

IBM Storage Protect Snapshot: Installation and User Guide

Db2

Version 8.2.0



Contents

List of Figures.....	5
List of Tables.....	6
Who should read this guide	9
Publications	9
New in IBM® Storage Protect Snapshot 8.2.0.....	10
New in IBM® Storage Protect Snapshot 8.1.24	10
New in IBM® Spectrum Protect Snapshot 8.1.11.....	10
New in IBM® Spectrum Protect Snapshot 8.1.4	11
.....	11
New and modified parameters or functions	11
Overview	12
Backup and restore methods with and snapshots	12
Snapshots and FlashCopies	12
Types of snapshot backups.....	13
Transferring snapshots to an IBM® Storage Protect server.....	13
Database cloning.....	13
Software components.....	14
Planning.....	16
IBM® Storage Protect Snapshot Prerequisite Checker.....	16
Capacity planning.....	16
Space requirement for global product installation	17
Space requirement for instances	17
Space requirement for snapshot copies	17
Required communication ports	18
Storage solutions.....	18
IBM® XIV® Storage System	18
SAN Volume Controller and Storwize® family storage systems	20
DS8000® storage system.....	24
Reconciliation of backups	25
Remote mirror integration.....	26
Remote mirroring and consistency groups.....	28
HyperSwap integration	29
Preparing applications that run on VMware or KVM.....	31
Checking the KVM setup	31
Installing and setting up IBM® Storage Protect Snapshot	32
Preparing for installing.....	32
Prerequisite checker for DB2®.....	33
Verify the DB2® system	35
Preparing the DB2® high-availability disaster recovery server or DB2® standby server	35
Preparing backup and cloning servers.....	37
Installing and uninstalling IBM® Storage Protect Snapshot for DB2	40
Installing IBM® Storage Protect Snapshot in interactive mode	40
Installing in silent mode.....	41
Uninstalling the software	41
Activating an instance.....	42
Setting up IBM® Storage Protect Snapshot separately on backup or clone servers.....	42
Setting up IBM® Storage Protect Snapshot on a backup server	42
Setting up IBM® Storage Protect Snapshot on a clone server.....	43
Upgrading IBM® Storage Protect Snapshot on a backup or clone server.....	44
Setting up a backup server for multiple production systems in a DB2® HADR environment.....	44
Configuring or reconfiguring IBM® Storage Protect Snapshot.....	45
Running the setup script for IBM® Storage Protect Snapshot for DB2®	46
Configuring storage environments	48
Backup and clone server assignment	58
Managing backups and clones with the DEVICE_CLASS parameter.....	59
Configuring for remote mirroring.....	61
Setting up daemons.....	63
Postinstallation and post-configuration tasks.....	64
Upgrading.....	64
Upgrading from IBM® Tivoli® Storage FlashCopy® Manager version 3.1.....	65
Protecting your data with IBM® Storage Protect Snapshot.....	66

Backing up data	66
DB2® backup commands	66
Backing up DB2® (HADR) standby servers	68
Snapshot backup of individual mirrors	69
Using Snapshot backup in a HyperSwap environment	73
Using Volume Group Snapshot in IBM FlashSystems	73
Restoring data	74
Restoring DB2® databases	74
Restoring data with remote mirroring	76
Using Snapshot restore in a HyperSwap environment	77
Cloning databases	79
Cloning and IBM® System Storage® SAN Volume Controller	79
Cloning databases with IBM® Storage Protect Snapshot	80
Database cloning preprocessing and postprocessing	80
Configuration files used for cloning	80
Usability states of snapshot backup operations	82
Usability state diagrams	83
Snapshot backup	83
Snapshot restore	84
Snapshot delete	84
Snapshot mount	85
Snapshot offload	85
Troubleshooting	87
General troubleshooting procedure	87
Logging and tracing files	88
Log files and trace files	88
Storage system log and trace files	91
CIM log and trace files	91
GPFS™ log files	92
IBM® Storage Protect for Enterprise Resource Planning log and trace files	92
Troubleshooting mirroring relationships	92
Troubleshooting storage solutions	93
Troubleshooting connectivity problems	93
When the production server and backup server are separated by a firewall, socket connections might time out	93
Troubleshooting DB2® setup problems	94
Setting up a backup server for multiple production systems in a DB2® HADR environment	94
Internet Protocol Version 6 (IPv6) support	94
Configuration files	96
Profile	96
GLOBAL	99
ACSD	100
CLIENT	101
CLONING	105
DEVICE_CLASS <i>device</i>	108
OFFLOAD	124
DB2STANDBY	128
OFFLOAD_DB2STANDBY	131
Changing profile parameters	131
Interdependency of LVM_FREEZE_THAW and TARGET_DATABASE_SUSPEND	132
Target set and target volumes files	132
Manage target volumes files for your storage system	133
DS8000® target volume parameter settings	135
SAN Volume Controller and Storwize® family target volume parameter settings	136
Target set handling for cloning	137
IBM® Storage Protect Snapshot password file	138
Commands and scripts	140
Backup, restore, cloning commands, and utilities	140
Commands that use DB2® utilities	140
Cloning commands	144
FlashCopy® cloning function-clauses	145
Deleting snapshot backups	146
Deleting a target volume or target set	146
Snapshot backup status in the repository	147
Administrative commands	147
Configuration commands	147
Background daemons	155
Mounting and unmounting snapshots on a secondary system	159
Integration with IBM® Storage Protect	164
IBM® Global Security Kit configuration	168
Enforcing SP800-131 compliant encryption	169
Uninstall GSKit	169

Examples	170
DB2® overall disk layout example.....	170
DB2® profile example.....	171
Target volumes file examples	172
SAN Volume Controller and Storwize® family target volumes file example	173
SAP in a multi-partition DB2® target volume file example	174
SAN Volume Controller and Storwize® family target volumes file example	175
DB2® HADR standby server profile example.....	177
Accessibility features for the IBM Storage® Protect product family.....	178
Overview	178
Keyboard navigation	178
Interface information	178
Vendor software	178
Related accessibility information.....	178
Notices	179
Trademarks.....	180
Terms and conditions for product documentation	180
Privacy policy considerations	181
Glossary	182
Index	183

List of Figures

Figure 1: IBM® Storage Protect Snapshot backup and restore environment	13
Figure 2: IBM® Storage Protect Snapshot and database cloning	14
Figure 3: IBM® Storage Protect Snapshot system components	15
Figure 4: Remote mirroring using Metro Mirror and Global Mirror sources	28
Figure 5: IBM® Storage Protect Snapshot in a HyperSwap environment	30
Figure 6: Protection of DB2® databases in a HADR environment with IBM® Storage Protect Snapshot	36
Figure 7: IBM® Storage Protect Snapshot host assignments for DB2®	59
Figure 8: Cross-site mirrored SAP database that is protected with IBM® Storage Protect Snapshot and an IBM® Storage Protect server	70
Figure 9: IBM® Storage Protect Snapshot in an LVM environment	73
Figure 10: Confirming site location	74
Figure 11: IBM® Storage Protect Snapshot after a HyperSwap restore	78
Figure 12: IBM® Storage Protect Snapshot after running the reconfiguration script	79
Figure 13: Usability States during snapshot backup	84
Figure 14: Usability states during snapshot restore	84
Figure 15: Usability states during snapshot delete	85
Figure 16: Usability states during snapshot mount	85
Figure 17: Usability states during snapshot offload	86
Figure 18: fcmcli command	141
Figure 19: fcmcli command functions	141
Figure 20: Protection of DB2® HADR standby nodes with IBM® Storage Protect Snapshot for Custom Applications	142
Figure 21: fcmcli command	143
Figure	144
Figure 22: fcmcli help	144
Figure	145
Figure	145
Figure 23: fcmcli command: -f password	151
Figure 24: acsd management agent	155
Figure 25: acsd management agent help	155
Figure 26: acsgen generic device agent	156
Figure 27: acsgen generic device agent help	156
Figure 28: fcmcli command	159
Figure 29: fcmcli command	160
Figure 30: -f mountfunction-clauses	161
Figure 31: -f mountfunction-clauses	161
Figure 32: -f unmountfunction-clause with force option	161
Figure 33: -f unmountfunction-clause with force option	161
Figure 34: -f mountand -f unmountfunction-clauses	162
Figure 35: -f unmountfunction-clause with force option	162
Figure 36: fcmcli command	164
Figure 37: -f tape_backupfunction-clause	165
Figure 38: -f update_statusfunction-clause	165
Figure 39: Example overall disk layout for a DB2® environment	170

