically on all RAC nodes when it is used to change the profile for example. All daemons must be stopped on all non-local RAC nodes before you start the setup script. You must also ensure to restart the daemons on all non-local RAC nodes when the setup script is finished.

3. Go to the Oracle installation directory, as follows:

```
INSTANCE_DIR: Oracle_instance_owner_${HOME}_directory
```

4. Issue one of the following commands to stop or start the daemons:

- `./setup_ora.sh -a stop -d Oracle_instance_owner_${HOME}_directory`
- `./setup_ora.sh -a start -d Oracle_instance_owner_${HOME}_directory`

Activating a database

The installer can activate the database or database instance on the backup or clone server.

About this task

If you were prompted to activate the database or database instance during the installation already, then you can skip the activation step and continue to configure the instance.

Procedure

1. Change to the root User ID, and change to the `FCM_INSTALL_DIR` directory with the following command

```
cd /opt/tivoli/tsfcm/acs_<versionAIXLinux>. For example, on a Linux system, this path is /opt/tivoli/tsfcm/acs_<versionAIXLinux>. For AIX the path is /usr/tivoli/tsfcm/acs_<versionAIXLinux>
```

2. Run the setup script with the option `-a install`, and `-d`, as follows.

```
./setup_ora.sh -a install -d Oracle_instance_owner_${HOME}_directory/
```

This command copies the binary files into the `INSTANCE_DIR` directory.

What to do next

Configure the database instance on the backup or clone server, "Preparing a database or database instance for configuration" on page 45.

Configuring storage environments

You must configure all storage devices that are storing backups from IBM Spectrum Protect Snapshot, but IBM System Storage DS8000 storage devices require more configuration to prepare for source and target volume relationships. Similarly, IBM System Storage SAN Volume Controller and IBM Storwize family must be configured when you use predefined target volumes.

The IBM Spectrum Protect Snapshot profile configuration file can contain one or more **DEVICE_CLASS** sections. This section is used to configure IBM Spectrum Protect Snapshot for use with a particular storage solution. The parameters do not depend on the database or custom application that is protected. Follow the steps in the appropriate procedure for your disk storage environment. For your disk storage subsystem, data files must be defined on volume groups that are separate from the volume groups where the control files and redo logs are defined.

For more information about volume group layout requirements, see Chapter 3, “Preparation for installation,” on page 27.

Configuring Storwize V7000 and SAN Volume Controller dynamic target allocation (SVCDTA)

To allow dynamic volume creation during backup operations, you must enable Secure Shell (SSH) remote access to the storage system command-line interface (CLI) with Secure Shell (SSH) keys. An SSH key pair must be created to authenticate users for a secure connection to SAN Volume Controller.

Before you begin

Verify that the OpenSSH client is installed on the production server, and the backup or clone server where IBM Spectrum Protect Snapshot is installed. The OpenSSH client is installed by default on most AIX and Linux distributions. If it is not installed on your system, consult your AIX or Linux installation documentation.

About this task

SSH is used to remotely enter commands on the SAN Volume Controller CLI. The following steps are required to enable CLI access with SSH keys:

- Generate a public and a private key pair
- Import the public key to the storage system
- Configure IBM Spectrum Protect Snapshot to authenticate with the private key.

The IBM Spectrum Protect Snapshot user must have a unique SSH key at the SAN Volume Controller. After you generate the key pair, import the public key and add a key file for the SAN Volume Controller user as specified in the IBM Spectrum Protect Snapshot profile. The parameters are **COPYSERVICES_USERNAME** and **COPYSERVICES_REMOTE_USERNAME**. The user ID at the remote site also needs a key file. The IBM Spectrum Protect Snapshot user owns the private key and has RW access to that key file.

The full path to the private key file is specified in the profile. By default, the path is `$HOME/.ssh/svc_sshkey`. The public counterpart of the private key file must be imported to the SAN Volume Controller and associated to the user ID.

Procedure

1. Generate an RSA key pair on the production server for the storage user name to access the storage system by entering the following command from the `$HOME/.ssh` directory. Ensure to enter the command as the database instance owner or application backup user from the `$HOME/.ssh` directory.

```
ssh-keygen -t rsa
```

This command generates two files, which you are prompted to name. If you select the name `svc_sshkey`, the private key is named `svc_sshkey`, and the public key is named `svc_sshkey.pub`.

Tip: Do not enter a passphrase for the file when prompted. For SVCDDTA dynamic target allocation, the passphrase must be empty.

2. If you do not remotely install the backup or cloning servers with SSH, you must copy the key pair to the backup and clone servers. Ensure that the key pair is stored in the same path as on the production server.
3. Upload the public key to the storage system for the SAN Volume Controller user that is specified by **COPYSERVICES_USERNAME** in the profile.

For instructions about how to upload to the storage system, see the documentation that is provided for your storage system. The documentation is available in IBM SAN Volume Controller Knowledge Center <http://www.ibm.com/support/knowledgecenter/STPVGU/welcome?lang=en>.

4. Run the IBM Spectrum Protect Snapshot for UNIX and Linux setup script in advanced mode by entering the following command:

```
./setup_gen.sh -advanced
```

Note: If you do not want to use an alternative SSH binary and the private key file is named `svc_sshkey` in the default path `$HOME/.ssh`, you can proceed to run the setup script in basic mode.

5. When prompted to specify a **SSH_DIR** path, enter the path where the Secure Shell protocols and executable files are installed. The default location is `/usr/bin`.
6. When prompted to specify a **SVC_SSHKEY_FULLPATH** path, enter the path and the file name for the private keyfile. The following example shows the default path and file name:

```
SVC_SSHKEY_FULLPATH    $HOME/.ssh/svc_sshkey
```

7. Continue configuring IBM Spectrum Protect Snapshot for SAN Volume Controller with the setup script for your component. When you are configuring SAN Volume Controller Dynamic Target Allocation, the profile that is created is saved with the necessary parameters.
8. Complete the process by restarting the IBM Spectrum Protect Snapshot daemons.

What to do next

If you are using SAN Volume Controller remote mirroring, the setup script asks if you want to create another SSH key to facilitate mirroring with the remote cluster. The key file **SVC_REMOTE_SSHKEY_FULLPATH** parameter specifies the private key file that is used for connecting to the secondary SAN Volume Controller site, and is specified by **COPYSERVICES_REMOTE_SERVERNAME**. The remote site `userId` is the one specified by the parameter **COPYSERVICES_REMOTE_USERNAME**.

Migrating from SVC with static target allocation to SVC with dynamic target allocation (SVCDTA)

You can change an existing configuration of IBM Spectrum Protect Snapshot for UNIX and Linux to use dynamic target allocation (COPYSERVICES_HARDWARE_TYPE: SVCDTA) without losing older backups. If the profile is using a device class that is configured for static target allocation (COPYSERVICES_HARDWARE_TYPE: SVC), you can create a new device class for SVCDTA and add it to the profile.

Before you begin

To start the configuration process, run the generic setup script with the following command: `./setup_gen.sh`

About this task

The following information demonstrates how to modify an existing IBM Spectrum Protect Snapshot configuration profile to use a new device class with dynamic target allocation. In this example, the Client profile is modified to change the device class from 'STANDARD' to a new device class called 'STANDARD_DTA'.

Procedure

1. Choose (m) when presented with the following options:
 - (c)reate a new profile
 - (r)euse the profile unchanged
 - (m)odify the profile
2. The profile parameters for the configuration that is being modified are displayed, in this case, for the 'CLIENT' section. Within this section, for the DEVICE_CLASS parameter, replace STANDARD with STANDARD_DTA.

Note: These steps are applicable for the 'CLONING' profile section also.

3. You are asked if you want to delete the device class that is being replaced. Device section STANDARD is no longer referenced. Do you want to delete it?[y|n]

Choose n to ensure that the existing device class is not deleted.

Important:

The existing device class, in this case 'STANDARD', **must be retained** to allow for any existing backup snapshots to be mounted or restored.

4. The profile parameters for the new device class 'STANDARD_DTA' are displayed, starting with the COPYSERVICES_HARDWARE_TYPE. Change this setting from SVC to SVCDTA.

Tip:

If the MAX_VERSIONS parameter is set to 'ADAPTIVE', you must return to the CLIENT profile section, and change the MAX_VERSIONS parameter from ADAPTIVE to a fixed number.

5. Enter the existing server information for the storage system host name (COPYSERVICES_SERVERNAME). Because you are using the same storage system server, but with a different storage adapter, a warning message is displayed. This message lists the restrictions that are associated with configuring both SVC and SVCDTA device classes on the same server.
6. You are asked if you want to proceed with the current configuration.

Enter (r) to retry or (i) to ignore and proceed.

- Choose i if you want to proceed with the configuration, acknowledging that some restrictions apply.
 - Choose r if you want to change the configuration, and use a different storage system server.
7. Enter the user name for the primary storage device (COPYSERVICES_USERNAME). The default name is superuser.
 8. Enter the path and the file name of the private SSH key file (SVC_SSHKEY_FULLPATH). For example:

```
SVC_SSHKEY_FULLPATH    $HOME/.ssh/svc_sshkey
```


, where \$HOME/.ssh/svc_sshkey is the default.
 9. Accept the defaults for the remaining parameters, or change where necessary. For example, change the FlashCopy type from NOCOPY to COPY.
 10. The profile is saved, and you are asked if you would like to specify a backup system or to quit the configuration.
Currently no backup system is setup. To configure a backup system please select option n.
 - Choose n if you want to specify a new backup system.
 - Choose q if you want to quit the configuration.

Results

IBM Spectrum Protect Snapshot for UNIX and Linux is now configured to use the SAN Volume Controller storage adapter with dynamic target allocation on the SAN Volume Controller storage server that was already in use for device type SVC.

Restriction:

If a configuration uses both device types 'SVC' and 'SVCDTA' on the same IBM Storwize v7000/IBM System Storage SAN Volume Controller server, the following limitations apply.

- No new backups can be created for the DEVICE_CLASS sections that use COPYSERVICES_HARDWARE_TYPE: SVC. If you attempt to create such a backup, a clear error message is displayed.
- Existing backups that were created with these DEVICE_CLASS sections can be mounted and restored, **but any newer backups are destroyed, even if they were created with the SVCDTA adapter.**
- Existing device classes that use COPYSERVICES_HARDWARE_TYPE: SVC must not be deleted until all backups that were created using this device class are expired and deleted from the IBM Spectrum Protect Snapshot repository, and also from the storage system.

Configuring the CIM adapter for SP 800-131A compliant encryption

CIM agents are provided by IBM System Storage SAN Volume Controller, IBM Storwize, and IBM System Storage DS8000 systems. IBM Spectrum Protect Snapshot for UNIX and Linux communicates with a CIM agent through the CIM interface. You must configure the CIM adapter to use the security standards, as defined in the National Institute of Standards and Technology (NIST) Special Publications (SP) 800-131A for encryption.

Before you begin

Ensure that the storage system is enabled for SP 800-131A standard encryption. For instructions about how to identify if the system is enabled, see the documentation that is provided for your storage system. For the new SVC adapter with dynamic target allocation (type SVCDTA), compliance with SP 800-131A is provided by the OpenSSH client version that is installed on the same host as the product.

Note: For IBM System Storage SAN Volume Controller and IBM Storwize family, this configuration applies only in the case of static target allocation (type SVC); the new SVC adapter with dynamic target allocation (type SVCDTA) uses the CLI interface via Secure Shell (SSH) rather than the CIMOM interface.

Procedure

1. Extract the Secure Sockets Layer (SSL) certificate from the IBM storage system cluster. The certificate must be in the Privacy Enhanced Mail (PEM) format. From any Linux or UNIX system with a LAN connection to the storage system, run the following shell command,

```
echo | openssl s_client -connect ibm_storage_cluster_ip:5989 2>&1  
| sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p'
```

where *ibm_storage_cluster_ip* specifies the IP address of the storage system, and *5989* specifies the port number for the HTTPS connection.

2. Save the output to a text file and place the file in a secure location on the production and backup servers.
3. Run the setup script in advanced mode by entering the following command:

```
./setup.sh -advanced
```
4. When prompted for the **COPYSERVICES_CERTIFICATEFILE** parameter for the storage system device class, enter the fully qualified path to the certificate file. For example:

```
COPYSERVICES_CERTIFICATEFILE ACS_DIR/truststore/svc_cluster.cert
```
5. Follow the setup script instructions to save the profile and restart the daemons.

Defining Logical Unit Numbers on DS8000 storage subsystems

Logical Unit Numbers (LUNs) must be defined for the DS8000 storage subsystem.

Before you begin

Before you start defining LUNs on the storage subsystem, verify that the following prerequisites are met:

- The LUNs are located where the production database or application is located.
- The size of the LUNs is dependent upon the size of the database or application.

- The size of the source volumes on the production server and size of the target volumes on the backup server must be the same.
- Both the source volume and target volume must be defined on the same storage subsystem.
- Assign the source volume to the DS8000 volume group that is associated with the production server.

Procedure

Perform these steps so that the correct LUNs are defined on both the production server and backup server.

1. Use the DS8000 Storage Manager to create two or more fixed block LUNs on the production server.

2. Use the DS8000 Storage Manager to create the same number of LUNs for the backup server as were created for the production server in the previous step.

Real-time manager (or Simulated manager)-> Configure storage -> Open systems -> Volumes-open systems

These LUNs must also be the same size as the LUNs created for the production server. Assign the target volume to the DS8000 volume group that is associated with the backup server.

3. Identify the serial numbers of the target LUNs by using the DS8000 Storage Manager.

Real-time manager (or Simulated manager)-> Configure storage -> Open systems -> Volumes-open systems

Select the target LUNs created on the backup server in Step 2. Identify the serial numbers with the matching size in the source LUNs. For example:

```
7501901
Nickname      Number Status Type GB
sandburr_3300 3300  Normal DS  2.0
sandburr_3400 3400  Normal DS  2.0
```

In this example, the serial numbers are 75019013300 and 75019013400.

4. Define the **TARGET_VOLUME** parameter in the target volumes file specified by the **VOLUMES_FILE** profile parameter with the appropriate serial numbers of the target LUN. For example:

```
TARGET_VOLUME 75019013300
TARGET_VOLUME 75019013400
```

This setting specifies the target volumes where the database or application is backed up.

Defining virtual disks on SAN Volume Controller and Storwize V7000

When you define virtual disks for the SAN Volume Controller and the Storwize V7000 storage devices, you can use either the graphical user interface or the command-line interface.

Before you begin

Before you start defining virtual disks, verify that the following prerequisites are met:

- A storage area network is available.

- Storage disks are attached and available in the SAN Volume Controller or Storwize V7000 environment.
- Subsystem Device Driver (SDD) or Subsystem Device Driver Path Control Module (SDDPCM) is installed and available on the host systems.
- A cluster is available in the SAN Volume Controller or Storwize V7000 environment.
- Each host has at least two paths to the SAN Volume Controller or Storwize V7000 storage device.

Procedure

To create virtual disks on the production server and backup server, complete the following steps.

1. From the graphical user interface, select **Work with Virtual Disks > Virtual Disks > Create Virtual Disks**. The virtual disks are created by using the managed disk group.
2. Map the virtual disk to the hosts that are created for the production server. To map the virtual disks to the backup server, in the IBM Spectrum Protect Snapshot profile file, configure the **BACKUP_HOST_NAME** parameter by assigning one of the following values:
 - Assign the value `PREASSIGNED_VOLUMES` to use a static predefined map.
 - Assign the `backup_server_hostname` to allow IBM Spectrum Protect Snapshot to dynamically map the target virtual disks when needed.

Note: The value `PREASSIGNED_VOLUMES` is not allowed if you select SAN Volume Controller and Storwize V7000 dynamic target allocation.

3. Define the **TARGET_VOLUME** parameter in the target volumes file (.fct). This name is specified by the **DEVICE_CLASS > TARGET_SETS > VOLUMES_FILE** parameter with the appropriate virtual disk names of the target LUNs in the profile configuration file. For example:

```
TARGET_VOLUME A01pro1_1_t1
TARGET_VOLUME A01pro1_2_t1
```

In this example, the source volume names are `A01pro1_1` and `A01pro1_2` with target set named 1.

Alternatively, you can define the target names by using the **TARGET_NAMING** parameter in the IBM Spectrum Protect Snapshot profile file.

Note: The parameter **TARGET_SETS** in the device class section is not allowed if you select SAN Volume Controller and Storwize V7000 dynamic target allocation.

Select the FLASHCOPY_TYPE

DS8000, SAN Volume Controller, and Storwize V7000 storage solutions support various FlashCopy types that provide different capabilities for your backup strategy.

Using different FlashCopy types for different backup generations is a valid strategy for IBM Spectrum Protect Snapshot. To implement such a backup strategy, define multiple **DEVICE_CLASS** sections in the profile, where each section specifies the same storage device. The only difference is that each section specifies a different FlashCopy type. These **DEVICE_CLASS** section definitions allow rules to be defined in the **CLIENT** profile section. The rules allow IBM Spectrum Protect Snapshot to

select the appropriate **DEVICE_CLASS** section for the next backup. For more information about the **DEVICE_CLASS** parameter, see the **CLIENT** section.

If the **FLASHCOPY_TYPE** is changed for one **DEVICE_CLASS**, complete the following steps:

1. Unmount the backup if it is mounted on a backup system.
2. Delete the backup with the delete force option.
3. Change the **FLASHCOPY_TYPE** in the **DEVICE_CLASS** and run a new backup with the new **FLASHCOPY_TYPE**.

Note: If you use SAN Volume Controller and Storwize V7000 dynamic target allocation you do not have to delete any old backups.

Table 4. Selecting the **FLASHCOPY_TYPE** for DS8000, SAN Volume Controller, and Storwize V7000

FLASHCOPY_TYPE	DS8000	SAN Volume Controller Storwize V7000
COPY	Can be used for backup and restore. Protects from physical failures of the source volumes when the background copy completes.	Can be used for backup and restore. Protects from physical failures of the source volumes when the background copy completes. For more information, see Note 1 in this table.
INCR	Same characteristics as COPY FLASHCOPY_TYPE but with fewer COPY activities in the background. DS8000 allows at most 1 incremental FlashCopy per source volume. In mirroring environments, this setting allows it to retain 1 backup generation per mirror. For DS8000, there must be only one target set specified in the target volumes file (.fct) for incremental snapshots. CIM errors might occur when more than 1 target set is specified.	Same characteristics as COPY FlashCopy but with fewer COPY activities in the background. For more information, see Notes [®] 1 and 2 in this table.
NOCOPY	Can be mounted remotely, but cannot be restored.	Can be mounted remotely and can be restored. Can be used to create a FlashCopy to a space-efficient target, but does not offer protection from physical failures to the source volume. Space-efficient target volumes can reach capacity limits in which case they go offline. In this scenario, you lose the current backup and all older backups that are not at FULL_COPY. You can choose to create space-efficient targets with the AUTOEXPAND option. In this scenario, the target is allocated more physical storage to prevent it going offline.

Note 1: If space-efficient source volumes are used in combination with space-efficient target volumes, IBM Spectrum Protect Snapshot can be configured to use **FLASHCOPY_TYPE** COPY, INCR, or NOCOPY. If fully allocated source volumes are used in combination with space-efficient target volumes, then IBM Spectrum Protect Snapshot can be configured to use **FLASHCOPY_TYPE** COPY, INCR, or NOCOPY. These options are available when the profile parameter **ALLOW_ALL_FLASHCOPY_TYPES** is set to YES. The default value of **ALLOW_ALL_FLASHCOPY_TYPES** is NO. When the default value is used, only **FLASHCOPY_TYPE** NOCOPY is possible.

Note 2: The information in Note 1 only applies if you use SAN Volume Controller and Storwize V7000 static target allocation. If you use SAN Volume Controller and Storwize V7000 dynamic target allocation, then **FLASHCOPY_TYPE** INCR and profile parameter **ALLOW_ALL_FLASHCOPY_TYPES** are not available.

The types of snapshots that are supported by IBM Spectrum Protect Snapshot, depending on the storage solution and operating system, are indicated in the following table.

Table 5. Supported storage subsystems and FlashCopy types

Device	COPY	INCR	NOCOPY	Space-efficient snapshots	Changes made to a mounted snapshot backup
IBM System Storage DS8000	Yes	Yes	Yes	N/A	Remains persistent and alters the content of the backup.
IBM System Storage SAN Volume Controller IBM Storwize family with static target allocation	Yes	Yes	Yes Includes space-efficient copies if configured so.	N/A	Remains persistent and alters the content of the backup.
IBM System Storage SAN Volume Controller IBM Storwize family with dynamic target allocation	Yes	No	Yes	N/A	Reverted during unmount and does not alter the backup.
IBM XIV Storage System	N/A	N/A	N/A	Yes	Reverted during unmount and does not alter the backup or remains persistent and alters the content of the backup.

Target set definitions

IBM Spectrum Protect Snapshot requires target sets to be defined for SAN Volume Controller, Storwize V7000, and DS8000.

Define targets by using target set definition files (SAN Volume Controller, Storwize V7000, and DS8000) or by using a naming convention (SAN Volume Controller and Storwize V7000 only). This convention determines the name of the target for both the source volume name and the target set name as specified for the current operation.

Tip: There is no requirement to define target volumes, if you select SAN Volume Controller and Storwize V7000 dynamic target allocation.

Target set definition files

A target set definition file contains a list of target volumes that are organized into target sets.

During the backup process, IBM Spectrum Protect Snapshot software matches source volumes to suitable targets within a target set. To determine source target relations, associate a source name with a target name in a target set definition file. In this scenario, the relationship between the source and target is required. Backup processing fails if one of the targets is unavailable for the specified source. For details on the target selection algorithms, see “Target set and target volumes” on page 153.

If IBM Spectrum Protect Snapshot attempts to mount the target set, the volumes within the target set must be assigned to a backup host. For example, the target set is mounted to create a backup to IBM Spectrum Protect. Because all target volumes within a single target are mounted to the same host, assign all target volumes

within a target set to the same host. When you use multiple backup servers within your environment, use multiple target set definition files.

For SAN Volume Controller and Storwize V7000 storage solutions, IBM Spectrum Protect Snapshot can assign the target volumes dynamically during the mount operation. In this case, you must not assign the target volumes in advance of the mount operation.

```
>>> TARGET_SET SET_1 # FCM determines a suitable target for every source
TARGET_VOLUME 40913158
TARGET_VOLUME 40A13158
TARGET_VOLUME 40B13158
<<<
>>> TARGET_SET SET_2 # For every source the target is mandated in the target set
# definiton (source name following target name)
TARGET_VOLUME 40C13158 40613158
TARGET_VOLUME 40D13158 40713158
TARGET_VOLUME 40E13158 40813158
<<<
```

Referring to target set definitions from the profile

The target set definition file must be specified in the `DEVICE_CLASS` section of the profile.

The following example is a section from an IBM Spectrum Protect Snapshot profile file that shows the association between `TARGET_SETS`, `VOLUMES_FILE`, and *name of target set definition file* parameters.

```
>>> DEVICE_CLASS STANDARD
COPYSERVICES_HARDWARE_TYPE DS8000
COPYSERVICES_PRIMARY_SERVERNAME <hostname> #
TARGET_SETS VOLUMES_FILE
VOLUMES_FILE name of target set definition file
FLASHCOPY_TYPE INCR
<<<
```

If multiple `DEVICE_CLASS` configuration sections are specified within the profile, each `DEVICE_CLASS` section must be associated with a unique target set definition file. The target set names must be unique across all target set definition files. If all target sets within the target set definition file are assigned to the same host and associated with one `DEVICE_CLASS`, they are mounted on the same host.

Target set definitions using the naming convention

Target set definitions can also be provided by using a naming convention on SAN Volume Controller and Storwize V7000.

IBM Spectrum Protect Snapshot supports using a naming convention, instead of a definition file, for target set definitions on SAN Volume Controller and Storwize V7000 storage systems. IBM Spectrum Protect Snapshot determines the target volume names from the name of the target set, used for the current backup, and the name of the source volume.

Target sets are specified directly in the `DEVICE_CLASS` configuration section of the profile for example, `TARGET_SETS 1 2 3`. The names are generated from `TARGET_SETS` and are sequentially numbered, 1, 2, 3, 1, 2, and so on. When target sets are defined in the profile, the target set name must be unique in the entire profile. For example, you cannot have the `TARGET_SETS` parameter, set to t1 for more than one device class. The following example shows multiple device classes that are named in the `DEVICE_CLASS` configuration section of the profile:

```

>>> Device_Class SVC_01
.
.
TARGET_SETS t1 t2
.
.
<<<
>>> Device_Class SVC_02
.
.
TARGET_SETS t3 t4
.
.
<<<
>>> Device_Class SVC_03
.
.
TARGET_SETS t5 t6
.
.
<<<

```

A TARGET_NAMING rule is also specified to determine the name of the target volume from the name of the source. For example, TARGET_NAMING %SOURCE_bt%TARGETSET. If the application is stored on a volume named *db_vol*, the targets required by IBM Spectrum Protect Snapshot are *db_vol_bt1*, *db_vol_bt2*, and *db_vol_bt3*. These targets depend on the target set that is selected for the current backup.

```

>>> DEVICE_CLASS STANDARD
COPYSERVICES_HARDWARE_TYPE SVC
COPYSERVICES_PRIMARY_SERVERNAME <hostname>
TARGET_SETS 1 2 3
TARGET_NAMING %SOURCE_bt%TARGETSET
FLASHCOPY_TYPE NOCOPY
<<<

```

The given TARGET_SETS or TARGET_NAMING definition results in the following target volume names:

```

    name of source volume_bt1
    name of source volume_bt2
    name of source volume_bt3

```

ASM failure group environment

In a Logical Volume Manager (LVM) mirroring on AIX®, and an Oracle Automatic Storage Management (ASM) failure group environment, multiple DEVICE_CLASS configuration sections are required. One section per storage subsystem or LVM mirror is required.

The **LVM_MIRRORING** parameter must be specified in the DEVICE_CLASS configuration section with a value of YES. This example shows the configuration,

```

>>> DEVICE_CLASS MIRR_1
COPYSERVICES_HARDWARE_TYPE DS8000
COPYSERVICES_PRIMARY_SERVERNAME DS8000_1
LVM_MIRRORING YES
TARGET_SETS VOLUMES_FILE
VOLUMES_FILE <name of target set definition file 1>
FLASHCOPY_TYPE INCR
<<<
>>> DEVICE_CLASS MIRR_2
COPYSERVICES_HARDWARE_TYPE DS8000

```

```

COPYSERVICES_PRIMARY_SERVERNAME DS8000_2
LVM_MIRRORING YES
TARGET_SETS VOLUMES_FILE
VOLUMES_FILE <name of target set definition file 2>
FLASHCOPY_TYPE INCR
<<<

```

LVM mirroring environments

In a Logical Volume Manager (LVM) mirroring on AIX, multiple `DEVICE_CLASS` configuration sections are required. One section per storage subsystem or LVM mirror is required.

The `LVM_MIRRORING` parameter must be specified in the `DEVICE_CLASS` configuration section with a value of `YES`. This example shows the configuration,

```

>>> DEVICE_CLASS MIRR_1
COPYSERVICES_HARDWARE_TYPE DS8000
COPYSERVICES_PRIMARY_SERVERNAME DS8000_1
LVM_MIRRORING YES
TARGET_SETS VOLUMES_FILE
VOLUMES_FILE <name of target set definition file 1>
FLASHCOPY_TYPE INCR
<<<
>>> DEVICE_CLASS MIRR_2
COPYSERVICES_HARDWARE_TYPE DS8000
COPYSERVICES_PRIMARY_SERVERNAME DS8000_2
LVM_MIRRORING YES
TARGET_SETS VOLUMES_FILE
VOLUMES_FILE <name of target set definition file 2>
FLASHCOPY_TYPE INCR
<<<

```

Backup and clone server assignment

With IBM Spectrum Protect Snapshot software, you can mount backup images and clone images. Each backup image and clone image is mounted on a server. However, you cannot mount a backup image or a clone image on more than one server at one time.

IBM Spectrum Protect Snapshot mount operation can be started by one of the following methods:

- By issuing a mount command from the command-line interface.
- By issuing a create or refresh clone command from the command-line interface.
- When IBM Spectrum Protect Snapshot is used with IBM Spectrum Protect and you offload backups to IBM Spectrum Protect.

The information that you enter during the installation and configuration of IBM Spectrum Protect Snapshot is used to create a profile configuration file. The `DEVICE_CLASS` section of this profile specifies the backup host name where the backup or clone images are mounted. There can be multiple `DEVICE_CLASS` sections. The `CLIENT` section specifies the `DEVICE_CLASS` to use for backup and offload operations. The `CLONING` section specifies the `DEVICE_CLASS` to use for cloning operations.

FlashCopy or snapshot target volumes are mounted and assigned to selected backup or clone server. Depending on the storage system and profile configuration the following assignments occur:

IBM XIV Storage Systems.

The assignment automatically occurs during the mount request.

SAN Volume Controller and Storwize V7000

If the **BACKUP_HOST_NAME** parameter is specified as *backup_server_hostname* in the **DEVICE_CLASS** section, the target volumes are mapped dynamically from the storage system to the backup and clone server.

DS8000, SAN Volume Controller, and Storwize V7000

If the **BACKUP_HOST_NAME** parameter is specified as *PREASSIGNED_VOLUMES* in the **DEVICE_CLASS** section, the target volumes must be preassigned to a specific backup or clone server before you issue a mount command. Ensure that the target volumes of all target sets associated with a specific **DEVICE_CLASS** are assigned to the same hosts. If target set definition files are used, assign all volumes within one target set definition file to the same host. This setting ensures that targets associated with a single device class are mounted from the same backup or clone server.

For all IBM Spectrum Protect Snapshot mount operations, there can be only one backup or clone server for each device class. If the identified servers have not mounted a backup or clone image, the mount request is propagated to those servers. The backup or clone is then mounted.

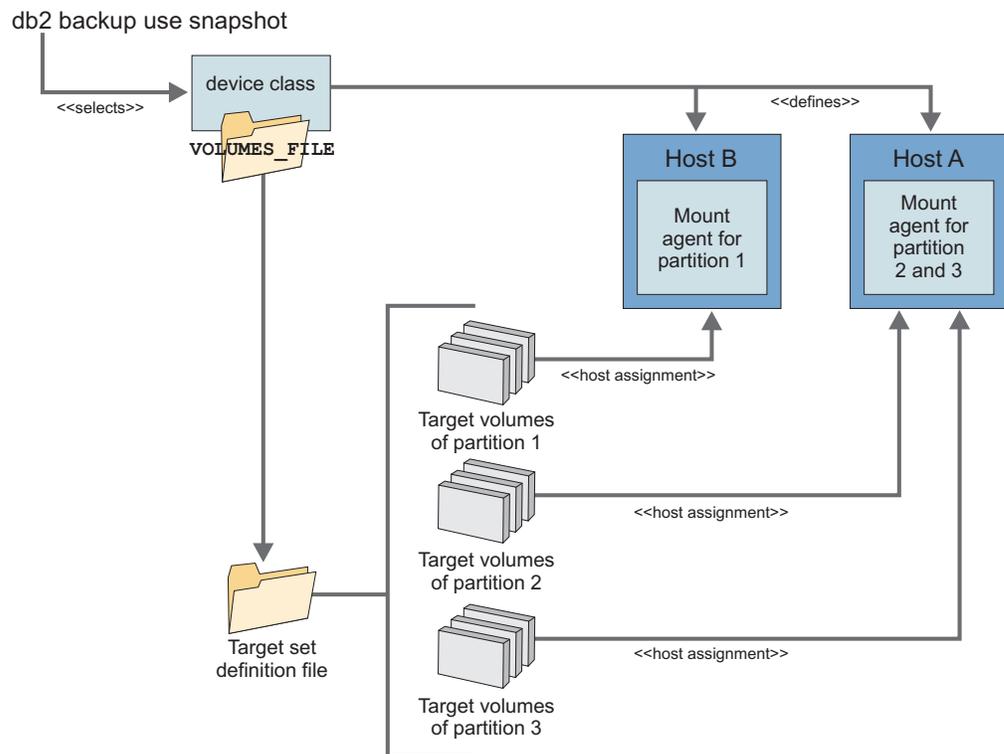


Figure 6. IBM Spectrum Protect Snapshot host assignments. This example shows a DB2 configuration.

Managing backups and clones with the `DEVICE_CLASS` parameter

Use the `DEVICE_CLASS` parameter in the `CLIENT` section of the IBM Spectrum Protect Snapshot profile file to select the storage device configuration for backups. In the `CLONING` section of the profile file, use this parameter to select the storage device configurations for cloning.

The IBM Spectrum Protect Snapshot `DEVICE_CLASS` profile parameter can be used as a filter to determine these backup criteria:

- Partition number
- Day of week
- Time of backup
- Cloning only: Clone database name

When used in this manner, the `DEVICE_CLASS` parameter provides access to a specific storage device. This device is identified by the copy services type, user name, and server name that is defined by the corresponding `DEVICE_CLASS` profile section. It also provides a backup policy that is device-specific. For example, this device-specific backup policy might be defined by these factors:

- List of target sets on DS8000, SAN Volume Controller, or Storwize V7000
- The type of FlashCopy backup to be completed (for example, incremental or copy)
- The mount location of the backup
- Whether a backup to IBM Spectrum Protect server storage is created from the snapshot

The `DEVICE_CLASS` parameter is specified in the client section of IBM Spectrum Protect Snapshot profile file. The settings for this parameter can be overridden with a command-line option during backup operations. Use the following command-line option:

From the Oracle in an SAP environment interface

`-S device class` in SAP BR*Tools configuration profile (.sap) `util_options` parameter.

From the Oracle environment

`-s device class` on the `acsora -f` backup command.

The `DEVICE_CLASS` parameter cannot be specified with the `restore`, `mount`, `unmount`, and `delete` commands. You can specify the backup ID, if it is not specified the latest backup is used. IBM Spectrum Protect Snapshot automatically uses the `DEVICE_CLASS` that was used for the selected backup at backup time.

Examples of how to use `DEVICE_CLASS` filters

This example creates alternating backups to each mirror. Device classes `MIRROR_1` and `MIRROR_2` refer to two separate storage clusters. Only those backups that are created to `MIRROR_2` are backed up to IBM Spectrum Protect server storage:

```
>>> CLIENT
TSM_BACKUP_LATEST_USE_FOR_MIRROR_2
DEVICE_CLASS MIRROR_1 MIRROR_2
[...]
```

```
<<<
```

This example creates backups to MIRROR_1 on Monday (1), Wednesday (3), and Friday (5). It creates backups to MIRROR_2 on Sunday (0), Tuesday (2), and Thursday (4), and Saturday (6). All backups are stored on IBM Spectrum Protect server storage:

```
>>> CLIENT
TSM_BACKUP LATEST
DEVICE_CLASS MIRROR_1 USE_AT Mon Wed Fri
DEVICE_CLASS MIRROR_2 USE_AT Sun Tue Thu Sat
[...]
<<<
```

This example creates disk only backups during the specified period of the day. These disk only backups are considered space-efficient. A full backup is also created at midnight that is stored on IBM Spectrum Protect server storage. Although the *DAYTIME* and *MIDNIGHT* device classes might have the same configuration, two different device classes are used. This setting is used even if both device classes point to the same SAN Volume Controller cluster:

```
>>> CLIENT
TSM_BACKUP LATEST USE_FOR MIDNIGHT
DEVICE_CLASS DAYTIME FROM 1:00 TO 23:59
DEVICE_CLASS MIDNIGHT FROM 0:00 TO 0:59
[...]
<<<
>>> DEVICE_CLASS DAYTIME
COPYSERVICES_HARDWARE_TYPE SVC
FLASHCOPY_TYPE NOCOPY
[...]
<<<
>>> DEVICE_CLASS MIDNIGHT
COPYSERVICES_HARDWARE_TYPE SVC
FLASHCOPY_TYPE INCR
SVC_COPY_RATE 80
[...]
<<<
```

Note: The time period that is specified cannot span midnight for a device class. If a device class time period is required to span midnight, you must specify two time periods for the device class. The first time period must end with a value 1 minute before midnight and the second time period must start at midnight. The following example shows how to specify a time period that spans midnight for a device class:

```
DEVICE_CLASS myClass FROM 20:00 TO 23:59
DEVICE_CLASS myClass FROM 00:00 TO 06:00
```

This example demonstrates how to create clone databases on different clone servers from the same production database. In this scenario, there are two clone servers, each one uses a different device class. The clone server host1 uses *DEVICE_CLASS CLONE1* and host2 uses *DEVICE_CLASS CLONE2*. When a clone request is started with clone database B01 selected, this clone is created with *DEVICE_CLASS CLONE1* and it is created on clone server host1.

```
>>> CLONING
DEVICE_CLASS CLONE1 USE_FOR_CLONING B01 C01
DEVICE_CLASS CLONE2 USE_FOR_CLONING B02 C02
<<<
>>> DEVICE_CLASS CLONE1
CLONE_DATABASE YES
...
<<<
```

```
>>> DEVICE_CLASS CLONE2
CLONE_DATABASE YES
...
<<<
```

Configuring for remote mirroring

When you configure IBM Spectrum Protect Snapshot, you can set the configuration parameters to create snapshots by using target volumes of remote mirroring relationships. These target volumes are used to create application consistent snapshot backups.