List of Tables

Table 1: Space requirements for a global product installation of IBM® Storage Protect Snapshot	17
Table 2: IBM® Storage Protect Snapshot for UNIX™ and Linux™ default port numbers	18
Table 3: Dynamic target volumes and predefined target volumes feature comparison.	20
Table 4	48
Table 5: Selecting the FLASHCOPY_TYPE for DS8000®, SAN Volume Controller, and Storwize® family.....	53
Table 6: Supported storage subsystems and FlashCopy® types	54
Table 7: Summary of backup commands for DB2® databases	66
Table 8: Summary of Restore Commands for DB2®	74
Table 9: Usability states	82
Table 10: Message prefixes used in the summary log file	87
Table 11: IBM® Storage Protect Snapshot log files	88
Table 12: IBM® Storage Protect Snapshot trace files	89
Table 13: IBM® Storage Protect Snapshot return codes.....	89
Table 14: IBM® Storage Protect Snapshot installer exit codes	90
Table 15: DB2® vendor reason codes	90
Table 16: Actions taken depending on values of LVM_FREEZE_THAW and TARGET_DATABASE_SUSPEND	132
Table 17: Managing target volume LUNs by storage system	133
Table 18: TARGET_VOLUME parameters	135
Table 19: TARGET_VOLUME parameters (SAN Volume Controller and Storwize® family)	136
Table 20: Options for starting the management agent, acsd, as a daemon process.....	155
Table 21: Options for starting the generic device agent, acsgen.....	157
Table 22: Options for starting the generic device agent, acsgen.....	158
Table 23: Options for starting the generic device agent, acsgen.....	158

Note:

Before you use this information and the product it supports, read the information in [“Notices” on page 179](#).

This edition applies to version 8, release 2 of IBM® Storage Protect Snapshot (product numbers 5725-X22, and 5608-AB8) and to all subsequent releases and modifications until otherwise indicated in new editions.

About this guide

This guide provides you with information about how to set up IBM® Storage Protect Snapshot for UNIX™ and Linux™. The information brings you through the steps from Planning, through to installing, configuring, administering, and operating the product for your particular setup.

IBM® Storage Protect Snapshot for Oracle® is provided as a single installation package for AIX® or Linux™. The product runs on the following storage systems:

- IBM® System Storage® DS8000®
- IBM® System Storage® SAN Volume Controller
- IBM® XIV® Storage System
- IBM® Storwize® family and IBM® Storwize® V7000 Unified
- IBM FlashSystem family

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

IBM® Storage 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® Storage 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® Storage 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® Storage Protect

Publications

The IBM Storage® Protect product family includes IBM Storage® Protect Plus, IBM Storage® Protect for Virtual Environments, IBM Storage® Protect for Databases, and several other storage management products from IBM®.

To view IBM® product documentation, see [IBM® Documentation](#).

What's new for Db2®

Learn about new features and enhancements in IBM® Storage Protect Snapshot 8.2.0.

New in IBM® Storage Protect Snapshot 8.2.0

IBM Storage Protect Snapshot now supports DB2 12.1.0 version

Support for DB2 12.1.0 version is now available in IBM® Storage Protect Snapshot.

Support for IBM FlashSystem coordinated HA snapshots

Beginning with version 8.2.0, IBM® Storage Protect Snapshot supports backup, restore, delete, and offload operations using Coordinated HA snapshot feature of Policy-Based High Availability (HA) topologies for the IBM FlashSystem.

Added the following new parameters under the SVCDDTA device class section to support IBM FlashSystem coordinated high availability snapshots

- **SVC_HA_SITE_SERVERNAME** specifies the second site of IBM FlashSystem in Policy-Based High Availability (HA) topology. Supported values for this parameter is the name of the second site as configured in the IBM FlashSystem.
- **SVC_HA_SITE_USERNAME** specifies the username for **SVC_HA_SITE_SERVERNAME**.
- **SVC_HA_SITE_SSHKEY_FULLPATH** specifies the ssh full key path that can be used to perform operations on **SVC_HA_SITE_SERVERNAME**.

Note: For enabling coordinated high availability feature, you must configure **COPYSERVICES_SERVERNAME**, **COPYSERVICES_USERNAME**, **SVC_SSHKEY_FULLPATH** for one of the FlashSystem and above newly added parameters for another FlashSystem in a Policy Based High Availability setup.

Note:

- IBM FlashSystem 2 Site HA configuration is supported.
- IBM FlashSystem 3 Site HA DR configuration is not supported.

New in IBM® Storage Protect Snapshot 8.1.24

Support for IBM FlashSystem volume group snapshots

Beginning with version 8.1.24, IBM® Storage Protect Snapshot supports backup, in-place restore, delete, and offload operations from single site and Policy-Based High Availability (HA) topologies for the IBM FlashSystem volume group snapshots.

Added a new parameter under the SVCDDTA device class section to enable IBM FlashSystem volume group snapshots

- The SVCDDTA device class section now contains an additional parameter named **SVC_ENABLE_VOLUMEGROUP_SNAPSHOT** that enables IBM FlashSystem volume group snapshots.
- Supported values for this parameter are **IGNORE_FLASHCOPY_MAPPINGS** and **ENFORCE_NO_FLASHCOPY_MAPPINGS**.

For more information, see [“Configuring Storwize family and SAN Volume Controller dynamic target allocation \(SVCDDTA\)”](#) on page 50.

New in IBM® Spectrum Protect Snapshot 8.1.11

Avail of a new level of data availability with IBM® Spectrum Protect Snapshot support for IBM HyperSwap technology, which enables mission-critical operations to continue without interruption during a complete site failure. HyperSwap systems support a dual-site, active-active solution that provides continuous availability of data during planned and unplanned outages.

For information on planning a HyperSwap system supported IBM® Spectrum Protect Snapshot, see the planning topic about [HyperSwap integration](#).

New in IBM® Spectrum Protect Snapshot 8.1.4

Specify different backup version retention values

For all storage devices, you can specify different backup retention periods for each device class that is configured. During the configuration process, you can define different values for each device class with the **MAX_VERSIONS** parameter in the profile. For more information about this feature, see [Backup version retention](#).

IBM SAN Volume Controller Dynamic Target Allocation (SVC DTA) incremental backups

Run incremental backups for SVC DTA with the **FLASHCOPY_TYPE** option INCR. After the initial FlashCopy backup, incremental FlashCopy backups are run to the same target volume. The ability to run incremental backups to the SVC, results in improvements in how long the backup takes and eases the load on the SVC. For more information about incremental flash copies, see [Incremental backups and MAX_VERSIONS](#).