Before you begin

Before you configure IBM Spectrum Protect Snapshot to use target volumes that are associated with remote mirroring one of the following technologies must be deployed:

- SAN Volume Controller or Storwize V7000 Global Mirror and Metro Mirror
- IBM XIV Storage System Synchronous Remote Mirroring and Asynchronous Remote Mirroring

About this task

To configure IBM Spectrum Protect Snapshot with SAN Volume Controller or Storwize V7000 Global Mirror and Metro Mirror, complete the following steps:

Procedure

1. On the SAN Volume Controller or Storwize V7000 system, create a partnership between the primary and secondary clusters. For example, you can run the following commands from the command-line interface:

```
ssh -i/dir/ssh-identity username@hostname or ip_primary_cluster
svctask mkpartnership -bandwidth bandwidth_in_mbps remote_cluster_name
or remote_cluster_id
```

2. Start the Global Mirror and Metro Mirror relationship by using either the graphical user interface or command-line interface. If you use the command-line interface, the following commands are provided as an example:

```
ssh -i/dir/ssh-identity username@hostname or ip_primary_cluster
svctask chpartnership -start remote_cluster_name or remote_cluster_id
```

3. Verify that the following information is true for the environment:
 - The production volumes are on the primary storage system.
 - The production volumes are in a remote mirror relationship with the remote volumes that are either in the secondary cluster, or in the same cluster.
 - All the remote mirror relationships are defined in a consistency group.
4. Run the setup script to configure a dedicated device class for the snapshot backups on the remote cluster. When you configure the new `DEVICE_CLASS` section with the setup script, look for the following prompt:

```
Is the FlashCopy/Snapshot taken from the mirror volumes {COPYSERVICES_REMOTE}.
```

Enter *yes*. The `COPYSERVICES_REMOTE_SERVERNAME`, `COPYSERVICES_REMOTE_USERNAME`, and `TAKEOVER_HOST_NAME` parameters are also required for remote mirroring.

5. The SSH parameter `SVC_SSHKEY_FULLPATH` specifies the path and the file name to the private SSH key file required for SAN Volume Controller. For remote mirroring, `SVC_REMOTE_SSHKEY_FULLPATH` specifies the second SSH key file to be

used for authentication on the remote site storage device. The key file is used to authenticate to the storage system with the user name specified for the **COPYSERVICES_REMOTE_USERNAME** parameter. If you do not want to create a new key pair for the remote site, one key can be shared for both storage sites.

6. If you are using SAN Volume Controller with static target allocation, you must allocate target volumes. On the remote cluster of the SAN Volume Controller or Storwize V7000, specify the corresponding snapshot target volumes for each source. To specify the snapshot target volumes, use one of the following options:
 - Parameter **TARGET_SETS** with **VOLUMES_FILE**. For example:

```
TARGET_SETS VOLUMES_FILE
VOLUMES_FILE /<component database>/DS0/acs/volumes/STANDARD_gm.fct
```
 - Parameter **TARGET_SETS** with **TARGET_NAMING**. For example:

```
TARGET_SETS dc2 dc3 dc4 dc5
TARGET_NAMING %SOURCEx%TARGETSET
```
7. At the end of the setup script configuration process, verify the user name and password. When you see the following prompt, enter *yes*:

Do you want to continue by specifying passwords for the defined sections?

Configure XIV remote mirroring

To configure IBM Spectrum Protect Snapshot with XIV Synchronous Remote Mirroring and Asynchronous Remote Mirroring, complete the following steps:

Procedure

1. Define a coupling between peer volumes on the master and subordinate XIV systems, which creates a mirror relationship between the two.
2. Activate the XIV remote mirror couplings.
3. Define a coupling between peer consistency groups on the master and subordinate XIV systems, which creates a mirror relationship between the two.
4. Add volume mirror couplings to the consistency group couplings.
5. Run the setup script to configure a dedicated device class for the snapshot backups on the remote cluster. When you configure the new **DEVICE_CLASS** section with the setup script, look for the following prompt:

Is the FlashCopy/Snapshot taken from the mirror volumes {COPYSERVICES_REMOTE}.

Enter *yes*. The **COPYSERVICES_REMOTE_SERVERNAME**, **COPYSERVICES_REMOTE_USERNAME**, and **TAKEOVER_HOST_NAME** parameters are also required for remote mirroring.

Example

The following information is provided as an example of how a team can complete asynchronous remote mirror configuration across two sites:

To configure IBM Spectrum Protect Snapshot with IBM XIV Storage System with Asynchronous Remote Mirroring at both sites, certain ports must be open within the firewalls:

- On the production system, the production host, backup host, and primary XIV system must have ports open within the firewall.
- On the takeover system, the takeover host, backup host, and secondary XIV system must have ports open within the firewall.

For both the primary and secondary sites, the following ports must be open within the firewall:

- TCP port 3260 (iSCSI) open within firewalls for iSCSI replication
- Ports: http, https, ssh, and telnet
- TCP/IP ports: 55697, 5997, 5998, and 7778

All ports must be bidirectional.

Setting up the daemons on the production and backup systems

Before manually starting the IBM Spectrum Protect Snapshot daemon processes, identify the daemons that must run on the production, backup, and cloning systems.

Procedure

You can manually set up the daemon processes. The following list specifies where the daemons can run.

- Run the following daemons on the production system only:
 - `INSTANCE_DIR/acsd` (management agent)
 - `INSTANCE_DIR/acsgen -D` (generic device agent)
- If offloaded backups are configured, run the following daemon on the production server:
`INSTANCE_DIR/fcmcli -D` (offload agent)
- Run the mount agent on all backup servers or cloning servers:
`INSTANCE_DIR/acsgen -D -M [-s deviceclass[,deviceclass]] [-H hostname]`

Postinstallation and configuration

After you install and configure IBM Spectrum Protect Snapshot, you can set up extra backup and clone servers.

You can use the setup script to update the profile and configure IBM Spectrum Protect Snapshot on multiple backup servers from the production server when you install Open Secure Shell (OpenSSH) to enable backup servers for remote installation and configuration from the production server. NFS shares between the production server and backup server are not required for this type of remote installation.

Upgrades and reconfiguration must be run only from the master production server node.

If OpenSSH is not available, follow the instructions for “Installing separately on backup or clone servers” on page 44 and run the setup script. Choose **On-site Backup server configuration** as the configuration type. Before you run the setup script on a backup or clone server, stop IBM Spectrum Protect Snapshot on the production server. For details about how to stop an activated IBM Spectrum Protect Snapshot instance, see IBM Spectrum Protect Snapshot commands and scripts.

Typically, it is not necessary to run the setup script on the backup server after the initial configuration. Exceptions to this rule include:

- The use of alternative storage hardware might require a reconfiguration of IBM Spectrum Protect Snapshot on the backup server.

- Changes to the scheduling policy for offloaded IBM Spectrum Protect backups might require you to configure the backup server again.
- If self-signed certificates are used, all changes to the certificates require a reconfiguration of the backup server.
- If OpenSSH is not used, you must copy the `fcmselcert.arm` file to the backup server before the setup script is run to configure the backup server again.

In these cases, stop IBM Spectrum Protect Snapshot on the production server before reconfiguration of the backup server. Otherwise, you are prompted to stop IBM Spectrum Protect Snapshot on the production server.

Chapter 7. Backing up data

Instructions about how to back up data and applications using IBM Spectrum Protect Snapshot are provided.

About this task

While IBM Spectrum Protect Snapshot focuses on snapshot backups, the software can be integrated with IBM Spectrum Protect clients for offloaded backups to IBM Spectrum Protect.

Backing up Oracle databases

IBM Spectrum Protect Snapshot integrates with multiple components when you back up an Oracle database.

The following table summarizes the commands that are associated with backing up an Oracle database.

Table 6. Summary of Backup Commands for Oracle

Snapshot backup (disk only)	Back up to IBM Spectrum Protect		
	From production database (tape only)	Integrated with snapshot	From existing snapshot
acsora -f backup	RMAN by using Data Protection for Oracle	acsora -f backup ¹	fcmlcli -f tape_backup ^{2, 3}

Note:

1. The IBM Spectrum Protect Snapshot profile parameter **TSM_BACKUP** is set to one of the following options: YES, MANDATE, or LATEST, and the offload agent (tsm4acs) is running in daemon mode on the production server.
2. The profile parameter **TSM_BACKUP** is set to one of the following options: YES, MANDATE, or LATEST, and the offload agent is not running in daemon mode.
3. The **fcmlcli -f tape_backup** operation must be issued from the production system.

You must use the Oracle instance owner ID and go to the **INSTANCE_DIR** directory before you enter any backup commands. IBM Spectrum Protect Snapshot backs up the database control file and the database profile on the production server to the IBM Spectrum Protect Snapshot repository. These files are required by IBM Spectrum Protect Snapshot on the backup server to start the database instance and to run the offloaded backup to the IBM Spectrum Protect server. Furthermore, the database control file can be optionally restored during a FlashCopy restore (profile parameter **DATABASE_CONTROL_FILE_RESTORE**).

By default, the IBM Spectrum Protect Snapshot software does not back up the transaction logs of the Oracle database. The database administrator is responsible for backing up the Oracle database transaction logs. After every full database online backup, back up the transaction logs to the IBM Spectrum Protect server.

The same commands can be used to back up Oracle databases that use Oracle Automatic Storage Management (ASM) and Oracle Real Application Clusters (RAC) based on Oracle ASM. An Oracle recovery database is required for IBM

Spectrum Protect Snapshot snapshot backups. This recovery database is a requirement even if a disk only backup is requested.

Backing up an Oracle database to IBM Spectrum Protect server

IBM Spectrum Protect Snapshot require specific software when you back up an Oracle database to IBM Spectrum Protect server.

Before you back up to IBM Spectrum Protect server, ensure that the following requirements are met:

- Data Protection for Oracle must be configured on the backup server.
- A RMAN backup script that is created by the user must contain the Data Protection for Oracle *TDPO_OPTFILE* environment variable. Specify the fully qualified path name to the *tdpo.opt* options file with the *TDPO_OPTFILE* environment variable. The *allocate channel* command and the **ENV** parameter must be specified on the same line. Also, the database command must be specified on a line separate from the backup command in the RMAN backup script. See “RMAN backup script example” on page 196.

Manually backing up Oracle databases

A manual backup performs a one time backup of an Oracle database.

Procedure

1. Log on to the production server and use the database instance owner user ID.
2. Change to the *INSTANCE_DIR* directory.
3. Enter the following command:

```
acsora -f backup
```

Table 7. Files used during a manual backup

File name	Description	Default location
acsora	IBM Spectrum Protect Snapshot production system executable file	<i>INSTANCE_DIR</i>
fccli	IBM Spectrum Protect Snapshot production and backup system executable file	<i>INSTANCE_DIR</i>
profile	IBM Spectrum Protect Snapshot profile	<i>ACS_DIR</i>

Note: Using some storage systems, the snapshot backup requires a certain amount of available space on the target storage pool, so that it can create the snapshot. If there is not enough storage space available, you can increase the capacity on the requested storage pool, or free up some items that are using existing capacity. Check the message for the exact amount of storage space that is required.

Automating Oracle database backups

A fully automated backup uses a server script to fully automate online backups of Oracle databases.

Synchronous and asynchronous offloaded backups are different. Synchronous means that the backup on the backup host starts as soon as a FlashCopy is started. In this case, the offload agent, tsm4acs, is started as a daemon by the init process. Asynchronous means that offloaded backups are triggered manually or triggered by another scheduler by starting the fcmcli agent. This type of backup is useful if the backup must be delayed until the required resources are available. During the configuration on the production server, you determine when to run the offloaded backups.

Oracle Data Guard standby database backup and restore

Run offline backups of Oracle Data Guard standby database as you would backup an Oracle database. The configuration, commands, and options are the same.

The Oracle Data Guard setup is composed of a primary database and one or more standby databases. Access to the primary database is read/write, while the standby databases that are always synchronized to the primary database, have read-only access. This setup is used as a high-availability setup, in which any of the standby databases can take over the role of primary database. In a failover or switchover scenario, the standby database becomes the primary database.

Related concepts:

“Backing up Oracle databases” on page 75

“Restore Oracle Data Guard standby or primary database” on page 92

Support of Oracle Automatic Storage Management failure groups

Oracle ASM organizes data in disk groups that consist of a collection of disk drives that are in the same loop as configured by the storage subsystem. IBM Spectrum Protect Snapshot uses an ASM instance to map these disk groups to physical disks. Each disk group can have multiple failure groups that are redundant copies of each other. These failure groups can be used as a technique to mirror storage volumes.

To use this technique, define disk groups with normal redundancy that are composed of two failure groups, and place the volumes for each of the failure groups on a dedicated storage cluster. Alternatively, define disk groups with high redundancy that are composed of three failure groups. Although the default Oracle System ID (SID) for the ASM instance is +ASM, other SIDs are supported.

In such a configuration, IBM Spectrum Protect Snapshot can create FlashCopy backups of an individual failure group for all of the following supported storage devices:

- DS8000
- Storwize V7000
- SAN Volume Controller
- IBM XIV Storage System
- IBM System Storage N series

FlashCopy backup of individual mirrors

IBM Spectrum Protect Snapshot supports mirroring.

Mirroring using the AIX logical volume manager (LVM mirroring)

IBM Spectrum Protect Snapshot provides LVM mirroring support for DS8000, IBM XIV Storage System, Storwize V7000, and SAN Volume Controller. For those devices, IBM Spectrum Protect Snapshot creates a FlashCopy backup where only one of the mirrors is copied during the backup. When LVM is used to mirror the database across sites, you can create offloaded tape backups on either site with IBM Spectrum Protect Snapshot. In this situation, you do not have to transfer the backup image across sites. To complete this task, a backup server is required on either site where backup images can be mounted to transfer them to secondary backup media. For DS8000, you can create at most one INCREMENTAL FlashCopy per source volume. However, in LVM environments, each source volume is mirrored. Therefore, IBM Spectrum Protect Snapshot can create two INCREMENTAL FlashCopy backups for DS8000.

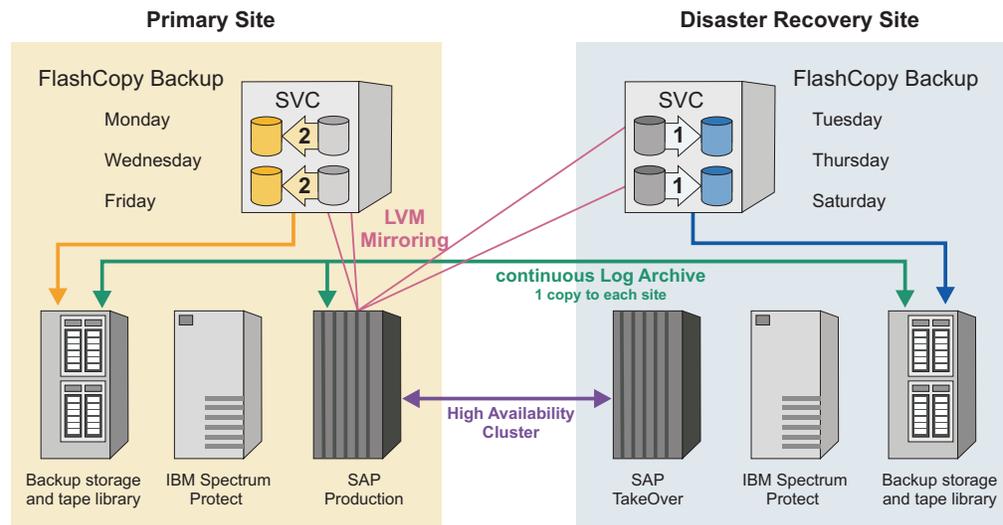


Figure 7. Cross-site mirrored SAP® database protected with IBM Spectrum Protect Snapshot and IBM Spectrum Protect.

Support of AIX enhanced concurrent capable volume groups

To support high-availability environments, IBM Spectrum Protect Snapshot supports enhanced concurrent capable volume groups.

Heterogeneous device mirroring

IBM Spectrum Protect Snapshot does not require the storage devices of different mirrors to be at the same version level.

Backing up data with remote mirroring

When you back up data with remote mirroring, you can create local and remote snapshot backups.

About this task

The local and remote snapshot backups can be created for Oracle and Oracle in an SAP environment databases.

These steps can be applied to the following scenarios:

- SAN Volume Controller snapshot backup at the auxiliary cluster with either Metro Mirror or Global Mirror.
- XIV system snapshot backup at the remote site with either Synchronous Remote Mirroring or Asynchronous Remote Mirroring.

To create local application-consistent snapshot backups with the source volumes of the system that is running remote mirroring, verify that one `DEVICE_CLASS` section is configured for the primary cluster. The production volumes are on the primary cluster. You can run the setup script to create or change `DEVICE_CLASS` sections. From the production host, start the local snapshot backup. There are no additional requirements.

To create application-consistent remote snapshot backups with the target volumes of the storage system that is running remote mirroring, complete the following steps. The first few steps do not include all details that are needed to complete the step. These steps are usually completed before you start the following procedure. The information is provided for your convenience. You can verify that you have the environment set up completely before the backup begins.

Procedure

1. Verify IBM Spectrum Protect Snapshot is installed in a supported environment. You must have a supported database that is running on the primary cluster. The primary cluster is mirrored to a remote cluster with the storage feature for remote mirroring.
2. Use the setup script wizard to configure IBM Spectrum Protect Snapshot for remote mirroring. When configuring for remote mirroring, the following parameters are set in the `DEVICE_CLASS` section:
 - `COPYSERVICES_REMOTE` YES
 - `COPYSERVICES_REMOTE_SERVERNAME` *SERVER_NAME*
 - `COPYSERVICES_REMOTE_USERNAME` *USER_NAME*
 - `TAKEOVER_HOST_NAME` *HOST_NAME*
3. At the end of the setup script wizard, the following question is displayed:
Do you want to continue by specifying passwords for the defined sections?
Enter *y* for yes.
4. Verify that the `DEVICE_CLASS` section, created for remote mirroring during the configuration process, is selected. To verify, go to the `CLIENT` section of the profile. In the `CLIENT` section, the `DEVICE_CLASS` to use is selected. When backing up data stored on Oracle in an SAP environment databases, the `DEVICE_CLASS` is specified in the SAP backup `.utl` file.
5. From the production host, start the remote snapshot backup by typing in the following command:

Oracle database, remote snapshot backup

```
acsora -f backup
```

or, Oracle in an SAP environment, remote snapshot backup

```
brbackup -p /oracle/SID/dir/initSID.sap -t online -d  
util_vol_online -c
```

When a snapshot backup is attempted, but the remote mirroring relationships are not synchronized, the backup fails and an error message is displayed. Before you can back up data, the mirroring relationships must be in the consistent synchronized state.

There is a snapshot consistency group created in the remote cluster. The target of the mirroring relationships is the source of this new snapshot consistency group.

Important: Using some storage systems, the snapshot backup requires a certain amount of available space on the target storage pool, so that it can create the snapshot. If there is not enough storage space available, you can increase the capacity on the requested storage pool, or free up some items that are using existing capacity. Check the message for the exact amount of storage space that is required.

6. To verify that the backup is complete, from a command prompt window, enter the following command:

```
fccli -f inquire_detail
```

What to do next

When you have completed the steps, you can mount and unmount the backup with the following commands:

- Mount the backup, from a command prompt window, by entering the following command: **fccli -f mount**
- Unmount the backup, from a command prompt window, by entering the following command: **fccli -f unmount**

Related reference:

“Mounting and unmounting snapshots on a secondary system” on page 185

Usability states of snapshot backup operations

To view the usability states of a snapshot backup, use the **-f inquire_detail** command option with the application-specific commands for example **fccli**, **acsora**, or **backint**.

Table 8. Usability states

Usability state value	Meaning
REMOTELY_MOUNTABLE	Backup data can be mounted from a remote system.
REPETITIVELY_RESTORABLE	Backup data can be restored. The image can be used multiple times.
DESTRUCTIVELY_RESTORABLE	Data can be restored. After the restore, other backups and possible the backup to be restored can potentially be deleted.
SWAP_RESTORABLE	Restore is possible by using the backup volumes directly rather than copying the data back to the source volumes.

Table 8. Usability states (continued)

Usability state value	Meaning
PHYSICAL_PROTECTION	The snapshot ensures protection from physical failures on the source volumes, there is no longer a dependency on the source volumes. This state does not necessarily mean that a FULL_COPY must be created with each snapshot. For example, block-level continuous data protection (CDP) mechanisms typically replicate the data only once, and then record changes only.
FULL_COPY	A full copy of the data was generated.
INCOMPLETE	A portion of the data that was backed up is deleted and can no longer be restored. This situation can happen, for example, after a partial restore of an old backup that is only DESTRUCTIVELY_RESTORABLE .
MOUNTING	A mount operation was requested on the backup server.
MOUNTED	This backup is mounted on a backup server.
DELETING	Indicates that a backup is marked for deletion. The deletion was requested.
DELETED	Indicates that the backup was deleted.
BACKGROUND_MONITOR_PENDING	Indicates that a required background copy process is not yet active or not yet finished. The device agent checks for backups with this state and monitors the associated volumes until the background copy is finished. This state is then replaced by FULL_COPY .
TAPE_BACKUP_PENDING	Indicates that a requested tape backup is not yet started or is not yet finished successfully. The offload agent checks for backups with this state, and runs the requested tape backup. After the tape backup finishes successfully, this state is reset. If the tape backup stops with an error, the TAPE_BACKUP_PENDING state remains set, TAPE_BACKUP_IN_PROGRESS is reset, and a <i>retry</i> counter is incremented.
TAPE_BACKUP_IN_PROGRESS	Indicates that the requested tape backup was started by the IBM Spectrum Protect Snapshot offload agent. If the backup fails, this state is reset.
TAPE_BACKUP_COMPLETE	Indicates that the requested tape backup is finished by the IBM Spectrum Protect Snapshot offload agent.
TAPE_BACKUP_FAILED	Indicates that the tape backup of the IBM Spectrum Protect Snapshot offload agent was not successful.
CLONE_DATABASE	Indicates that an IBM Spectrum Protect Snapshot cloning operation was run.
RESTORING	Indicates that an IBM Spectrum Protect Snapshot restore operation was run.

Usability state diagrams

The following usability state diagrams show the state changes during different operations. The green arrows are used for actions that you can start. The blue arrows are used for actions that are done automatically by IBM Spectrum Protect Snapshot. The black arrows indicate IBM Spectrum Protect Snapshot operations that you can use to change usability states.

Snapshot backup

The first state diagram shows the usability state changes during an IBM Spectrum Protect Snapshot snapshot backup operation. Depending on the storage system (DS8000, SAN Volume Controller, and XIV system, some states differ).

For example, on XIV system, the snapshot backup is immediately restorable and the restore can be repeated multiple times. On DS8000 and SAN Volume Controller the snapshot backup requires a background monitoring operation (**acsngen -D**) that removes the **BACKGROUND_MONITOR_PENDING** state and instead sets the **FULL_COPY** and **PHYSICAL_PROTECTION** state. This requirement depends on the FlashCopy type that was used for the snapshot backup. Background monitoring operations (**acsngen -D**) are automatically running.

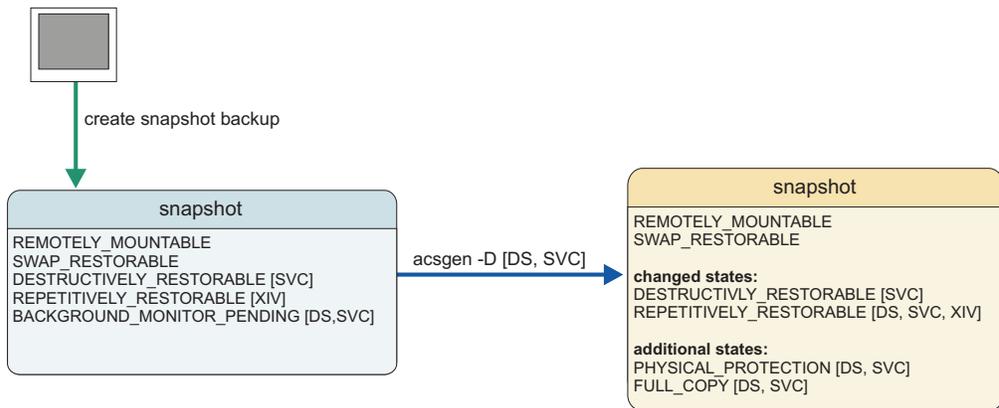


Figure 8. Usability States during snapshot backup

Snapshot restore

The second state diagram shows the usability state changes during an IBM Spectrum Protect Snapshot snapshot restore operation. On the DS8000 and SAN Volume Controller storage systems, the usability states change during a snapshot restore operation.

For DS8000 and SAN Volume Controller systems, the **BACKGROUND_MONITOR_PENDING** state is turned on and in a **RESTORING** state. The background monitor process (**acsngen -D**) resets both states when the copy process in the storage system finishes. Background monitoring operations (**acsngen -D**) are automatically running.

For XIV system there is no usability state change.

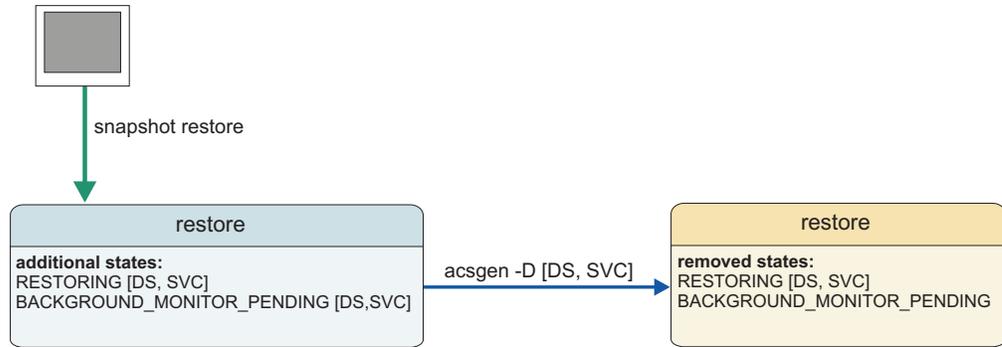


Figure 9. Usability states during snapshot restore

Snapshot delete

The next state diagram shows the usability state changes during an IBM Spectrum Protect Snapshot snapshot delete operation. There are two types of delete operations: delete and delete with force option. For both types, the snapshot backup is marked with the **DELETING** state and a background monitoring operations (**acsgen -D**), which is running automatically in background, switches the states to **DELETED**. On the XIV system, the snapshot in the XIV system, is deleted by the background monitor agent and the snapshot backup is also deleted from the IBM Spectrum Protect Snapshot repository.

On the DS8000 and SAN Volume Controller storage systems, the FlashCopy relations are not deleted by the background monitor operation unless the delete force option was used on the delete command. On the DS8000 and SAN Volume Controller systems, the snapshot backup is not deleted from the IBM Spectrum Protect Snapshot repository. Instead, a deleted snapshot backup can be reused by a new creation of a snapshot backup.

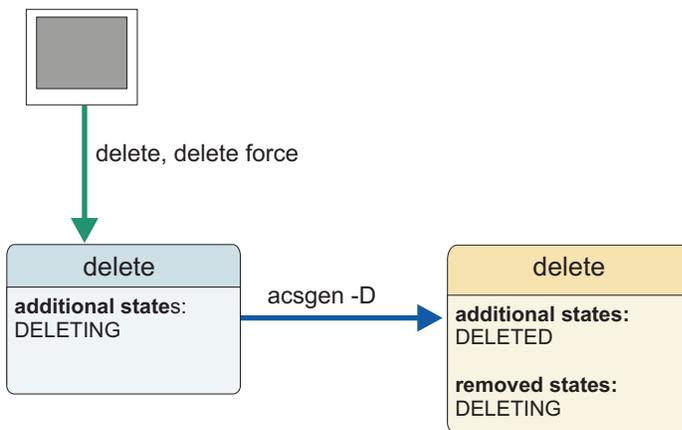


Figure 10. Usability states during snapshot delete

Snapshot mount

The next state diagram shows the usability state changes during an IBM Spectrum Protect Snapshot snapshot mount operation. You can start a snapshot mount operation by using the mount function of the command-line interface or start it automatically during the creation of a snapshot backup. In the latter case, it is named a forced mount operation. In either case, the mount operation first changes

the state to **MOUNTING**. If the mount operation finishes successfully, the state changes from **MOUNTING** to **MOUNTED**. If the mount operation fails, the state stays **MOUNTING**. The only operation that is allowed to remove a **MOUNTING** or **MOUNTED** state is a successful IBM Spectrum Protect Snapshot unmount operation. If the unmount operation finishes successfully, the **MOUNTING** or **MOUNTED** state is removed. If the unmount operation fails, the state remains as **MOUNTING** or **MOUNTED**. An unmount force operation is not needed for unmounting unless an offloaded tape backup is in progress.

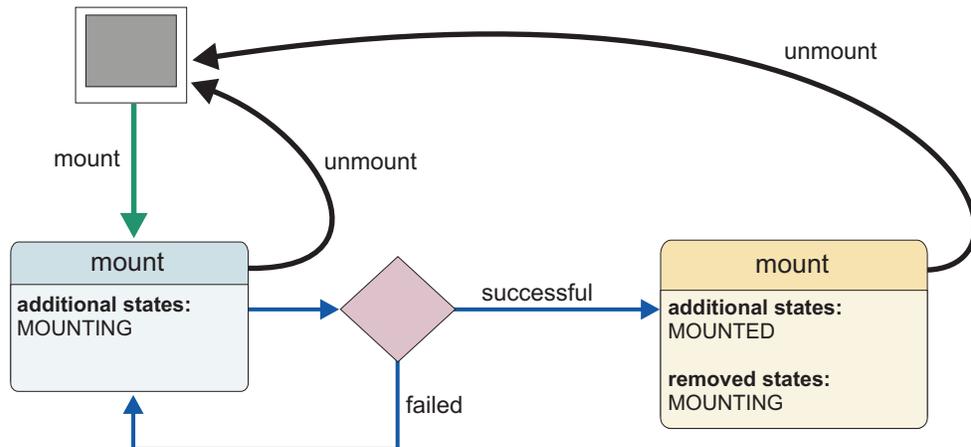


Figure 11. Usability states during snapshot mount

Snapshot offload

The last state diagram shows the usability state change during an IBM Spectrum Protect Snapshot snapshot offload operation. You can start a snapshot offload operation with the **tape_backup** function of the command-line interface. Alternatively, run it automatically with the offload agent that is running in the background (**fcml i -D**). If the snapshot backup is not already mounted successfully, a mount operation is started automatically. The mount operation changes the state first to **MOUNTING** and then to **MOUNTED**. After that or in case that the snapshot backup was already mounted, the offload operation adds the state **TAPE_BACKUP_IN_PROGRESS** and runs the offloaded tape backup. If this operation is successful, the state switches from **TAPE_BACKUP_IN_PROGRESS** to **TAPE_BACKUP_COMPLETE**. Otherwise, the **TAPE_BACKUP_IN_PROGRESS** state switches to a **TAPE_BACKUP_FAILED** state and the **TAPE_BACKUP_PENDING** state persists. In either case, the automatic unmount operation is started and the **MOUNTED** state is removed when the operation completes successfully. If the mount operation fails, or the tape backup operation stops then the **MOUNTED** or **MOUNTING** state remains. The only operation that can remove these states is a successful IBM Spectrum Protect Snapshot unmount operation. If the unmount operation finishes successfully, the **MOUNTED** or **MOUNTING** state is removed. If the unmount operation fails, the states are not removed. An unmount force operation is only needed for unmounting when an offloaded tape backup is in progress (**TAPE_BACKUP_IN_PROGRESS** is still set). The unmount force operation resets the **TAPE_BACKUP_IN_PROGRESS** state when it successfully completes the unmount operation.

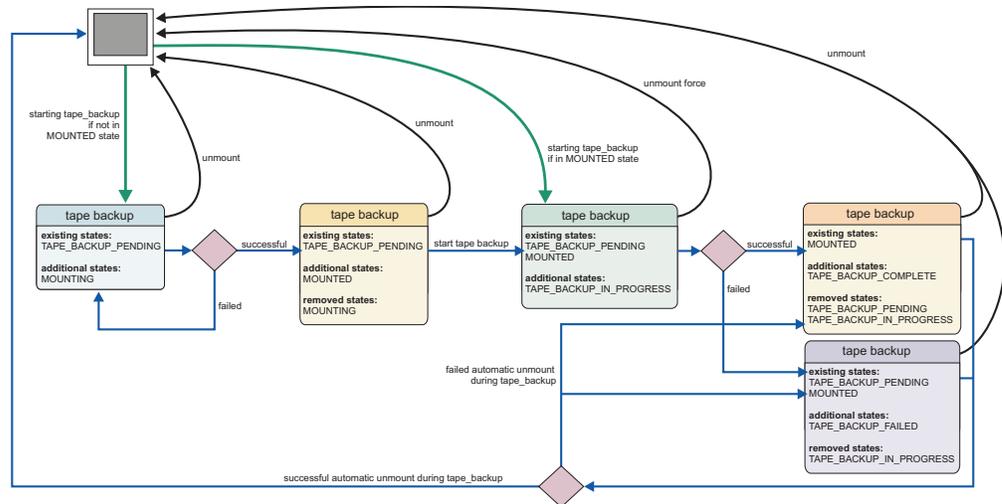


Figure 12. Usability states during snapshot offload

The usability state **TAPE_BACKUP_PENDING** can be removed by using the IBM Spectrum Protect Snapshot function **update_status** with the option **-S TSM_BACKUP=NO**. This state is also removed by starting a new snapshot backup with the option **TSM_BACKUP[_FROM_SNAPSHOT] LATEST**. This option automatically removes the usability state **TAPE_BACKUP_PENDING** from all snapshot backups that exist in the IBM Spectrum Protect Snapshot repository.

Chapter 8. Restoring Oracle databases

Restore databases with IBM Spectrum Protect Snapshot for Oracle by restoring from a snapshot on the storage subsystem, or restoring data from IBM Spectrum Protect. Specific command entries are used when restoring Oracle databases.

The following table summarizes the command entries according to the type of restore:

Table 9. Summary of Restore Commands for Native Oracle

Snapshot Restore	Restore from IBM Spectrum Protect
acsora -f restore [-b backup_ID]	Using Data Protection for Oracle, RMAN.

This section describes how to restore your Oracle database using the snapshot restore feature.

Restoring from a snapshot backup

Follow the steps to restore an Oracle database. In a case where a control file copy was stored in the data volume group or data disk group, follow the steps to recover the current copy of the control file.

Before you begin

Ensure that the following requirements are met:

- The redo logs for the database are in a disk group that is not shared with any Oracle data files.
- The Oracle RAC data files are stored on a supported storage system.

About this task

After the snapshot restore process completes, you can start the recovery of the restored RAC database. In the example that is used, the database is called MYDB.

Procedure

1. Stop the database that is to be restored. Log on to the production server with the Oracle RAC instance owner user ID.

```
srvctl stop database -d MYDB
```

To find an existing backup ID, type the following command: `acsora -f inquire`.

2. Restore a snapshot backup. Choose one of the following methods:
 - Restore the last snapshot backup by specifying the command, `acsora -f restore`
 - To restore an older snapshot backup, specify the backup ID of that backup as follows:
`acsora -f restore -b backup ID`

where the *backup ID* is the ID of the snapshot backup that you want to restore.

3. After the snapshot restore process completes, you must recover the database. Choose one of the following recovery types: The recovery type is specified in the profile file. For more information about recovery types, see **DATABASE_CONTROL_FILE_RESTORE** in "ORACLE" on page 124.
 - If **DATABASE_CONTROL_FILE_RESTORE YES** is specified in the IBM Spectrum Protect Snapshot profile file, an incomplete recovery is needed.
 - If **DATABASE_CONTROL_FILE_RESTORE NO** is specified in the IBM Spectrum Protect Snapshot profile file, a complete recovery is needed.
- 4.

Tip: The steps for advanced recovery follow:

If a control file exists in the data volume group or data disk group, and it was not restored, there will be inconsistent copies of the control file after the snapshot restore operation finishes. To restore the control file copy that was overwritten by the restore operation from a current copy, follow these steps:

- a. Start the instance without mounting the database by entering the following command:


```
SQL> startup nomount
```
- b. Use RMAN to restore a current copy of the control file that is not part of the data disk group, and is not restored.


```
RMAN> restore controlfile from '<control_file_copy_name>'
```
- c. Mount the database, and recover the database for a complete recovery.


```
SQL> alter database mount;
SQL> recover database;
```
- d. Open the database if you want to continue working with the fully recovered database.


```
SQL> alter database open;
```

Results

If you restored the control file from a current copy, the control file in the data disk group is overwritten to that current version.

Important: On some storage system types, the snapshot restore requires sufficient available space on the target storage pool in order to restore the necessary volume. If there is not enough storage space available, you can increase the capacity on the requested storage pool, or free up some capacity. Check the message for the exact amount of storage space that is required.

Restoring Oracle Real Application Clusters database snapshots

Follow the steps to restore an Oracle Restoring Oracle Real Application Clusters (RAC) database. In the example that is used in this task, an Oracle RAC database snapshot with two instances is restored.

Before you begin

Ensure that the following requirements are met:

- The redo logs for the database are in a disk group that is not shared with any Oracle data files.
- The Oracle RAC data files are stored on a supported storage system.

About this task

After the snapshot restore process completes, you can start the recovery of the restored RAC database. In the example that is used, the database is called MYDB, with two instances that are called MYDB1 and MYDB2. The procedure is for a complete recovery. For incomplete recoveries, refer to the Oracle documentation.

Procedure

1. Stop the database that is to be restored. Log on to the production server with the Oracle RAC instance owner user ID.

```
srvctl stop database -d MYDB
```
2. Restore a snapshot backup. Choose one of the following methods:
 - Restore the last snapshot backup by specifying the command,

```
acsora -f restore
```
 - To restore an older snapshot backup, specify the backup ID of that backup as follows:

```
acsora -f restore -b backup ID
```

where the *backup ID* is the ID of the snapshot backup that you want to restore.

To find an existing backup ID, type the following command: `acsora -f inquire`.

3. After the snapshot restore processing completes, you must recover the database. The recovery type is specified in the profile file. For more information about recovery types, see **DATABASE_CONTROL_FILE_RESTORE** in "ORACLE" on page 124.
 - If **DATABASE_CONTROL_FILE_RESTORE YES** is specified in the IBM Spectrum Protect Snapshot profile file, an incomplete recovery is needed.
 - If **DATABASE_CONTROL_FILE_RESTORE NO** is specified in the IBM Spectrum Protect Snapshot profile file, a complete recovery is needed.