New and modified parameters or functions

The following parameters are modified for IBM® Spectrum Protect Snapshot 8.1.4:

MAX_VERSIONS for all device classes

Use this parameter in the CLIENT section of the profile for each **DEVICE_CLASS** to create incremental FlashCopies. When **FLASHCOPY_TYPE** is set to INCR, the incremental FlashCopy is refreshed depending on the **MAX_VERSIONS** value set. Specify the maximum number of snapshot backup versions to be kept before the oldest backup is deleted by using **MAX_VERSIONS** to specify the maximum. For SVC DTA, when the version-delete deletes the oldest snapshot backup with an incremental FlashCopy relation, it reuses the target volumes of the expired backup and refreshes the FlashCopy relationship instead of deleting it. For more information about **MAX_VERSIONS**, see [“CLIENT” on page 101](#).

Overview

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

The following list identifies the applications that can be protected and cloned with IBM® Storage 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 in an SAP environment
- Oracle with Automatic Storage Management (ASM)
- Oracle in a RAC environment.
- Custom Applications (without cloning)

IBM® Storage Protect Snapshot can back up DB2® databases that are on snapshot-oriented storage systems or file systems.

IBM® Storage Protect Snapshot supports AIX® and Linux™ operating systems.

When you use IBM® Storage Protect Snapshot with other IBM® Storage Protect products, snapshots can be sent to the server.

- Transfer snapshots by using IBM® Storage Protect for Enterprise Resource Planning, IBM® DB2® Storage Protect client for DB2® environments, or the IBM® Storage Protect backup-archive client.
- To send snapshot backups to IBM® Storage Protect, you must configure a backup server or cluster.
- For DB2® in SAP environments, use IBM® Storage Protect for Enterprise Resource Planning.
- For DB2® environments, use the IBM® DB2® Storage Protect client for DB2® environments.

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

- IBM® XIV® Storage System
- IBM® Storwize® family
- IBM® System Storage® SAN Volume Controller
- IBM® System Storage® DS8000®

Backup and restore methods with and snapshots

The terms *snapshot* or *FlashCopy* are used differently depending on your hardware storage. Both denote a logical point-in-time copy where the target volume represents an exact copy of the data on a source volume. The term *snapshot* is used generically to apply to all hardware types.

Snapshots and FlashCopies

Depending on the storage hardware you use, the terms *snapshot* or *FlashCopy* are used. Both denote a logical point-in-time copy, where the target volume represents an exact copy of the data on a source volume. Data is transferred to the target volume as the source volume is modified. This action is called copy-on-write, or redirect-on-write. The logical copy can be transformed into a physical full copy on the target volume. When target volumes must be provided in advance, 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. When data is restored, it is copied from the target to the source volume. The term *snapshot* is used to signify snapshot and FlashCopy.

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 that are written on the source volume after the snapshot was created are copied to the target volume.

Transferring snapshots to an IBM® Storage Protect server

When you use IBM® Storage Protect Snapshot with IBM® Storage Protect products, you can transfer snapshots to the IBM® Storage Protect server. To send these snapshot backups to the IBM® Storage 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® Storage 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® Storage Protect Snapshot uses the FlashCopy® or snapshot function of the storage hardware for database cloning. Cloned databases are required in 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 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® Storage 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® Storage Protect Snapshot to mount a cloned database.

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




Figure 2: IBM® Storage Protect Snapshot and database cloning

Software components

IBM® Storage Protect Snapshot is composed of several software components.




Figure 3: IBM® Storage 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® Storage 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 specific adapters for each storage device and the management 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 specific agents for each device type 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® family, and SAN Volume Controller with static target allocation.
- The storage adapter for SVCDDTA communicates with the CLI interfaces of SAN Volume Controller and the Storwize® family storage systems with dynamic target allocation via Secure Shell.
- The XIV storage adapter is used with the generic device agent `acsgen`. This adapter communicates with the `acsgen` agent and issues commands to the XIV® Storage System by using the command line interface `XCLI`.
- Third-party adapters that can communicate with non-IBM devices are available. Third-party adapters are not part of the IBM® Storage Protect Snapshot product.

Offload agent

The offload agent `fcmcli` is used to send an existing snapshot to an IBM® Storage Protect server. This agent also calls the generic device agent for mount and unmount operations on a backup system.

IBM® Storage Protect Snapshot command line interface

The command line interface `fcmcli`, is used to issue various commands.

Planning

Before you install IBM® Storage 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® Storage 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 *Pre-installation Checklist* documents the requirements that can be verified automatically by running the IBM® Storage Protect Snapshot Prerequisite Checker.

The *Pre-installation Checklist* is published here: <http://www.ibm.com/support/docview.wss?uid=swg214276>

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 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® Storage 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® Storage 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® Storage Protect Snapshot for DB2®.

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

The hardware and software requirements for IBM® Storage 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® Storage Protect Snapshot.

The storage space that is required for IBM® Storage 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 with IBM® Storage Protect Snapshot.
- Space that is required on the storage 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® Storage 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® Storage Protect Snapshot		
Operating system	Installation path	Space required (MB)
AIX®	/usr/tivoli/tsfcm/acs_version_number	1250
Linux™	/opt/tivoli/tsfcm/acs_version_number	500

Space requirement for instances

IBM® Storage Protect Snapshot must also be installed on each 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 an instance-specific directory. The space that is required for each IBM® Storage Protect Snapshot enabled application is equal to the amount of space that is required for the global product installation.

The same amount of space is required for any backup or clone instance.

Extra space is required for IBM® Storage Protect Snapshot log files. Log files are written continuously by IBM® Storage 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 require the most space. The space that is required depends on the following factors:

- The total size of all volumes in the storage system that are part of the volume groups that 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 remote mirroring with any of the following storage systems, each backup copy uses space on the remote site storage and on the local site until it is deleted.

XIV®

SAN Volume Controller

IBM® Storwize® family

Use the **MAX_VERSIONS** parameter in the IBM® Storage Protect Snapshot profile to limit the number of snapshots that are stored on a storage 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.

Volume group snapshots are supported on IBM® SAN Volume Controller and IBM® Storwize® family.

Required communication ports

IBM® Storage 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® Storage Protect Snapshot for UNIX™ and Linux™.

To determine the default port numbers that are used for IBM® Storage Protect Snapshot for UNIX™ and Linux™ see the following table:

Communication port numbers

Table 2: IBM® Storage Protect Snapshot for UNIX™ and Linux™ default 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]	Production server and backup/cloning server	SAN Volume Controller CIM agent, with static target allocation Storwize® family cluster CIM agent, with static target allocation
22	Production server and backup/cloning server	SSH port on SAN Volume Controller or Storwize® family cluster, with dynamic target allocation
6989 (HTTPS port) ^[1] 6988 (HTTP port) ^[1]	Production server and backup/cloning server	DS8000® DS8000® CIM Agent
7778	Production server and backup/cloning server	XIV® XIV® CLI
[1] The protocol is specified in the COPYSERVICES_COMMPROTOCOL parameter of the IBM® Storage Protect Snapshot profile.		

Storage solutions

Before you install and configure IBM® Storage Protect Snapshot software, review the storage solution setup. The storage device and its storage volumes must be accessible from all backup and clone servers in the environment.

IBM® XIV® Storage System

When IBM® Storage 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. By default, this snapshot is a space-efficient read-only copy of the application.

If you set the **USE_WRITABLE_SNAPSHOTS** parameter to NO, the snapshots are not mounted directly on a backup host. Instead, IBM® Storage Protect Snapshot creates duplicates from the snapshots as part of the mount procedure, and these duplicates are removed when the backup is unmounted. The duplicate is a space-efficient logical copy of the snapshot, and this copy is writable. The duplicate is effectively another image, so 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 operations of that backup. A subsequent mount operation presents the image as created when the snapshot occurred.

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. For more information, see [“LVM mirroring environments” on page 57](#). A typical IBM® Storage Protect Snapshot profile section for IBM® XIV® Storage System 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
<<<

```

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

Dependent software packages

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

Support for LVM mirroring (AIX® only)

If AIX® Logical Volume Manager (LVM) mirroring is used in the environment, IBM® Storage Protect Snapshot can create separate snapshots of either mirror.

IBM® Storage 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® Storage Protect Snapshot support virtual I/O with n-port ID virtualization. On the production server, IBM® Storage 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® Storage Protect Snapshot supports virtual I/O with NPIV only.

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

You can use IBM® XIV® 11.6 Real-time Compression™ with IBM® Storage Protect Snapshot. The usage of IBM® Storage Protect Snapshot with compressed volumes is not changed. However, when you transform volumes that are managed by IBM® Storage 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® Storage Protect Snapshot operations (for example, back up, restore, and mount) fail. The XIV® adapter returns the **FMM18137E** message. Run the volume transformation at a time that does not overlap with scheduled backups or other IBM® Storage Protect Snapshot actions that run 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) by 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® Storage Protect Snapshot 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 that completes 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 the restore to the volume that is seen by the host appears to be successful, but the source volume visible to the host system is unchanged.

When you use IBM® Storage Protect Snapshot to protect volumes to be transformed, delete the existing snapshot backups, regardless of the `delete_source` option setting.

Related information

[Remote mirror integration](#)

SAN Volume Controller and Storwize® family storage systems

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

SAN Volume Controller storage adapter device types

IBM® Storage Protect Snapshot for UNIX™ and Linux™ offers two backup solutions with Storwize® family and SAN Volume Controller storage systems.

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

SVCDTA

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

SVC

Storwize® family 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.

Restriction: Both **SVC** and **SVCDTA** values are considered to be different hardware types, so limitations apply when they are used on the same storage system.

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.

Dynamic target volumes and predefined target volumes features

Table 3: Dynamic target volumes and predefined target volumes feature comparison.		
Feature	Dynamic target volumes	Static target volumes
Command line interface	Storwize® family or SAN Volume Controller command-line interface (CLI)	Common Information Model (CIM) interface

Feature	Dynamic target volumes	Static target volumes
Number of snapshot images retained	Specify a value for each device class MAX_VERSIONS parameter. Click here for information about the CLIENT section of the profile and values for 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)

If AIX® Logical Volume Manager (LVM) mirroring is used in the environment, IBM® Storage Protect Snapshot can create separate FlashCopy® images of either mirror. Each mirror must be located in a different storage system.

Support for virtual I/O (AIX® only)

SAN Volume Controller, and Storwize® family 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® Storage 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® Storage 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® Storage 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® Storage Protect Snapshot profile.

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® Storage 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

[Running the setup script for IBM Storage Protect Snapshot for DB2](#)
[IBM Storage Protect Snapshot - All Requirements Document](#)

Incremental backups and MAX_VERSIONS for SVCDDTA

When you configure SAN Volume Controller Dynamic target allocation, you can choose to run incremental backups. When the maximum number of backups as defined by **MAX_VERSIONS** is reached for a device class

with **FLASHCOPY_TYPEINCR**, the oldest backup is deleted just before the new backup is taken. This new backup refreshes the INCRFlashCopy relation of the previous deleted backup.

For more information about device class settings, see [Device class backup version retention](#).

Related information

[Configuring Storwize family and SAN Volume Controller dynamic target allocation \(SVCDTA\)](#)

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® Storage Protect Snapshot uses the Storwize® family 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: Using the **MAX_VERSIONS** parameter in the **CLIENT** section of the profile, specify the number of backups to retain. Use a specific number of backups to retain or use the **ADAPTIVE** option for each **DEVICE_CLASS** for **DEVICE_CLASS<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® Storage Protect Snapshot profile when the **COPYSERVICES_REMOTE** is **YES**.

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

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

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

Backup operations and cloning operations

An IBM® Storage 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® Storage Protect Snapshot profile. When the snapshot 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 snapshot backup in a multiple target set configuration. This situation can also happen when a new snapshot cloning operation is started with the force option (-F).

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

Backup operations

An IBM® Storage 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® Storage 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.

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

Restore operation

An IBM® Storage 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® family 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.

Static target allocation

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

When you restore snapshot 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 snapshot target from becoming fully allocated. When you use either SAN Volume Controller or Storwize® family, and IBM® Storage 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® family support the creation of automatically expanding target volumes. If you create target volumes that automatically expand, more storage is assigned to the target when the level of unused real volume capacity decreases. This additional storage ensures that sufficient capacity is available.

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

FlashCopy® relationships

During a restore, IBM® Storage 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® Storage 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® family dynamic target allocation. With SVC DTA, no backups are deleted during a restore operation.

Target sets

IBM® Storage Protect Snapshot cannot reuse a target set for a new snapshot backup unless it corresponds to the last snapshot mapping in a cascaded snapshot relationship. This scenario implies that when IBM®

Storage 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® Storage 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® family dynamic target allocation.