Important: For all scenarios using snapshot technology with IBM Spectrum Protect Snapshot and using certain storage systems, the snapshot restore requires available space on the target storage pool so that it can restore the necessary volume. If there is not enough storage space available, you can increase the capacity on the requested storage pool, or free up some items that are using existing capacity. Check the message for the exact amount of storage space that is required.

For a complete recovery of the database, follow these steps:

- a. Mount the database on the database instance where you want to run the recovery.

```
srvctl start instance -d MYDB -i MYDB001 -o mount
```

If you want an incomplete recovery, do not proceed to the next sub step, see your Oracle documentation for information about incomplete recoveries.
- b. Recover and open the database to complete the recovery operation.

```
SQL> recover database; SQL> alter database open;
```
- c. Start all other database instances by entering the following command:

```
srvctl start instance -d MYDB -i MYDB002
```
- d. If you used normal or high redundancy, synchronize the failure groups that were not included in the restore operation.

Restoring Oracle databases from IBM Spectrum Protect

IBM Spectrum Protect backups are restored as an entire database (Restore Method One) or with datafile granularity (Restore Method Two). RMAN must be used to perform restore procedures.

Restoring the entire database

Follow these tasks to restore IBM Spectrum Protect backups as an entire database.

Procedure

To restore an entire database backup, complete the following steps:

1. If the database is running, enter the following command to stop the database:
shutdown;
2. Enter the following command to mount the database:
startup mount;
3. Enter the following command to start RMAN and connect to the target database and the recovery catalog:
rman target username/password rcvcat username/password@connect_string
4. Issue the RMAN **run** command by specifying the allocation of channels and the restoration of the database. The following example is from an AIX installation:

```
run
{
allocate channel t1 type 'sbt_tape' parms
'ENV=(TDPO_OPTFILE=/usr/tivoli/tsm/client/oracle/bin/tdpo.opt)';
allocate channel t2 type 'sbt_tape' parms
'ENV=(TDPO_OPTFILE=/usr/tivoli/tsm/client/oracle/bin/tdpo.opt)';
allocate channel t3 type 'sbt_tape' parms
'ENV=(TDPO_OPTFILE=/usr/tivoli/tsm/client/oracle/bin/tdpo.opt)';
allocate channel t4 type 'sbt_tape' parms
'ENV=(TDPO_OPTFILE=/usr/tivoli/tsm/client/oracle/bin/tdpo.opt)';
restore database;
}
```

The previous example also applies to Linux, Solaris, and HP-UX, except that the path to the `tdpo.opt` file differs. On Linux, Solaris and HP-UX, the path is likely to start with `/opt/tivoli`.

5. Recover the database by entering the following command to connect to the target database:
recover database;

What to do next

If your restore is not successful and you receive an error message, see the default `tdpoerror.log` error log file for assistance.

Restoring files

To restore IBM Spectrum Protect Snapshot backups by restoring specific files, complete the following steps.

Procedure

1. If the database is running, enter the following command to stop the database:

```
shutdown;
```

2. Enter the following command to mount the database:

```
startup mount;
```

3. Enter the following command to start RMAN and connect to the target database and the recovery catalog. Enter the command on one line.

```
rman target username/password rcvcat username  
/password@connect_string
```

The RMAN command in the preceding example is divided to accommodate page formatting. The actual RMAN command string is on one line.

4. Issue a RMAN **run** command by specifying the allocation of channels and the restoration of the data file *n*, where *n* is the number of the data file. The following example is from an AIX installation:

```
run  
{  
  allocate channel t1 type 'sbt_tape' parms  
  'ENV=(TDPO_OPTFILE=/usr/tivoli/tsm/client/oracle/bin/tdpo.opt)';  
  allocate channel t2 type 'sbt_tape' parms  
  'ENV=(TDPO_OPTFILE=/usr/tivoli/tsm/client/oracle/bin/tdpo.opt)';  
  allocate channel t3 type 'sbt_tape' parms  
  'ENV=(TDPO_OPTFILE=/usr/tivoli/tsm/client/oracle/bin/tdpo.opt)';  
  allocate channel t4 type 'sbt_tape' parms  
  'ENV=(TDPO_OPTFILE=/usr/tivoli/tsm/client/oracle/bin/tdpo.opt)';  
  restore datafile n;  
}
```

This example can also apply to Linux, Solaris and HP-UX platforms, except that the path to the `tdpo.opt` file differs. On Linux, Solaris and HP-UX platforms, the path is likely to start with `/opt/tivoli`.

5. Enter the following SQL command to bring the data file online. The *n* variable is the number of the data file.

```
alter database datafile n online;
```

6. Recover the data file by connecting to the target database and entering the following command:

```
recover datafile n;
```

What to do next

If your restore is not successful and you receive an error message, see the error log file for assistance.

Restore Oracle Data Guard standby or primary database

When you are restoring an Oracle Data Guard database, there are two scenarios to choose from that depend on if a failover or switchover between the primary and the standby database took place. After a switchover or failover, when the standby database is changed to the primary database, the value of the **DATABASE_CONTROL_FILE_RESTORE** allows you to choose from two recovery settings.

In the first scenario, the standby database remains in the standby database role. When that standby database is restored it results in a standby database with the point in time of the IBM Spectrum Protect Snapshot backup. The synchronization with the Oracle Data Guard primary database can be restarted. In order for the synchronization to succeed, all the required archive logs must be available on the primary database.

In the second scenario, if the standby database was changed to the primary database, choose one of the following settings:

- **DATABASE_CONTROL_FILE_RESTORE** parameter set to NO. A database recovery is required after the restore operation. In this case, the archived logs are required. This action results in an Oracle Data Guard primary database recovered to a specified point in time.
- **DATABASE_CONTROL_FILE_RESTORE** parameter set to YES. A database recovery is not possible after the restore operation. The result of the restore operation is an Oracle Data Guard standby database with the point in time of the IBM Spectrum Protect Snapshot backup.

Related concepts:

“Oracle Data Guard standby database backup and restore” on page 77

Related reference:

Chapter 8, “Restoring Oracle databases,” on page 87

Restoring data with remote mirroring

Restore data on a remote site with IBM Spectrum Protect Snapshot.

Before you begin

The restore operations for the remote site must meet the following environment conditions:

- Data is successfully backed up and the backup copy of data is accessible on the remote site.
- A takeover host is running with the same operating system level as the production host.
- The takeover host is configured on the remote side.
- The database instance is created on the takeover host.
- IBM Spectrum Protect Snapshot software is installed on the takeover host. The software level on the production host and on the takeover host are the same.

Note: Never edit the existing **DEVICE_CLASS** parameters in the profile. For takeover operations, always add new **DEVICE_CLASSES** for the new local and new remote sites.

Restoring Oracle data with remote mirroring

About this task

The takeover operation is complete, and the reversal of roles and remote relationships are already in place. If not already included in the takeover operation, stop the acsd daemon on the primary production host, and transfer all the repository files from the primary production host to the takeover host. The repository files are in the directory defined by the parameter **ACS_REPOSITORY** in the ACS section of the profile.

Note: The snapshot restore operation requires sufficient available space on the target storage pool so that it can restore the necessary volume. Increase the capacity on the requested storage pool or free up some items that are using existing capacity in cases where there is insufficient space.

The database is restored to the takeover host by using the IBM Spectrum Protect Snapshot snapshot local repository after the remote backup. Complete the following steps:

Procedure

1. Update the IBM Spectrum Protect Snapshot configuration parameters with the setup script wizard. Specifically, set the **ACSD** parameter to use the acsd on the takeover host in the GLOBAL section. Do not use the acsd of the production host.
2. Start the IBM Spectrum Protect Snapshot acsd daemon on the takeover host.
3. From the backups that are displayed, select the remote backup to use for the restore. The backups are displayed after you enter the query command on the takeover host. For example, **acsora -f inquire_detail**.
4. Start the restore by entering the following command:

```
acsora -f restore -b <backup_id>
```

Results

The remote mirroring relationships are stopped. The volume groups with the file systems that contain the table spaces are restored from the FlashCopy targets to the remote mirroring targets. The file systems that contain the table spaces are mounted.

You must restart the remote relationships before taking another snapshot of remote mirroring targets. For IBM XIV Storage System, the remote relationships are removed. You must re-create the remote relationships before taking another snapshot of remote mirroring targets.

Chapter 9. Cloning databases

IBM Spectrum Protect Snapshot uses the FlashCopy or snapshot function of the storage solutions for database cloning. This method eliminates downtime and minimizes the impact on the production database.

For FlashCopy backup, the physical volume identification numbers (PVIDs) are not changed. For FlashCopy cloning, the PVIDs of the FlashCopy target disks are automatically changed by the IBM Spectrum Protect Snapshot software. You can have several cloned databases of one source database that are running on one host. Each clone target requires a unique database name.

With IBM Spectrum Protect Snapshot, a cloning process can be started with an online or offline source database. For online IBM Spectrum Protect Snapshot cloning, the source database is suspended for a short time. The suspension occurs when the storage system creates its FlashCopy or snapshot of the source database.

The cloned database (target database) can have the same database name as the source database. The cloned database can also be renamed to any valid database name during the IBM Spectrum Protect Snapshot cloning process. Each clone target requires a unique database name. IBM Spectrum Protect Snapshot requires the cloned database to be created on a different database server than the source database server regardless of whether the clone database name is changed.

Archiving logs on cloned databases

Note: Do not run any queries or operations on the Oracle database and ASM instance on the clone system while IBM Spectrum Protect Snapshot clone operations are running.

When the source database cloned is in ARCHIVELOG mode, the resulting cloned database will not be configured as ARCHIVELOG but as NOARCHIVELOG mode. If you want to archive redo logs on the cloned database you must configure the database to enable archive logging, and specify a valid path to store the logs.

Cloning and IBM System Storage SAN Volume Controller

When you clone databases and use IBM System Storage SAN Volume Controller, the space-efficient disks can be used as a target for FlashCopy cloning operations. However, when you use SAN Volume Controller space-efficient disks as a target for FlashCopy cloning, there are restrictions on the FlashCopy backups. You can complete cloning operations from the cloning source volumes. If you want to complete FlashCopy backup and FlashCopy cloning from the same source disks, use full target disks.

To use SAN Volume Controller space-efficient disks, in the **DEVICE_CLASS** that is used for cloning operations, set the **ALLOW_NOCOPY_FLASHCOPY** parameter to YES.

Cloning databases with IBM Spectrum Protect Snapshot

Create a database clone with IBM Spectrum Protect Snapshot using the `fccli -f create_clone` command or the `fccli -f refresh_clone` command.

When you enter one of the commands to create or refresh a clone, the following processing occurs:

1. The selected preprocessing scripts are run, including stopping the clone database. This step only occurs when using the `refresh_clone` command with the `-X pre-processing_configuration_file` option.
2. The FlashCopy clone is unmounted on the clone system. This step occurs only when using `refresh_clone` function.
3. A new FlashCopy clone is created, including the suspension and resumption of the source database, and mounted on the clone system.
4. The cloned database is recovered.
5. The cloned database is renamed to the target database name.
6. IBM Spectrum Protect Snapshot starts the cloned database.
7. The selected postprocessing scripts are run to clean up the clone database. This step occurs only when the `-Y post-processing_configuration_file` option is used.

Cloning Oracle RAC databases

After you share the FlashCopy repository, configuration and binary files, an Oracle RAC database clone can be created from every node of the Oracle RAC cluster. IBM Spectrum Protect Snapshot can operate on the clone from all RAC nodes, so that a clone created on one RAC node can be inquired, refreshed, or deleted from any other node within the RAC cluster.

Before you begin

For information about how to share the FlashCopy repository, configuration and binary files, see “Preparing Oracle RAC” on page 35.

About this task

There are multiple options for the clone system in an Oracle RAC multi-node architecture. While the production system is a multi-node RAC database, the installation options for the clone system include the following options:

- A single-node non-RAC instance
- A single-node RAC instance

To use one of the production system cluster nodes as a clone target system is not allowed. The ASM disk groups of the production system instance use the same names as the disk groups in the snapshots. This naming clash leads to a conflict.

Use the following procedure to clone Oracle RAC databases.

Procedure

1. Choose to the preprocessing scripts for cloning when you are creating or refreshing a clone.
2. Type in one of the following commands to clone the database:
 - `-f create_clone` if you are creating a clone.

- -f refresh_clone if you are refreshing a clone.
3. Type in one of the following commands to work with the clone:
- -f inquire_clone if you are querying a clone.
 - -f delete_clone if you are deleting a clone.

What to do next

IBM Spectrum Protect Snapshot uses a PFILE for initialization on the single-node RAC clone target server. If you intend to extend the clone target server by more nodes, you must create a shared SPFILE for all nodes of the clone target RAC cluster.

Note, in an Oracle RAC cluster, all instances must use the same SPFILE. Each local `init<ORACLE_SID>.ora` file contains only one entry that points to a shared server parameter file. The following instance shows an example:

```
SPFILE='+disk_group_name/dbunique_name/spfiledbname.ora'
```

For more information about adding RAC nodes to the created clone database, see the Oracle information here http://docs.oracle.com/cd/E11882_01/rac.112/e41960/adddelunix.htm.

Related concepts:

“Configuration files used for cloning” on page 98

Related reference:

“Cloning commands” on page 169

Oracle instance names on the clone server for Oracle RAC

IBM Spectrum Protect Snapshot always creates an Oracle instance numbered 1 on the clone target server. The Oracle instance number from where the command was entered does not change this numbering system.

Examples

In the following case, a database that is named ORR was cloned from an instance that is named ORR3, to a target database named CLO. IBM Spectrum Protect Snapshot creates an Oracle instance with SID CLO1 on the clone target system.

For a policy managed RAC database that is named ORR, with an instance named ORR_3, IBM Spectrum Protect Snapshot creates an Oracle instance with SID CLO_1 on the clone target system.

Cloning an Oracle Data Guard standby database

IBM Spectrum Protect Snapshot can clone an Oracle Data Guard standby database.

To clone an Oracle Data Guard standby database, on the Oracle Data Guard standby server complete the following steps:

- Shut down the Oracle Data Guard standby instance.
- Log on to the Oracle Data Guard standby server and use the Oracle instance user ID.
- To create a clone, enter the following command:

```
./fcmcli -f create_clone -C CLONE_ORACLE_SID -u CLONE_ORACLE_USER
```

This clone is a stand-alone Oracle database that can be opened in read and write mode. You can use the cloning command optional function `-X preprocessing-configuration-filename` and `-Y postprocessing-configuration-filename` to run preprocessing and postprocessing scripts on the cloned target system. You can use shell or sql scripts.

Tip: If the IBM Spectrum Protect Snapshot profile parameter `OVERWRITE_DATABASE_PARAMETER_FILE` is set to YES, the `initDBSID.ora` file is copied from the Oracle Data Guard standby server to the clone target system. This file contains Data Guard configuration information that is not required by the clone instance and can cause problems. Therefore, create a customized `initDBSID.ora` file for the Oracle clone instance and set the IBM Spectrum Protect Snapshot `OVERWRITE_DATABASE_PARAMETER_FILE` profile parameter to NO. You can copy the `initDBSID.ora` file and remove any specific Data Guard configuration information.

Typically Oracle temporary tablespace files are not present in the Data Guard standby server. If the temporary tablespace files are on the same file systems as the Oracle data files, no additional configuration is required by IBM Spectrum Protect Snapshot. If the temporary tablespace files are on a dedicated file system or volume group, use the `FLASH_DIR_LIST` parameter in the cloning section of IBM Spectrum Protect Snapshot profile to include these files. Use the `FLASH_DIR_LIST` parameter to specify a fully qualified directory name and file name. This file contains the mount points where the Oracle temporary files are located. Use a separate line for each mount point. Then, IBM Spectrum Protect Snapshot includes the mount points and the corresponding volume groups in the FlashCopy cloning operation.

Database cloning preprocessing and postprocessing

Repetitive processing steps that occur before and after database cloning can be automated by scripts.

The required functions in the automated scripts depend on the cloning environment. Because all possible environments cannot be covered by one package, preprocessing and postprocessing must be considered outside the scope of IBM Spectrum Protect Snapshot cloning.

IBM Spectrum Protect Snapshot provides a framework in which you can run shell scripts and component scripts on the clone system. Run the shell scripts before a clone database is unmounted and after a new clone database is created. Then, you can fully automate the cloning process.

Configuration files used for cloning

IBM Spectrum Protect Snapshot uses preprocessing and postprocessing configuration files during cloning operations. The functions that are provided by the processing scripts depend on the cloning environment where they are issued.

All processing configuration files and the scripts that are defined in the configuration files must meet the following requirements:

- Files and scripts are stored on the clone system.
- Files and scripts have permissions for read and write access for the clone database instance owner. The preprocessing and postprocessing scripts have permissions for read and write access for the user who updates and runs the

scripts. If the scripts are run by any user registered on the system, the scripts are owned by the root user. The root user has permission to read and write for the User, Group, and World user groups.

- Files and scripts have permission for read access for the production database instance owner.

Attention: If a write access level for the World user group is given, there is a security risk.

An example of a preprocessing configuration file for Oracle is: `/oracle/P01/acs/preprocessing.ini`. When adding processing configuration files, place each script on a separate line as shown in the following example:

```
/oracle/P01/acs/scripts/PreProcessing_stopsap.sh
/oracle/P01/acs/scripts/PreProcessing_stopdb.sh
```

Both processing configuration files support embedded user comments. A comment line in the configuration file is denoted by the number sign character: `#`. The scripts are specified with fully qualified file names. Each line of the processing configuration file represents one processing script. The IBM Spectrum Protect Snapshot Offload Agent, `tsm4acs`, uses these arguments and their values when calling the scripts:

DBNAME_PROD

The database name on the production system.

DBNAME_CLONE

The database name on the cloning system.

DBHOST_PROD

The host name of the production system.

DBHOST_CLONE

The host name of the cloning system.

CLONE_TIMESTAMP

The timestamp when the clone was created. This entry is also the time when the production database is suspended and the FlashCopy operation begins. The timestamp format is `YYYYMMDDhhmmss`. During preprocessing, the timestamp identifies when the previous FlashCopy clone is created. During postprocessing, the timestamp identifies when the current FlashCopy clone was created.

SCHEMA

The database schema of the production database as specified by the profile parameter **DATABASE_SCHEMA**. Depending on SAP® Kernel release, this schema is `SAPR3` or `SAPDBname`.

You can use the following processing scripts:

- SQL scripts with the extension `.sql`.
- Shell scripts with the extension `.sh`. Shell scripts can be started by a database user who is different from the clone database user. For example, when installing the SAP license for the cloned SAP system, start the postprocessing shell script as the SAP administration user `sidadm`:

```
scripts/PostProcessing_saplicense.sh:c01adm
```

By adding `:c01adm` to the script file name, the script runs as user `c01adm` instead of user `orac01`. This addition requires that the owner of the script be the same as the user who is the intended operator of the script. In this example, `c01adm` is

the owner of the script. There is one exception. If a preprocessing or postprocessing script is owned by the root user, the script can be run by any user registered on the system.

The processing scripts that are defined in the processing configuration files run sequentially. The return code of each script is validated. The following values are used:

- RC=0** Processing ends successfully. If this script is the last script to be run, continue cloning. If this script is not the last script, continue with the next script.
- RC=1** Processing ends successfully with warning. If this script is the last script to be run, continue cloning. If this script is not the last script, continue with the next script.
- RC=2** Processing ends with an error. Cloning immediately stops. No additional scripts run.

The return code for each script is written to the cloning log files. The output is written to dedicated log files with the following file names:

```
clone_preproc.<timestamp>  
clone_postproc.<timestamp>
```

Cloning processing example

An example of a cloning configuration file, showing the production database named P01, and the clone database named C01.

```
./fcmcli -f preproc_clone -u oraclerc01 -C C01 -X /oracle/C01/acs/preprocessing.ini  
./fcmcli -f postproc_clone -u oraclerc01 -C C01 -Y /oracle/C01/acs/postprocessing.ini
```

If a processing script needs extra command-line options, add these options to each line of the configuration file. In this example, the additional command-line argument LC01 is added to the script entry in the configuration file:

```
/oracle/C01/acs/scripts/PostProcessing_startListener.sh LC01
```

The IBM Spectrum Protect Snapshot command-line interface issues a call to the processing script with the six default arguments. After these arguments are provided, extra command-line options are sent. In this example, the additional command-line argument LC01 is passed to the PostProcessing_startListener.sh script as the seventh argument:

```
#!/bin/ksh  
# FOLLOWING ACTIONS ARE PERFORMED -----  
# start the Oracle Listener  
  
DBNAME_PROD=$1  
DBNAME_CLONE=$2  
DBHOST_PROD=$3  
DBHOST_CLONE=$4  
CLONE_TIMESTAMP=$5  
  
# ${SCHEMA} is schema owner (for SAP Kernel > 6.10, for userid other than SAPR3)  
SCHEMA=$6  
SCHEMA=$(echo ${SCHEMA} | tr [a-z] [A-Z])  
  
# ${LISTENER} is the name of the listener to be started (taken from listener.ora)  
LISTENER=$7  
  
lsnrctl start ${LISTENER}
```

Chapter 10. Troubleshooting

There are multiple resources for support.

The following list identifies the various ways that you can find information online:

- IBM Spectrum Protect Snapshot wiki on the developerWorks® site.
- Service Management Connect site.
- IBM Spectrum Protect Snapshot product support. Enter the search term, such as an authorized program analysis report (APAR) number, release level, or operating system to narrow the search criteria for your support need.

General troubleshooting procedure

This procedure is valid for all IBM Spectrum Protect Snapshot applications.

The starting point for problem determination is the summary log file located in the <ACS_DIR>/logs directory. The summary log file name is summary.<timestamp>.log where <timestamp> is an entry that represents the four-digit year, month, and day (for example, summary.20090817.log). A new log file is created each day. This file contains a list of all operations and the most important messages. Each line begins with one of these prefixes to indicate the type of operation:

Table 10. Message prefixes used in the summary log file

Prefix	Operation
GEN	Generic message
DB	Database backup or restore; inquire or delete of FlashCopy backups
MON	Monitoring of the background copy that is performed by the storage device
TSM	Off-loaded backup to IBM Spectrum Protect
MNT	Mount and unmount services
CLO	FlashCopy cloning operations

The summary log file only contains the information about operations that were performed and whether they completed successfully. Error messages are also logged when they occur. A dedicated log file is created for each operation in the <ACS_DIR>/logs/details. These files should be checked for detailed information when an error occurs.

This summary log file example shows a FlashCopy backup of a database. Messages with the DB prefix are issued by the database client. This is the application that requests the backup operation.

```
GEN 00:10:00 (70a)
=====
New backup operation started for database instance db2h51, database H51.

=====
DB 00:10:00 (70a) FMM1510I New connection received.
DB 00:10:00 (70a) FMM1513I *****> Database client connected: db2s95, database S95,
partition NODE0000
DB 00:10:00 (70a) FMM1574I Backup for db2s95.S95.DEVICE_CLASS:STANDARD.NODE0000 is
created using DEVICE_CLASS
DEVICE_CLASS:STANDARD.
```

```

DB 00:10:01 (80c) FMM1510I New connection received.
DB 00:10:01 (80c) FMM1514I *****> Device client connected.
DB 00:10:01 (80c) FMM6219I Backup to TSM: NO
DB 00:10:01 (80c) FMM1582I The target set 1 will be used for the current backup.
DB 00:10:44 (70a) FMM1014I Operation backup completed successful.
GEN 00:12:28 (70e)
=====

```

Logging and tracing files

Log and trace files are updated during IBM Spectrum Protect Snapshot operations.

Log and trace files are written to during backup and restore processing by these products:

- IBM Spectrum Protect Snapshot
- Storage system
- CIM
- General Parallel File System (GPFS) for IBM Spectrum Protect Snapshot for Custom Applications.
- IBM Spectrum Protect for ERP
- Operating system

Log files and trace files

Refer to these examples of the log and trace files that are maintained by IBM Spectrum Protect Snapshot.

IBM Spectrum Protect Snapshot document each operation in log files. In addition, trace files can be requested with the TRACE parameter in the profile. Do not activate tracing unless requested by IBM Support. If TRACE is set to YES, each IBM Spectrum Protect Snapshot component creates an extra trace file in the log directory.

Tip: Ensure to look for, and manage the amount of free space of the file system that contains the ACS_DIR/logs directory.

The following tables list the log and trace files that are maintained by IBM Spectrum Protect Snapshot. These files are in ACS_DIR/logs.

Table 11. IBM Spectrum Protect Snapshot log files

Purpose	File
Overview of operations and their result.	summary.timestamp.log
Overview about the monitoring of the background copy that is done by the storage device.	monitor.timestamp.log
Detailed log of a particular operation.	details/function.longtimestamp
<p>Note:</p> <ul style="list-style-type: none"> • <i>timestamp</i> is the date (yyyymmdd) • <i>longtimestamp</i> is the date and time (yyyymmddHHMMSS) • <i>function</i> is a value of backup, restore, inquire, delete, mount, unmount, tsm, or clone <p>The summary log file is always used as an entry point. All major events, such as the start of a new operation or errors, are recorded in this file. A new summary log file is created for every day and records all operations of one day within a single file.</p>	

Table 12. IBM Spectrum Protect Snapshot trace files.

Component	File
Management Agent (acsd)	acsd. <i>id</i> .trace
Application client (for DB2, the Snapshot Backup Library)	client. <i>instance.db name.node.id</i> .trace
Generic Device Agent (acsgen)	acsgen. <i>hostname.device class.node num.id</i> .trace acsgen. <i>hostname.function.id</i> .trace acsgend. <i>hostname.id</i> .trace
Device Agent for IBM XIV [®] Storage System Devices	xivadapter_ <i>id</i> _function.trace
Device Agent for CIM Devices (DS8000, SAN Volume Controller, Storwize V7000)	fmcima. <i>hostname.function.id</i> .trace fmcima. <i>hostname.device class.node num.id</i> .trace
Offload Agent (tsm4acs)	tsm4acs. <i>host.id</i> .trace
fcmcli	fcmcli. <i>host.id</i> .trace
RMAN (when started by IBM Spectrum Protect Snapshot)	rman. <i>SID.id</i> .log
<p>Notes:</p> <ul style="list-style-type: none"> • Names ending in -d are daemon processes (started with -D option). • <i>id</i> is the date (yyyymmdd) for log files written by daemon processes, date, and process ID (yyyymmdd.xxxxxx) for trace files written by daemon processes or a timestamp (yyyymmddHHMMSS) for log and trace files for other processes. • <i>device class</i> can be a device class specified in the profile or all if no command-line parameter -s device class was specified for the device agent. It can also be omitted for traces of the device agent. • <i>instance</i> and <i>db hostname</i> can be <i>undef</i> for query and delete requests that are started with db2acsutil. • <i>node num</i> is the DB2 partition number in the case of DB2 and SAP with DB2. It is 0 for Oracle and SAP with Oracle or it can also be omitted for Oracle and SAP with Oracle. • <i>function</i> is backup, delete, restore, mount, unmount, or reconcile. 	

Table 13. IBM Spectrum Protect Snapshot return codes.

Reason code	Explanation	User response
0	Operation is successful	None
1	Operation terminated successfully with warnings	The IBM Spectrum Protect Snapshot operation was successful but warning messages were reported. Check the IBM Spectrum Protect Snapshot summary log file and the therein referenced detail log files for more information.
2	Operation terminated with error	The IBM Spectrum Protect Snapshot operation failed. Check the IBM Spectrum Protect Snapshot summary log file and the therein referenced detail log files for more information.

Table 14. IBM Spectrum Protect Snapshot installer exit codes.

Exit Code	Explanation	User Response
0	The operation completed successfully	The installation completed successfully without any warnings or errors.
1	The operation completed successfully with warnings.	The installation completed successfully, but one or more of the actions from the installation sequence caused a warning or a non-fatal error. See the IBM Spectrum Protect Snapshot installer log file installation.log in the installation directory for details.
-1	The operation terminated with error	One or more of the actions from the installation sequence caused a unrecoverable error. See the IBM Spectrum Protect Snapshot installer log file installation.log in the installation directory for details.
>=1000	The operation terminated with error Note: There more error codes with numbers greater than or equal to 1000 which all mean that some kind of error occurred.	One or more of the actions from the installation sequence caused a unrecoverable error. See the IBM Spectrum Protect Snapshot installer log file installation.log in the installation directory for details.

Table 15. DB2 vendor reason codes.

Reason Code	Explanation	User Response
0	The operation is successful.	None
2	Communication error with device	The IBM Spectrum Protect Snapshot operation failed. Check the db2diag.log and the IBM Spectrum Protect Snapshot summary log file for details.
3	The DB2 and vendor products are incompatible	The IBM Spectrum Protect Snapshot operation failed during initialization of the IBM Spectrum Protect Snapshot vendor library. The DB2 API version does not match the IBM Spectrum Protect Snapshot vendor library version. Check the db2diag.log for details.
6	Object specified cannot be found	The IBM Spectrum Protect Snapshot operation failed because the requested object cannot be found in the IBM Spectrum Protect Snapshot repository. Check the db2diag.log and the IBM Spectrum Protect Snapshot summary log file for details.
8	Invalid user ID specified	The IBM Spectrum Protect Snapshot operation failed because an invalid user ID was specified on the db2 command line. Check the db2diag.log.

Table 15. DB2 vendor reason codes (continued).

Reason Code	Explanation	User Response
9	Invalid password provided	The IBM Spectrum Protect Snapshot operation failed because an invalid password was specified on the db2 command line. Check the db2diag.log.
10	Invalid options specified	The IBM Spectrum Protect Snapshot operation failed because an invalid db2 command-line option was specified. Check the db2diag.log.
11	Initialization failed	The IBM Spectrum Protect Snapshot operation failed because the IBM Spectrum Protect Snapshot vendor library cannot be initialized. Check the db2diag.log and the IBM Spectrum Protect Snapshot summary log file for details.
14	End of data reached	Not an error condition.
18	Device error	The IBM Spectrum Protect Snapshot operation failed. Check the IBM Spectrum Protect Snapshot summary log file for details.
19	Warning	The IBM Spectrum Protect Snapshot operation is successful with warning messages. Check the IBM Spectrum Protect Snapshot summary log file for details.
21	More data to come	Not an error condition.
26	Delete object fails	The IBM Spectrum Protect Snapshot delete operation failed. Check the IBM Spectrum Protect Snapshot summary log file for details.
29	Abort request failed	The IBM Spectrum Protect Snapshot abort request failed. Check the IBM Spectrum Protect Snapshot summary log file for details.
30	Unexpected Error	The IBM Spectrum Protect Snapshot operation failed. Check the IBM Spectrum Protect Snapshot summary log file for details.
31	No data has been returned	Not an error condition.
32	Object not under Backup Adapter control	The IBM Spectrum Protect Snapshot operation failed because the object specified for a restore or query is not under the control of IBM Spectrum Protect Snapshot. It might be under control of IBM Spectrum Protect for ERP, for example. Check the db2diag.log and the IBM Spectrum Protect Snapshot summary log file for details.

Table 15. DB2 vendor reason codes (continued).

Reason Code	Explanation	User Response
34	Another database or application is using the same storage groups	The IBM Spectrum Protect Snapshot snapshot backup operation failed because another database or application is using the same storage group. Check the db2diag.log and the IBM Spectrum Protect Snapshot summary log file for details.

Storage system log and trace files

Storage system log and trace files are updated during IBM Spectrum Protect Snapshot operations.

Consult the documentation for the configured storage system.

CIM log and trace files

CIM log and trace files are updated during IBM Spectrum Protect Snapshot operations.

For more information about log and trace files for CIM, see the CIM documentation. The DS8000 Open API, SAN Volume Controller, and Storwize V7000 master console produce log and trace output.

IBM Spectrum Protect for ERP log and trace files

IBM Spectrum Protect for ERP log and trace files are updated during backup and restore operations.

See the section *How to find files containing message output (log files)* in the IBM Spectrum Protect for ERP *Installation and User's Guide* for details concerning logs and traces within IBM Spectrum Protect for ERP.

Important: A trace file can be requested by specifying the TRACEFILE parameter in the IBM Spectrum Protect for ERP profile. However, do not place this file on NFS, because this might cause network problems due to the high volume of trace entries being written.

Troubleshooting mirroring relationships

There are some questions that might arise when implementing IBM Spectrum Protect Snapshot and storage systems with mirroring technologies. The following information is provided to help you answer questions unique to your environment.

Question

Why are some remote mirroring relationships missing?

Answer

The target volumes that are referenced in this solution are part of the remote mirror relationship. The target volumes are used as the source for the snapshot operation.

Before you start the snapshot backup that uses the target volumes, verify that the remote mirroring relationships are established. You can verify the

relationships by using either the graphical user interface or the command-line interface. For example, if using SAN Volume Controller global mirror, you can enter the following command to verify the mirroring relationship:

```
ssh -i/<dir>/ssh-identity <username>@<hostname>  
svctask mkrcrelationship -master <vdiskname local> -aux <vdiskname remote>  
-cluster <clusterid> -name <relation name> -consistgrp <consgrp name>  
-global
```

Question

The remote mirroring relationships are not in the state `consistent_synchronized`. How does the state for remote mirroring relationship get updated?

Answer

Go to the storage solution. Synchronize the consistency groups. For more information about synchronizing consistency groups, see the documentation that is provided with the storage hardware.

Question

(SAN Volume Controller only) One or more of the FlashCopy target volumes for the remote site are missing. Where is the FlashCopy target volume?

Answer

Use either the graphical user interface or command-line interface to start the Metro Mirror or Global Mirror consistency group. For example, you can enter the following command from the command-line interface:

```
ssh -i/<di>ssh-identity <username>@<hostname of the cluster> svctask  
startrcconsistgrp consist group id
```

Question

(XIV only) One of the following issues exists.

- The remote mirroring is not operational.
- For XIV system synchronous mirroring, the state of the consistency group is not `consistent_synchronized`.
- For XIV system asynchronous mirroring, the state of the consistency group is not `RPO_OK`.

How are these issues resolved?

Answer

Verify that the consistency groups meet the following requirements:

- Consistency groups need to be enabled and synchronized.
- The volumes that are assigned to the consistency groups need to be correctly identified and enabled.

One consistency group per database partition is needed.

Troubleshooting storage solutions

There are some common problems that might occur when using IBM Spectrum Protect Snapshot and storage solutions. These problems and the solutions are provided to help you complete problem determination activities.

Question

During the backup or cloning on a storage solution running a supported AIX operating system, the mount of one or more file systems fails on the auxiliary host with the following message:

```
FMM0644E Error on running command: mount: 0506-334
/oracle/C21/mirrlog2 is not a known file system.
```

How can this error be resolved?

Answer

When the storage solution running a supported AIX operating system imports a volume group, use the label of the logical volume for the new mount point. Check the production system to determine the labels of the logical volumes that support backup and clone operations. The fields **mount point** and **label** should have identical values. For example:

```
# ls1v lvDS1data1
LOGICAL VOLUME: lvDS1data1 VOLUME GROUP: DS1data1vg
...
MOUNT POINT: /db2/DS1/db2ds1/NODE0001 LABEL: /db2/DS1/db2ds1/NODE0001
```

Troubleshooting connectivity problems

This information covers a problem that can occur with connectivity. The problem and the solution are provided to help you complete problem determination activities.

When the production server and backup server are separated by a firewall, socket connections might time out

Question

After a successful snapshot backup operation, why is it not possible to mount or unmount this snapshot backup on a backup or clone server?

Answer

The socket connection failure can result from a mismatch between the firewalls connection timeout setting and the operating systems frequency of sending keep alive network packets. When a firewall or other network devices such as a router or switch exists between the production and backup server, the daemon connection can time out. A similar situation can exist between a production and clone server. To prevent connections from timing out, the management agent `acsd` on the production server, requests that the operating system sends out network packets. These packets keep the connection between the servers alive.

The `tcp_keepidle` operating system parameter specifies the interval of inactivity. Depending on the operating system, this parameter might vary. After this interval of inactivity, the TCP generates a keep alive transmission for the application that requests it. This interval is measured in half seconds. For AIX operating systems, the keep alive default value for this parameter is 14400 (2 hours). This frequency is sufficient for many environments. Decrease this value when the following conditions exist:

- A firewall or other network device exists between the production and backup or clone server.
- If the device connection timeout is less than 2 hours.

For AIX operating systems, issue the following network command to reduce the **tcp_keepidle** parameter value and send a keep alive transmission every 5 minutes:

```
no -o tcp_keepidle=600
```

This change remains in effect until you restart the production server. To permanently modify this parameter, add the command to the `/etc/rc.net` file.

Troubleshooting tips for IBM Spectrum Protect Snapshot for Oracle

Resolving problems encountered when using IBM Spectrum Protect Snapshot requires tasks specific to the native Oracle database environment.

If an error condition occurs during an IBM Spectrum Protect Snapshot event, there are several sources of information you can view to help determine what the problem might be. Be aware of the following information:

- Make sure to increase the size of the following two Oracle options located in the `$ORACLE_HOME/dbs/init(database_name).ora` file:

```
sort_area_size = 10000000
sort_area_retained_size = 10000000
```

- When using IBM Spectrum Protect Snapshot to back up an Oracle database, the target database being backed up *cannot* reside on the same volume group as the file system containing `$ORACLE_HOME`. Make sure that the Oracle Server does not share a volume group with the target database.
- When performing a full offline backup of a database, the target database on the production server must be in “startup mount” state at the time **acsora** is issued. Otherwise it will not be possible to restore the resulting backup without performing recovery.