Recommendations for setting up the environment with static target volumes

When you set up the SAN Volume Controller and Storwize® family environments for use with IBM® Storage 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® Storage 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® Storage 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® Storage Protect backup at 00:15.

If the value for the profile parameter **MAX_VERSIONS** is set to **ADAPTIVE**, all disk-only backups taken before that point in time are deleted. Otherwise, the version policy causes the dual backup to fail if **MAX_VERSIONS** 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® Storage Protect Snapshot profile.

You can mount images on a remote server and back up the images to an IBM® Storage Protect server 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® Storage Protect Snapshot requires that FlashCopy® backup target volumes be created in advance on DS8000®. To provide a target set definition to IBM® Storage Protect Snapshot, organize target volumes into target sets, where each target set represents one backup generation.

IBM® Storage 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)

If AIX® Logical Volume Manager (LVM) mirroring is used in the environment, IBM® Storage Protect Snapshot can create separate FlashCopy® images of either mirror.

DS8000® allows one incremental FlashCopy® per source volume. When production volumes are mirrored by using Logical Volume Manager (LVM) mirroring, 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® Storage 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® Storage 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® Storage 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® Storage 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® Storage 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® Storage 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 the IBM® Storage Protect Snapshot repository.

The reconciliation process removes IBM® Storage Protect Snapshot backups when the following events occur 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.

The reconciliation process removes IBM® Storage Protect Snapshot backups when the following events occur on the IBM® XIV® Storage System

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

The reconciliation process removes IBM® Storage Protect Snapshot backups when the following events occur on IBM® System Storage® SAN Volume Controller and IBM® Storwize® family storage systems with static target allocation

- When a FlashCopy® backup becomes invalid because it was created after the creation of the original backup that was later restored. This case 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.
- When FlashCopy® mappings of target volumes are used by the storage system for FlashCopy® backups. When they are used in a specific FlashCopy® backup, then previous FlashCopy® backups can become invalid if they depend on the same mapping. This case applies to backups with space efficient target volumes or if the background copy process is not finished.

The reconciliation process removes IBM® Storage Protect Snapshot backups when the following event occurs on IBM® System Storage® DS8000®

When a source target relationship is withdrawn backups are removed. This process does not happen automatically.

Remote mirror integration

When you use storage solutions with mirror technologies in combination with IBM® Storage Protect Snapshot, certain criteria must be met by the environment to integrate backup, restore, and cloning operations. For IBM® System Storage® SAN Volume Controller and IBM® System Storage® DS8000® series, 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® Storage 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® Storage Protect Snapshot backs 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 snapshot backup. When you use IBM® Storage Protect Snapshot in this environment, do not use snapshot targets as Global Mirror and Metro Mirror sources.

IBM® XIV® Storage System

IBM® Storage Protect Snapshot can back up application data consistently on XIV® 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® Storage 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® Storage Protect Snapshot. The following list identifies what you must configure when you are setting up the cluster partnership:

- IBM® Storage Protect Snapshot is installed on the production and backup host on the local site (primary cluster).
- IBM® Storage 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® Synchronous Remote Mirroring and Asynchronous Remote Mirroring, configure either the source or the targets as a source for the snapshot backup.
- IBM® Storage Protect Snapshot uses a consistency group on the SAN Volume Controller and XIV® storage solutions for the FlashCopy or snapshot. A consistency group is a group of volumes that are associated with a snapshot pair, which is a snapshot group of two corresponding instant copies of data, that is, point-in-time copies of a volume. For the snapshot pair, the logically related data must be kept consistent across the volumes. The snapshot 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® Storage Protect Snapshot:

SAN Volume Controller

- A consistency group contains a list of snapshot or Remote Copy relationships.
- The IBM® Storage Protect Snapshot software creates a snapshot consistency group on the secondary site to build a consistency unit between the source and target of the snapshot.
- 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®

- 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 you start 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® synchronous mirroring, the state of the consistency group must be consistently synchronized.
- For XIV® 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.

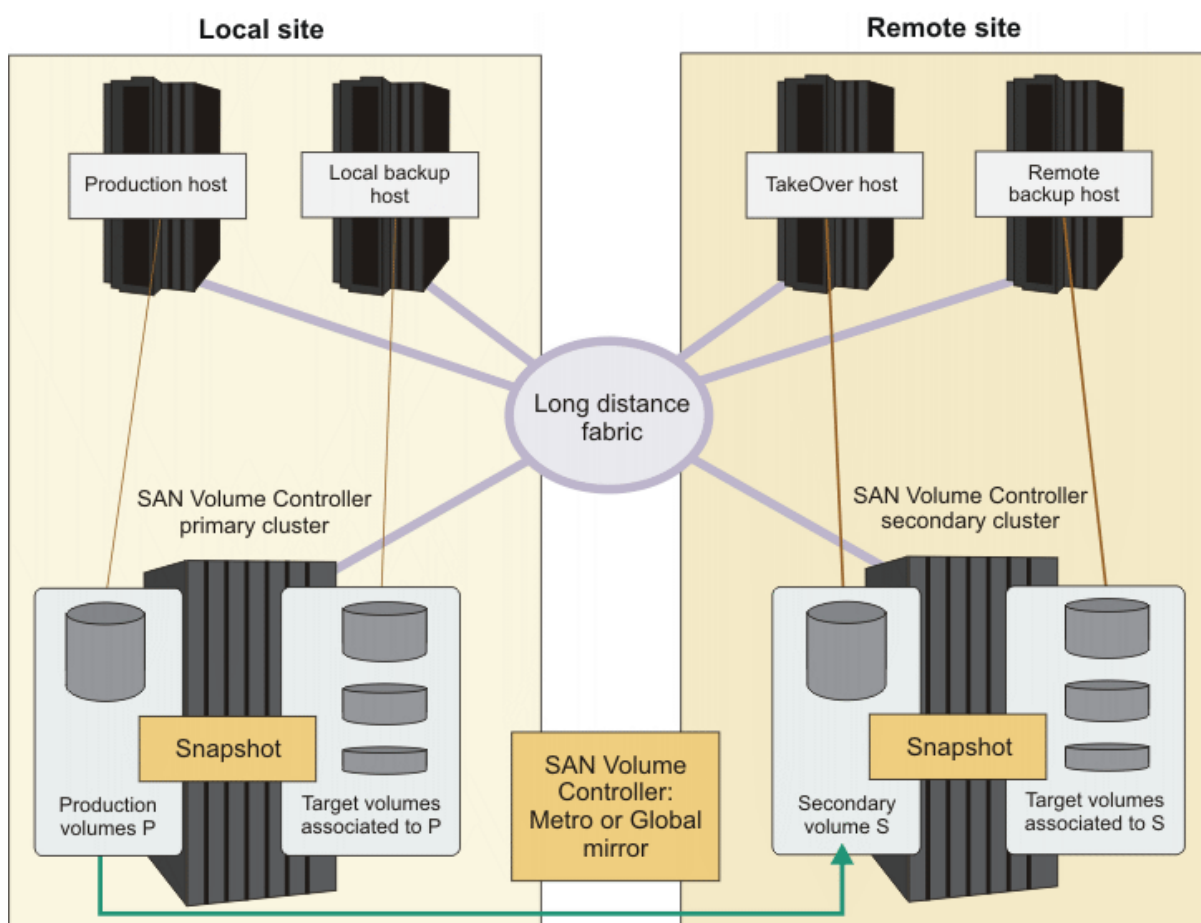


Figure 4: Remote mirroring using Metro Mirror and Global Mirror sources

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® Storage 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® Storage 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® Storage Protect Snapshot FlashCopy® backup time. IBM® Storage 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® Storage 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.

HyperSwap integration

HyperSwap systems support a dual-site, active-active solution that provides continuous availability of data during planned and unplanned outages.

A fully independent copy of the data is maintained at each site. When data is written by hosts at either site, both copies are synchronously updated before the write operation is completed. If storage at either site goes offline, HyperSwap will automatically failover storage access to the system at the surviving site. The HyperSwap function automatically re-synchronizes the two copies of the data as soon as possible after the outage.

The HyperSwap solution requires one IBM FlashSystem 9200 control enclosure at each site, or an equivalent system that can support more than one I/O group such as Storwize, V7000 or SAN Volume Controller. IBM FlashSystem 9200 systems are virtualized, Non-Volatile Memory Express (NVMe), all-flash software-defined storage solution. IBM FlashSystem 9200 runs IBM Storage Virtualize, which enables rapid deployment of cloud storage for disaster recovery.

A third FlashSystem 9200 site or an SAN Volume Controller system such as V7000 acts as a tie-breaking quorum device that provides an automatic tie-break in case of a link failure between the two main sites. The tie-breaker site can be implemented as a fiber channel storage area network (FC SAN) or an IP-based quorum application.

The HyperSwap function works with the standard multipathing drivers that are available on a wide variety of host types, with no additional host support required to access the highly available volume. Where multipathing drivers support Asymmetric Logical Unit Access (ALUA), the storage system tells the multipathing driver which nodes are closest to it, and can be used to minimize I/O latency. The HyperSwap topology uses additional system resources to support a full independent cache on each site, allowing full performance even if one site is lost. You simply indicate to the storage system which site a host is connected to, and the storage system configures the optimal host path.