This RMAN script template will restore the database backed up offline as described in the previous paragraph. It restores control files, datafiles, and opens the database *without* any application of logs. This script must be started with the target database in a “startup mount” state:

```
run
{
  allocate channel ch1 type 'SBT_TAPE' parms
  'ENV=(TDPO_OPTFILE=<full path of tdpo.opt file>)';
  set until scn = <Ckp SCN for backup being restored>;
  restore control file to '<full path of 1st control file>';
  restore control file to '<full path of 2nd control file>';
  restore control file to '<full path of 3rd control file>';
  alter database mount;
  restore
  (database);
  sql 'alter database open RESETLOGS';
  release channel ch1;
}
```

The database will in an open state and in a new incarnation after this script completes. All that remains is to issue the **reset database** command to RMAN

and back up the database again since the previous backups are now rendered unusable since the database is in a new incarnation.

The <Ckp SCN for backup being restored> value is the Checkpoint SCN listed for the backup being restored in the RMAN **list backup** command. For example, the Checkpoint SCN is 32024 in the following list:

```
List of Backup Sets
Key Recid Stamp LV Set Stamp Set Count Completion Time
-----
26081 4 469212393 0 469212319 5 06-AUG-02

List of Backup Pieces
Key Pc# Cp# Status Completion Time Piece Name
-----
26082 1 1 AVAILABLE 06-AUG-02 05dvf74v_1_1

List of Datafiles Included
File Name LV Type Ckp SCN Ckp Time
-----
1 /dev/rmyfilelv 0 Full 32024 06-AUG-02
2 /dev/rmyrollbk1v 0 Full 32024 06-AUG-02
3 /dev/rmytemplv 0 Full 32024 06-AUG-02
4 /dev/rmyuserlv 0 Full 32024 06-AUG-02
```

Note that for an offline backup, the Checkpoint SCN should be the same for all of the datafiles.

Guidelines for Oracle variables

When the SQL*Plus or Oracle Net configuration files do not reside in their default location, set the **TNS_ADMIN** environment variable.

To run offloaded backups of Oracle databases, a recovery catalog database is needed. The database must be accessible by RMAN from the production host and the backup host. On the production host, use the following command to verify if the connection can be established. Before entering the command, log on as the database instance owner.

```
rman target / catalog catalog_db_user/catalog_user_password@catalog_db_connect_string
```

On the backup host, log on as the root user and enter the following command:

```
su - oracle_instance_owner -c rman target / catalog catalog_db_user/
catalog_user_password@catalog_db_connect_string
```

If you receive errors that say RMAN is unable to connect to the catalog database, verify the configuration of the `tnsnames.ora` on the host where the command was run. In addition, verify the listener configuration on the host where the catalog database runs. For details about the setup of the catalog database and the listener configuration, see documentation provided by Oracle.

IBM Spectrum Protect Snapshot for Oracle miscellaneous errors

Certain unique errors might occur when you use IBM Spectrum Protect Snapshot for native Oracle.

If you receive the following errors:

IBM Spectrum Protect Snapshot fails on the backup server in DBCS locales when the datafile or the path to the datafile contains a DBCS name.

This error is an Oracle problem that was reported to the Oracle development team. The Oracle Technical Assistance Request (TAR) number for this problem is 2367962.999.

The following procedure provides a workaround until the problem is resolved by Oracle:

1. Take the table space that contains the DBCS name in its datafile or the path to its datafile offline.
2. If the DBCS name is in the datafile, rename the DBCS datafile to an English name. If the DBCS name is in the path to the datafile, move the datafile to a path with an English name.
3. Log in to the Server Manager and issue the following command:

```
ALTER TABLESPACE <dbcs_tablespace_name> RENAME DATAFILE  
'dbcs_path/dbcs_datafile' TO 'english_path/english_datafile';
```
4. Bring the table space online.
5. Delete the DBCS datafile if necessary.

Although IBM Spectrum Protect Snapshot supports table spaces that are named with DBCS, datafiles or paths to the datafiles that contain DBCS must be renamed to English before running IBM Spectrum Protect Snapshot.

Internet Protocol Version 6 (IPv6) support

The IBM Spectrum Protect Snapshot for UNIX and Linux software operates in IPv4, IPv6, and mixed environments.

The network configuration determines which protocol is used by the IBM Spectrum Protect Snapshot software. The acsd service listens for IPv4 and IPv6 connection requests. Connection requests to the acsd service are made for the addresses that are returned by the system for the respective port on the local host. Connection requests to other systems are made for the addresses that are specified by the user. When TCP/IP addresses are set from a command-line interface, or when you are setting configuration parameters with the setup script, IPv6 addresses are supported. When an IP address and a port are specified in the following format:

<IPv4 address>:<service or port>

the format needs to be changed for IPv environments only:

<service or port>@<IP address>

In pure IPv4 environments, the traditional format can be used.

Appendix A. Configuration files

When you complete the setup script, the information you enter is used to configure IBM Spectrum Protect Snapshot.

IBM Spectrum Protect Snapshot uses the following configuration files:

- Profile
- Target volumes
- Password
- (Oracle in an SAP environment) Backint for SAP database
- (Oracle in an SAP environment) BR*Tools
- (Native Oracle) IBM Spectrum Protect options

The parameter and option information provided for the IBM Spectrum Protect Snapshot configuration files is for reference only. Do not edit these configuration files. The configuration files are updated when you use the setup script.

Profile configuration file

When you complete the setup script, the information you enter is used to create the profile configuration file. Each section of the profile includes parameters and options that determine how the IBM Spectrum Protect Snapshot backs up and restores data in your environment. For references, the following information explains the various parameters and options.

In the IBM Spectrum Protect Snapshot executable files, the profile is identified by the value specified for option `-p`.

The profile is divided into the following sections:

- GLOBAL
- ACSD
- CLIENT
- DEVICE_CLASS *device*
- OFFLOAD
- ORACLE
- CLONING

There can be multiple DEVICE_CLASS sections. Each DEVICE_CLASS section must have a unique *device* instance name.

The profile must be available on all database nodes and on the system where the management agent, `acsd`, is running. In addition, the GLOBAL section of the profile is required on the host where the clone databases reside.

GLOBAL The GLOBAL section contains information that is required and used by all IBM Spectrum Protect Snapshot components. The section is used by all database nodes, and the management, device, and offload agents. The components reference the information in the GLOBAL section during the start up process. Changes to this section require a restart of IBM Spectrum Protect Snapshot.

IBM Spectrum Protect Snapshot can be installed on multiple systems within an environment. For example, when a database is distributed among multiple application hosts or when a backup server is used to transfer snapshot backups to IBM Spectrum Protect. When IBM Spectrum Protect Snapshot is installed on multiple systems within an environment, there is only one active management agent. The location of this management agent is specified in GLOBAL section using the **ACSD** parameter.

Other parameters in the GLOBAL section specify the location for logging, tracing, and password files. On the backup server, the only section of profile that is referenced is GLOBAL.

ACSD The ACSD section contains information that is used exclusively by the management agent, acsd. This section includes the **ACS_REPOSITORY** parameter. The **ACS_REPOSITORY** parameter specifies the directory where the management agent stores its backup repository. This repository is the most important collection of IBM Spectrum Protect Snapshot data. If the repository is lost, any previously created backup cannot be restored.

CLIENT The CLIENT section contains all parameters relating to backup operations, including parameters for database applications, the number of backup versions, whether an IBM Spectrum Protect backup is to be created from the snapshot, how many snapshot backup generations to retain, and which **DEVICE_CLASS** section is used during snapshot creation. The CLIENT section is used by the snapshot backup library that is loaded to start backup or restore processing.

DEVICE_CLASS *device*

The **DEVICE_CLASS** section contains parameters that are related to the storage solution or file system (file system snapshots). At least one **DEVICE_CLASS** section is required for the configuration of the management agent. A **DEVICE_CLASS** section describes the characteristics of a storage device or file system that can be used to create a snapshot backup. The parameters and options that are used in the **DEVICE_CLASS** section depend on the storage solution.

Each storage solution that is used in the environment must have a **DEVICE_CLASS** section and must have a unique *device* instance name.

The **DEVICE_CLASS** section that is used is determined by the value of the **DEVICE_CLASS** parameter in the CLIENT section of the profile for backup operation. For cloning operations, this value is determined by the **DEVICE_CLASS** parameter in the CLONING section of the profile. If the same value is specified for the **DEVICE_CLASS** parameter in both the CLIENT and CLONING sections, an error is reported.

The value of **DEVICE_CLASS** *device* is recorded in the IBM Spectrum Protect Snapshot repository to identify the appropriate **DEVICE_CLASS** section during the restore process. Therefore, use caution when you delete or rename **DEVICE_CLASS** sections. If the appropriate section cannot be found, then the data that is backed up cannot be restored.

For each **DEVICE_CLASS** section, a password is required and can be set by running the setup script without the **-a** action option. For example:

```
setup_type.sh -d <Instance owner $HOME directory>
```

The password can be set in a batch processing mode using the following **fcml** command: **fcml -f password**

These passwords are used by IBM Spectrum Protect Snapshot to authenticate to the storage solution represented by the `DEVICE_CLASS` section.

OFFLOAD

The parameters and options in the OFFLOAD section determine how a snapshot is transferred to IBM Spectrum Protect. The information is sent to the offload agent, `tsm4acs`.

When the offload agent is started, it connects to the management agent and queries for snapshot backups that have been backed up with the `TSM_BACKUP` parameter that is set to YES. For Oracle in an SAP environment systems, the `TSM_BACKUP_FROM_SNAPSHOT` parameter is used instead of the `TSM_BACKUP` parameter. If this parameter and option is found, the offload agent mounts the snapshot and initiates an IBM Spectrum Protect backup using one of the following applications:

- (Oracle in an SAP environment) IBM Spectrum Protect for Enterprise Resource Planning (IBM Spectrum Protect for ERP)
- (Oracle) Oracle RMAN and Data Protection for Oracle

If one of the following conditions exists, the OFFLOAD section is required. If neither condition exists, the OFFLOAD section is optional.

- (Oracle in an SAP environment) IBM Spectrum Protect for ERP is used for an offload tape backup. In this scenario, the OFFLOAD section includes at least the `PROFILE` parameter.
- One or more of the default values must be overridden.

CLONING

The CLONING section contains the parameters used for cloning operations. The section is ignored for all other operations.

CLONING

The CLONING section contains the parameters used for cloning operations. The section is ignored for all other operations.

Examples

All parameters in a section are indicated by a section start notation, `>>>` *section_name*, and a section end notation, `<<<` *section_name*. The name is optional on the section end notation. Comments can be used at any place within the profile. Comments start with a `#` character and extend to the end of the line. Tab characters are permitted. The following example provides an example of the profile configuration file:

```
# Global section
>>> GLOBAL
parametername1 value1
parametername2 value1 value2
....
<<<
# ACSD section
>>> ACSD
parametername1 value1
parametername2 value1 value2
....
<<<
# CLIENT section
>>> CLIENT
parametername1 value1
parametername2 value1 value2
....
```

```

<<<
# DEVICE_CLASS device section
>>> DEVICE_CLASS device
parametername1 value1
parametername2 value1 value2
....
<<<
# DEVICE_CLASS device2 section
>>> DEVICE_CLASS device2device2
parametername1 value1
parametername2 value1 value2
....
<<<
# OFFLOAD section
>>> OFFLOAD
parametername1 value1
parametername2 value1 value2
....
<<<
# ORACLE section
>>> ORACLE
parametername1 value1
parametername2 value1 value2
....
<<<
# CLONING section
>>> CLONING
parametername1 value1
parametername2 value1 value2
....
<<<

```

GLOBAL

The profile parameters in the GLOBAL section contain basic configuration information. Examples of the type of information that is specified by the parameters are the port that is used by IBM Spectrum Protect Snapshot and the location of log files. The parameters are independent of the storage solution, database application, and custom application.

The following list provides the parameters, a description of each parameter, and default values for the GLOBAL section of the profile configuration file.

ACS_DIR

Path to the IBM Spectrum Protect Snapshot configuration directory. This parameter is required. The following subdirectories are included in this directory:

logs The subdirectory contains all log and trace information for IBM Spectrum Protect Snapshot.

shared The subdirectory contains information that is shared among all IBM Spectrum Protect Snapshot components.

When the subdirectory is initially created, the only file that is stored in the directory is the password file: `pwd.acsd`. This file contains the passwords for all devices that are specified within the profile. The file also contains a master password that is used from all components for authentication when they are connecting to the management agent. When you run remote configuration tasks from the production system with SSH, the information in these directories is promoted to all systems that belong to the instance where IBM Spectrum Protect

Snapshot is configured. When you run configuration tasks separately, you must promote the information manually.

Default

user_home/acs

Advanced mode only

Yes

ACSD

The host name and port of the system where the management agent is running. The following format is used for **ACSD**: *hostname port*

This parameter must be identical on all systems where IBM Spectrum Protect Snapshot is installed for a database instance. While the parameter must be identical, each database instance can be managed by an individual management agent.

Default

hostname 57328

Advanced mode only

Yes

ENFORCE_TLS12

IBM Spectrum Protect Snapshot uses the security suite, IBM Global Security Kit (GSKit) for Secure Socket Layer / Transport Layer Security (SSL/TLS) TCP/IP connections. GSKit is able to provide SP800-131 compliant encryption by using the TLS protocol V1.2. To enforce the use of this protocol, select the option YES, otherwise the TLS version 1.0 and 1.1 is enabled by default.

Default

NO

Advanced mode only

Yes

TRACE

There are two options for **TRACE**: YES and NO. YES means that tracing is enabled. NO means that tracing is not enabled.

This parameter can also be set in the `.utl` file.

Default

NO

Advanced mode only

Yes

ACSD

Except where noted, the profile parameters in the ACSD section are independent of the storage device or application.

ACS_REPOSITORY

This parameter sets the path to the IBM Spectrum Protect Snapshot repository. This directory is used during restore operations and must be in a secure location. If the repository is lost, all backups are not available.

The directory that is referenced by the **ACS_REPOSITORY** parameter cannot be in a file system that participates in snapshot backup operations. If the directory is part of a file system that is used for snapshot backup operations, IBM Spectrum Protect Snapshot reports a failure. The IBM Spectrum Protect Snapshot repository cannot be in the main IBM Spectrum Protect Snapshot

directory that is specified by the **ACS_DIR** parameter. Ideally, the **ACS_REPOSITORY** directory is a subdirectory of the **ACS_DIR** directory. For example:

```
<ACS_DIR>/acsrepository
```

Before you configure IBM Spectrum Protect Snapshot, the path to the **ACS_REPOSITORY** is set, but the directory does not exist. The **ACS_REPOSITORY** directory is created during the configuration process. If the directory specified for the **ACS_REPOSITORY** parameter exists, the setup script, used to configure IBM Spectrum Protect Snapshot, reports an error.

Default

user_home/acs/acsrepository.

Advanced mode only

Yes.

ADMIN_ASSISTANT

In non-SAP environments, this parameter is ignored. If IBM Spectrum Protect for ERP and the Administration Assistant component are installed, when this parameter is set, IBM Spectrum Protect Snapshot sends backup and restore information to the Administration Assistant.

<server> <port>

Server and port where the IBM Spectrum Protect for ERP Administration Assistant server component is listening.

NO Do not send data to the Administration Assistant.

Default

NO

Advanced mode only

Yes.

REPOSITORY_LABEL

When this parameter is set, a prefix is added to each volume name on the IBM XIV Storage System. The prefix contains 3 characters in one of the following ranges:

[a-z]
[A-Z]
[0-9]

Note: If the repository label changes, backups that are created with the prior repository label are excluded from reconciliation.

Default

TSM

Advanced mode only

Yes.

SYNCHRONOUS_RECONCILE

This parameter is used to configure IBM Spectrum Protect Snapshot to synchronously reconcile and delete snapshot backups. If the **RESTORE_AND_DELETE** option is specified for this parameter, a delete and restore operation also starts a synchronous delete and reconcile operation. This process can be useful for storage systems that can delete snapshot backups during an IBM Spectrum Protect Snapshot snapshot restore process. Deletion can occur on Storwize V7000, or SAN Volume Controller storage systems. Also, this

process is useful if you manually delete snapshot backups and use the force option (-f) on DS8000, SAN Volume Controller, or Storwize V7000 storage systems.

If the YES option is specified in addition to the delete and restore operation, a backup operation also starts a synchronous delete and reconcile process. This process can be useful for storage systems that delete snapshot backups during an IBM Spectrum Protect Snapshot backup or cloning operation. Deletion can occur on SAN Volume Controller or Storwize V7000 storage systems. The following list identifies the possible options:

NO Use this option not to start a synchronous delete and reconcile operation.

YES

Use this option to start a synchronous delete and reconcile process as part of a backup, restore, and delete operation.

RESTORE_AND_DELETE

Use this option to start a synchronous delete and reconcile process as part of a restore and delete operation.

Default

RESTORE_AND_DELETE

Advanced mode only

YES

CLIENT

IBM Spectrum Protect Snapshot uses the Oracle specific parameters to configure Oracle backup and restore operations. These parameters are defined in the client section of the IBM Spectrum Protect Snapshot profile configuration file.

Oracle databases

When using Oracle databases, in addition to the CLIENT section, an ORACLE section is included in the profile configuration file. For more information about the ORACLE section, see "ORACLE" on page 124.

APPLICATION_TYPE

This parameter specifies the environment. There is only one option:

ORACLE

Use as an Oracle system.

DEVICE_CLASS

This parameter specifies the device classes to use. The following sample identifies the syntax that can be used with the **DEVICE_CLASS** parameter:

`DEVICE_CLASS list_of_device_classes [conditions]`

When a list of device classes is specified, the software determines which device class matches the device class in the environment. When multiple device classes are specified, separate the device classes names with a space. The condition statement is optional. When you use the condition statement, use the following syntax:

`[USE_AT days of week] [FROM time TO time]`

Note: The time period that is specified cannot span midnight for a device class. If a device class time period is required to span midnight, you must specify two time periods for the device class. The first time period must end

with a value 1 minute before midnight and the second time period must start at midnight. The following example shows how to specify a time period that spans midnight for a device class:

```
DEVICE_CLASS myClass FROM 20:00 TO 23:59
DEVICE_CLASS myClass FROM 00:00 TO 06:00
```

When there are different devices, multiple sections can be used. Each section provides information about a particular device. To select a particular section, use the **DEVICE_CLASS** parameter. When the software restores data, the software uses the **DEVICE_CLASS** value that is specified when the data was backed up.

The configuration wizard (the setup script) automatically adds **DEVICE_CLASS** sections to the IBM Spectrum Protect Snapshot profile when you add more instances of the **DEVICE_CLASS** parameter to the CLIENT section of the profile.

Default

STANDARD

Advanced mode only

No

ENHANCED_PARTITIONING

The **ENHANCED_PARTITIONING** parameter is used to control processing of the application file systems during the backup or restore operation. IBM Spectrum Protect Snapshot fails, when a file system contains symbolic links that point to a file system on a different volume group that is not part of the FlashCopy operation. Set the **ENHANCED_PARTITIONING** parameter to NO to ensure that symbolic links if present are not processed. With this setting, there is no check for additional files that are not associated with the application. If you use this setting, the run time of the backup operation is likely to decrease depending on the file system structure. The following list identifies the possible options:

YES

Use this option to ensure that IBM Spectrum Protect Snapshot processes all symbolic links of files or directories.

NO Use this option to ensure that IBM Spectrum Protect Snapshot does not process symbolic links of files or directories.

Default

YES

Advanced mode only

Yes.

GLOBAL_SYSTEM_IDENTIFIER

This parameter specifies a string to be used in the IBM Spectrum Protect for Enterprise Resource Planning Administration Assistant that uniquely identifies an Oracle database in the system landscape. This parameter is only valid when the **ADMIN_ASSISTANT** parameter is specified in the ACSD section of the profile.

Default

The default value is *ORA_<DBname>*.

Advanced mode only

Yes

LVM_FREEZE_THAW

This parameter specifies when to enable file system freeze and thaw actions. The following list identifies the possible options:

YES

Enable file system freeze before the snapshot operation and the thaw after

the snapshot operation. For AIX, the YES value can be used only when all file systems included in the backup are JFS2 file systems.

NO Do not freeze the file system. To set this parameter to NO, a licensed version of IBM Spectrum Protect Snapshot is needed and a backup server is required for mounting the snapshot to ensure file system consistency.

The value NO is required if at least one file system that does not support freeze or thaw actions, such as JFS, is involved.

AUTO

If the **TARGET_DATABASE_SUSPEND** parameter is YES, then this parameter is set with the following option: **LVM_FREEZE_THAW YES**. If the file system does not support freeze actions, the AUTO value is NO.

For more information, see “Interdependency of **LVM_FREEZE_THAW** and **TARGET_DATABASE_SUSPEND**” on page 152.

Default

AUTO

Advanced mode only

Yes

MAX_VERSIONS

This parameter specifies the number of snapshot versions to store. The following list identifies the possible options:

ADAPTIVE

The maximum number varies depending on the available space. IBM Spectrum Protect Snapshot reuses the oldest target set as the target for the current backup.

n Where *n* is the maximum number of snapshot versions to be stored. The amount of space that is required depends on the following factors:

- The number of snapshots.
- For each snapshot, the number of changes to the file system content since the snapshot was taken.

When this limit is reached, the oldest version is deleted.

Default

ADAPTIVE

Advanced mode only

No

NEGATIVE_LIST

This parameter is used to control file processing. This processing occurs when files that are not associated with the database are stored within the same file system that is used for the backup and restore operations. This parameter is required. The following list identifies the possible options:

NO_CHECK

This is the default value, and it means that there are no checks for extra files. The operation ignores any additional files that are identified. When you use the default value and data is restored, all files on the file system or volume group are overwritten.

WARN

Use this option to receive a warning message for each file that is identified on the volume, but not part of the FlashCopy operation. The processing

continues. When you use this option and data is restored, all files on the file system or volume group are overwritten.

ERROR

Use this option to receive an error message for each file that is discovered on the volume, but not part of the FlashCopy operation. The processing ends.

filename

Where *filename* is a name of a file that contains a list of fully qualified names of files and directories, each name requires a new line. Only files or directories that are not associated with the database but are stored within the file system that is used for backup operations are listed. Any file that is identified by IBM Spectrum Protect Snapshot that is not part of the database files or is not in the **NEGATIVE_LIST** file, causes processing to end. Any directory that is listed in the **NEGATIVE_LIST** file is processed recursively. For example, all files within the directory, including subdirectories, are processed during a backup or restore request.

When you are restoring data with remote mirroring, the value of this parameter is forced to **NO_CHECK**. This value is used because at the time after the takeover operation there are no file systems mounted on the takeover host.

Default

NO_CHECK

Advanced mode only

Yes

TARGET_DATABASE_SUSPEND

This parameter determines if the activity is suspended on the target database until the FlashCopy operation completes. The following list identifies the possible options:

Yes

This option suspends the target database until the FlashCopy operation completes. When there are many transactions processing, use this option.

NO This option means that the target database is available while FlashCopy operations run.

OFFLINE

All backups must be offline for the FlashCopy operations to run. When either *acsora* or *acsutil* is issued, the target database on the production system must be in a *startup mount* state. If the state is not *startup mount*, recovery is necessary to restore the database.

Default

Yes.

Advanced mode only

Yes.

For more information about the **TARGET_DATABASE_SUSPEND**, see “Interdependency of **LVM_FREEZE_THAW** and **TARGET_DATABASE_SUSPEND**” on page 152.

TIMEOUT_FLASH

This parameter specifies the maximum time, in seconds, that the database agent waits for a response to the management agent call during the *flash* phase. If the database agent does not receive a response within the specified time, an

error message is displayed. This parameter allows the maximum time to be specified for a database to be suspended. This parameter also implies the maximum time when JFS2 file systems can be frozen. When the timeout is reached, the file systems thaw, the database is resumed, and the backup operation ends with an error. If the **LVM_FREEZE_THAW** parameter is set to either **AUTO** or **YES**, the minimal value for **TIMEOUT_FLASH** is 5 seconds. In other scenarios, the minimal value is 1 second.

Default

The default value is 120 seconds.

Advanced mode only

Yes

TIMEOUT_PHASE

This parameter specifies the maximum time, in seconds, that the database agent waits for a response to the management agent call during a specific operation phase. If the database agent does not receive a response within the specified time, either the backup or restore operation ends and an error message is shown.

Specify one of the following phase values for a FlashCopy backup:

- **PARTITION**
- **PREPARE**
- **FLASH** (this parameter has a separate description)
- **VERIFY**
- **CLOSE**

For example, **TIMEOUT_PREPARE**.

Specify one of the following phase values for a FlashCopy restore:

- **PREPARESTORE**
- **FLASHRESTORE**
- **COMPLETERESTORE**
- **CLOSE**

For example, **TIMEOUT_FLASHRESTORE**.

Default

The default value is 3600 seconds.

Advanced mode only

Yes

TSM_BACKUP

This parameter specifies whether to create an IBM Spectrum Protect backup from a snapshot. IBM Spectrum Protect Snapshot must be installed on a backup server. When the **TSM_BACKUP** parameter is set to **YES**, **MANDATE**, or **LATEST**, and after the offload agent runs, an IBM Spectrum Protect backup is created from the snapshot. The following list identifies the possible options:

YES

This option creates an IBM Spectrum Protect backup from a snapshot. If the IBM Spectrum Protect backup operation does not successfully complete, the target set can be reused.

MANDATE

This option creates an IBM Spectrum Protect backup from a snapshot. However, the target set cannot be reused until the IBM Spectrum Protect backup successfully completes.

LATEST

This option removes a backup request to IBM Spectrum Protect from a previous backup. When a new snapshot with **TSM_BACKUP** set to LATEST, YES, or MANDATE is created, IBM Spectrum Protect Snapshot removes any unsuccessful backup request that were previously created with the **TSM_BACKUP** option set to LATEST. This option prevents backup requests to IBM Spectrum Protect from queuing if they are not completed in time.

NO Keeps the snapshot backup but the snapshot is not used as a source for a subsequent tape backup operation.

TSM_ONLY

After the IBM Spectrum Protect backup is completed, during the unmount operation, the backup is automatically marked for deletion. This action occurs regardless of whether the backup is successful or not.

USE_FOR *list of device classes*

To create an IBM Spectrum Protect backup from snapshots that are run with particular device classes, as specified in the profile, combine this attribute with other options. When you list device classes, separate device classes with the space character. There is no limit of the number of device classes.

Default

None

Advanced mode only

No

VOLUME_MGR

The following list identifies the possible options:

ASM

When this option is selected, the option that is set for the **LVM_FREEZE_THAW** parameter is ignored. When this parameter is ignored, there is no file system, and the wizard does not query for data.

LVM

When this option is selected, the ASM-related options in the **DEVICE_CLASS** section is ignored and not queried by the wizard software.

Default

LVM

Advanced mode only

No.

ORACLE

The ORACLE section is an extension to the CLIENT and CLONING sections for ORACLE environments. The parameters do not depend on the storage device.

CATALOG_DATABASE_CONNECT_STRING

The recovery catalog connect string. This parameter specifies the connect string of the Recovery catalog database that is used to catalog backup information. This value corresponds to the value defined in the *\$ORACLE_HOME/network/admin/tnsnames.ora* file.

Default

There is no default value. This parameter is specified by the user.

Advanced mode only

No.

CATALOG_DATABASE_USERNAME

This parameter sets the user name that has the Oracle system database administrator privileges on the Recovery catalog database.

Default

There is no default value. This parameter is specified by the user.

Advanced mode only

No.

TARGET_DATABASE_PARAMETER_FILE

This parameter specifies the fully resolved path and file name of the Oracle parameter file for the target database. This file is a text-based Oracle parameter file (PFILE) and not a binary Oracle Server Parameter File (SPFILE).

Default

`${ORACLE_HOME}/dbs/init${ORACLE_SID}.ora`

Advanced mode only

Yes.

DATABASE_BACKUP_SCRIPT_FILE

Name of the RMAN backup script that contains the Data Protection for Oracle environment variables.

Default

There is no default value. For offload configuration, this parameter is specified.

Advanced mode only

No.

DATABASE_CONTROL_FILE_RESTORE

This parameter specifies whether to restore Oracle control files after snapshot restore processing completes. There are two options:

- YES: restores Oracle control files and you complete the incomplete recovery up to the point when the control files were backed up.
- NO: does not restore Oracle control files. A full snapshot recovery up to the current image is completed, using the existing control files in the system.

Default

NO

Advanced mode only

Yes.

ASM_INSTANCE_USER

This parameter specifies the user that owns the Oracle Grid Infrastructure installation. A setting of AUTO, refers to the default value grid.

Default

grid

Advanced mode only

Yes

ASM_INSTANCE_ID

This parameter is deprecated. The ORACLE_SID of the ASM instance is automatically determined from the environment of the user that is specified in the **ASM_INSTANCE_USER** value.

Default

AUTO

Advanced mode only

Yes

ASM_ROLE

This parameter specifies the role that is used when connecting to the ASM instance. There are two options:

- `sysasm`: This option is the default role for connections to the ASM instance.
- `sysdba`: This role is supported for Oracle 11gR1, but is deprecated.

Default

`sysasm`

Advanced mode only

Yes.

CLONING

The CLONING section of the IBM Spectrum Protect Snapshot profile contains parameters that are used for cloning operations. The parameters are independent of the storage device or application.

The following lists provide the parameters, a description of each parameter, and default values for the CLONING section.

The following parameters apply to Oracle, and Oracle in an SAP environment databases:

CLONE_TARGET_DATABASE_TYPE

This parameter identifies the installation type that is on the clone system.

There are three options as follows:

AUTO

With this value, it is assumed that the clone system is of the same type as the production system. For example, if the production system is configured for Oracle RAC then *AUTO* for this parameter tells IBM Spectrum Protect Snapshot that the clone system is an Oracle RAC installation. *AUTO* is the default value for this parameter.

RAC

The clone system is configured for Oracle RAC with one node even if the production system is not configured for RAC.

NON_RAC

The clone system is not configured for RAC. Use this setting if you clone a RAC production database to a clone system that was not configured for RAC.

You cannot clone from non-RAC Oracle to a RAC configuration.

Default

AUTO.

Advanced mode only

Yes.

DEVICE_CLASS

This required parameter identifies the device class to use when you are cloning a database. The following code sample provides an example of how to specify options for this parameter:

```
DEVICE_CLASS device class USE_FOR_CLONING list of clone database names
```

There is an optional *conditions* statement that can be used. The following code sample includes an example of how to use the *conditions* statement, which is optional. When you use the condition statement, use the following syntax:

```
[USE_AT days of week] [FROM time TO time]
```

The time period that is specified cannot span midnight for a device class. If a device class time period is required to span midnight, you must specify two time periods for the device class. The first time period must end with a value 1 minute before midnight and the second time period must start at midnight. The following example shows how to specify a time period that spans midnight for a device class:

```
DEVICE_CLASS myClass1 USE_FOR_CLONING CL1 FROM 20:00 TO 23:59  
DEVICE_CLASS myClass2 USE_FOR_CLONING CL2 FROM 00:00 TO 06:00
```

If multiple **DEVICE_CLASS** statements are used, ensure that a unique 1-to-1 relation between the clone database name and the device class exists.

Default

Not applicable.

Advanced mode only

No.

ENHANCED_PARTITIONING

The **ENHANCED_PARTITIONING** parameter is used to control processing when extra file systems that are not database files are specified by the **FLASH_DIR_LIST** parameter in a cloning operation. When **VOLUME_MGR** is set to **ASM**, the **ENHANCED_PARTITIONING** parameter is not evaluated by IBM Spectrum Protect Snapshot, and the default setting applies. IBM Spectrum Protect Snapshot fails, when a file system that is specified by the **FLASH_DIR_LIST** parameter contains symbolic links that point to a file system on a different volume group that is not part of the FlashCopy operation. Set the **ENHANCED_PARTITIONING** parameter to **NO** to ensure that symbolic links if present are not processed. You must manually add this parameter to the IBM Spectrum Protect Snapshot profile file. The following list identifies the possible options:

YES

Use this option to ensure that IBM Spectrum Protect Snapshot processes all symbolic links of files or directories that are specified in the **FLASH_DIR_LIST** profile parameter.

NO Use this option to ensure that IBM Spectrum Protect Snapshot does not process symbolic links of files or directories that are specified in the **FLASH_DIR_LIST** profile parameter.

Default

YES

Advanced mode only

Yes.

FLASH_DIR_LIST

This parameter is used to include files systems that are not part of the database files in the FlashCopy operation. For example, when you are cloning an SAP Advanced Business Application Programming and Oracle Java™ system, the Java instance is not part of the database files. A clone of the Java instance is created along with the clone of the database. In this scenario, use the **FLASH_DIR_LIST** parameter to include the Java instance directories. If **VOLUME_MGR** is set to **ASM**, the **FLASH_DIR_LIST** is ignored.

Specify a fully qualified directory name and file name. For example:

[ON DBPARTITIONNUM *list of partitions*] *fully qualified file name*

Inside the file, specify one fully qualified file or directory on each line. IBM Spectrum Protect Snapshot uses the FlashCopy function to FlashCopy the complete volume groups where the specified files or directories are located.

The default value is an empty list. This value prevents extra files or directories from participating in the FlashCopy operation.

Default

By default, no file name is specified.

Advanced mode only

Yes.

DATABASE_SCHEMA

When a clone database is created from the production database, the database schema does not change. The clone database uses the database schema that is used by the production database. The **DATABASE_SCHEMA** parameter is used to specify the database schema. For Oracle databases, the **DATABASE_SCHEMA** parameter is required. For Oracle in an SAP environment databases, the **DATABASE_SCHEMA** parameter is optional. The default database schema is determined by the `db5_ora_schema` environment variable. This environment variable is set on the production database instance owner environment.

If these environment variables are not set, the default database schema value is `SAPR3`. When the **DATABASE_SCHEMA** parameter is used for a Oracle in an SAP environment database, the specified database schema value overrides all default database schema values. The **DATABASE_SCHEMA** parameter is evaluated when the following conditions exist:

- A processing script is used with the **preproc_clone** or **postproc_clone** command.
- The **refresh_clone** command is entered with the `-X` or `-Y` cloning parameter.
- The **create_clone** command is entered with the `-Y` cloning parameter.

Default

The default value is determined by environment variables.

Advanced mode only

Yes.

NEGATIVE_LIST

The **NEGATIVE_LIST** parameter is used to control processing when files not associated with the database are stored within the same file system that is used for the backup and restore operations. This parameter is required. The following list identifies the options:

NO_CHECK

Use this option to not check for extra files. The operation ignores any additional files that are identified. When you use this option and data is restored, all files that are on the file system or volume group are overwritten.

WARN

Use this option to receive a warning message for each file that is identified on the volume, but not part of the FlashCopy operation. The processing continues. When you use this option and data is restored, all files that are on the file system or volume group are overwritten.

ERROR

Use this option to receive an error message for each file that is discovered on the volume, but not part of the FlashCopy operation. The processing ends.

filename

Use this option to back up and restore files that are not part of the database tablespace files. Using this option includes files in the FlashCopy operations. When you use this option, specify the fully qualified names of the files and directories. Use one line for each entry. When these files are identified, processing continues. When other files are identified, but not part of the database tablespace files or identified in the **NEGATIVE_LIST** file, processing ends. Any directory that is listed in the **NEGATIVE_LIST** file is processed recursively. For example, all files within the directory, including subdirectories, are processed during a backup or restore request.

Default

There is no default for this required parameter.

Advanced mode only

Yes.

GLOBAL_SYSTEM_IDENTIFIER

Use this parameter to specify a string to be used in the IBM Spectrum Protect for Enterprise Resource Planning Administration Assistant that uniquely identifies an Oracle database in the system landscape. This parameter is only valid when the **ADMIN_ASSISTANT** parameter is specified in the ACSD section of the profile.

Default

The default value is *ORA_<DBname>*.

Advanced mode only

Yes.

TIMEOUT_FLASH

This parameter specifies the maximum time, in seconds, that the database agent waits for a response to the management agent call during the *flash* phase. If the database agent does not receive a response within the specified time, an error message is displayed. This parameter allows the maximum time to be specified for a database to be suspended. This parameter also implies the maximum time when JFS2 file systems can be frozen. When the timeout is reached, the file systems thaw, the database is resumed, and the backup operation ends with an error. The minimum value for **TIMEOUT_FLASH** is 5 seconds.

Default

The default value is *120* seconds.

Advanced mode only

Yes

TIMEOUT_<PHASE>

Specify the maximum time (in seconds) that the database agent waits for a response to the management agent call during the *<phase>* phase. If the database agent does not receive a response within the specified time, the backup or restore operation ends and an error message is displayed. The default value is *3600* seconds.

You can specify one of these phase values for a FlashCopy backup. For example: **TIMEOUT_PREPARE**

- **PARTITION**
- **PREPARE**
- **VERIFY**
- **CLOSE**

You can specify one of these phase values for a FlashCopy restore. For example: **TIMEOUT_FLASHRESTORE**

- **PREPARERESTORE**
- **FLASHRESTORE**
- **COMPLETERESTORE**
- **CLOSE**

Advanced mode only

Yes.

The following parameters apply to Oracle and Oracle in an SAP environment databases:

OVERWRITE_DATABASE_PARAMETER_FILE

This parameter is used only with Oracle and Oracle in an SAP environment databases. This parameter is also included in the OFFLOAD section of the configuration profile. The parameter specifies whether the database configuration file on the clone server is overwritten with the file from the production server. The parameter value in the OFFLOAD section is not applicable to cloning operations and is ignored.

To copy the database configuration file from the production system to the clone system, specify this parameter in the CLONING section of the configuration profile. The IBM Spectrum Protect Snapshot software requires two database configuration files to be available in the clone instance on the clone system. The default database configuration file name is `${ORACLE_HOME}/dbs/init${ORACLE_SID}.ora`

In the scenario where the production system is `${ORACLE_SID}=P01` and the clone system is `${ORACLE_SID}=C01`, the following database configuration files are required:

```
/oracle/C01/102_64/dbs/initP01.ora
/oracle/C01/102_64/dbs/initC01.ora
```

The `initP01.ora` file is used during the cloning process to recover the database that is used in the FlashCopy operation on the clone system. The `initC01.ora` file is used to rename and start the clone database.

Specify one of the following values:

YES

Copy the database configuration file from the production system to the clone system. The following processes occur:

- The clone database configuration file `initP01.ora` is copied on the clone system. The existing file, `/oracle/C01/102_64/dbs/initP01.ora`, is overwritten. If the production database uses a binary Oracle Server Parameter file (SPFILE) then it is dumped to a temporary pfile and copied to the clone system.
- If the production database is configured to use a binary Oracle SPFILE, do not specify a value for the **TARGET_DATABASE_PARAMETER_FILE** parameter. In this scenario, the default database configuration file name `${ORACLE_HOME}/dbs/init${ORACLE_SID}.ora` is used on the clone system.

- The clone database configuration file `/oracle/C01/102_64/dbs/initP01.ora` is copied to `/oracle/C01/102_64/dbs/initC01.ora`. The existing file, `/oracle/C01/102_64/dbs/initC01.ora`, is overwritten. All occurrences of P01 in this file are renamed to C01.

NO Do not copy the database configuration file from the production system to the clone system. This value requires that the database configuration files `/oracle/C01/102_64/dbs/initP01.ora` and `/oracle/C01/102_64/dbs/initC01.ora` are available on the clone system. You must verify that these files are available and are valid.

If the name of the database configuration file on the production database is not the default file name, `${ORACLE_HOME}/dbs/init${ORACLE_SID}.ora`, use the **TARGET_DATABASE_PARAMETER_FILE** parameter, in the CLONING section, to specify the correct name. In this scenario, the clone database configuration file name is created by replacing the `${ORACLE_SID}` value of the production database with the name of the clone database.

Default

YES

Advanced mode only

Yes.

TARGET_DATABASE_PARAMETER_FILE

This parameter is used only with Oracle and Oracle in an SAP environment databases. This parameter is also included in the ORACLE section of the configuration profile.

The parameter specifies the database configuration file name. The parameter value in the ORACLE section is not applicable to cloning operations and is ignored.

Specify this parameter in the CLONING section of the profile to identify the name of the Oracle parameter file for the production database. Enter the fully resolved path and file name of the Oracle parameter file for the production database. By default, the file name of the Oracle parameter file for the production database is `initSID.ora`. This file must be a text-based Oracle parameter file (PFILE) and not a binary Oracle SPFILE. The default value is `${ORACLE_HOME}/dbs/init${ORACLE_SID}.ora`. If the production database is configured to use a binary Oracle SPFILE, this file is automatically detected and this parameter must not be specified.

Default

`${ORACLE_HOME}/dbs/init${ORACLE_SID}.ora`

Advanced mode only

Yes.

VOLUME_MGR

The following list identifies the possible options:

ASM

When this option is selected, the option that is set for the **LVM_FREEZE_THAW** parameter is ignored. When this parameter is ignored, there is no file system, and the wizard does not query for data.

LVM

When this option is selected, the ASM-related options in the **DEVICE_CLASS** section is ignored and not queried by the wizard software.

Default

LVM

Advanced mode only

No.

DEVICE_CLASS *device*

The IBM Spectrum Protect Snapshot profile configuration file can contain one or more **DEVICE_CLASS** sections. The device class section configures IBM Spectrum Protect Snapshot for use with a particular storage or file system solution. The parameters do not depend on the database or custom application that is protected.

Use care when you rename or delete a **DEVICE_CLASS** section from the profile, as you cannot access backups that were taken with the original **DEVICE_CLASS** section. Therefore, first remove backups and clones that are associated with the **DEVICE_CLASS** before you rename or delete the **DEVICE_CLASS** section.

A *device* refers to supported IBM XIV Storage System , IBM Storwize family, IBM System Storage SAN Volume Controller, and IBM System Storage DS8000 series.

Updating DEVICE_CLASS device for mirroring

To use the mirroring technologies, a **DEVICE_CLASS** section specific to the storage solution used for mirroring needs to be added to the profile configuration file. There is one exception to this statement: If remote backups are run, the existing **DEVICE_CLASS** section for the device is sufficient. No additional **DEVICE_CLASS** section is needed.

About this task

When creating a **DEVICE_CLASS** section for the storage solution used for mirroring, the section includes the same parameters as the device class for the local site, specific vales for the remote site, and the parameters that are required to connect and send requests to the remote cluster. The parameters required to connect and send requests to the remote cluster are identified in the following list:

COPYSERVICES_REMOTE

The option set for this parameter determines if the backup is taken at the remote site. The options are YES and NO. The default option is set to NO.

COPYSERVICES_REMOTE_SERVERNAME

This parameter specifies the IP address or hostname for the secondary cluster.

If the **COPYSERVICES_REMOTE** parameter is set to YES, the parameter is required.

If the **COPYSERVICES_REMOTE** parameter is set to NO, the

COPYSERVICES_REMOTE_SERVERNAME parameter cannot be used. If the parameter is used, an error occurs.

COPYSERVICES_REMOTE_USERNAME

This parameter specifies the user name used to connect to the secondary cluster. The default option is superuser. If the **COPYSERVICES_REMOTE** parameter is set to NO, the **COPYSERVICES_REMOTE_SERVERNAME** parameter cannot be used. If the parameter is used, an error occurs.

TAKEOVER_HOST_NAME

This parameter is required when restoring a remote mirroring backup after a takeover procedure on the remote side. The value for this parameter is the host name of the takeover host and is only used in combination with the secondary cluster defined by the **COPYSERVICES_REMOTE_SERVERNAME** parameter. The value

specified for this parameter needs to match the value defined in the storage system. If the values do not match, an error occurs.

The following `DEVICE_CLASS` parameters need to be common to both clusters:

- `COPYSERVICES_COMMPROTOCOL`
- `COPYSERVICES_CERTIFICATEFILE`
- `COPYSERVICES_SERVERPORT`

DEVICE_CLASS XIV system Storage System parameters

The parameters that are defined in the device class section of the IBM Spectrum Protect Snapshot profile file, configure IBM Spectrum Protect Snapshot for use with the IBM XIV Storage System.

BACKUP_HOST_NAME

This parameter specifies the name of the backup host that is used during offloaded tape backups only. The following list identifies the possible options:

backup_server_hostname

Enter the host name or cluster name of the backup server as configured on the XIV system Storage System.

None

This option is used if you do not have a backup server.

Default

None

Advanced mode only

No.

CLONE_DATABASE

This parameter is preset by the setup script. If you use the setup script for configuration, it is not necessary to manually update any parameters. The following list identifies the possible options:

YES

Use the device class for cloning. When the parameter is set to YES, the device class is unavailable for non-cloning backup or restore operations. The device class is ignored during backup expiration and reconciliation processing.

NO Do not use the device class for cloning. When the parameter is set to NO, any cloning request fails with an error message and return code 2.

The following example shows the `CLONE_DATABASE` parameter that is specified in the `DEVICE_CLASS device` section of the profile:

```
>>> DEVICE_CLASS STANDARD
CLONE_DATABASE YES
COPYSERVICES_HARDWARE_TYPE XIV
PATH_TO_XCLI /home/xivtest/XCLI
COPYSERVICES_SERVERNAME nextra
COPYSERVICES_USERNAME admin
# RECON_INTERVAL 12
# USE_WRITABLE_SNAPSHOTS AUTO
BACKUP_HOST_NAME acsback5
<<<
```

Default

This parameter is not explicitly set. The setup script sets the value, depending on if the device class is specified in the `CLIENT` or `CLONING` section.

Advanced mode only

No.

COPYSERVICES_HARDWARE_TYPE

This parameter is required. Only one device can be specified.

XIV

Specify the XIV option, when the database is stored on the XIV system Storage System.

On the console, any notifications that refer to IBM XIV Storage System operations and **COPYSERVICES_HARDWARE_TYPE** are displayed as **COPYSERVICES_HARDWARE_TYPE=GENERIC**. Similarly, when you view the log or trace files in the ACS_DIR/logs directory, any references that are related to the **COPYSERVICES_HARDWARE_TYPE** for the XIV system Storage System are displayed as **COPYSERVICES_HARDWARE_TYPE=GENERIC**.

Default

Not available.

Advanced mode only

No.

COPYSERVICES_SERVERNAME

This parameter identifies the TCP/IP host name of the storage system where the data to protect is located.

Default

None

Advanced mode only

No.

COPYSERVICES_USERNAME

This parameter identifies the user name. Use the *XIV user* name that you use log on to the XIV system Storage System.

Default

superuser

Advanced mode only

No.

LVM_MIRRORING

Set this parameter to YES if your volume groups use AIX Logical Volume Manager mirroring.

Default

No.

Advanced mode only

Yes.

RECON_INTERVAL

This parameter specifies the interval, in hours, between two subsequent reconciliation operations. The options are whole numbers between 0 and 24 inclusive.

Default

12

Advanced mode only

Yes.

PATH_TO_XCLI

This parameter specifies the path where the XIV command-line interface, XCLI, is installed. There is no default value. This parameter is only valid when **COPYSERVICES_HARDWARE_TYPE** specifies XIV.

Default

None.

Advanced mode only

No.

USE_WRITABLE_SNAPSHOTS

This parameter determines whether writable snapshots are used. Writable snapshots are required in LVM mirrored environments. The following list identifies the options:

YES Writable snapshots are used.

NO Writable snapshots are not used.

AUTO Based on the environment, the value is automatically selected.

Default

AUTO

Advanced mode only

Yes

Storwize family and SAN Volume Controller Storage System parameters**DEVICE_CLASS parameters for static target allocation:**

The parameters that are defined in the device class section of the IBM Spectrum Protect Snapshot profile file, configure IBM Spectrum Protect Snapshot for use with the IBM Storwize family or IBM System Storage SAN Volume Controller storage systems.

CLONE_DATABASE

This parameter is preset by the setup script. If you use the setup script for configuration, it is not necessary to manually update any parameters. The following list identifies the possible options:

YES Use the device class for cloning. When the parameter is set to YES, the device class is unavailable for non-cloning backup or restore operations. The device class is ignored during backup expiration and reconciliation processing.

NO Do not use the device class for cloning. When the parameter is set to NO, any cloning request fails with an error message and return code 2.

Default

This parameter is not explicitly set. The setup script sets the value, depending on if the device class is specified in the CLIENT or CLONING section.

Advanced mode only

No

COPYSERVICES_HARDWARE_TYPE

This parameter is required. Only one device can be specified.

SVC

Specify the *SVC* option, when the database is stored on either the SAN Volume Controller or the Storwize V7000 storage system.

Tip: You must manually create backup target volumes in advance on the storage system.

Default

Not available

Advanced mode only

No

COPYSERVICES_USERNAME

This parameter identifies the user name. Use the *SVC user* name that you use to log on to the SAN Volume Controller master console or cluster. For Storwize V7000, use the *Storwize V7000 user* name that you use to log on to the Storwize V7000.

Default

superuser

Advanced mode only

No

RECON_INTERVAL

This parameter specifies the interval, in hours, between two subsequent reconciliation operations. The options are whole numbers between 0 and 24 inclusive.

Default

12

Advanced mode only

Yes

LVM_MIRRORING

Set this parameter to YES if your volume groups use AIX Logical Volume Manager mirroring.

Default

No.

Advanced mode only

Yes.

COPYSERVICES_COMMPROTOCOL

This parameter identifies the protocol to be used for communication with the CIM Agent. The options are HTTP, for communication in a non-secure mode, and HTTPS, for communication in a secure mode.

Default

HTTPS

Advanced mode only

Yes

COPYSERVICES_CERTIFICATEFILE

When **COPYSERVICES_COMMPROTOCOL** is set to HTTPS, there are two options:

certificate_filename

Name of a certificate file that is created for secure communication between the CIM Client and the CIM Agent.

NO_CERTIFICATE

Select for null trust provider mode.

By default, the CIM Agent for DS8000, which is preinstalled on the HMC, requires communication in secure mode. For this scenario, clients such as IBM Spectrum Protect Snapshot must connect by using HTTPS instead of HTTP. This connection requires that the CIM Client obtain the public key that is used for encryption from the *truststore* certificate in the CIM Agent. After the client obtains the public key, the CIM Client is authenticated by using the user name and password.

To enable the HTTPS protocol, the IBM Spectrum Protect Snapshot profile parameter **COPYSERVICES_COMMPROTOCOL** must specify HTTPS. For this scenario, the **COPYSERVICES_CERTIFICATEFILE** parameter can define a certificate file name, and IBM Spectrum Protect Snapshot exports the certificate by using this file.

The CIM Agent also provides another communication mode that is known as *null trust provider*. In this scenario, the CIM Agent does not verify that the certificate passed by the client matches a known certificate. Rather, it accepts any certificate from the client, including a null string for the file name. To enable this mode, the value of **COPYSERVICES_CERTIFICATEFILE** must be **NO_CERTIFICATE**. This mode is used only if the production and backup systems, and the storage system, are protected by a firewall. If **NO_CERTIFICATE** is used, the `cimom.properties` parameter **DigestAuthentication** must be set to `false`.

Default

NO_CERTIFICATE

Advanced mode only

Yes

COPYSERVICES_PRIMARY_SERVERNAME

This parameter identifies the server name or address that defines the TCP/IP address of the host that is running the CIM Agent for DS Open API. This host manages the SAN Volume Controller master console and the embedded CIM Agent in the Storwize V7000 storage system. For SAN Volume Controller, the **COPYSERVICES_PRIMARY_SERVERNAME** parameter, if specified, must point directly to the SAN Volume Controller cluster with the embedded CIM server. For Storwize V7000, the **COPYSERVICES_PRIMARY_SERVERNAME** parameter must point to the Storwize V7000 cluster.

Default

localhost

Advanced mode only

No

COPYSERVICES_SERVERPORT

This parameter identifies the server port number on the CIM Agent for DS Open API. This information is used to manage the primary and secondary Copy Services servers of the SAN Volume Controller master console or the embedded CIM Agent on the Storwize V7000 storage system.

Default

The default port number depends on the settings of **COPYSERVICES_HARDWARE_TYPE** and **COPYSERVICES_COMMPROTOCOL**:

COPYSERVICES_HARDWARE_TYPE	COPYSERVICES_COMMPROTOCOL	Default Port
SVC	HTTPS	5989
	HTTP	5988

Advanced mode only

Yes

COPYSERVICES_TIMEOUT

This parameter identifies the maximum length of time, in minutes, that the CIM Client waits for a response to a call put to the CIMOM (CIM Agent). If the CIM Client does not receive a response within this time, an error message is displayed.

Default

6

Advanced mode only

Yes

FLASHCOPY_TYPE

This parameter specifies whether the storage solution does a bit-level copy of data from one logical volume to another. This parameter applies to any FlashCopy storage system. The following options are available:

COPY Directs the storage system to run a bit-level copy of the data from one physical volume to another. Specify this value when the following conditions are true:

- A fast snapshot restore of a backed-up database is required.
- A complete copy of the database data on the target volume is required.

NOCOPY Directs the storage system to run a bit-level copy of a track if the data is modified after the initial FlashCopy request. This technique is typically referred as copy-on-write. This option applies only to FlashCopy devices. Specify this value when the following conditions are true:

- A complete copy of the source volumes that contain the database files is not required on the target volumes.
- Backup time constraints are a concern.

INCR This option is similar to the COPY option but the INCR option copies only those tracks that were modified since the previous incremental FlashCopy was created. This option applies only to FlashCopy devices. Specify this value when the following conditions are true:

- IBM Spectrum Protect backups are taken from disk copies. This type of backup creates less burden on the storage system than for the COPY option.
- A snapshot restore operation of the backed up database is to be completed.
- More frequent backups for the database are scheduled.

The **SVC_COPY_RATE** parameter is forced to 0 when the **FLASHCOPY_TYPE** parameter is specified as NOCOPY.

Default

COPY

Advanced mode only

No

RESTORE_FORCE

This parameter specifies whether to force a restore. During a rerun of a snapshot restore, the message FMM0200E can be generated. This problem occurs if the background copy process of the previous snapshot restore is still running and the **RESTORE_FORCE** parameter is set to NO. There are two ways to resolve the issue that is identified by the message:

- Wait until the background copy process ends.
- Set the **RESTORE_FORCE** parameter to YES in the profile configuration file and try the snapshot restore again. This option withdraws all existing source and target relationships, and creates new source and target relationships. A full copy is completed. If you want to set **RESTORE_FORCE** to YES for a specific restore, you can create a temporary profile configuration file.

Default

NO

Advanced mode only

Yes

TARGET_SETS

This parameter specifies the target volumes to be used in the FlashCopy operation. The following list identifies the possible options:

VOLUMES_FILE

The name of the target volumes file (.fct).

list_of target_set_names

A list of target set names. For example: TARGET_SETS 1 2 3

To define the naming convention for the target volumes, specify the **TARGET_NAMING** parameter. For example: TARGET_NAMING *string_with_wildcards_%SOURCE_and_%TARGETSET*

This parameter and option define the naming convention for target volumes. When a backup volume is required, IBM Spectrum Protect Snapshot determines the name of the target set for the operation and the name of the source volume to be backed up. The name of the target volume that stores the backup is the name that is specified after the following strings are replaced with the respective values in the operation: *%SOURCE_and_%TARGETSET*.

Default

None

Advanced mode only

No

VOLUMES_FILE

This parameter specifies the name of the target volumes file (.fct).

Default

None

Advanced mode only

No

ALLOW_NOCOPY_FLASHCOPY

Use this parameter with the **CLONE_DATABASE** parameter. The following list identifies the possible options:

YES Create an IBM Spectrum Protect Snapshot clone on space-efficient targets. For this device class, use space-efficient targets and set **FLASHCOPY_TYPE** to NOCOPY. FlashCopy backups cannot be stored on the same source volumes.

NO Do not create an IBM Spectrum Protect Snapshot clone on space-efficient targets. If both backup and cloning must be completed

on the same source volumes, cloning is completed to full targets and the **ALLOW_NOCOPY_FLASHCOPY** parameter is set to NO.

Default

NO

Advanced mode only

Yes

ALLOW_ALL_FLASHCOPY_TYPES

Use this parameter when IBM Spectrum Protect Snapshot is configured with **FLASHCOPY_TYPE FULL**, or **FLASHCOPY_TYPE INCR**. Use the parameter when the source volumes are fully allocated and the target volumes are space efficient. The following list identifies the available options:

YES Allows IBM Spectrum Protect Snapshot to be configured to use **FLASHCOPY_TYPE FULL**, or **FLASHCOPY_TYPE INCR** when the source volumes are fully allocated and the target volumes are space efficient.

NO If the source volumes are fully allocated and the target volumes are space efficient, you can set the parameter **FLASHCOPY_TYPE** to **NOCOPY** only.

Default

NO

Advanced mode only

Yes

SVC_CLEAN_RATE

This parameter specifies the cleaning rate for the FlashCopy mapping. A value from 1 to 100 can be entered.

Default

None

Advanced mode only

Yes

SVC_COPY_RATE

This parameter specifies the priority that the SAN Volume Controller or Storwize V7000 gives to the FlashCopy background process for the current backup or restore. A value from 0 to 100 can be entered.

A value of 100 indicates the highest priority, but places the greatest burden on the responsiveness of the storage system. A value of 0 indicates the lowest priority, but suppresses the background copy process and forces the **FLASHCOPY_TYPE** parameter to have the **NOCOPY** option.

Default

50

Advanced mode only

No

SVC_GRAIN_SIZE

This parameter specifies the grain size, in KB, for FlashCopy mapping for space-efficient virtual disks on SAN Volume Controller or Storwize V7000. The grain size of the space-efficient virtual disk must match the grain size of the FlashCopy. The options for this parameter are 64, and 256.

After the parameter is set, the value cannot be changed until the backup is deleted with the option -F to remove the mappings.

Default

256

Advanced mode only

Yes

DEVICE_CLASS parameters for dynamic target allocation:

The parameters that are defined in the device class section of the IBM Spectrum Protect Snapshot profile file, configure IBM Spectrum Protect Snapshot for use with IBM Storwize family or IBM System Storage SAN Volume Controller storage systems.

CLONE_DATABASE

This parameter is preset by the setup script. If you use the setup script for configuration, it is not necessary to manually update any parameters. The following list identifies the possible options:

YES Use the device class for cloning. When the parameter is set to YES, the device class is unavailable for non-cloning backup or restore operations. The device class is ignored during backup expiration and reconciliation processing.

NO Do not use the device class for cloning. When the parameter is set to NO, any cloning request fails with an error message and return code 2.

Default

This parameter is not explicitly set. The setup script sets the value, depending on if the device class is specified in the CLIENT or CLONING section.

Advanced mode only

No

COPYSERVICES_HARDWARE_TYPE

This parameter is required. Only one device can be specified.

SVCDDTA

Specify the SVCDDTA option when the storage system is SAN Volume Controller or Storwize V7000 and you require the target volumes to be dynamically allocated during the backup process.

Default

None

Advanced mode only

No

COPYSERVICES_SERVERNAME

Defines the TCP/IP host name of the storage system where the application data to protect is allocated.

Default

None

Advanced mode only

No

COPYSERVICES_USERNAME

Identifies the user name. Specify the user name that is used to log on to the SAN Volume Controller cluster. For Storwize V7000, specify the Storwize V7000 user name.

Default

superuser

Advanced mode only

No

SVC_SSHKEY_FULLPATH

Specifies the path and the file name to the private SSH key file. The key file is used to authenticate to the storage system with the user name specified for the **COPYSERVICES_USERNAME** parameter. In an Oracle SAP environment, both the Oracle and `sidadm` users can start Snapshot operations. The private ssh key must be located in the home directory `$HOME/.ssh` for both users.

Default

\$HOME/.ssh/svc_sshkey

Advanced mode only

Yes

SVC_REMOTE_SSHKEY_FULLPATH

This parameter specifies the second SSH key file to be used for authentication on the remote site storage device. The key file is used to authenticate to the storage system with the user name specified for the **COPYSERVICES_REMOTE_USERNAME** parameter. If you do not want to create a new key pair for the remote site, one key can be shared for both storage sites.

Default

\$HOME/.ssh/svc_sshkey

Advanced mode only

Yes

SSH_DIR

Specifies the path to the Secure Shell protocols and executable files.

Default

/usr/bin

Advanced mode only

Yes

SVC_COPY_RATE

Specifies the priority that the storage system gives to the FlashCopy background process for the current backup or restore operation. Enter a value from the range 1 - 100.

The **SVC_COPY_RATE** parameter only applies for full copy backups (FLASHCOPY_TYPE COPY). For space-efficient backups (FLASHCOPY_TYPE NOCOPY), the copy rate is implicitly set to 0.

Default

0

Advanced mode only

Yes

LVM_MIRRORING

Set this parameter to YES if your volume groups use AIX Logical Volume Manager mirroring.

Default

No.

Advanced mode only

Yes.

FLASHCOPY_TYPE

Specifies whether the storage solution does a bit-level copy of data from one logical volume to another. This parameter applies to any FlashCopy storage system. The following options are available:

COPY Directs the storage system to run a bit-level copy of the data from one physical volume to another. Specify this value when the following conditions are true:

- A fast snapshot restore of a backed-up database is required.
- A complete copy of the database data on the target volume is required.

NOCOPY Directs the storage system to run a bit-level copy of a track if the data is modified after the initial FlashCopy request. This technique is typically referred as copy-on-write. Specify this value when the following conditions are true:

- A complete copy of the source volumes that contain the database files is not required on the target volumes.
- A fast snapshot restore of a backed-up database is required.
- Backup time constraints are a concern.

Default

NOCOPY

Advanced mode only

No

SVC_GRAIN_SIZE

Specifies the grain size, in KB, for FlashCopy mapping for space-efficient virtual disks on SAN Volume Controller or Storwize V7000. The grain size of the space-efficient virtual disk must match the grain size of the FlashCopy. The options for this parameter are 64, and 256.

After the parameter is set, the value cannot be changed until the backup is deleted with the option -F to remove the mappings.

Note: When you are migrating from the SVC adapter with static target allocation, you must ensure that the grain size for the new SVCDDTA device classes is set to the same value as it was for the device classes for SVC.

Default

256

Advanced mode only

Yes

SVC_POOLNAME

This parameter specifies the name of the storage pool that is used to create target volumes for the FlashCopy backups. A value must be assigned if a source volume has two copies in the SVC, and these copies are in two different

storage pools. If the `DEVICE_CLASS` is configured for remote site backup `COPYSERVICES_REMOTE YES`, the specified pool name is related to the remote site storage device.

Default

Name of the storage pool where the source volume is located.

Advanced mode only

Yes

SVC_IOGROUP

Specifies the name of the input and output (IO) group, which is used to create target volumes for the FlashCopy backups. If the `DEVICE_CLASS` is configured for remote site backup `COPYSERVICES_REMOTE YES`, the specified IO group is related to the remote site storage device.

Default

Name of the IO group on the source volume where the FlashCopy relationship is established.

Advanced mode only

Yes

SVC_MOUNT_POOLNAME

Specifies the name of the storage pool that is used to create temporary duplicates of the target volumes of a FlashCopy backup, which then mounts to a host. If the `DEVICE_CLASS` is configured for remote site backup `COPYSERVICES_REMOTE YES`, the specified pool name is related to the remote site storage device.

Default

Name of the storage pool on the target volume that is used to create duplicate volumes for the mount operation.

Advanced mode only

Yes

SVC_MOUNT_IOGROUP

Specifies the name of the IO group, which is used to create duplicate volumes for the mount operation. If the `DEVICE_CLASS` is configured for remote site backup `COPYSERVICES_REMOTE YES`, the specified IO group is related to the remote site storage device.

Default

Name of the IO group on the target volume that is used to create duplicate volume for the mount operation.

Advanced mode only

Yes

SVC_TARGET_VOLUME_REAL_SIZE

Specify the percentage of the source volume size to allocate, which is used to create the actual target volumes during the backup operation.

The `SVC_TARGET_VOLUME_REAL_SIZE` parameter only applies to `FLASHCOPY_TYPE NOCOPY`

Default

10

Advanced mode only

Yes

RECON_INTERVAL

This parameter specifies the interval, in hours, between two subsequent reconciliation operations. The options are whole numbers between 0 and 24 inclusive.

Default

12

Advanced mode only

Yes

DEVICE_CLASS DS8000 Storage System parameters

The parameters that are defined in the device class section of the IBM Spectrum Protect Snapshot profile file, configure IBM Spectrum Protect Snapshot for use with the IBM System Storage DS8000.

BACKUP_HOST_NAME

This parameter specifies the name of the backup host that is used during offloaded tape backups only. The following list identifies the possible options:

PREASSIGNED_VOLUMES

Specify this option when the target volumes are preassigned to a specific backup server.

None

This option is used if you do not have a backup server.

Default

None.

Advanced mode only

No.

CLONE_DATABASE

This parameter is preset by the setup script. If you use the setup script for configuration, it is not necessary to manually update any parameters. The following list identifies the possible options:

YES Use the device class for cloning. When the parameter is set to YES, the device class is unavailable for non-cloning backup or restore operations. The device class is ignored during backup expiration and reconciliation processing.

NO Do not use the device class for cloning. When the parameter is set to NO, any cloning request fails with an error message and return code 2.

Default

This parameter is not explicitly set. The setup script sets the value, depending on if the device class is specified in the CLIENT or CLONING section.

Advanced mode only

No

COPYSERVICES_HARDWARE_TYPE

This parameter is required. Only one device can be specified.

DS8000

Specify the DS8000 option, when the database is stored on one of the following storage systems:

- IBM DS8100
- IBM DS8300

- IBM DS8700
- IBM DS8800
- IBM DS8870

Default

None.

Advanced mode only

No.

COPYSERVICES_USERNAME

This parameter identifies the user name, use the *cim user* of the CIM Agent for DS Open API. The CIM Agent for DS Open API manages the primary and secondary copy services servers of the DS8000 cluster.

Default

superuser

Advanced mode only

No.

RECON_INTERVAL

This parameter specifies the interval, in hours, between two subsequent reconciliation operations. The options are whole numbers between 0 and 24 inclusive.

Default

12

Advanced mode only

Yes

LVM_MIRRORING

Set this parameter to YES if your volume groups use AIX Logical Volume Manager mirroring.

Default

No.

Advanced mode only

Yes.

COPYSERVICES_COMMPROTOCOL

This parameter identifies the protocol to be used for communication with the CIM Agent. The options are HTTP, for communication in a non-secure mode, and HTTPS, for communication in a secure mode.

Default

HTTPS

Advanced mode only

Yes.

COPYSERVICES_CERTIFICATEFILE

When **COPYSERVICES_COMMPROTOCOL** is set to HTTPS, there are two options:

certificate_filename

Name of a certificate file that is created for secure communication between the CIM Client and the CIM Agent.

NO_CERTIFICATE

Select for null trust provider mode.

By default, the CIM Agent for DS8000, which is preinstalled on the HMC, requires communication in secure mode. For this scenario, clients such as IBM Spectrum Protect Snapshot must connect by using HTTPS instead of HTTP. This connection requires that the CIM Client obtain the public key that is used for encryption from the *truststore* certificate in the CIM Agent. After the client obtains the public key, the CIM Client is authenticated by using the user name and password.

To enable the HTTPS protocol, the IBM Spectrum Protect Snapshot profile parameter **COPYSERVICES_COMMPROTOCOL** must specify HTTPS. For this scenario, the **COPYSERVICES_CERTIFICATEFILE** parameter can define a certificate file name, and IBM Spectrum Protect Snapshot exports the certificate by using this file.

The CIM Agent also provides another communication mode that is known as *null trust provider*. In this scenario, the CIM Agent does not verify that the certificate passed by the client matches a known certificate. Rather, it accepts any certificate from the client, including a null string for the file name. To enable this mode, the value of **COPYSERVICES_CERTIFICATEFILE** must be **NO_CERTIFICATE**. This mode is used only if the production and backup systems, and the storage system, are protected by a firewall. If **NO_CERTIFICATE** is used, the `cimom.properties` parameter **DigestAuthentication** must be set to `false`.

Default

`NO_CERTIFICATE`

Advanced mode only

Yes.

COPYSERVICES_PRIMARY_SERVERNAME

This parameter identifies the server name or address that defines the TCP/IP address of the host that is running the CIM Agent for DS Open API. This host manages the primary and secondary copy services servers of the DS8000 cluster.

Default

`localhost`

Advanced mode only

No.

COPYSERVICES_SECONDARY_SERVERNAME

This parameter identifies the name of the backup Copy Services server that is located within a snapshot devices cluster. Specify either the IP address or the server DNS name. This parameter can be used only in environments with DS8000 in combination with the proxy CIM Agent.

Default

`None`

Advanced mode only

Yes.

COPYSERVICES_SERVERPORT

This parameter identifies the server port number of the host that is running the CIM Agent for DS Open API.

Default

The default port number depends on the settings of **COPYSERVICES_HARDWARE_TYPE** and **COPYSERVICES_COMMPROTOCOL**:

COPYSERVICES_HARDWARE_TYPE	COPYSERVICES_COMMPROTOCOL	Default Port
DS8000	HTTPS	6989
	HTTP	6988

Advanced mode only

Yes.

COPYSERVICES_TIMEOUT

This parameter identifies the maximum length of time, in minutes, that the CIM Client waits for a response to a call sent to the CIMOM (CIM Agent). If the CIM Client does not receive a response within this time, an error message is sent.

Default

6

Advanced mode only

Yes.

FLASHCOPY_TYPE

This parameter specifies whether the storage solution does a bit-level copy of data from one logical volume to another. This parameter applies to any FlashCopy storage system. The following options are available:

COPY Directs the storage system to run a bit-level copy of the data from one physical volume to another. Specify this value when the following conditions are true:

- A fast snapshot restore of a backed-up database is required.
- A complete copy of the database data on the target volume is required.

NOCOPY Directs the storage system to run a bit-level copy of a track if the data is modified after the initial FlashCopy request. This technique is typically referred as copy-on-write. This option applies only to FlashCopy devices. Specify this value when the following conditions are true:

- A complete copy of the source volumes that contain the database files is not required on the target volumes.
- Backup time constraints are a concern.

INCR This option is similar to the COPY option but the INCR option copies only those tracks that were modified since the previous incremental FlashCopy was created. This option applies only to FlashCopy devices. Specify this value when the following conditions are true:

- IBM Spectrum Protect backups are taken from disk copies. This type of backup creates less burden on the storage system than for the COPY option.
- A snapshot restore operation of the backed up database is to be completed.
- More frequent backups for the database are scheduled.

There must be only one target set specified in the target volumes file (.fct) for incremental snapshots. CIM errors might occur when more than one target set is specified. A successful backup of the database to the IBM Spectrum Protect server is possible even if the parameter is set to NOCOPY.

Default

COPY

Advanced mode only

No.

RESTORE_FORCE

This parameter specifies whether to force a restore. During a rerun of a snapshot restore, the message FMM0200E can be generated. This problem occurs if the background copy process of the previous snapshot restore is still running and the **RESTORE_FORCE** parameter is set to NO. There are two ways to resolve the issue that is identified by the message:

- Wait until the background copy process ends.
- Set the **RESTORE_FORCE** parameter to YES in the profile configuration file and try the snapshot restore again. This option withdraws all existing source and target relationships, and creates new source and target relationships. A full copy is completed. If you want to set **RESTORE_FORCE** to YES for a specific restore, you can create a temporary profile configuration file.

Default

NO

Advanced mode only

Yes

TARGET_SETS

This parameter specifies the target volumes to be used in the FlashCopy operation. The following list identifies the possible options:

VOLUMES_FILE

The name of the target volumes file (.fct).

Default

None.

Advanced mode only

No.

VOLUMES_FILE

This parameter specifies the name of the target volumes file (.fct).

Default

None.

Advanced mode only

No.

OFFLOAD

The OFFLOAD section of the profile configuration contains information that is related to IBM Spectrum Protect backups from a snapshot.

File names that are specified in the offload section typically point to files that are on a backup server. The offload section is optional and can exist for Oracle, and Oracle in an SAP environment. The parameters do not depend on the storage device.

The following list provides the parameters, a description of each parameter, and default values applicable for Oracle, and Oracle in an SAP environment.

BACKUP_METHOD

This parameter is preset by the setup script (the profile configuration wizard). The setup script value depends on the environment where the setup script is running.

- Oracle - ORACLE
- Oracle in an SAP environment - BACKINT.

Default

Preset by the setup script, according to the environment.

Advanced mode only

Yes.

OFFLOAD Oracle parameters

The following list provides the parameters, a description of each parameter, and default values that are applicable in Oracle environments.

OVERWRITE_DATABASE_PARAMETER_FILE

This parameter is used with Oracle and Oracle in an SAP environment databases only. The parameter specifies if the database configuration file on the backup server is overwritten with the file from the production server. The database configuration file on the backup server is used to mount the database so that it can be offloaded to IBM Spectrum Protect. Specify one of the following values:

YES

Automatically copies the database configuration file to the backup system with the version defined on the production system.

NO Do not copy the production system database configuration file to the backup system.

If the production database is configured to use a text-based Oracle parameter file (PFILE), this setting requires that the default database configuration file `{ORACLE_HOME}/dbs/init${ORACLE_SID}.ora` is available and valid on the backup system. If the name of the database configuration file on the production database is not the default file name, `{ORACLE_HOME}/dbs/init${ORACLE_SID}.ora`, use the **TARGET_DATABASE_PARAMETER_FILE** parameter, in the ORACLE section, to specify the correct name of the configuration file. If the production database is configured to use a binary Oracle SPFILE, this setting requires that the database configuration file `fileOracle_instance_owner_${HOME_directory}/acs/tempfiles/init${ORACLE_SID}.ora_fromSPfile` is available and valid on the backup system. If you need a modified version of the SPFILE from the production system, set **OVERWRITE_DATABASE_PARAMETER_FILE YES** to create this file automatically with a first backup and offload, then switch **OVERWRITE_DATABASE_PARAMETER_FILE NO** and modify this file for subsequent backups and offloads.

Default

YES

Advanced mode only

Yes.

DATABASE_BACKUP_INCREMENTAL_LEVEL

This parameter specifies the level of backup. Any numerical value can be entered. The following conditions apply:

- To complete a full backup, use option 0. A full backup must be completed before an incremental backup can be run.
- To complete an incremental backup, enter a numerical value greater than 0. Incremental backups are progressive. For example, a level 0 backup must complete before a level 1 backup can start. A level 1 backup must be complete before a level 2 backup can occur.

Default

0

Advanced mode only

No.

ASM_INSTANCE_USER

This parameter is used for the backup server. If this parameter is not specified for the OFFLOAD section, the value of this parameter, as specified in the ORACLE section, is used for the backup server.

user name

Specify the user name of the ASM instance owner. Use this parameter when the target database and the ASM instance are running under different user IDs. The ASM instance has one of the following permissions: sysdba, sysasm, or sysadm.

AUTO

When this parameter is set to AUTO, the database user who is running the process is used.

Default

There is no default value.

Advanced mode only

Yes.

ASM_INSTANCE_ID

This parameter specifies the SID of the ASM instance. This parameter is used for the backup server. If this parameter is not specified for the OFFLOAD section, the value of this parameter, as specified in the ORACLE section, is used for the backup server.

You can have a SID for the ASM instance other than +ASM. In this scenario, this profile parameter specifies the ASM instance SID.

Default

0

Advanced mode only

Yes.

ASM_ROLE

Specify the role that is used when you are connecting to the ASM instance. When you use Oracle 11g, specify the sysasm role.

Default

0

Advanced mode only

Yes.

OFFLOAD parameters for Oracle in an SAP environment

The following list provides the parameters, a description of each parameter, and default values applicable in an Oracle in an SAP environment:

PROFILE

This required parameter is used to specify the name of the external SAP backint profile (init<SID>.utl).