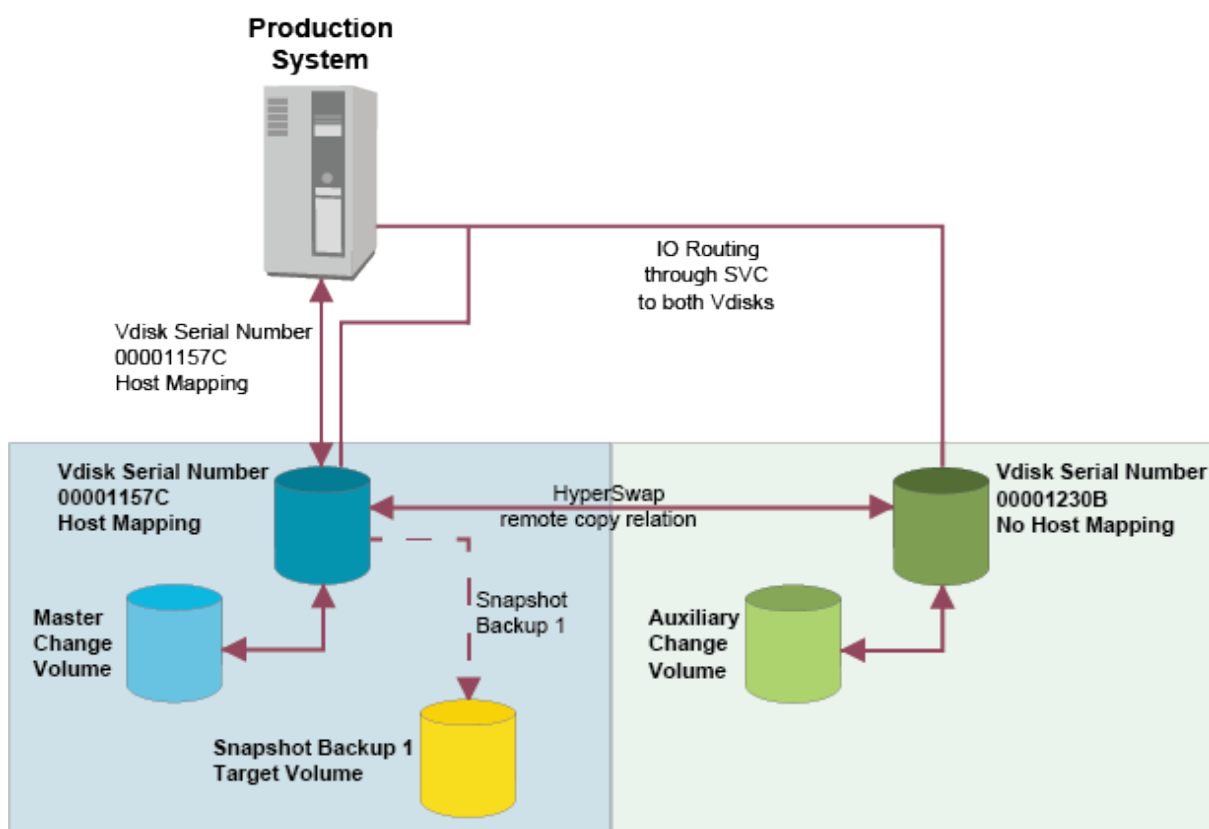


Figure 5: IBM® Storage Protect Snapshot in a HyperSwap environment

A HyperSwap volume is a single logical entity. From an IBM® Storage Protect Snapshot perspective, an SVC-based HyperSwap cluster looks like a normal cluster. Similarly, from the operating system perspective the HyperSwap LUNs that are mapped to a host look like normal SVC vDisks. To the operating system, the LUNs present as a single LUN serial number or UID.

In practice, a HyperSwap volume is realized by using four physical volumes, a set of FlashCopy maps and a remote copy active-active relationship. The master and auxiliary volumes are 2 vDisks residing in storage pools at different sites. The vDisks have different LUN serial numbers or UIDs. Internally the vDisks have different vDisk names. To the SVC GUI the vDisks present themselves with the same vDisk name with the site name added in parentheses. When using the SVC CLI, however, you can see the real vDisk name of the disk on the auxiliary site.


A typical HyperSwap-based production server performs IO routing to the two vDisks, using SVC capabilities. Host mapping is set up between a production system and a vDisk using a LUN serial number. The primary vDisk in turn synchronizes changes with a Master Change Volume.

On the other side of the HyperSwap remote copy relationship, the LUN serial number of the vDisk has no host mapping. This vDisk synchronizes changes with an Auxiliary Change Volume.

SVC handles all internal HyperSwap operations transparently for the operating system and the application, that is using the data on the HyperSwap volumes. If for example one site in SVC has IO problems, then SVC automatically passes all read and write activity over to the vDisks on the second site.

From the operating system perspective nothing has changed. The mapped LUN serial number is the same, even if all IO activity is now handled by the vDisk on site2 (which internally has a different LUN serial number). The only place where such a failover can be seen is the SVC GUI or the SVC CLI when you inspect the Primary Volume column. The column shows the value **Auxiliary** instead of **Master** after a failover.

Note: Because all HyperSwap volumes belonging to an application are in a remote copy consistency group, all volumes failover at the same time, even if the a problem is with a single disk. In the SVC GUI

you can also see a small symbol  next to the state **Consistent Synchronized**, indicating that the

Master and Auxiliary are switched. When you hover over the graphic, a pop-up displays the text **Master and Auxiliary are Switched**. This will be relevant when we look at restore scenarios.

Preparing applications that run on VMware or KVM

Before you install IBM® Storage 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® Storage Protect Snapshot configuration requirements. For more information about application-specific requirements, see [“Planning” on page 16](#).

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® Storage 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® Storage 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® Storage 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.

- 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, `360050768018205de4000000000001949` 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`.

Installing and setting up IBM® Storage Protect Snapshot

To install IBM® Storage Protect Snapshot you must follow the installation steps, run the setup script for your component, activate the installation for every application you want to protect, and configure the product. The first step is to install IBM® Storage Protect Snapshot on the production server. If you choose to, you can upgrade your system from a previous version of IBM® Storage Protect Snapshot to Version 8.1.4.

About this task

Depending on your environment, a separate installation of IBM® Storage Protect Snapshot can be required on a backup or clone server. The following set of tasks are required to complete the installation process.

Procedure

1. Install IBM® Storage Protect Snapshot on the production server. During the installation, the product is installed to a global installation directory.
The production server is the server where the application to be protected by IBM® Storage Protect Snapshot is located.
2. Activate the installation for every database instance you want to protect with IBM® Storage Protect Snapshot.
During the activation, all the necessary files are copied from the installation directory, to the database instance installation directory.
3. Configure IBM® Storage Protect Snapshot by using a dedicated configuration script for the application that you want to protect, and optionally install and configure the product on a backup or clone server. If Open Secure Shell (OpenSSH) is configured between the production and the backup or clone servers, IBM® Storage Protect Snapshot can be activated and configured on the backup or clone server by using the dedicated configuration script from the production server. Otherwise, a separate installation is required.
4. If backup and clone servers are not automatically activated and configured by the setup script, set up IBM® Storage Protect Snapshot on the required servers.
Backup servers or clone servers are auxiliary hosts that are required by IBM® Storage Protect Snapshot to mount backup images and clone databases. A backup server also is required to offload backups to an IBM® Storage Protect server.

Result

The following files and directories are created on the production server, and optionally on the backup or clone servers, during the configuration process:

- When the ACS_DIR directory is not identical to the instance directory, an ACS_DIR configuration directory is created. The path for the ACS_DIR directory is specified in the IBM® Storage Protect Snapshot profile.
- A profile in the ACS_DIR configuration directory.
- A symbolic link from the <instance directory>/profile that points to the ACS_DIR/profile when the two directories are not identical is created.
- A password file in the ACS_DIR/shared directory.

IBM® Storage Protect Snapshot daemon processes are stopped and restarted if requested.

Preparing for installing

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

The hardware and software requirements for IBM® Storage 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 prepare your environment for IBM® Storage Protect Snapshot for UNIX™ and Linux™, run the Prerequisite Checker tool.

Related information

[Prerequisite checker for DB2](#)

Prerequisite checker for DB2®

Check your system by running the Prerequisite checker tool before you install IBM® Storage Protect Snapshot for DB2®.

You must complete the *Pre-installation Checklist* checklist before you install IBM® Storage 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 preparation activities, install and run the Prerequisite Checker tool before you install or upgrade to a new version of IBM® Storage 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® Storage Protect Snapshot.

Before you begin

In a distributed DB2® data partitioned environment, the Prerequisite Checker can be installed and run on any system that is hosting an active DB2® partition. Verify that your environment meets the hardware and software requirements of the IBM® Storage Protect Snapshot Prerequisite Checker.

Procedure

1. Download the IBM® Storage Protect Prerequisite Checker installation file for your operating system from the download website. 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:
AIX®: `<VERSION>-FCM-PREREQ-AIX.bin [-i console | -i swing]`
Linux™: `<VERSION>-FCM-PREREQ-Linux.bin [-i console | -i 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. This method is the default.
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® Storage Protect Snapshot, with the root user ID. Check the following requirements:

- The database is activated. Verify that the database is activated.
- The default environment of the instance owner must contain all the environment settings necessary for interaction with the database.

- The instance 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:

```
./fcmprereqchecker.sh -u <instance_owner> -s <storage_management_IP_address>
-p <storage_management_port> [-o <output_path>] [-d <database_name>]
```

Where,

`instance_user` is the name of the instance 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.

For DB2®, the `database_name` must be specified when multiple databases are in the instance that is owned by `dbusername`.

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 failures, you must modify your system before you proceed to install IBM® Storage Protect Snapshot for UNIX™ and Linux™.

About this task

The Summary reports the overall result of the checks run; the status is either `Failed`, `Warning`, or `Passed`. The machine name, Operating System, and serial number are listed with the results for each check.

- Find the `result.html` file and open it in your browser.
The `result.html` file is stored in the `<output_path>` of the Prerequisite Checker. The default output path is `<checker_path>/logs`. If you specified a different output path with the `-o` option, the `result.html` and log files are stored there. There is a text file version of the results stored there called, `result.txt`. For information about fails and warnings, including message information, review the log file in the same directory.
- 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 your environment before you install the product.
- 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® Storage Protect Snapshot product.

Procedure

1. Log on with the root user ID.
2. Enter the following command:

```
<checker_path>/uninstall/uninstall.bin [-i console | -i 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.

Result

The Prerequisite Checker executable files are removed from your system.

Verify the DB2® system

Ensure that the following structure is in place before you install and use IBM® Storage Protect Snapshot.

The DB2® configuration must have the following structure:

- At least one volume group for each database partition that contains table spaces and the local database directory.
- At least one volume group for each database partition that contains database log files.
- Data must be on file systems that are supported.

IBM® Storage Protect Snapshot supports DB2® Database Partitioning Feature (DPF) environments, where data of multiple logical partitions are within a single volume group. This setting is referred to as volume sharing, see the **PARTITION_GROUP** parameter for more details. With volume sharing, the partitions sharing volume groups cannot be restored individually. In this case, the granularity of the restore operation is defined by **PARTITION_GROUP**.

Any other data that is stored on the volume groups is processed and included in the backup image. When you restore a backup image, the whole backup image is restored. Therefore, any data on the file system that is updated or created after the backup is taken is overwritten. To ensure that nothing is overwritten unintentionally, do not store any other data in the data and log volume groups.

Extra files that are not associated with the database backup operation, but are stored in the volume groups, can cause the backup to fail. Use the **NEGATIVE_LIST** parameter in the profile to control the processing of these files. For information about this parameter, see the “**CLIENT**” on page 101 profile section for details.

IBM® Storage Protect Snapshot processes table spaces, the local database directory, and log files. The volume group layout for DB2® is detailed in the *Pre-installation Checklist*.

The hardware and software requirements for IBM® Storage 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® Storage Protect Snapshot does not support a volume and storage layout where a single database partition is spread across multiple storage devices. In an AIX® logical volume manager mirroring environment, each mirror must be located within a separate storage cluster.

To offload backups from this environment to the server, IBM® Storage Protect Snapshot requires DB2® and the DB2® instance to be installed on a backup server. This server must be configured to ensure that the topology of the production and backup server are the same. The number of DB2® partitions and DB2® partition numbers must be identical on the production and backup server. However, the DB2® partitions on the backup server can be consolidated into one backup server even if the production database is spread across multiple hosts.

Preparing the DB2® high-availability disaster recovery server or DB2® standby server

Before you install IBM® Storage Protect Snapshot, verify the configuration of the DB2® high-availability disaster recovery (HADR) or DB2® standby server environment. Preparing a DB2® (HADR) server is not applicable to a partitioned DB2® database.

A DB2® (HADR) server in primary mode is protected by IBM® Storage Protect Snapshot in the same way as for standard non-HADR DB2® servers. DB2® (HADR) servers in standby mode can be protected with IBM® Storage Protect Snapshot but use different commands independent from the standard DB2® commands.

Protecting a DB2® (HADR) server in standby mode with IBM® Storage Protect Snapshot assumes that the corresponding server in primary mode is also protected by IBM® Storage Protect Snapshot. After a DB2® (HADR) takeover operation, the former DB2® (HADR) standby server takes the role as the primary server. In this case, IBM® Storage Protect Snapshot for DB2® continues protecting the former standby server as the primary.

The storage layout for a DB2® (HADR) primary and standby server environment must be identical to the storage layout described in “Verify the DB2 system” on page 35. In addition, the names of files, file systems, and mount points must be identical for both the DB2® primary and the DB2® standby environments. When you use DB2® automatic storage, then the table space container IDs and names can be different between the primary and standby systems. However, the automatic storage paths must be identical.

Remember: In a typical environment, different servers in a DB2® HADR environment can act as the primary or the standby system. To support such configurations, you can create a single IBM® Storage Protect Snapshot profile at each site. This profile can be used regardless of whether the system is the active or inactive copy of DB2®. For more information about DB2 database protection in the HADR environment, see “Figure: Protection of DB2® databases in a HADR environment with IBM® Storage Protect Snapshot” on page 36.

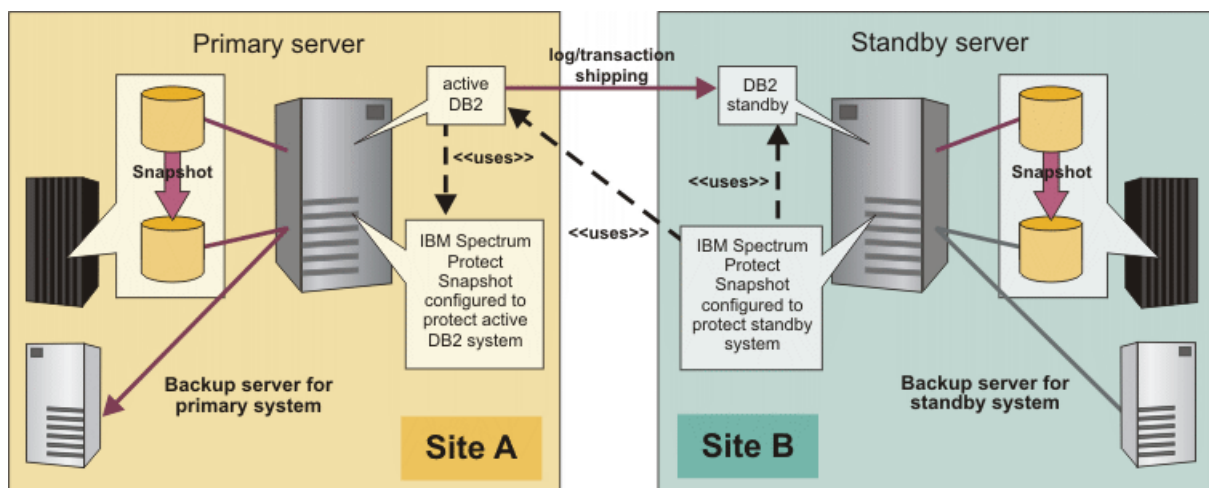


Figure 6: Protection of DB2® databases in a HADR environment with IBM® Storage Protect Snapshot

Although a single IBM® Storage Protect Snapshot profile can be used on both sites, you must use different commands for backing up DB2®. For the primary server, the command is `db2 backup use snapshot`. For the standby server, use the command `fcmlcli -f backup_db2standby`. In an environment where IBM® Storage Protect is configured, this configuration also affects the way IBM® Storage Protect Snapshot triggers the offload backup.