Default

There is no default value.

Advanced mode only

Yes.

Changing profile parameters

Except for the GLOBAL and ACSD sections, changes to the profile take effect immediately and do not require restarting IBM Spectrum Protect Snapshot. Updates to the GLOBAL and ACSD sections require a restart of IBM Spectrum Protect Snapshot.

About this task

To change the GLOBAL and ACSD sections, complete the following steps:

Procedure

1. For each system where IBM Spectrum Protect Snapshot is installed, enter the following command to stop IBM Spectrum Protect Snapshot:
`setup_type.sh -a stop`
2. Start the setup script by entering the appropriate command for your database environment: `./setup_ora.sh`
To use the advanced mode, use the `-advanced` option with the appropriate setup script command. In the advanced mode, you can specify more parameters.
3. Follow the setup script instructions that are displayed.
4. For each system where IBM Spectrum Protect Snapshot is installed, enter the following command to start IBM Spectrum Protect Snapshot:
`setup_type.sh -a start -d Instance_owner_${HOME} directory`

Interdependency of LVM_FREEZE_THAW and TARGET_DATABASE_SUSPEND

The **LVM_FREEZE_THAW** and **TARGET_DATABASE_SUSPEND** parameters are interdependent.

These two IBM Spectrum Protect Snapshot profile parameters are interdependent in the following manner:

- If **LVM_FREEZE_THAW** is set to YES, the database must be suspended. Otherwise, write operations to the database might time out and leave the database in an inconsistent state. A specified value of YES for **TARGET_DATABASE_SUSPEND** prevents this situation.
- If **LVM_FREEZE_THAW** is set to NO, the user might want to suspend the database without freezing the file system. Also, if JFS is used, freeze and thaw are not supported.
- If **LVM_FREEZE_THAW** is set to AUTO, and the file systems support the freeze function, the effect of AUTO is described in the following table. If the file systems do not support the freeze function, the AUTO value resolves to NO.

For Oracle ASM environments, **TARGET_DATABASE_SUSPEND** is independent of **LVM_FREEZE_THAW**, and **LVM_FREEZE_THAW** is not allowed for ASM.

The following table summarizes the actions taken depending on the values of the two parameters:

Table 16. Actions taken depending on values of LVM_FREEZE_THAW and TARGET_DATABASE_SUSPEND

Value of LVM_FREEZE_THAW	Value of TARGET_DATABASE_SUSPEND		
	YES	NO	OFFLINE
YES	Suspend and freeze	Terminate with an appropriate error message. Conflicting parameters.	Offline with freeze
NO	Suspend, no freeze	No suspend, no freeze	Offline without freeze
AUTO	Treat as LVM_FREEZE_THAW YES	Treat as LVM_FREEZE_THAW NO	Offline with freeze

Target set and target volumes

FlashCopy backups on DS8000, SAN Volume Controller, and Storwize V7000, require a target set for each set of source volumes to be backed up. The target set is a set of target volumes, and several target sets can be defined for use in different FlashCopy backups. The target volumes file, with extension .fct, identifies the target volumes to be used for an IBM Spectrum Protect Snapshot backup.

The volumes in each target set that are used in a backup, must be specified in a separate target set. These target sets are specified in a target volumes file, the .fct file. The target set section name begins with the prefix **TARGET_SET** and is appended with a target set name. The target set name differentiates different target set sections. The target set name can be any alphanumeric value.

In the **TARGET_SET**, use the **TARGET_VOLUME** parameter for every target volume in the target set as shown in the following example:

```
>>> TARGET_SET 1
TARGET_VOLUME ...
.
.
TARGET_VOLUME ...
<<<
```

To specify multiple target sets in the target volumes file, add the next target set section with a unique target set ID as shown in this example:

```
>>> TARGET_SET 2
TARGET_VOLUME ...
.
.
TARGET_VOLUME ...
<<<
```

Comments can be entered before the first target set section only, and are indicated by a “#” character in the first column of each line. Tab characters can be entered.

When **VOLUMES_FILE** is specified in the profile, the target volumes file can have any file name and does not conform to any naming convention.

Target set definitions are not required for XIV system.

Related concepts:

Appendix C, “Examples,” on page 193

Manage target volumes files for your storage system

Different storage systems require different methods of target volume mapping. Use the **VOLUMES_FILE** parameter to share a target volume file between multiple device classes.

DS8000 and SAN Volume Controller, and Storwize V7000 storage systems, need the **TARGET_SETS** parameter to specify the target volumes file, **VOLUMES_FILE**. For XIV system, target LUNs are created automatically without the target volumes files, as shown in the following table:

Table 17. Managing target volume LUNs by storage system

DS8000	SAN Volume Controller and Storwize V7000	XIV system
Manual target LUN creation with the target volumes file (.fct) that defines the VOLUMES_FILE parameter.	Manual target LUN creation with the target volumes file (.fct) that defines the VOLUMES_FILE parameter. Or, Naming convention that defines the TARGET_NAMING parameter.	Automatic target LUN creation without using target volumes file (.fct).

For IBM Spectrum Protect Snapshot to associate a target volume to a source volume, the following criteria must be met:

- The source volume and target volume must be in the same storage system.
- The source volume and target volume must be the same size.

A target volume is selected for validation as a suitable target volume for the source volume depending on the value of the parameter **TARGET_SETS**.

VOLUMES_FILE

The **VOLUMES_FILE** parameter is used to share a target volume file between multiple device classes by restricting a target set to a specific **DEVICE_CLASS**. The target volume is validated as suitable for the source volume based on the value of the **TARGET_SETS** parameter. The following criteria must be in place for a valid target volume:

- A target volumes file, .fct, must be specified.
- A list of target volumes must be specified in the target volumes file. The source volumes and the size are optional.

This example shows the syntax of target volumes files that are specified by the **VOLUMES_FILE** parameter:

```
>>> TARGET_SET <target set name>

DEVICE_CLASS <device class name> # this parameter is optional and allows to
                                # restrict the use of this target set to a
                                # specific device class

>>> PARTITION <name of partition> # e.g. NODE0000 for partition 0 or NODE0001 for
#partition 1, ...
TARGET_VOLUME <target> [<source>] [<size>]
[...]
```

[...]

<<<

[...]

If no source is specified in the **TARGET_SETS** parameter and a FlashCopy relation exists between target volumes and a source volume, IBM Spectrum Protect for Advanced Copy Services checks for each of the specified target volumes. If a FlashCopy relation exists, it is reused for the next FlashCopy backup. However, if no FlashCopy relation exists to a source volume, a new relation between one source volume and the target is created with the next FlashCopy backup. In this case, the created source-target pairs are unpredictable because they depend on the order of the target volumes as listed in the target volumes file. There is also a dependency on the order of the source volumes as they occur in the operating system. If you want predefined source-target pairs, you must specify the dedicated source volume for each of the target volumes in the target volumes file. Alternatively you can ensure that all FlashCopy relations exist in the storage system before the start of the FlashCopy backup.

Related reference:

“DS8000 target volume parameter settings”

“SAN Volume Controller and Storwize V7000 target volume parameter settings” on page 156

“Target set handling for cloning” on page 158

DS8000 target volume parameter settings

Each target volume that is planned for use must be specified by its serial number for a DS8000 configuration.

A snapshot backup operation looks for either a source volume and target volume correlation, or a target-volume-only specification. A target set definition file contains a list of target volumes that are organized into target sets. IBM Spectrum Protect Snapshot attempts to match source volumes to suitable targets within a target set during backup.

Table 18. Parameters of the 'VOLUMES_SET_x' Topic (DS8000)

Parameter Name	Value
TARGET_VOLUME <target volume serial number> <source volume serial number> <source volume size>	<p>Specify a source serial number with a target serial number in the target set definition file. This action determines source target relations. The relation between the source and target is required. Backup processing fails if one of the targets is unavailable for the specified source.</p> <p>This example shows a configuration where the DS8000 source volume with serial 75924811011 must be used in a FlashCopy with the target volume with serial number 75924811001.</p> <pre>TARGET_VOLUME 75924811001 75924811011 Size=2.0_GB</pre> <p>The source serial number and the size can be omitted or dashes can be entered in both fields as placeholders, as shown in the following example:</p> <pre>TARGET_VOLUME 75924811001 - -</pre> <p>Target volumes must meet the following requirements:</p> <ul style="list-style-type: none"> • The size of the target volume must be the same as the size of the source volume • The source and target volumes that are listed in one TARGET_SET must be in the same storage system • The order of the parameters, target volume serial number, source volume serial number, and size of source volume must not be changed.

Use the **FLASHCOPY_TYPE** parameter for DS8000 and SAN Volume Controller, and Storwize V7000. The following actions are possible:

- Change the **FLASHCOPY_TYPE** value of an existing target set
- Remove a target volume from an existing target set
- Remove a complete target set.

You must use the sequence of commands that are described in “Deleting snapshot backups” on page 171 with the force option.

Related concepts:

“Target set and target volumes” on page 153

Related reference:

“DS8000 target volumes file example” on page 197

SAN Volume Controller and Storwize V7000 target volume parameter settings

Each target volume that is used, must be specified by the corresponding virtual disk name. A snapshot backup operation looks for either a source volume and target volume correlation, or a target-volume-only specification.

A target set definition file contains a list of target volumes that are organized into target sets. During the backup process, the IBM Spectrum Protect Snapshot software attempts to match source volumes to suitable targets within a target set.

Table 19. Parameters of the 'VOLUMES_SET_X' topic (SAN Volume Controller and Storwize V7000)

Parameter Name	Value
TARGET_VOLUME <target volume virtual disk name> <source volume virtual disk name> <source volume size>	<p>Specify a source virtual disk name with a target virtual disk name in the target set definition file. This action determines source target relations. The relation between the source and target is required, backup processing fails if one of the targets is unavailable for the specified source.</p> <p>This example shows a configuration where the SAN Volume Controller source volume with virtual disk name <i>svdfsrc4</i> must be used in a FlashCopy with the target volume with virtual disk name <i>svdftgt4</i>. TARGET_VOLUME svdftgt4 svdfsrc4 Size=2.0_GB</p> <p>The source virtual disk name and the size can be omitted or dashes can be entered in both fields as placeholders, as shown in the following example: TARGET_VOLUME svdftgt4 - -</p> <p>Target volumes must meet the following requirements:</p> <ul style="list-style-type: none"> • The size of the target volume must be the same or greater than the size of the source volume. • The source and target volumes listed in one TARGET_SET must be in the same SAN Volume Controller cluster. • The order of the parameters must not be changed.

For more information about the criteria that are used to associate a target volume to a source volume, see “Target set and target volumes” on page 153.

Use the **FLASHCOPY_TYPE** parameter for DS8000, SAN Volume Controller, and Storwize V7000. The following actions are possible:

- Change the **FLASHCOPY_TYPE** value of an existing target set
- Remove a target volume from an existing target set
- Remove a complete target set.

To complete these types of changes, use the sequence of commands that are described in “Deleting snapshot backups” on page 171 with the force option.

For SAN Volume Controller 6.1 or later and Storwize V7000, with IBM Spectrum Protect Snapshot software you can delete FlashCopy mappings that are not dependent on other FlashCopy mappings. Only the source and target FlashCopy mappings of the oldest backup can be deleted. If multiple backup generations are used and you want to delete a backup that is not the oldest backed up version, the background operation that deletes the mappings is delayed until all older backups are deleted or are reused by a new backup request.

The following example presents a typical Multiple Target FlashCopy (MTFC) cascade:

S->T4->T3->T2->T1

S = Source volume

T1-T4 = Snapshots taken at t1, t2, t3, t4 where T1 is the oldest,
T4 the most recent snapshot

T1 depends on T2,T3,T4,S
T2 depends on T3,T4,S
and so on...

Following the path from *S* to *T4* is called *downstream*. The opposite direction is called *upstream*.

Example 1: T2 is restored

All upstream snapshot mappings are stopped: *T3,T4*

Example 2: T2 is overwritten by a new backup

All downstream snapshot mappings are stopped: *T1*

Related reference:

“SAN Volume Controller and Storwize V7000 target volumes file example” on page 198

Target set handling for cloning

Cloning operations require specific settings for target sets.

The TARGET_SETS profile parameter identifies the target volumes to be used in the FlashCopy operation. This parameter must be specified in the device class section of the profile. You can specify one of these values with cloning operations:

VOLUMES_FILE *name of the target volumes file (.fct)*

Specify the name of the target volumes file (.fct). The USE_FOR_CLONING *list of clone database names* statement identifies the correct target set to use for a specific clone database name. When more than one clone database name is specified in the list, the referenced target set is used for all specified clone database names. Each name that is specified in the list must be separated by a space. In this situation, the target set must be used by those clone databases only that are identified in the list. The USE_FOR_CLONING list of clone database names must be specified in the target volumes file.

TARGET_NAMING *string with wildcards %SOURCE USE_FOR_CLONING list of clone database names*

Available for SAN Volume Controller only. Specify the naming convention for target volumes. When a backup volume is required at backup time, IBM Spectrum Protect Snapshot determines the name of the target set for the current operation and the name of the source volume to be backed up. The name of the volume that stores the backup is the name that is specified when the string %SOURCE is replaced with the respective value in the current operation. The required USE_FOR_CLONING *list of clone database names* statement identifies the correct target set to use for a specific clone database name. When more than one clone database name is specified in the list, the referenced target set is used for all specified clone database names. Each name that is specified in the list must be separated by a space. In this situation, only the target set must be used by those clone databases that are identified in the list. The USE_FOR_CLONING list of clone database names must be specified with the TARGET_NAMING parameter itself. It is possible to have multiple TARGET_NAMING entries in the device class where each represents a different clone database name.

Restriction: For SAN Volume Controller and Storwize V7000, when a new backup is started on a target volume that is not the oldest in the chain, SAN Volume Controller stops all mappings to older target volumes. When a restore is requested from a target volume that is not the youngest in the chain, SAN Volume Controller

stops all mappings to newer target volumes. When a mapping to a target volume stops in either of these situations, this target volume immediately goes offline if any of these conditions exist:

- The target volume is a space-efficient volume.
- The mapping was for an incremental copy that was ongoing.
- The mapping was for a full copy that was ongoing.

As a result, the target volumes for the production database to be cloned, and the target volumes for the FlashCopy backup of the same database, must not be on the same SAN Volume Controller or Storwize V7000 cluster. If you are cloning databases in an AIX Logical Volume Mirroring (LVM) environment, use FlashCopy cloning on one of the SAN Volume Controller or Storwize V7000 clusters and FlashCopy backup on the other SAN Volume Controller or Storwize V7000 cluster. Avoid space-efficient target volumes for cloning. If space-efficient target volumes are used, the profile parameter `ALLOW_NOCOPY_FLASHCOPY YES` must be specified in the cloning device class section of the profile.

Target volumes file (.fct) cloning examples

The target volumes file (specified by the `VOLUMES_FILE` parameter) must have the following syntax for Oracle:

```
>>> TARGET_SET target set name
DEVICE_CLASS <device classes> USE_FOR_CLONING <list of clone database names>
# this parameter is mandatory for FlashCopy Cloning and allows to
# restrict the use of this target set to a specific device class
# and to a specific clone database name
TARGET_VOLUME target [source] [size]
[...]
```

IBM Spectrum Protect Snapshot password file

To access the storage system where the database volumes are stored, IBM Spectrum Protect Snapshot requires a password file.

The password file contains a *master password* that is required by the agents such as application agents or offload agents, when they are authenticating or connecting to the Management Agent. When IBM Spectrum Protect Snapshot agents are running in a distributed environment across multiple servers, separate password file instances can be used for different nodes. In a distributed environment, you must ensure that each local password file instance contains all the passwords that are needed by the agents that are running on the node. The master password must be included in all instances. Use the SSH for the setup to ensure that the password files are replicated to all nodes automatically.

The master password is only prompted for in advanced mode, and is only needed when you are installing IBM Spectrum Protect Snapshot separately on the backup servers or cloning servers without using SSH. In this case, you must know the password. The password is defined when you configure the production server; for backup and cloning servers this password must be used so that the servers can connect to the management agent on the production server. When you use SSH for remote deployment to the backup and cloning servers, the password file is copied to the servers automatically.

A password file is created during the IBM Spectrum Protect Snapshot configuration process. The setup script that is used for the configuration also updates information that is stored in the `/etc/inittab` directory. An example of the path to the password file follows:

```
<ACS_DIR>/shared/pwd.acsd
```

where `<ACS_DIR>` is the value of the `ACS_DIR` parameter in the profile. In basic mode, the password is not prompted as it is generated automatically if it is not set earlier. A generated password is available as the default password in advanced mode.

The minimum length of the master password is 8 characters. The password must contain at least one number and one letter. The use of special symbols increases the strength of the password.

IBM Global Security Kit configuration

IBM Spectrum Protect Snapshot uses the security suite IBM Global Security Kit (GSKit), for Secure Socket Layer (SSL) and Transport Layer Security (TLS) TCP/IP connections. GSKit supports Federal Information Processing Standards (FIPS140-2) and also incorporates the new security standards as defined in the Special Publications 800131 (SP 800-131). GSKit is automatically installed by IBM Spectrum Protect Snapshot.

This security standard requires longer key lengths, stronger cryptographic algorithms, and incorporates TLS Protocol version 1.2.

During the installation, IBM Spectrum Protect Snapshot automatically creates a new key pair and a self-signed certificate if no default certificate exists. The key pair is stored in the local key database file. The self-signed certificate is created from the key pair and automatically distributed to all backup and cloning servers through the existing SSH remote deployment mechanisms.

If you do not use the SSH remote deployment capabilities of IBM Spectrum Protect Snapshot, you must complete the following steps:

1. Manually copy the self-signed certificate `fcmsel fcert.arm` file to the IBM Spectrum Protect Snapshot `INSTANCE_DIR` directory on the backup and cloning servers. The manually copied self-signed certificate is imported automatically when the setup routine is rerun on the backup or cloning servers.
2. Globally install GSKit on each server by running the setup script as root user on the backup or cloning server. The required installation files are available in the `gskit_install` subdirectory of the IBM Spectrum Protect Snapshot `INSTANCE_DIR` directory. The files are visible to the backup and cloning servers. To install GSKit, enter the command, `./setup_ora.sh -a install_gskit -d instance_directory`

If manually copying the self-signed certificate file to the backup and cloning servers is not feasible, as an alternative, use a signed certificate. The signed certificate can be from an internal or external certificate authority (CA). When SP800-131 encryption is enforced, the signed certificate must comply with the standard as defined by the National Institute of Standards and Technology (NIST) SP800-131 standard encryption. This standard requires a minimum key size = 2048 bits and a signature algorithm = RSA with SHA-224 or higher. Import the CA signed certificate to the key database on the production server.

If you use a standard CA-signed certificate, you do not need to handle `fcmselcert.arm` files. You must import the CA-signed certificate manually into the production server key ring. Use the GSKit command-line utilities to import the certificate to the production server. If the CA-signed certificate is not a standard certificate that GSKit has a root certificate for, you must import the certificate to all sites. No further action is necessary on the auxiliary server.

The following GSKit files are installed by IBM Spectrum Protect Snapshot:

- A key database file, `fcncert.kdb`, is in the **INSTANCE_DIR** directory.
The KDB file on the production server contains a new key pair and a self-signed certificate. On the backup and cloning servers, the KDB file contains the public part of the self-signed certificate.
- A request database file, `fcncert.rdb`, is in the **INSTANCE_DIR** directory.
The request database file is used to store certificate requests that are associated with the key database. This file is automatically created when IBM Spectrum Protect Snapshot creates a key database file. This file is created with the same name as the key database file, but with a `.rdb` extension.
- An encrypted stash file, `fcncert.sth`.
The password that is protecting the key database file is generated automatically and is stored in the encrypted stash file.
- An ASCII encoded binary file, `fcmselcert.arm`.
This file is used to export the public part of the self-signed certificate. It is also used to import the public part of the self-sign certificate to the backup and cloning servers.
When you install backup and clone servers separately without the use of SSH, the installation process installs and sets up IBM GSKit. In this scenario, after IBM GSKit installation, manually copy the self-signed certificate to the backup and cloning servers.
- A certificate revocation list file, `fcncert.crl`.
This file contains a list of revoked certificates.

The `.kdb`, `.rdb`, `.crl`, and the `.sth` files contain critical security parameters and these parameters must be protected against unauthorized access by the operating system. It is advisable to back up the key database files regularly, especially if you are using a CA signed certificate.

If you are working with the self-signed certificates that are created by the setup script, you need to ensure that the `.arm` file is integrated on the auxiliary server. To do this, run the setup script on the production server through OpenSSL, or manually copy it to the auxiliary server and run the setup script there.

If you are using a CA signed certificate, you must use the GSKit command-line utilities to import the certificate to the production server. If the CA signed certificate is not a standard certificate that GSKit has a root certificate for, you must import the certificate to all sites.

Enforcing SP800-131 compliant encryption

The files that are needed for IBM GSKit are automatically installed during the installation. To enforce SP800-131 compliant encryption, during the configuration of IBM Spectrum Protect Snapshot, you must set the **ENFORCE_TLS12** parameter to YES in the IBM Spectrum Protect Snapshot profile file. You must use the advanced

mode during the configuration to specify this parameter. Otherwise, TLS Protocol version 1.0 and 1.1 is enabled as the default value for the **ENFORCE_TLS12** parameter is NO.

Any existing self-signed certificates that were created by a previous version of IBM Spectrum Protect Snapshot must be deleted to allow IBM Spectrum Protect Snapshot to create new self-signed certificates. To remove any existing self-signed certificates, go to the IBM Spectrum Protect Snapshot installation (*INSTANCE_DIR*) directory and enter the following command:

```
rm fmcert.*
```

Note: It is not required to delete existing external certificate authority (CA) signed certificates. However, if the CA signed certificate does not meet the minimum SP800-131 criteria, you must manually replace it with a new one.

Uninstall GSKit

GSKit must not be uninstalled unless you are sure that no product on the system is using it. By uninstalling GSKit you are removing the global GSKit installation from the system entirely.

If required, you can globally uninstall GSKit on each server by running the setup script on the backup or cloning server.

```
For Oracle in an SAP environment or Oracle: ./setup_ora.sh -a  
uninstall_gskit -d instance_directory
```

Option files used by Data Protection for Oracle

When you are using Data Protection for Oracle, the following IBM Spectrum Protect option files are used:

- Client system options (*dsm.sys*)
- Client user options (*dsm.opt*)
- Data Protection for Oracle options (*tdpo.opt*)
- RMAN backup script

Example 1 shows you how to configure the system options file (*dsm.sys*) to point to the same IBM Spectrum Protect.

In the following examples, the client user options file *dsm.opt*, stored in the */usr/tivoli/tsm/client/ba/bin* and */usr/tivoli/tsm/client/api/bin* directories, includes a server with the following TCP/IP address: *arrow.la.xyzcompany.com*. The *servername* option in the *dsm.opt* and *dsm.sys* files define the server stanza names only. The *tcpserveraddress* option indicates which server is contacted.

For the *dsm.opt* file that is stored in the */usr/tivoli/tsm/client/ba/bin* directory, use the following example:

```
servername tdphdw
```

For the *dsm.sys* file that is stored in the */usr/tivoli/tsm/client/ba/bin* directory, use the following example:

```
servername          tdphdw  
commethod           tcpip  
tcpport             1500
```

```
tcpserveraddress arrow.la.xyzcompany.com
passwordaccess generate
schedmode prompted
nodename hdworc1
```

For the `dsm.opt` file that is stored in the `/usr/tivoli/tsm/client/api/bin` directory, use the following example:

```
servername tdporc
```

For the `dsm.sys` file that is stored in the `/usr/tivoli/tsm/client/api/bin` directory, use the following example:

```
servername tdporc
commmethod tcpip
tcpport 1500
tcpserveraddress arrow.la.xyzcompany.com
passwordaccess prompt
nodename hdworc1
```

Example 2 shows you how to configure multiple server stanzas in the system options file (`dsm.sys`).

To configure multiple server stanzas in the system options file (`dsm.sys`), copy the option settings from the IBM Spectrum Protect for Databases: Data Protection for Oracle `dsm.sys` file to the IBM Spectrum Protect Snapshot `dsm.sys` file. For example, a combined `dsm.sys` file for a server with the name `arrow`:

```
servername tdphdw
commmethod tcpip
tcpport 1500
tcpserveraddress arrow.la.xyzcompany.com
passwordaccess generate
schedmode prompted

servername tdporc
commmethod tcpip
tcpport 1500
tcpserveraddress arrow.la.xyzcompany.com
passwordaccess prompt
```

Appendix B. Commands and scripts

A list of various commands and scripts that are used with IBM Spectrum Protect Snapshot operations is provided.

About this task

You can issue various commands for example to trigger a snapshot backup or snapshot restore. In addition, administrative tasks such as to start or stop IBM Spectrum Protect Snapshot can be issued from the command line.

Procedure

- Use `-d database-name` to specify the database name when you use the `-F` option. There is no default value.
- Use `-i instance-name` to specify the instance name that applies to the command you are running. This option is must be specified in the `-F` option. There is no default value.

Example

`-d` is always `ORCL__`
`-i` is the `<SID>`

Backup, restore, cloning commands, and utilities

You can issue commands to trigger a snapshot backup or snapshot restore, and to inquire and delete snapshot backups in the IBM Spectrum Protect Snapshot repository. You can create and manage database clones from the command-line interface.

The acsora user interface for Oracle

Use the production system user interface, `acsora`, to issue commands in an Oracle database environment.

The `acsora` syntax follows:

```
acsora [-p profile] -f function [-B backupID] [-F] [-s DEVICE_CLASS ]
```

where `<function>` is one of:

```
backup
restore
delete
inquire
inquire_detail
```

Table 20. Parameters for Oracle databases

Option	Description	Default
<code>-p profile</code>	Full path and name of the profile that is used by the Management Agent The Management Agent uses the GLOBAL and ACS sections of the profile.	<code><ACS_DIR>/profile</code>

Table 20. Parameters for Oracle databases (continued)

Option	Description	Default
-B backupID	The backup ID for restore, delete, inquire functions	
-f backup	Backup database	Function backup if -f is not specified.
-f restore	Restore database	
-f delete	Delete snapshot backup	
-f inquire	List snapshot backups	
-f inquire_detail	List snapshot backups	
-F	When specified with the -f delete or -f inquire options, the -F option withdraws source and target relationships.	
-s DEVICE_CLASS	The name of the DEVICE_CLASS section in the profile that is used for the backup operation.	As specified in the profile.

The return code of acsora is 0 if it finishes the request without an error. The return code 1 indicates one or more issues occurred during the process. Although not considered critical, resolve these minor issues to prevent potential critical issues later. Return code 2 indicates that an error occurred during command processing and that the operation did not complete successfully.

-f backup

This command backs up the Oracle database according to the profile settings.

Example (backup database):

```
acsora -f backup -s STANDARD
```

-f restore

This command restores the Oracle database from the backup that is specified by the backup ID, or the latest backup.

The Oracle database is available for immediate use after a snapshot restore and a rollforward recovery. However, for DS8000 and SAN Volume Controller, background copy processing from the target volumes to the source volumes might require more time to complete, especially if **FLASHCOPY_TYPE COPY** is specified. Although the database is available, you cannot perform another IBM Spectrum Protect Snapshot backup or restore until background copy processing completes.

Example (restore specified backup):

```
acsora -f restore -B A0FZ36AY8G
```

See “Restoring Oracle databases from IBM Spectrum Protect” on page 90 and Chapter 8, “Restoring Oracle databases,” on page 87 for detailed instructions on how to restore your Oracle database.

-f inquire

This command lists the details for the snapshot backup that is denoted by the backup ID, if specified by using the **-B** option, or all backups.

Example (list all backups):

```
acsora -f inquire
```

-f inquire_detail

This command queries the backup repository and lists all available backups.

This command is like the **-f inquire** command, but it prints more information like usability states and background copy progress, if available, for each backup.

-f delete

This command deletes the snapshot backup that is denoted by the entered backup ID.

Example (deleted specified backup):

```
acsora -f delete -B A0FZ36AY8G
```

acsutil Snapshot Object Manager for Oracle

The Snapshot Object Manager for Oracle, *acsutil*, provides a snapshot backup query and restore interface for Oracle and Oracle in an SAP environment environments.

Functions of the acsutil command

The Snapshot Object Manager for Oracle, *acsutil*, provides an interface for *acsora* to show available backups, run restore operations, and delete unwanted backups. It communicates with *acsora* through input and output files.

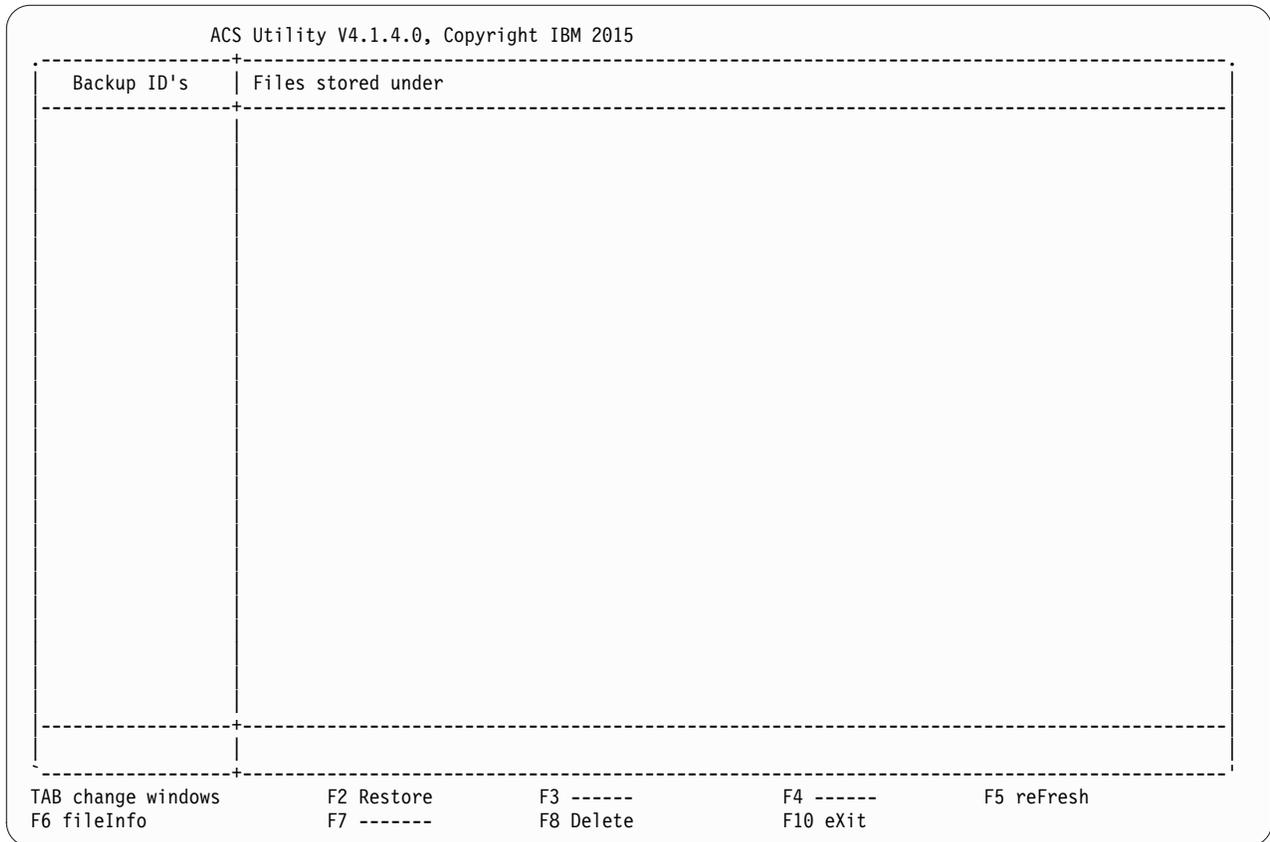
Syntax of the acsutil command

```
acsutil [-p profile]
```

Where **-p profile** is the path and name of the IBM Spectrum Protect Snapshot profile. The default value is *ACS_DIR/profile*.

The Snapshot Object Manager user interface consists of a split window, which is character-based.

The first step is an automatic inquire operation for all backup IDs. The following figure shows the screen layout for the list of backup IDs found by the Snapshot Object Manager when the inquiry is complete.



If you mark the backup ID you are interested in and then press the Tab key, all file names that belong to the marked backup ID are displayed.

Tab - Switch window side

Move the cursor between the sides of the window.

F2 - Restore

Restore the marked backup ID.

F5 - Refresh

Refresh the list of backup IDs and file names.

F6 - Fileinfo

Opens a separate window to display file information.

For backup IDs, the sequence number (backup version count) is shown.

F8 - Delete

Delete the selected backup ID and all corresponding files.

F10 - Exit

Exit from Snapshot Object Manager

ENTER - Mark/unmark backup ID

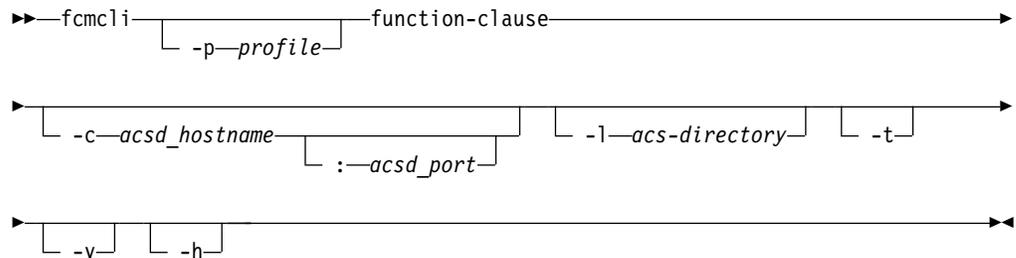
Mark or unmark the backup ID.

The Snapshot Object Manager can delete backup IDs with all included files. It is not possible to delete single files within a backup ID. To delete a backup ID, it must be highlighted. After pressing F8 you must confirm the deletion operation. The backup ID and all included files are deleted.

For each restore, a log file is created.

Cloning commands

You can use the IBM Spectrum Protect Snapshot command-line interface, `fccli`, to create and manage clones of component databases.



Where:

-p *profile*

Full profile name. Default value: `INSTANCE_DIR/profile`

-c *acsd_hostname*

Name of the server where the management agent (acsd) is running. Default value: `localhost`.

acsd-port

TCP/IP port number or service name on which the management agent (acsd) is listening. Default value: `57328`.

-l *acs-directory*

Directory where the logs and shared directories are located. Default value: `ACS_DIR`.

-t Start trace on. Default value: Trace off.

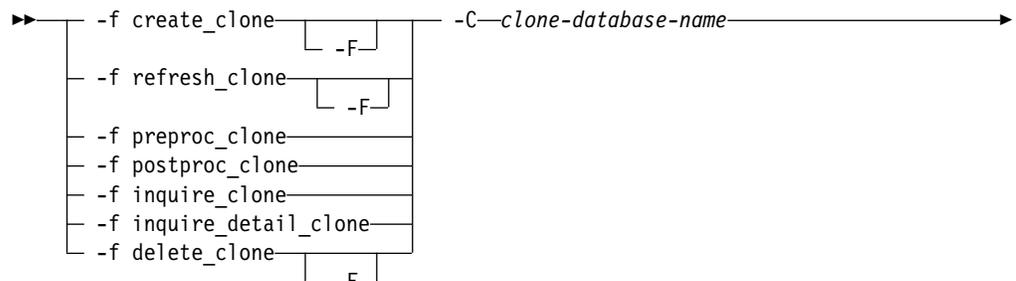
-v Show version.

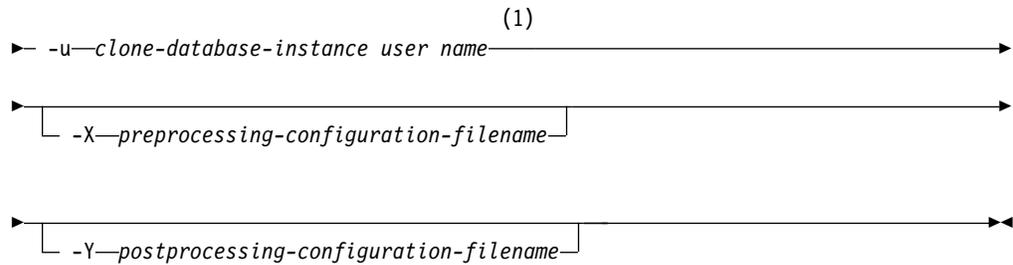
-h Show help text.

The values for the `function-clause` parameter are described in the following sections.

FlashCopy cloning function-clauses

The following functions are supported by the `fccli` command option `-f 'function'` for FlashCopy cloning operations:





Notes:

- 1 This option is not required for **-f inquire_clone** or **-f inquire_detail_clone** commands.

Where:

-F The force command option is optional. Depending on the cloning command it can have the following results:

- **delete_clone**: The force option causes the clone to be unmounted, marked as deleted, and also deletes the FlashCopy relationships. Otherwise, the **delete_clone** function unmounts the clone and marks it as deleted in the FlashCopy Manager repository.
- **create_clone, refresh_clone**: The force option deletes all backup versions that are older than the clone targets that are reused for the new or refreshed clone. Otherwise, a failure can occur when there are backup versions older than the clone targets that are reused for the new or refreshed clone.

This option is valid for Storwize V7000 and SAN Volume Controller Version 5.1, or later.

-C clone-database-name

The name of the cloned database on the clone system. This option must be specified for all cloning functions.

Specify a valid database name. The *clone database name* can be the name of the production database or you can specify an alternative name.

-u clone-database-instance user name

Specify the user name of the clone instance owner. This option is required when the following functions are issued:

- **create_clone**
- **delete_clone**
- **refresh_clone**
- **preproc_clone**
- **postproc_clone**

-X preprocessing-configuration-filename

The name of the configuration file to be used with the preprocessing script. The preprocessing configuration file must be on the clone server.

-Y postprocessing-configuration-filename.

The name of the configuration file to be used with the postprocessing script. The postprocessing configuration file must be on the clone server.

The return code of the **fcmcli** command is 0 if it finishes the request without an error or if there were no candidates for the request. The return code is 1 if one or

more minor issues occur which are not critical but must be checked to prevent major issues later. Return code 2 indicates that an error occurred during the command execution.

Issue cloning-related commands on the production system as the production database instance owner. The cloning commands must be issued from the `INSTANCE_DIR` directory where the IBM Spectrum Protect Snapshot production files are located. The `fccli` command identifies the name of the production database in the following order:

- For Oracle databases, the value of the `ORACLE_SID` environment variable is used to identify the production database name.

Deleting snapshot backups

IBM Spectrum Protect Snapshot snapshot backups can be deleted from the snapshot repository.

Before you begin

Optionally, you can delete snapshot backups on DS8000 and SAN Volume Controller storage subsystems that contain a dedicated set of target volumes in one or more target sets. With IBM XIV Storage System solutions you can create as many snapshot backups as needed, and old backups are manually deleted. Old backups can also be deleted automatically by using the `MAX_VERSIONS` (`MAX_SNAPSHOT_VERSIONS`) parameter.

About this task

Manually delete an IBM Spectrum Protect Snapshot snapshot backup by following the procedure.

Procedure

1. Run the following command to unmount the file systems and export the volume groups on a backup system. This method is used when the backup that is using this target set is currently mounted. This step can be omitted if the backup is not currently mounted.

```
fccli -f unmount [-B <backupID>]
```

2. Based on the use of this target set, any existing source, and target FlashCopy relationships (such as INCR or NOCOPY) must be withdrawn. Run the following command:

```
(Oracle) acsora -f delete -B <backupID>
```

```
(Oracle in an SAP environment) backint -f delete -B <backupID>
```

Results

Note: For IBM XIV Storage System, these commands delete the snapshot backup in the IBM Spectrum Protect Snapshot snapshot repository, and the snapshot on the storage system is also deleted.

Note: (DS8000 or SAN Volume Controller): These commands delete the snapshot backup in the IBM Spectrum Protect Snapshot snapshot repository only. The source and target relations on DS8000 or SAN Volume Controller are not withdrawn.

Deleting a target volume or target set

To remove a target volume from a target set or to remove a complete target set, run the following steps to free up the target volumes:

Procedure

1. Run the following command to unmount the file systems and export the volume groups on a backup system. If the backup is not mounted, do not run this step.

```
fccli -f unmount [-B <backupID>]
```

This method is used when the backup that is using this target set is mounted

2. Based on the use of this target set, any existing source, and target FlashCopy relationships (such as INCR or NOCOPY) must be withdrawn. Run the following command:

```
For Oracle, acsora -f delete -F -B <backupID>
```

```
or, for Oracle in an SAP environment, backint -f delete -F -B <backupID>
```

Results

The withdrawal of the source and target FlashCopy relationship is done by the IBM Spectrum Protect Snapshot generic device agent, acsgen, as a background operation. This process can take up to 10 minutes. Do not try to reuse the target volumes before the actual process completes successfully.

Snapshot backup status in the repository

Ensure that you routinely check the status of the IBM Spectrum Protect Snapshot repository.

To check the status of snapshot backups in the IBM Spectrum Protect Snapshot repository, use one of the following commands:

```
For Oracle, acsora -f inquire[_detail]
```

```
or, acsutil
```

```
For Oracle in an SAP environment, backint -f inquire[_detail] -t  
volume|file -p <SAP Backint profile (.utl)>
```

```
or, acsutil
```

When using the inquire_detail command with the appropriate tool, output similar to the following displays:

```
Type Partition Backup-ID TSM Backup-ID State  
DevClass TargetSet Background Copy BytestobeFlashcopied  
#BACKUP NODE0000 C01__A0FY303K6B IN-PROGRESS MIRROR1 1 3.000GB of 3.000GB  
3.000GB  
UsabilityStates :  
REMOTELY_MOUNTABLE,REPETITIVELY_RESTORABLE,SWAP-RESTORABLE,PHYSICAL_PROTECTION,  
FULL_COPY,TAPE_BACKUP_PENDING
```

Administrative commands

You can use commands to administer IBM Spectrum Protect Snapshot.

Administrative commands are available for you to do the following tasks:

- Start, stop, or configure IBM Spectrum Protect Snapshot.
- Mount or unmount a snapshot backup on a secondary system.
- Create a backup to IBM Spectrum Protect from a snapshot if you have IBM Spectrum Protect configured in your environment

To use the commands to automate operations for IBM Spectrum Protect Snapshot, add entries to the cron table (crontab) file. Because there are so many ways to implement IBM Spectrum Protect Snapshot software, there are no templates. To automate operations, either specify the commands in the crontab file, or create scripts and add the scripts to the crontab file.

Configuration commands

Use configuration commands to run the setup script, maintain IBM Spectrum Protect Snapshot passwords, and query the amount of storage space that is used for backups.

Installation setup script

The setup script provides instructions for configuration. The setup script is used by the IBM Spectrum Protect Snapshot installation program. The setup script can also be used to manually set up IBM Spectrum Protect Snapshot, and to complete a basic configuration.

The setup script uses the following command syntax:

```
setup_type.sh -a action -d Instance_owner_$HOME directory
```

For the *type* parameter, in the setup script name, the following values can be specified:

- setup_ora.sh

You can use the setup script for the following purposes:

- Activation or upgrade of IBM Spectrum Protect Snapshot for one instance-specific installation, as root user:

```
setup_type.sh -a install -d Instance_owner_$HOME_directory
```

The setup script is run from the *FCM_INSTALL_DIR* directory.

- Initial configuration and reconfiguration:

```
setup_type.sh
```

The setup script must be run as the database instance owner.

For custom applications, run the script as the application backup user. Run the script from the *INSTANCE_DIR* directory.

- Initial configuration and reconfiguration in advanced mode:

```
setup_type.sh -advanced
```

- Stopping an activated instance:

```
setup_type.sh -a stop -d Instance_owner_$HOME directory
```

The command must run as the database instance owner.

For custom applications, run the command as the application backup user. The command must be run from the *INSTANCE_DIR* directory.

- Starting an activated instance:

```
setup_type.sh -a start -d Instance_owner_$HOME_directory
```

The command must be run as the database instance owner.

For custom applications, run the command as application backup user. The command must be run from the *INSTANCE_DIR* directory.

- Disabling a stopped instance:

```
setup_type.sh -a disable -d Instance_owner_$HOME_directory
```

The command must be run as the database instance owner.

For custom applications, run the command as the application backup user. The command must be run from the *INSTANCE_DIR* directory. This command completely removes the entries from the */etc/inittab*.

For a typical configuration, these commands are run on a production system. There are some scenarios where these commands need to be run on a backup system. If you are running the commands on both systems, when you stop or disable IBM Spectrum Protect Snapshot, run the command on the production system before the backup system.

The setup script can be used to install IBM Spectrum Protect Snapshot on multiple backup nodes from the production server. As a prerequisite, Open Secure Shell (OpenSSH) must be installed on all of the nodes in the backup server. NFS shares between the production server and backup server nodes are not required for this type of remote installation. OpenSSH is the preferred method for IBM Spectrum Protect Snapshot.

The script must be run from the database instance-specific installation directory:

- (Oracle) *Instance owner \$HOME_directory/acs/*

The default action, *setup*, is performed and the instance is configured.

If the script is called without parameters, it can be issued as the instance owner. The script creates a profile or changes an existing profile, and updates the daemon jobs according to the current profile (production system) or user preference (backup system).

If IBM Spectrum Protect Snapshot cannot be stopped, stop IBM Spectrum Protect Snapshot on the production system before you run the script with the *-a install -d Instance_owner_\$HOME_directory* options.

Setup script values

The following values are available for *setup_type.sh*.

Use *setup_ora.sh* to configure IBM Spectrum Protect Snapshot for Oracle.

Or,

Use *setup_ora.sh* to configure IBM Spectrum Protect Snapshot for Oracle in an SAP environment.

The following values are available for action: The instance directory name (*-d* option) is required for all explicit actions.

disable

This call can be issued as the root or instance owner. It stops IBM Spectrum Protect Snapshot and removes all daemon jobs. To reactivate IBM Spectrum Protect Snapshot, call the script without parameters.

If IBM Spectrum Protect Snapshot cannot be stopped, stop IBM Spectrum Protect Snapshot on the production system before running `setup_type.sh -a install -d <Instance owner $HOME directory>`.

install

This call needs to be issued with the root user ID. When issued, the following actions are completed:

1. Stops IBM Spectrum Protect Snapshot (`setup_type.sh -a stop -d <Instance owner $HOME directory>`) For DB2 databases, change `<INSTANCE owner $HOME directory>` to `<INSTANCE owner $HOME directory>/sqllib`.
2. Copies all binary files from the IBM Spectrum Protect Snapshot installation directory to the instance-specific installation directory (`INSTANCE_DIR`)
3. Sets the appropriate access rights for the binary files.
4. Restarts IBM Spectrum Protect Snapshot (`setup_type.sh -a start -d <Instance owner $HOME directory>`).

The steps to start and stop IBM Spectrum Protect Snapshot are skipped if it is not configured.

If IBM Spectrum Protect Snapshot cannot be stopped, stop IBM Spectrum Protect Snapshot on the production system before running `setup_type.sh -a install -d <Instance owner $HOME directory>`.

start

This call can be issued as the root or instance owner. The call starts a previously installed and configured version of IBM Spectrum Protect Snapshot. This call starts the configured daemon jobs.

stop

This call can be issued as the root or instance owner. It stops the version of IBM Spectrum Protect Snapshot that is currently running. This call updates the configured daemon jobs and checks that IBM Spectrum Protect Snapshot is stopped successfully (a write lock can be acquired for the `.lock` file that is located in the instance-specific installation directory).

This call fails on the backup system in environments where the instance-specific installation directory is shared between the production and backup systems, if IBM Spectrum Protect Snapshot is running on the production system. To successfully stop IBM Spectrum Protect Snapshot in those environments, stop IBM Spectrum Protect Snapshot on the production system.

This option is not required for the default setup function.

Setting or changing passwords with the setup script

You can set or change passwords by issuing the setup script without the `-a` action option.

Use the command in this example:

```
setup_type.sh
```

Running the setup script without the `-a` action option proceeds through several tasks that are similar to the tasks described in Chapter 5, “Installing and upgrading,” on page 41.

When this command is issued, the profile wizard starts. You can use the profile wizard to edit the profile, and to set or change passwords. Using this wizard to administer passwords is preferred because the wizard updates changed passwords on the backup systems. To update passwords on the backup system, specify `YES` at the following prompt:

Select the backup system to update or delete:

- 1) acsback5
- n) configure a new backup system
- b) return to previous menu
- q) quit configuration

Select one of the options.

1

The selected backup system is acsback5

The backup system on acsback5 is configured with the device class(es) `DISK_ONLY`.

Select the action you want to take on the backup system acsback5:

- 1) update IBM Spectrum Protect Snapshot installation
- 2) start IBM Spectrum Protect Snapshot services
- 3) stop IBM Spectrum Protect Snapshot
- 4) uninstall IBM Spectrum Protect Snapshot
- 5) setup the SSH key authentication
- b) return to backup system selection
- q) quit the configuration

Select one of the options.

1

Do you want to update the Backup System installation on acsback5? [y|n] [y]

Password administration

You can use the `setup.sh` script or the `fcmlcli -f password` command to change the IBM Spectrum Protect Snapshot passwords.

The `fcmlcli -f password` command supports an interactive and a non-interactive mode. To use the interactive mode, do not enter a password when you issue the command and you are prompted to enter the following passwords:

- The master password, which is the password of the acsd management agent. By default, a 32 character password is automatically generated. However, you can enter an alternative password.
- The password for the ORACLE section if defined in the specified profile.
- The password for the DB2STANDBY section if defined in the specified profile.
- The passwords for the disk storage subsystems that are referenced by the `DEVICE_CLASS` sections in the specified profile.

If the specified profile contains multiple `DEVICE_CLASS` sections that reference the same physical disk storage subsystem, the password is queried one time by combining these `DEVICE_CLASS` sections.

The interactive mode is the preferred method for setting passwords. Using this method, the passwords are verified by testing the connections to the corresponding

Spectrum Protect Snapshot to auto-generate a password. For example, issue the following command to auto-generate the master password:

```
./fcmcli -f password :auto
```

GSKit commands

If you are not using SSH for remote installation and configuration of IBM Spectrum Protect Snapshot on backup and cloning systems, use GSKit commands to manually import a self-signed certificate. If you decide to use a CA signed certificate, use GSKit commands to complete a manual setup.

Manually importing the self-signed certificate

The self-signed certificate is automatically created by IBM Spectrum Protect Snapshot. When the IBM Spectrum Protect Snapshot setup script is run on the production server, it automatically creates the `fcmselcert.arm` file. It is stored on the production server in the default installation path. The `fcmselcert.arm` file is automatically imported on the backup and cloning servers from the production server with the SSH remote deployment mechanisms of the setup script. When remote deployment is not used and you separately run the setup script on the backup or cloning server, the `fcmselcert.arm` file if present is automatically imported to the local key database and then deleted. To use this automation, copy the `fcmselcert.arm` file from the production server to either the backup or cloning server before you start the setup routines on the backup or cloning server.

Alternatively, you can import the self-signed certificate by using the following GSKit command. However, in most scenarios this step is not necessary as the file is automatically imported as part of the IBM Spectrum Protect Snapshot setup process.

```
gsk8capicmd_64 -cert -add -db fcmcert.kdb -stashed -label "FCM server certificate" -file <path to fcmselcert.arm> -format ascii
```

This command fails if the key database already contains a certificate with the label FCM server certificate. To remove the certificate with the label FCM server certificate, you can use the following command:

```
gsk8capicmd_64 -cert -delete -db fcmcert.kdb -stashed -label "FCM server certificate"
```

CA Certificate

You can use a CA signed certificate for IBM Spectrum Protect Snapshot. If the certificate that is assigned by a CA has no built-in GSKit support, import the CA root certificate into the key database file (`fcmcert.kdb`). Use the GSKit command-line utilities to update the file on the production system, the backup system, and the cloning system. The root certificate of a trusted CA certificate is in the key database. GSKit has the following trusted root certificates:

- Entrust.net Global Secure Server Certification Authority
- Entrust.net Global Client Certification Authority
- Entrust.net Client Certification Authority
- Entrust.net Certification Authority (2048)
- Entrust.net Secure Server Certification Authority
- VeriSign Class 3 Public Primary Certification Authority
- VeriSign Class 2 Public Primary Certification Authority
- VeriSign Class 1 Public Primary Certification Authority
- VeriSign Class 4 Public Primary Certification Authority - G2

- VeriSign Class 3 Public Primary Certification Authority - G2
- VeriSign Class 2 Public Primary Certification Authority - G2
- VeriSign Class 1 Public Primary Certification Authority - G2
- VeriSign Class 4 Public Primary Certification Authority - G3
- VeriSign Class 3 Public Primary Certification Authority - G3
- VeriSign Class 2 Public Primary Certification Authority - G3
- VeriSign Class 1 Public Primary Certification Authority - G3
- Thawte Personal Premium CA
- Thawte Personal Freemail CA
- Thawte Personal Basic CA
- Thawte Premium Server CA
- Thawte Server CA
- RSA Secure Server Certification Authority
- Secure Server Certification Authority

The following example shows the command to request that a CA signed certificate is included:

```
gsk8capicmd_64 -certreq -create -db fmcert.kdb -stashed -label "FCM server certificate request" -dn dist_name -target fcmservercertreq.arm
```

For SP800-131 compliance, when the **ENFORCE_TLS12** parameter is set to YES in the IBM Spectrum Protect Snapshot profile, ensure that the certificate meets the minimum requirement by adding the following two options:

- `-size 2048` (or higher)
- `-sigalg sha224` (or higher)

Note: IBM Spectrum Protect Snapshot creates a self-signed certificate that is signed with SHA512, and the size is 4086 bits.

The **label** parameter can have any value except FCM server certificate. This value is already used by the self-signed certificate in the key database.

When you use a certificate that is signed by a CA that has no built-in GSKit support, you must import the CA root certificate. This task must be completed before the certificate is received or imported. The CA root certificate must be imported into the key database (KDB) files on the production system. The CA root certificate must also be imported into the KDB files on the backup and cloning servers. Issue the following command to import the root certificate:

```
gsk8capicmd_64 -cert -add -db fmcert.kdb -stashed -label "FCM server certificate request" -file path to CARootCertificate.arm
```

Issue the following command to import a signed certificate when it is received from a CA:

```
gsk8capicmd_64 -cert -receive -file fcmservercertsigned.arm -db fmcert.kdb -stashed
```

Rename the CA signed certificate label to FCM server certificate. Usually, the key database still contains the self-signed certificate, it must be deleted before the CA signed certificate can be renamed. To remove the self-signed certificate from the key database, issue the following command:

```
gsk8capicmd_64 -cert -delete -db fmcert.kdb -stashed -label "FCM server certificate"
```

To rename the CA signed certificate issue the following command:

```
gsk8capicmd_64 -cert -rename -db fmcert.kdb -stashed -label  
"FCM server certificate request" -new_label "FCM server certificate"
```

The file `fcmselfcert.arm` is used to export the self-signed certificate. When you use a CA certificate, the `.arm` file is obsolete and must be deleted on the production system. The self-signed certificate is automatically removed from the key database on the backup or cloning system during the next remote update with the setup script. If remote deployment is not used, you can manually remove the self-signed certificate from the key database on the backup and cloning servers. To remove the self-signed certificate, issue the following command:

```
gsk8capicmd_64 -cert -delete -db fmcert.kdb -stashed -label "FCM server  
certificate"
```

Monitoring the expiry date of certificates

When a self-signed certificate is created, an expiry date can be specified. The expiration time of the certificate is specified in days. The default is 365 days. The duration is 1-7300 days (20 years). The IBM Spectrum Protect Snapshot setup script creates the self-signed certificate for the production, backup, and cloning servers. The expiration time of all self-signed certificates that is generated by the setup script is 20 years. If you are using CA signed certificates, the expiration date is set by the certificate authority. You must monitor certificates for expiry and remove any expired certificates. If the key database does not contain a valid certificate with the label `FCM server certificate` and the setup script is rerun, a new self-signed certificate is generated. The `.kdb`, `.rdb`, `.arm` and `.sth` files are rewritten.

Related information:

 [ftp://ftp.software.ibm.com/software/webserver/appserv/library/v80/
GSK_CapiCmd_UserGuide.pdf](ftp://ftp.software.ibm.com/software/webserver/appserv/library/v80/GSK_CapiCmd_UserGuide.pdf)

Query managed capacity

Use the **managed_capacity** command to display information about IBM Spectrum Protect Snapshot managed capacity and licensing.

When you run the **managed_capacity** command, an XML managed capacity and licensing report is printed to the ACS directory or to another specified directory:

The report that is generated lists the capacity value that is calculated from source disks that are protected by IBM Spectrum Protect Snapshot for which a FlashCopy or snapshot backup was created. If a volume contains multiple backups, that volume is counted once during the query. Identify the repository from which to list backups by specifying the profile that is associated with the source volume. The output displays the total managed capacity for all source volumes.

The **fccli -f managed_capacity** syntax is as follows:

```
fccli -f managed_capacity [-p profile] [-c] [-o <output_directory>]
```

- p** Specify the name of the IBM Spectrum Protect Snapshot profile that is associated with the backups on the volume.
- c** Specify this option to display the output as comma-separated values.
- o** Specify this option to print the report to a specified directory as an XML report to view in your browser. When you do not specify a **-o** directory, the report is printed to `ACS_DIR/capacity_reports`.

Tip: Ensure to regularly delete old copies of managed capacity reports from the output directory.

Example output

This command displays capacity for the profile in /orc/S01/acs:

```
fcmlcli -f managed_capacity -p /orc/S01/acs/profile
```

Output:

```
FMM0461I Created tracefile '/orc/S01/acs/logs/fmquery.trace' for process ID
'31634'.
FMM1498I Front-End Capacity Report: Total protected size: 108.723 MB
FMM1497I Front-End Capacity Report: Number of managed objects: 1
FMM1496I Back-End Capacity Report: Total protected size: 217.445 MB
FMM1493I Back-End Capacity Report: Number of managed objects: 2
FMM1495I Logical Unit (LUN) Capacity Report: Total protected size: 768.000 MB
FMM1494I Logical Unit (LUN) Capacity Report: Number of managed objects: 2
```

This command displays all volumes for the profile that is in /orc/S01/acs as comma-separated values:

```
fcmlcli -f managed_capacity -p /orc/S01/acs/profile -c
```

Output:

```
...
tsm_sur_capacity,0
tsm_sur_objects,0
fcm_be_capacity,0
fcm_be_objects,0
fcm_lun_capacity,8589934592
fcm_lun_objects,4
tsm,no
```

For more information about front-end and back-end capacity and how to measure them, see the latest User's Guide at this site ftp://public.dhe.ibm.com/storage/tivoli-storage-management/front_end_capacity_measurement_tools/

Background daemons

For IBM Spectrum Protect Snapshot to work, some background daemon processes are required. Background daemon processes are not started directly. Instead, they are usually added to the /etc/inittab through the **setup_*.sh** commands.

To support high availability environments where the /etc/inittab cannot be used, you can instruct the setup_*.sh scripts to provide you with the exact commands that must be added to your high availability scripts instead of adding entries to /etc/inittab.

Management agent: acsd

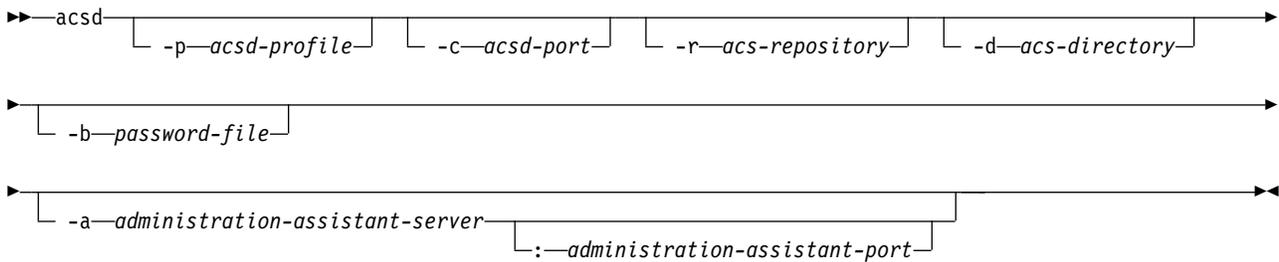
The management agent, acsd, coordinates the snapshot backup operation. It is a background daemon process that starts automatically.

The management agent, acsd, controls the backup flow and mediates between the other agents. The acsd agent provides access to the snapshot backup repository, which contains information about the valid snapshot backups and their relationships to snapshot capable storage devices.

(DB2) acsd must be started as the DB2 instance owner.

If you must deviate from the standard installation, the management agent offers the following command options for customization:

acsd management agent



Syntax for obtaining version or help information:

acsd management agent help

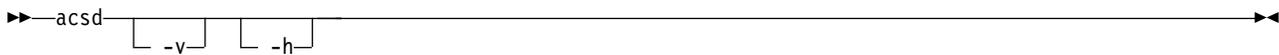


Table 21. Options for starting the management agent, acsd, as a daemon process

Option	Description	Default	Overrides profile parameter
-p acsd-profile	Full path and name of the profile that is used by the management agent. The management agent uses the GLOBAL and acsd sections of the configuration profile.	INSTANCE_DIR/profile	
-c acsd-port	TCP/IP port number or service name on which the management agent is listening	57328	ACSD (port number or service name)
-r acs-repository	Directory name where the snapshot backup repository is located	None	ACS_REPOSITORY
-d acs-directory	Name of IBM Spectrum Protect Snapshot directory	ACS_DIR	
-b password-file	File in which the IBM Spectrum Protect Snapshot management agent password is stored (in encrypted form). See notes.	ACS_DIR/shared/pwd.acsd	No corresponding profile parameter.
-a administration-assistant-server	(SAP) Host name of the server on which the Administration Assistant is running	None	ADMIN_ASSISTANT (hostname)
administration-assistant-port	(SAP) TCP/IP port on which the Administration Assistant is listening	None	ADMIN_ASSISTANT (port number)
-v	Display version and help information	None	N/A

Table 21. Options for starting the management agent, acsd, as a daemon process (continued)

Option	Description	Default	Overrides profile parameter
-h	Display help information only	None	N/A

All parameters override the values that are specified in the acsd profile or the corresponding default values. The shared and logs directories are automatically created in ACS_DIR. If no parameters are entered, acsd starts with the default profile and uses the default parameter values where applicable, or an error message is shown if this profile does not exist.

(DB2) When a user installs DB2 and creates a DB2 instance, the acsd management agent, is copied to the *DB2 instance directory/acs* directory. To activate IBM Spectrum Protect Snapshot, the user must start the setup script as the DB2 instance owner from this same directory. This script creates two entries in the /etc/inittab directory. The management agent, acsd, starts automatically from the /etc/inittab directory without any command-line arguments. The default values are used for configuring the management agent, acsd. The default values can be overridden by providing a profile. By default, this profile is in the directory *DB2 instance directory/acs*.

When acsd is started for the first time, or with a new **ACS_DIR** parameter, the following actions occur:

- Create the subdirectories shared and logs
- Create a password file pwd.acsd in the shared subdirectory
- Generate a master password

When the snapshot backup library uses the same ACS_DIR, it can authenticate itself to acsd with the password provided in the pwd.acsd file. If the snapshot backup library uses a different ACS_DIR, the default password file pwd.acsd must be copied to that directory so that they can read the master password from that directory.

Note: The minimum length of the master password is 8 characters. It must contain at least one number and one letter. The use of special symbols increases the strength of the password.

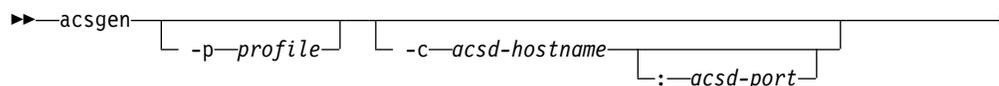
Generic device agent: acsgen

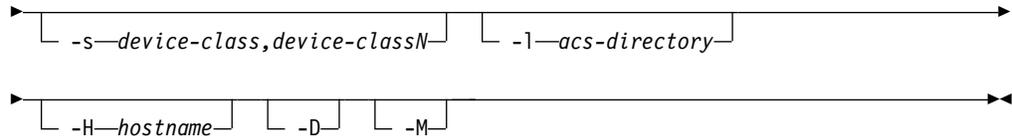
The generic device agent, acsgen, is the component that uses adapters to start snapshot commands on snapshot-compatible devices.

The generic device agent, acsgen, is started as a background daemon so you are not required to manually start it.

If you must deviate from the standard installation, the generic device agent, acsgen, offers the following command options for customization:

acsgen generic device agent





Syntax for obtaining version or help information:

acsgen generic device agent help



Table 22. Options for starting the generic device agent, `acsgen`. Description of `acsgen` options with default values if applicable.

Option	Description	Default
<code>-p profile</code>	Full profile name.	<code><INSTANCE_DIR>/profile</code>
<code>-c acsd-hostname</code>	Name of the server where the management agent, <code>acsd</code> , is running.	<code>localhost</code>
<code>acsd-port</code>	TCP/IP port number or service name on which the management agent, <code>acsd</code> , is listening.	<code>57328</code>
<code>-s device-class</code>	Section in the profile that pertains to the device class. Specify multiple device classes by separating each device class by a space.	<code>STANDARD</code>
<code>-l acs-directory</code>	Directory where the logs and shared directories can be found.	<code><ACS_DIR></code>
<code>-D</code>	Start as daemon. The <code>-a</code> option defines the usability states that the device agent responds to. Valid only when started from the following path: <code>/etc/inittab</code>	Run and end.
<code>-H hostname</code>	The host name where the process is running. The primary use is by the launchpad component to check its partitions in a DB2 multi-partition environment.	The system host name that is displayed by the <code>hostname</code> command.

Table 22. Options for starting the generic device agent, *acsngen* (continued). Description of *acsngen* options with default values if applicable.

Option	Description	Default
-M	Start the device agent as a mount agent. This agent is called for mounting or unmounting the target volumes on the backup system when any of the following situations exist: <ul style="list-style-type: none"> • An offloaded backup to IBM Spectrum Protect is requested • Database files on JFS file systems • Database files on AIX LVM mirrored volumes • The database is not suspended A mount verifies the consistency of the associated file systems.	Start as the monitoring agent.
-v	Display version and help information.	None
-h	Display help information only.	None

Mounting and unmounting snapshots on a secondary system

IBM Spectrum Protect Snapshot commands are available to mount or unmount a snapshot backup on a secondary system.

fcmdi command

```

>> fcmdi [ -p profile ] function-clause
      [ -c acsd_hostname ] [ -l acs-directory ] [ -t ]
      [ -v ] [ -h ]
  
```

Where:

-p profile

Full profile name. Default value: *INSTANCE_DIR/profile*

-c acsd-hostname

Name of the server where the management agent (acsd) is running. Default value: *localhost*

acsd-port

TCP/IP port number or service name on which the management agent (acsd) is listening. Default value: *57328*

-l *acs-directory*

Directory where the logs and shared directories are located. Default value:
ACS_DIR

-t Start with trace on. Default value: off.

-v Show version.

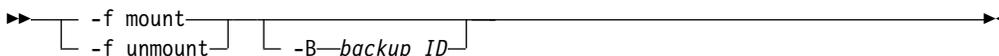
-h Show help text.

The return code of the **fccli** command is *0* if it finishes the request without an error or if there were no candidates for the request. Return code *1* indicates one or more minor issues occurred that are not critical but can be checked to prevent major issues later. Return code *2* indicates that an error occurred during the command execution.

FlashCopy administrative operations

The following functions are supported by the **fccli** command option **-f** 'function' for mount and unmount:

-f mount and -f unmount function-clauses



Where:

-f mount

Mount snapshot target set.

-f unmount

Unmount snapshot target set.

-B backup ID

The Backup ID as displayed by **fccli -f inquire [_detail]** command.

The following functions are supported by the **fccli** command option **-f** 'function' for forced unmount:

-f unmount function-clause with force option



Where:

-f unmount

Unmount snapshot target set.

-F Force a reset of **TAPE_BACKUP_IN_PROGRESS** usability states for the specified snapshot backup during the unmount force function. This parameter also requires the following **-B backup-id** argument.

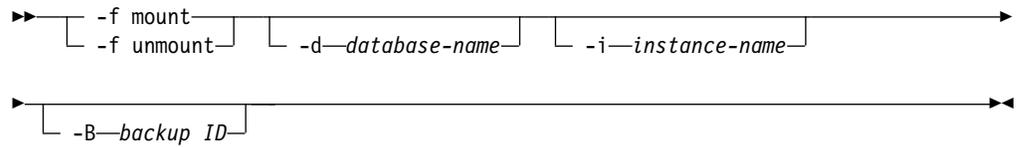
-B backup ID

The Backup ID as displayed by **fccli -f inquire [_detail]** command.

The functions **mount**, **unmount**, or **tape_backup** cannot run in parallel on the same backup server.

The following functions are supported by the **fcmlcli** command option **-f** 'function' for mount and unmount:

-f mount and -f unmount function-clauses



Where:

-f mount

Mount snapshot target set.

-f unmount

Unmount snapshot target set.

-d database-name

Database name.

-i instance-name

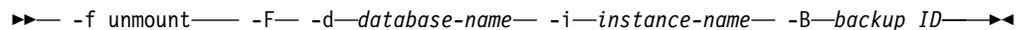
Instance name to apply to the command. There are no limitations.

-B backup ID

The Backup ID as displayed by **fcmlcli -f inquire [_detail]** or **db2acsutil** command.

The following functions are supported by the **fcmlcli** command option **-f** 'function' for forced unmount:

-f unmount function-clause with force option



Where:

-f unmount

Unmount snapshot target set.

-F Force a reset of **TAPE_BACKUP_IN_PROGRESS** usability states for the specified snapshot backup during the unmount force function. This parameter also requires the following arguments:

- **-d database-name**
- **-i instance-name**
- **-B backup-id**

-d database-name

Database name.

-i instance-name

Instance name to apply to the command. There are no limitations.

-B backup ID

The Backup ID as displayed by **fcmlcli -f inquire [_detail]** or **db2acsutil** command.

The functions **mount**, **unmount**, or **tape_backup** cannot run in parallel on the same backup server.

-f mount

This command mounts a snapshot backup on a backup system.

Mounting a backup means the following occurs:

1. Configure the target volumes, which might need to be assigned to the offload system (see the profile parameter **BACKUP_HOST_NAME** in “DEVICE_CLASS *device*” on page 132 for details).
2. Import the volume groups from the target volumes.
3. Mount all file systems within the volume groups.

The mount is done by one mount agent for each backup server. As a result, a mount agent is started by the launchpad daemon that runs on the respective backup server. By specifying additional options (filter arguments) such as

```
-i instance-name  
-d database-name  
-B backup-id
```

a specific snapshot backup can be selected for mounting on the offload system.

If no backup with the usability state **TAPE_BACKUP_PENDING** exists, the parameters **-i**, **-d**, and **-B** are mandatory. Here are two examples. The first one is generic:

```
fccli -f mount -d <database-name> -i <instance-name> -B <backup-id>.
```

Here is a specific example for Oracle:

```
fccli -f mount -d ORCL__ -i OR1 -B A0IFZH20KH
```

.

Here is a generic example for Oracle in an SAP environment:

```
fccli -f mount -d <backupidprefix> -i <database-name> -B <backup-id>
```

Here is a specific example:

```
fccli -f mount -d OAS__ -i OAS -B A0I9V162X5
```

Note: If the option **-B** is omitted, the oldest backup still in state *tape_backup_pending* is selected implicitly.

To reflect whether a snapshot backup is being mounted or is mounted, the usability states **MOUNTING** and **MOUNTED**, are set for those backups in the snapshot backup repository. These two state values prevent a duplicate mount request for a backup that is being mounted, or is already mounted, on the backup system. If multiple snapshot backups of a database are candidates to be mounted, IBM Spectrum Protect Snapshot picks the one with the most recent snapshot backup ID.

-f unmount

This command releases all resources on the offload server that were used by the mount command.

For *normal mode*, the unmount is completed by one mount agent for each backup server. A mount agent is started by the launchpad daemon that runs on the respective backup server. The following steps are completed by the software:

1. Unmount the file system that belongs to the target volumes.
2. Export the assigned volume group.
3. Remove the devices, `vpath/hdisk`, from the offload system.

When extra options, which are known as filter arguments, are specified, a specific snapshot backup can be selected for unmounting from the offload system. The following list identifies filter arguments:

```
-i instance-name  
-d database-name  
-B backup-id
```

If the unmount does not succeed because of problems that are related to the device agent, the usability state of the backup remains **MOUNTED** in the snapshot backup repository. After resolving the problems on the backup system, the **fcmlcli unmount** command must be issued again. The command is issued again to finalize the unmount of the file systems and update the usability state of the backup in the snapshot backup repository. If an off-loaded tape backup is running, the usability state **TAPE_BACKUP_IN_PROGRESS** is set and those backups are not be picked by IBM Spectrum Protect Snapshot for unmounting.

For *force mode*, unexpected system failures with offloaded tape backups can lead to an incorrect state of the backup reflected in the snapshot backup repository. The state **TAPE_BACKUP_IN_PROGRESS** is set. A built-in force option, **-F**, for the **fcmlcli unmount** function is provided to return the system to a usable state. Besides the normal unmount function, the unmount force option picks backups in the **TAPE_BACKUP_IN_PROGRESS** state as candidates to be unmounted and to reset the **TAPE_BACKUP_IN_PROGRESS** usability state for those backups. The **-d**, **-i**, and **-B** options are specified to uniquely identify the backup that is involved.

Integration with IBM Spectrum Protect

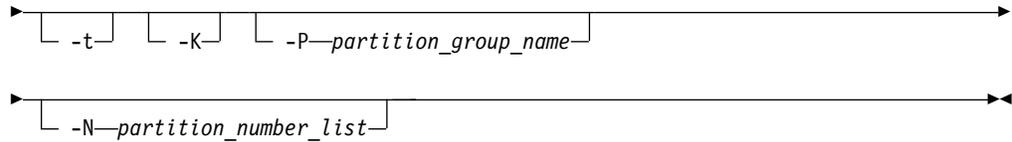
If IBM Spectrum Protect is set up and configured in your environment, you can create a backup to IBM Spectrum Protect from a snapshot.

The fcmlcli offload agent

The offload agent is a daemon process that manages offloaded backups to IBM Spectrum Protect. The agent also provides a command line interface offering functions for managing IBM Spectrum Protect backups.

fcmlcli command

```
➤ fcmlcli [ -p profile ] function-clause  
➤ [ -c acsd_hostname ] [ : acsd_port ] [ -l acs-directory ] [ -D ]
```



Where:

-p profile

Full profile name. Default value: *INSTANCE_DIR/profile*

-c acsd_hostname

Name of the server where the management agent (acsd) is running. Default value: *localhost*.

acsd-port

TCP/IP port number or service name on which the management agent (acsd) is listening. Default value: *57328*.

-l acs-directory

Directory where the logs and shared directories are located. Default value: *ACS_DIR*.

-D Run as daemon process. Valid only when started from */etc/inittab*. Default value: Run and end.

-t Start trace on. Default value: Trace off.

-K In a multi-partition environment, the partitions remain mounted when all participating partitions are successfully offloaded to IBM Spectrum Protect. The offload agent unmounts all partitions after the last partition is successfully offloaded. Default value: Off. The unmount operation is part of every IBM Spectrum Protect backup operation.

-P partition_group_name

The name of a partition group as specified in the profile with the **PARTITION_GROUP** parameter.

-N partition_number_list

A single number or list of numbers that are separated by a comma that specifies the partitions to apply the action against. When not specified, the action is applied to all partitions.

The values for the function-clause parameter are described in the following sections.

-f tape_backup:

This offload agent command backs up data to tape storage.

Note: IBM Spectrum Protect for Enterprise Resource Planning must be installed on the production and backup server if you use IBM Spectrum Protect Snapshot in an SAP environment with Oracle or DB2. IBM Spectrum Protect for Databases: Data Protection for Oracle as well as Oracle must be installed on the production and backup server if you use IBM Spectrum Protect Snapshot to protect an Oracle non-SAP environment. If IBM Spectrum Protect Snapshot for Custom Applications is used, the IBM Spectrum Protect backup-archive client must be installed on the backup server.

To create a snapshot backup with a subsequent tape backup, **TSM_BACKUP** or **TAPE_BACKUP_FROM_SNAPSHOT** must be specified either as part of the backup command or as a profile parameter, thus applying to all backups. The management agent updates the usability state with **TAPE_BACKUP_PENDING**. The IBM Spectrum Protect Snapshot offload agent then picks up all snapshot backups in the state **TAPE_BACKUP_PENDING** and backs them up to tape. The `fcmlcli -f` backup operation must be issued from the production system.

To start the offload backup to tape, enter the command:

```
fcmlcli -f tape_backup
```

By specifying additional options or filter arguments such as

```
-i instance-name  
-d database-name
```

the appropriate backup for the given instance and or database can be selected for offloading to tape. The `-B backup-id` option cannot be specified in conjunction with `-f tape_backup`. The backups should be processed in chronological order. The `tsm4acs` backs up the oldest snapshot eligible for transfer to IBM Spectrum Protect.

By specifying the `-D` option for the offload agent, it acts as a daemon process that periodically checks for outstanding tape backup requests. Furthermore, the offload agent, running as a daemon, tries to offload a snapshot backup to tape only one time. If the first attempt fails for some reason, the snapshot backup is marked accordingly and is not be picked a second time by the `tsm4acs` daemon for offloading to tape. This type of backup must be offloaded to tape manually by issuing the following command:

```
fcmlcli -f tape_backup filter_arguments
```

If multiple snapshot backups of a database are candidates for offloading to tape, the IBM Spectrum Protect Snapshot offload agent (whether as a daemon or with the `-f tape_backup` function) always selects the one with the oldest snapshot backup ID. This selection ensures that the IBM Spectrum Protect backups are created in the appropriate sequential order.

Tip: Whenever a new snapshot backup with **TSM_BACKUP** set to YES, MANDATE, or LATEST is created, IBM Spectrum Protect Snapshot sets the **TAPE_BACKUP_PENDING** status to NO for all snapshot backups that were previously created with **TSM_BACKUP** set to LATEST. This prevents backup requests to IBM Spectrum Protect from queuing if they cannot be completed in time.

The `tsm4acs` **tape_backup** function internally does the following steps:

1. Mount the file systems on the offload system if they were not previously mounted using `fcmlcli` with the 'mount' function or by a forced mount request. If all necessary file systems were mounted, this step is skipped.
2. Update the usability state to **TAPE_BACKUP_IN_PROGRESS** for all partitions that have the usability state **TAPE_BACKUP_PENDING** set.
3. Back up these partitions to tape.
4. Update usability states: For those partitions for which the backup succeeded, reset the usability state **TAPE_BACKUP_PENDING** and set **TAPE_BACKUP_COMPLETE**. For those partitions where the backup failed, set the usability state **TAPE_BACKUP_FAILED**. For all participating partitions, reset the usability state **TAPE_BACKUP_IN_PROGRESS**.
5. Unmount the file systems from the offload system.

When the usability state for a partition is **TAPE_BACKUP_IN_PROGRESS**, any request to restart the offload of that partition to tape is refused.

If a backup to IBM Spectrum Protect fails, the IBM Spectrum Protect Snapshot software can try the backup operation again.

-f update_status:

This offload agent command updates the usability state of a specified snapshot backup.

The usability state of a specified snapshot backup can be updated to either offload a snapshot to IBM Spectrum Protect (TSM_BACKUP=yes), or to not offload a snapshot (TSM_BACKUP=no). It is possible to offload a snapshot backup to IBM Spectrum Protect even though the TSM_BACKUP or TSM_BACKUP_FROM_SNAPSHOT profile parameter was deactivated during the snapshot backup operation. If there is no longer a need to offload the snapshot backup that was run with the parameter TSM_BACKUP or TSM_BACKUP_FROM_SNAPSHOT activated, the usability state can be reset.

To identify the backup whose state is to be modified, these parameters must also be specified using the **-f update_status** command:

-d database-name
-i instance-name
-B backup-id

Appendix C. Examples

Refer to these IBM Spectrum Protect Snapshot examples when you are configuring, updating, or following product tasks.

Oracle overall disk layout example

Refer to this example when configuring the disk layout in a Oracle environment.

The following figure shows file systems involved in a sample disk layout.

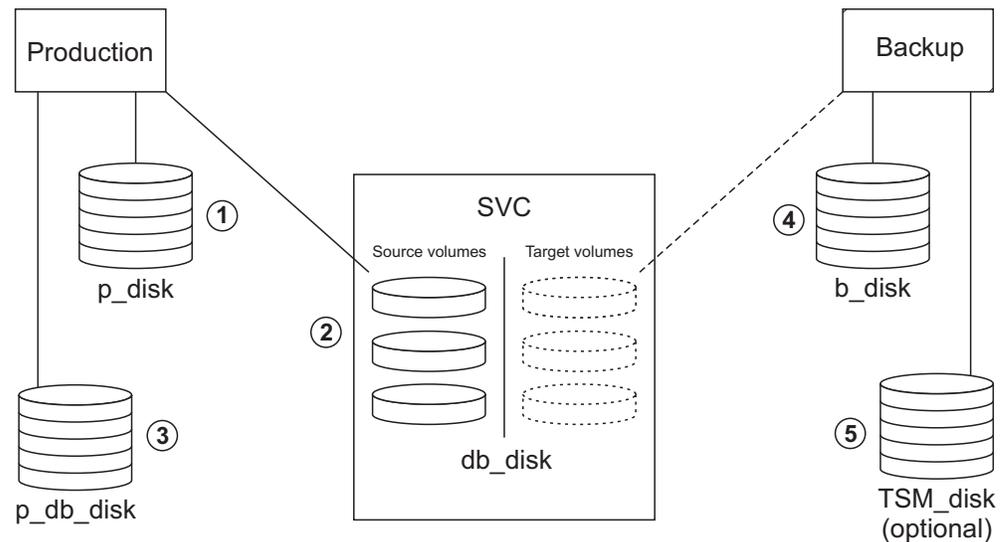


Figure 13. Example overall disk layout for a Oracle environment

The respective disk categories contain the following disk types that are used for the various file systems:

1. Local disks on the production system (p_disk category) for the file systems

```
/oracle/A01 part of VG ora_main
/oracle/A01/acs (ACS_DIR)
```

The Oracle database binary files are located in the directory that is defined by the ORACLE_HOME environment variable.

2. Source volume disks on the production system (db_disk category) for the file systems

```
/oracle/A01/oradata/system part of VG ora_d1
/oracle/A01/oradata/temp part of VG ora_d2
/oracle/A01/oradata/custom part of VG ora_d3
```

```
/oracle/A01/origlogA part of VG ora_11
/oracle/A01/origlogB part of VG ora_11
```

```
/oracle/A01/mirrlogA part of VG ora_12
/oracle/A01/mirrlogB part of VG ora_12
```

The file systems for the control files must not reside on volume groups that are part of the FlashCopy backup. For FlashCopy cloning at least one of the three Oracle control files must reside on a volume group that is part of the FlashCopy cloning operation.

```

/oracle/A01/cntrl/cntrlA01.dbf part of VG ora_main
/oracle/A01/misc/cntrlA01.dbf part of VG ora_misc
/oracle/A01/arch/cntrlA01.dbf part of VG ora_arch

```

3. The p_db_disk category is not used for this setup.
4. Local disks on the backup system (b_disk category) for the file system

```

/oracle/A01
/oracle/A01/acs (ACS_DIR)

```

There is no need to create separate volume groups for logs and control files.

Oracle profile example

The profile file provides parameters that customize how IBM Spectrum Protect Snapshot works within a particular environment. Use this example to verify that the configuration of the profile is correct for your Oracle environment.

The following sample profile is provided for when IBM Spectrum Protect Snapshot software backups up Oracle databases:

```

>>> GLOBAL
ACS_DIR /oracle/A10/acs          # directory for logs, password file, etc.
ACSD prodsrv 57328              # <server> <port>
# TRACE_NO                      # YES | NO
<<<<
>>> CLIENT
APPLICATION_TYPE ORACLE        #
TARGET_DATABASE_SUSPEND YES    # YES | NO | OFFLINE
# TSM_BACKUP_NO                # YES | NO
# MAX_VERSIONS ADAPTIVE        # num | ADAPTIVE
# LVM_FREEZE_THAW AUTO         # AUTO | YES | NO
# NEGATIVE_LIST NO_CHECK       # NO_CHECK | WARN | ERROR | <path to negative list file>
# DEVICE_CLASS STANDARD        #
<<<<
>>> OFFLOAD
BACKUP_METHOD ORACLE           #
# OVERWRITE_DATABASE_PARAMETER_FILE YES # YES | NO
# DATABASE_BACKUP_INCREMENTAL_LEVEL 0   #
<<<<
>>> ORACLE
CATALOG_DATABASE_CONNECT_STRING cat_db # *mandatory parameter*
CATALOG_DATABASE_USERNAME rman         # *mandatory parameter*
DATABASE_BACKUP_SCRIPT_FILE /oracle/A10/acs/tsm_backup.scr
# *mandatory parameter*
TARGET_DATABASE_PARAMETER_FILE /oracle/A10/10gr2/dbs/initA10.ora
# *mandatory parameter*
# DATABASE_CONTROL_FILE_RESTORE NO      # YES | NO
DATABASE_CONTROL_FILE_RESTORE NO        # YES | NO
<<<<
>>> ACSD
ACS_REPOSITORY /oracle/A10/acs/acsrepository # *mandatory parameter*
# ADMIN_ASSISTANT NO                      # NO | <server> <port>
<<<<
>>> DEVICE_CLASS STANDARD
COPYSERVICES_HARDWARE_TYPE SVC           # *mandatory parameter* SVC | DS8000 | XIV
COPYSERVICES_PRIMARY_SERVERNAME cim_srv  #
TARGET_SETS_VOLUMES_FILE # *mandatory parameter*
VOLUMES_FILE /oracle/A01/acs/volumes/volumes_file.fct
# COPYSERVICES_SECONDARY_SERVERNAME     #
# COPYSERVICES_USERNAME superuser       #
# SVC_COPY_RATE 80                      # num
# COPYSERVICES_COMMPROTOCOL HTTPS        # HTTP | HTTPS
# COPYSERVICES_CERTIFICATEFILE NO_CERTIFICATE # NO_CERTIFICATE | <certificate file>
COPYSERVICES_SERVERPORT 5999             # *mandatory parameter*
# FLASHCOPY_TYPE COPY                   # COPY | INCR | NOCOPY
# COPYSERVICES_TIMEOUT 6                 # num
# RESTORE_FORCE NO                       # YES | NO
BACKUP_HOST_NAME NONE
<<<<

```

Oracle ASM profile example

The profile file provides parameters that customize how IBM Spectrum Protect Snapshot works within a particular environment. Use this example to verify that the configuration of the profile is correct for your Oracle with ASM environment.

The following depicts a sample profile:

```
>>> GLOBAL
ACS_DIR /oracle/A10/acs          # directory for logs, password file, etc.
ACS prodsvr 57328                # <server> <port>
# TRACE NO                       # YES | NO
<<<<
>>> CLIENT
APPLICATION_TYPE ORACLE         #
TARGET_DATABASE_SUSPEND YES     # YES | NO | OFFLINE
# TSM_BACKUP NO                 # YES | NO
# MAX_VERSIONS ADAPTIVE         # num | ADAPTIVE
# LVM_FREEZE_THAW AUTO          # AUTO | YES | NO
# NEGATIVE_LIST NO_CHECK       # NO_CHECK | WARN | ERROR | <path to negative list file>
# DEVICE_CLASS STANDARD        #
VOLUME_MGR ASM                  # ASM|LVM
<<<<
>>> OFFLOAD
BACKUP_METHOD ORACLE           #
# OVERWRITE_DATABASE_PARAMETER_FILE YES # YES | NO
# DATABASE_BACKUP_INCREMENTAL_LEVEL 0   #
#ASM_INSTANCE_USER AUTO         # User name | AUTO
#ASM_INSTANCE_ID +ASM           # ASM instance SID | +ASM
#ASM_ROLE SYSDBA                # SYSDBA or SYSASM | SYSDBA
<<<<
>>> ORACLE
CATALOG_DATABASE_CONNECT_STRING cat_db # *mandatory parameter*
CATALOG_DATABASE_USERNAME rman         # *mandatory parameter*
DATABASE_BACKUP_SCRIPT_FILE /oracle/A10/acs/tsm_backup.scr
# *mandatory parameter*
TARGET_DATABASE_PARAMETER_FILE /oracle/A10/10gr2/dbs/initA10.ora
# *mandatory parameter*
DATABASE_CONTROL_FILE_RESTORE NO       # YES | NO
#ASM_INSTANCE_USER AUTO                 # User name | AUTO
#ASM_INSTANCE_ID +ASM                   # ASM instance SID | +ASM
#ASM_ROLE SYSDBA                        # SYSDBA or SYSASM | SYSDBA
<<<<
>>> ACSD
ACS_REPOSITORY /oracle/A10/acs/acsrepository # *mandatory parameter*
# ADMIN_ASSISTANT NO                       # NO | <server> <port>
<<<<
>>> DEVICE_CLASS STANDARD
COPYSERVICES_HARDWARE_TYPE SVC           # *mandatory parameter* SVC | DS8000 | XIV
COPYSERVICES_PRIMARY_SERVERNAME cim_srv  #
TARGET_SETS VOLUMES_FILE # *mandatory parameter*
VOLUMES_FILE /oracle/A01/acs/volumes/volumes_file.fct
# COPYSERVICES_SECONDARY_SERVERNAME     #
# COPYSERVICES_USERNAME superuser       #
# SVC_COPY_RATE 80                       # num
# COPYSERVICES_COMMPROTOCOL HTTPS        # HTTP | HTTPS
# COPYSERVICES_CERTIFICATEFILE NO_CERTIFICATE # NO_CERTIFICATE | <certificate file>
COPYSERVICES_SERVERPORT 5999            # *mandatory parameter*
# FLASHCOPY_TYPE COPY                    # COPY | INCR | NOCOPY
# COPYSERVICES_TIMEOUT 6                 # num
# RESTORE_FORCE NO                       # YES | NO
BACKUP_HOST_NAME NONE
<<<<
```

The **ASM_INSTANCE_USER**, **ASM_INSTANCE_ID**, and **ASM_ROLE SYSDBA** parameters in the OFFLOAD section contain values for the backup server. They can be commented out if the same parameters in the ORACLE section are also valid for the backup server.

RMAN backup script example

Refer to this example when you are configuring Data Protection for Oracle on the backup server.

The RMAN backup script must be specified in the profile with the **DATABASE_BACKUP_SCRIPT_FILE** parameter. The example shows the required syntax of the RMAN backup script. The keyword backup must be entered on a single line of text without any other text as shown in this example. There must be a line break after the backup keyword.

```
run
{
  allocate channel 'c1' type 'sbt_tape' parms 'ENV=(TDPO_OPTFILE=/home/oracle/tdpo.opt)';
  backup
  (database);
  release channel c1;
}
```

Oracle in an SAP environment (disk only) profile example

The profile file provides parameters that customize how IBM Spectrum Protect Snapshot works within a particular environment. Use this example to verify that the configuration of the profile is correct for your Oracle in an SAP environment.

The following depicts a sample profile:

```
>>> GLOBAL
# ACS_DIR /oracle/A01/acs
ACSD acsprod5 57328
# TRACE NO
<<<
>>> ACSD
ACS_REPOSITORY /oracle/A01/acs/pmtest
ADMIN_ASSISTANT no
# REPOSITORY_LABEL TSM
<<<
>>> CLIENT
# BACKUPIDPREFIX SAP_
APPLICATION_TYPE SAP_ORACLE
TARGET_DATABASE_SUSPEND YES
TSM_BACKUP YES
# MAX_VERSIONS ADAPTIVE
LVM_FREEZE_THAW 120
# TIMEOUT_FLASH 120
GLOBAL_SYSTEM_IDENTIFIER A01
# DEVICE_CLASS STANDARD
<<<
>>> DEVICE_CLASS STANDARD
COPYSERVICES_HARDWARE_TYPE XIV
# LVM_MIRRORING NO
PATH_TO_XCLI /home/xivtest/XCLI
COPYSERVICES_SERVERNAME nextra
COPYSERVICES_USERNAME admin
# RECON_INTERVAL 12
# USE_WRITABLE_SNAPSHOTS AUTO
BACKUP_HOST_NAME acsback5
<<<
```

DS8000 target volumes file example

Refer to this example when you are editing the target volumes file for a DS8000 storage subsystem configuration.

The following file is an example of a VOLUMES_FILE .fct file that includes the target set configuration that is used for cloning:

```
#
#***** First sample *****#
#
#=====#

>>> TARGET_SET 1
>>> PARTITION NODE0000
TARGET_VOLUME 13ABCTA0111 - -
TARGET_VOLUME 13ABCTA0112 - -
TARGET_VOLUME 13ABCTA0113 - -
<<<
<<<

>>> TARGET_SET 3

DEVICE_CLASS CLONE USE_FOR_CLONING D98
>>> PARTITION NODE0000
TARGET_VOLUME 13ABCTA011D - -
TARGET_VOLUME 13ABCTA011E - -
TARGET_VOLUME 13ABCTA011F - -
<<<
<<<

#=====#
```

The following file shows a VOLUMES_FILE .fct file for DB2 EEE configurations:

```
#
#***** Second sample *****#
#
#=====#

>>> TARGET_SET 1

DEVICE_CLASS CLONE USE_FOR_CLONING S98
  >>> PARTITION NODE0000
    TARGET_VOLUME S97p5d1_t1 - -
    TARGET_VOLUME S97p5d2_t1 - -
  <<<
  >>> PARTITION NODE0001
    TARGET_VOLUME S97p5d3_t1 - -
    TARGET_VOLUME S97p5d4_t1 - -
  <<<
  >>> PARTITION NODE0002
    TARGET_VOLUME S97p511_t1 - -
    TARGET_VOLUME S97p512_t1 - -
  <<<
  >>> PARTITION NODE0003
    TARGET_VOLUME S97p513_t1 - -
    TARGET_VOLUME S97p514_t1 - -
  <<<
<<<

#=====#
```

SAN Volume Controller and Storwize V7000 target volumes file example

Refer to this example when you are editing the target volumes file for an SAN Volume Controller or Storwize V7000 storage system configuration.

```
***** First sample *****#
#
#=====#
>>> TARGET_SET VOLUMES_SET_1
TARGET_VOLUME svdftgt1 svdrsrc2 -
TARGET_VOLUME svdftgt2 svdfsrc3 -
TARGET_VOLUME svdftgt3 svdfsrc4 -
TARGET_VOLUME svdftgt4 svdfsrc5 -
TARGET_VOLUME svdftgt5 svdfsrc6 -
<<<
#=====#
```

The following sample profile is an example of a profile in a non-mirrored environment. Create three space-efficient disk-only backups and one dual backup, at midnight, per day.

```
>>> CLIENT
...
TSM_BACKUP LATEST USE_FOR DISK_TSM
DEVICE_CLASS DISK_ONLY FROM 5:30 TO 23:59
DEVICE_CLASS DISK_TSM FROM 0:00 TO 05:29
<<<
>>> DEVICE_CLASS DISK_ONLY
COPYSERVICES_HARDWARE_TYPE SVC
FLASHCOPY_TYPE NOCOPY # space efficient targets
TARGET_SETS 1 2 3
TARGET_NAMING %SOURCE_%TARGETSET
...
<<<
>>> DEVICE_CLASS DISK_TSM
COPYSERVICES_HARDWARE_TYPE SVC
FLASHCOPY_TYPE NOCOPY # space efficient targets
TARGET_SETS DUAL
TARGET_NAMING %SOURCE_%TARGETSET
...
<<<
```

This scenario illustrates a profile in a mirrored environment. On MIRROR_1, two space-efficient FlashCopy backups are created on Monday, Wednesday, and Friday. The backup that is created at midnight is copied to IBM Spectrum Protect. The backup that is created at noon is retained only on disk. The backup that is created on Monday is retained until the target sets are reused on Wednesday. On MIRROR_2, only one incremental FlashCopy backup was created on Sunday, Tuesday, Thursday, and Saturday. This backup is also copied to IBM Spectrum Protect. The backup is retained until the next incremental backup is started.

```
>>> CLIENT
...
TSM_BACKUP LATEST USE_FOR MIRROR_1_DISK_TSM MIRROR_2
DEVICE_CLASS MIRROR_1_DISK_ONLY USE_AT Mon Wed Fri FROM 5:30 TO 23:59
DEVICE_CLASS MIRROR_1_DISK_TSM USE_AT Mon Wed Fri FROM 0:00 TO 05:29
DEVICE_CLASS MIRROR_2_USE_AT SUN Tue Thu Sat
<<<
>>> DEVICE_CLASS MIRROR_1_DISK_ONLY
COPYSERVICES_HARDWARE_TYPE SVC
FLASHCOPY_TYPE NOCOPY # space efficient targets
```

```

TARGET_SETS DO
TARGET_NAMING %SOURCE_%TARGETSET
...
<<<
>>> DEVICE_CLASS MIRROR_1_DISK_TSM
COPYSERVICES_HARDWARE_TYPE SVC
FLASHCOPY_TYPE NOCOPY # space efficient targets
TARGET_SETS DT
TARGET_NAMING %SOURCE_%TARGETSET
...
<<<
>>> DEVICE_CLASS MIRROR_2
COPYSERVICES_HARDWARE_TYPE SVC
FLASHCOPY_TYPE INCR
TARGET_SETS 1
TARGET_NAMING %SOURCE_%TARGETSET
...
<<<

```

This example is like the previous example, but the example does not create IBM Spectrum Protect backups from MIRROR_1. Rather, the example retains the space-efficient FlashCopy images for one week (same schedule).

```

>>> CLIENT
...
TSM_BACKUP LATEST USE_FOR MIRROR_1_DISK_TSM MIRROR_2
DEVICE_CLASS MIRROR_1_DISK_ONLY USE_AT Mon Wed Fri
DEVICE_CLASS MIRROR_2_USE_AT Sun Tue Thu Sat
<<<
>>> DEVICE_CLASS MIRROR_1_DISK_ONLY
COPYSERVICES_HARDWARE_TYPE SVC
FLASHCOPY_TYPE NOCOPY # space efficient targets
TARGET_SETS 1A 1B 3A 3B 5A 5B
TARGET_NAMING %SOURCE_%TARGETSET
...
<<<
>>> DEVICE_CLASS MIRROR_2
COPYSERVICES_HARDWARE_TYPE SVC
FLASHCOPY_TYPE INCR
TARGET_SETS 1
TARGET_NAMING %SOURCE_%TARGETSET
...
<<<

```

Appendix D. Accessibility features for the IBM Spectrum Protect product family

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Spectrum Protect family of products includes the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Spectrum Protect family of products uses the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the product.

The product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/about/releasenotes.html?view=kc#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

User interfaces do not have content that flashes 2 - 55 times per second.

Web user interfaces rely on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

Web user interfaces include WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Spectrum Protect product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those

websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

SoftLayer[®] is a registered trademark of SoftLayer, Inc., an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Glossary

A glossary is available with terms and definitions for the IBM Spectrum Protect family of products.

See the IBM Spectrum Protect glossary.

To view glossaries for other IBM products, see IBM Terminology.

Index

A

- accessibility features 201
- ACS_DIR 116
- ACS_REPOSITORY 117
- acsd 181
- ACSD 116, 117
- acsgen 183
- acsora 165
- acsutil 167
- activate a database 54
- adding instance
 - RAC 43
- ADMIN_ASSISTANT 117
- administrative commands 173
- APAR 101
- application agents
 - acsora 4
 - backint 4
 - fccli 4
- APPLICATION_TYPE 119
- ASM
 - DS8000 environment 33
 - failure groups 77
 - DEVICE_CLASS 65
 - planning 32
- ASM cloning 33
- ASM failure group 34
- ASM_INSTANCE_ID 124, 149
- ASM_INSTANCE_USER 124, 149
- ASM_ROLE 124, 149
- ASNODENAME 149
- Asynchronous Remote Mirroring 20
- audience ix

B

- background daemons
 - description 181
- backup
 - to IBM Spectrum Protect 189
- backup commands 165
- backup environment
 - backup servers
 - assignment 66
 - clone servers
 - assignment 66
 - preparing
 - backup servers 37
 - clone servers 37
- backup procedure
 - fully automated backup 77
 - Oracle 75, 76, 77
 - to IBM Spectrum Protect server 76
- backup server
 - virtual machine 40
- BACKUP_DESTINATION 124
- BACKUP_HOST_NAME 132, 133, 145
- BACKUP_METHOD 149
- BACKUPIDPREFIX 119
- block device mapping 26

- BUFFER_SIZE 149

C

- CA certificate 178
- capacity planning 8
- CATALOG_DATABASE_USERNAME 124
- CATALOG_DATABASE_CONNECT_STRING 124
- changing passwords 176
- CIM
 - acsgen 183
 - generic device agent 183
 - log files 106
 - trace files 106
- CLIENT 119
- clone devices
 - DS8000 17
 - SAN Volume Controller 12
 - Storwize V7000 12
- clone server
 - virtual machine 40
- CLONE_DATABASE 132, 135, 145
- cloning
 - commands 169
 - configuration files 98
 - target volumes file (.fct)
 - parameter settings 158
- command line interface
 - mount 188
 - unmount 189
- commands 165
 - administrative 173
 - backup 165
 - cloning 169
 - delete 165
 - inquire 165
 - restore 165
- communication ports 9
- configuration 49
- Configuration 49
- configuration files
 - configuration file 113
 - password file 159
- configuring 50
 - remote mirroring 70
- configuring database instance 45
- configuring db instance 45
- configuring instance 45
- Configuring SVC 55
- consistency groups 22
- control files
 - cloning files
 - description 98
 - profile
 - structure 113
- target volumes file (.fct)
 - cloning parameter settings 158
 - description 153
 - DS8000 parameter settings 155
 - SAN Volume Controller parameter settings 156
 - Storwize V7000 parameter settings 156

COPYSERVICES_HARDWARE_TYPE parameter
 description 141
 COPYSERVICES_CERTIFICATEFILE 132, 135, 145
 COPYSERVICES_CERTIFICATEFILE parameter
 description 141
 COPYSERVICES_COMMPROTOCOL 132, 135, 145
 COPYSERVICES_COMMPROTOCOL parameter
 description 141
 COPYSERVICES_HARDWARE_TYPE 132, 135, 145
 CLONE_DATABASE 133
 COPYSERVICES_PRIMARY_SERVERNAME 132, 135, 145
 COPYSERVICES_PRIMARY_SERVERNAME parameter
 description 141
 COPYSERVICES_SECONDARY_SERVERNAME 132, 135,
 145
 COPYSERVICES_SERVERNAME 132, 133
 COPYSERVICES_SERVERNAME parameter
 description 141
 COPYSERVICES_SERVERPORT 132, 135, 145
 COPYSERVICES_SERVERPORT parameter
 description 141
 COPYSERVICES_TIMEOUT 132, 135, 145
 COPYSERVICES_TIMEOUT parameter
 description 141
 COPYSERVICES_USERNAME 132, 133, 135, 145
 COPYSERVICES_USERNAME parameter
 description 141

D

database cloning
 commands 96
 description 3
 Oracle Data Guard 97
 postprocessing 98
 preprocessing 98
 process 95
 DATABASE_BACKUP_INCREMENTAL_LEVEL 149
 DATABASE_BACKUP_MSGLOG_FILE 124
 DATABASE_BACKUP_SCRIPT_FILE 124
 DATABASE_CONTROL_FILE_RESTORE 124
 DATABASE_MEMORY 149
 DBPARTITIONNUM 149
 defining LUNs
 DS8000 59
 defining target sets
 naming convention 64
 profile 64
 single partition 63
 defining virtual disks
 SAN Volume Controller 60
 Storwize V7000 60
 delete commands 165
 deleting
 snapshot backup 171
 developerWorks wiki 101
 device agents
 acsd 4
 acsgen 4
 CIM adapter 4
 command-line interface 4
 offload agent 4
 query capacity 4
 volume group takeover script 4
 DEVICE_CLASS 119
 device 65, 66, 132
 parameters 65, 66

devices
 DS8000 17
 IBM XIV Storage System
 preparing 10
 SAN Volume Controller 12, 15
 storage
 log files 106
 setting up 55
 trace files 106
 Storwize V7000 12, 15
 disability 201
 disk layout
 examples
 Oracle environment 193
 DMS_OPT 149
 DS8000
 defining LUNs 59
 examples
 target volumes file 197
 planning 17
 target volumes file (.fct)
 parameter settings 155
 DSM_DIR 149
 DSM_LOG 149
 Dynamic Target Allocation 55

E

encryption
 SP 800-131A 59
 ENHANCED_PARTITIONING 119
 environment
 ASM 33
 backup servers 37
 assignment 66
 clone servers 37
 assignment 66
 Oracle 31
 examples 193
 Oracle ASM environment
 profile 195
 Oracle environment
 disk layout 193
 profile 194
 Oracle in an SAP environment
 profile for disk only 196
 RMAN backup script 196
 target volumes file
 DS8000 197
 SAN Volume Controller 198
 Storwize V7000 198

F

fcmlcli 189
 FlashCopy 2
 FlashCopy backup
 reconciliation 19
 FLASHCOPY_TYPE 132, 135, 145
 FLASHCOPY_TYPE parameter
 description 141
 fmquery
 description 180
 fully automated backup 77

G

- generic device agent 183
- GLOBAL 116
- Global Mirror 20
- GLOBAL_SYSTEM_IDENTIFIER 119
- GSKit
 - configuring 160
 - FIPS 160
 - installing 160
- guest operating system 25

H

- high availability daemons 53

I

- IBM Knowledge Center x
- IBM Spectrum Protect
 - integration 189
- IBM Spectrum Protect for ERP
 - log files 106
 - trace files 106
- IBM Spectrum Protect Snapshot 57
 - snapshot restore 87, 88
- individual mirrors 77, 78
- inquire commands 165
- installation
 - description 41
 - prerequisites 27, 72
- installation planning 49
- instance ID 43
- instance names
 - cloning 97
 - Oracle RAC 97
- Internet Protocol Version 6 111
- introduction 1

K

- keyboard 201
- Knowledge Center x
- KVM 26

L

- log files
 - CIM 106
 - IBM Spectrum Protect for ERP 106
 - IBM Spectrum Protect Snapshot 102
 - storage subsystems 106
 - summary 102
- LVM_FREEZE_THAW 119
- TARGET_DATABASE_SUSPEND 152
- LVM_MIRRORING 132, 133, 135, 145

M

- management agent 181
- MAX_VERSIONS 119
- Metro Mirror 20
- migrating 47

- mirroring
 - individual 77, 78
 - LVM 23, 77, 78
 - DEVICE_CLASS 65, 66
 - Oracle ASM 65
 - overview 23
- MODE 149
- multipath
 - KVM check 26

N

- N Series deprecation xi
- naming convention 64
- NEGATIVE_LIST 119
- NetApp deprecation xi
- new xi
- NUM_BUFFERS 149
- NUM_SESSIONS 149
- NUMBER_BACKUPS_IN_PARALLEL 149

O

- OFFLOAD 149
- offload agent
 - description 189
 - tape_backup 190
 - update_status 192
- option files
 - IBM Spectrum Protect and Data Protection for Oracle 162
- OPTIONS 149
- Oracle 49, 92
 - ASM environment 33
 - profile example 195
 - environment 31
 - disk layout example 193
 - profile example 194
 - SAP environment
 - profile example for disk only 196
 - troubleshooting 109
 - miscellaneous errors 110
 - variables
 - troubleshooting 110
- ORACLE 124
- Oracle ASM
 - failure groups 77
 - DEVICE_CLASS 65
 - preparing 32
- Oracle configuration 50
- Oracle Data Guard
 - restore 92
- Oracle RAC
 - cloning 96
 - high availability 53
 - installation setup 35
 - preparation 35
- Oracle SAP 49
- overview 1
- OVERWRITE_DATABASE_PARAMETER_FILE 149

P

- PARALLEL_BACKUP 149
- PARALLEL_RESTORE 149
- PARALLELISM 149

- parameters
 - DEVICE_CLASS 68
- PARTITION_GROUP 119
- password administration 176
- password file 159
- passwords
 - changing 176
 - setup script 176
- PATH_TO_XCLI 132
- planning 7, 29
 - capacity 8
 - checklist 7
- preinstallation checklist 49
- preinstallation planning 7
- preparing db instance 45
- prerequisite checker 7, 29
 - installing 28
- Prerequisite checker
 - results 29
- Prerequisite Checker tool
 - Uninstall 30
- prerequisites
 - hardware 27, 72
 - software 27, 72
- problem determination
 - general procedure 101
 - Oracle environment 109
 - miscellaneous errors 110
 - Oracle variables 110
- product support 101
- profile
 - example for disk only
 - Oracle in an SAP environment 196
 - examples
 - Oracle ASM environment 195
 - Oracle environment 194
 - structure 113
 - target set definitions 64
 - naming convention 64
- PROFILE 149
- profile parameter sections
 - DEVICE_CLASS device 141
- publications x

Q

- query managed capacity (fmquery)
 - description 180

R

- RECON_INTERVAL 132, 133, 135, 145
- RECON_INTERVAL parameter
 - description 141
- reconciliation
 - FlashCopy backup 19
- Redundancy levels 32
- remote mirror 92
- remote mirroring xi, 22
- repository
 - snapshot backup
 - status 172
- REPOSITORY_LABEL 117
- restore commands 165
- restore procedure
 - Oracle database 87

- restore procedure (*continued*)
 - Oracle RAC 88
 - snapshot restore 87, 88
- RESTORE_FORCE 132, 135, 145
- restoring 92
 - Oracle databases 90, 91
- RMAN backup script
 - example 196

S

- SAN Volume Controller
 - examples
 - target volumes file 198
 - planning 12
 - target volumes file (.fct)
 - parameter settings 156
- Service Management Console 101
- setup script 49
 - description 173
 - values 174
- shared file system
 - RAC 35
- snapshot 2
 - mounting 185
 - unmounting 185
- snapshot backup 80
 - deleting 171
 - repository
 - status 172
- snapshot devices
 - DS8000 17
 - SAN Volume Controller 12
 - Storwize V7000 12
- Snapshot Object Manager, acsutil 167
- SSH key file 55
- status
 - repository 172
- storage solutions
 - preparing 10
- storage subsystems
 - DS8000
 - defining LUNs 59
 - log files 106
 - preparing
 - IBM XIV Storage System 10
 - SAN Volume Controller 15
 - Storwize V7000 15
 - SAN Volume Controller
 - defining virtual disks 60
 - setting up 55
 - Storwize V7000
 - defining virtual disks 60
 - trace files 106
- Storwize V7000
 - examples
 - target volumes file 198
 - planning 12
 - target volumes file (.fct)
 - parameter settings 156
- SVC 20, 22, 55
- SVC dynamic target allocation xi
- SVC_CLEAN_RATE 132, 135
- SVC_COPY_RATE 132, 135
- SVC_COPY_RATE parameter
 - description 141
- SVC_GRAIN_SIZE 135

SVCDDTA 20
SVCDDTA SVC Migrating to new adapter 57
Synchronous Remote Mirroring 20

T

target set definitions 63
files 63
naming convention 63, 64
profile 64
target volumes
storage systems 154
target volumes file
examples
DS8000 197
SAN Volume Controller 198
Storwize V7000 198
target volumes file (.fct)
description 153
parameter settings
cloning 158
DS8000 155
SAN Volume Controller 156
Storwize V7000 156
TARGET_DATABASE_HOME 124
TARGET_DATABASE_SID 124
TARGET_DATABASE_USERNAME 124
TARGET_DATABASE_PARAMETER_FILE 124
TARGET_DATABASE_PASSWORD_FILE 124
TARGET_DATABASE_SUSPEND 119
TARGET_NAMING 132, 135
TARGET_NAMING parameter
description 141
TARGET_SETS 132, 135, 145
TARGET_SETS parameter
description 141
TIMEOUT_<PHASE> 119
TIMEOUT_FLASH 119
TRACE 116
trace files
CIM 106
IBM Spectrum Protect for ERP 106
IBM Spectrum Protect Snapshot 102
storage subsystems 106
summary 102
troubleshooting 108
general procedure 101
Oracle environment 109
miscellaneous errors 110
Oracle variables 110
TSM_BACKUP 119

U

uninstalling 47
upgrade
production server 46
upgrading
migrating 47
process 41
usability states 80
USE_WRITABLE_SNAPSHOTS 132

V

V8.1.0 xi
VENDOR_BIN 149
VENDOR_LIB 149
VENDOR_PATH 149
VENDOR_PROFILE 149
virtual machine 25
VIRTUALFSNAME 149
VOLUME_MGR 119
VOLUMES_DIR 132
VOLUMES_FILE 132, 135, 145
VOLUMES_FILE parameter
description 141



Product Number: 5725-X22
5608-AB8

Printed in USA