- While DB2® is active, IBM® Storage Protect Snapshot issues a `db2 backup` command for the backup to IBM® Storage Protect.
- If DB2® is in standby mode, IBM® Storage Protect Snapshot creates the backup by using the IBM® Storage Protect backup-archive client.

Restriction: IBM® Storage Protect Snapshot communicates with the active site to query DB2® for a list of files to back up. Therefore, the names of mount points and file systems that contain DB2® data must be identical in site A and site B.

Remember: Although you can create a single profile on either site that can be used regardless of whether the site is active or passive, you must install IBM® Storage Protect Snapshot separately on each site. There is no installation option that installs IBM® Storage Protect Snapshot on either site simultaneously.

If the primary and standby servers are in the same site, you must copy the self-signed certificates to the standby server installation directory. For more information about how to do this, see [“Setting up a backup server for multiple production systems in a DB2 HADR environment” on page 44](#)

Preparing backup and cloning servers

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

A backup server or clone server is used to offload the backup image to an IBM® Storage Protect server. The sending of data to the IBM® Storage Protect server happens from the backup server and not 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® Storage 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® Storage 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.

When a backup server is needed

Backup servers are auxiliary hosts where IBM® Storage Protect Snapshot can mount backups. The number of required IBM® Storage Protect Snapshot backup servers is determined by the number of servers that are used to access backup images.

For example, to access backup images on either site of a disaster recovery environment, at least two backup servers are needed. A backup server can also simultaneously be used for multiple applications and multiple production servers. IBM® Storage Protect Snapshot can mount a backup image on a backup server. For the following scenarios, at least one backup server is required.

- Mount backup images on another server.
- When IBM® Storage Protect Snapshot is used with other products for example, IBM® Storage Protect for Enterprise Resource Planning to offload backups to an IBM® Storage Protect server.
- When IBM® Storage Protect Snapshot requires a mount during a backup operation because one of the following conditions exists:
 - The database is running in an LVM mirrored environment on AIX.®
 - Conditions that require a so called IBM® Storage Protect Snapshot forced mount operation for the different storage subsystem environments:

SAN Volume Controller, Storwize® family, and DS8000®

A forced mount is required if the option `PREASSIGNED_VOLUMES` is set for the profile parameter `BACKUP_HOST_NAME` and the operating system is Linux™.

DS8000®

A forced mount is required on AIX if the option `PREASSIGNED_VOLUMES` is set for the profile parameter `BACKUP_HOST_NAME` and a freeze and thaw action was not executed for the file systems.

For a physically partitioned DB2® database, multiple backup servers can be used. The physical partitioning of the production server is not required to be identical to the physical partitioning of the backup server. However, the number of DB2® partitions and the DB2® partition numbers must be identical on the production and backup server.

Ensure that the amount of DB2® partitions and DB2® partition numbers are identical on the production database and on the clone server. All partitions of the clone database must be on a single clone server.

Installation prerequisites for backup and clone servers

For hosts that are used as backup or clone servers, 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 user ID (UID) for the clone instance must be different from the owner of the production database when the name of the clone database instance is different from the name of the production instance.
- A database instance with the same version as the database instance on the production server must be installed on the backup server.

When IBM® Storage Protect Snapshot is used in an environment with IBM® Storage 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 an IBM® Storage Protect server.

The following IBM® Storage Protect Data Protection clients are used by IBM® Storage Protect Snapshot to start a backup to the IBM® Storage Protect server, and must be installed and configured on the backup server.

- IBM® DB2® Storage Protect client for DB2® environments.
- IBM® Storage Protect for Enterprise Resource Planning agent is optional for DB2® in an SAP environment.
- IBM® Storage Protect backup-archive client, which is used to offload backups from the DB2® standby server.

The IBM® Storage Protect backup-archive client is used by these clients and must be installed and configured on both the production and the backup servers.

Update the IBM® Storage Protect backup-archive client node password on the production server and all backup servers whenever it changes. When IBM® Storage Protect is configured to use the **PASSWORDACCESS GENERATE** parameter, the password can change without notification.

- If the IBM® Storage Protect backup-archive client is configured to use the **PASSWORDACCESS GENERATE** parameter, use the IBM® Storage Protect proxy-node capability to avoid authentication errors when the password is reset.
- Create one data node on the IBM® Storage Protect server where all IBM® Storage Protect 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.

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

IBM® Storage Protect Snapshot requires the following settings for cloning:

DB2® and DB2® in an SAP environment

Two database instances must be installed on the clone server. These database instances must be the same version as the database instance on the production server. One of the two database instances is assigned the same name as the production instance (on the production server). The other database instance is assigned the clone instance name. If the clone database name is the same name as the production database name, only one database instance with this common name is required on the clone server.

On the clone server, the database instance with the same name as production database instance must be assigned a user name and group name. The user name and group name must be the same names that are assigned to the database instance on the production server. The same user ID (UID) and group ID (GID) must be used. If the name of the clone database is different from the name of the production database, the user name and UID for the clone instance must be different from the owner name of the production database.

Backup server prerequisites

When IBM® Storage Protect Snapshot is used in an environment with IBM® Storage 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 an IBM® Storage Protect server. Ensure that the required settings are in place before you install IBM® Storage Protect Snapshot on the backup server.

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

- The user name and user ID (UID) for the clone instance must be different from the owner of the production database when the name of the clone database instance is different from the name of the production instance.
- A database instance with the same version as the database instance on the production server must be installed on the backup server.

The following IBM® Storage Protect Data Protection clients are used by IBM® Storage Protect Snapshot to start a backup to the IBM® Storage Protect server, and must be installed and configured on the backup server.

- IBM® DB2® Storage Protect client for DB2® environments.
- IBM® Storage Protect for Enterprise Resource Planning agent is optional for DB2® in an SAP environment.
- IBM® Storage Protect backup-archive client, which is used to offload backups from the DB2® standby server.

The IBM® Storage Protect backup-archive client is used by these clients and must be installed and configured on both the production and the backup servers.

Update the IBM® Storage Protect backup-archive client node password on the production server and all backup servers whenever it changes. When IBM® Storage Protect is configured to use the **PASSWORDACCESS GENERATE** parameter, the password can change without notification.

- If the IBM® Storage Protect backup-archive client is configured to use the **PASSWORDACCESS GENERATE** parameter, use the IBM® Storage Protect proxy-node capability to avoid authentication errors when the password is reset.
- Create one data node on the IBM® Storage Protect server where all IBM® Storage Protect 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.

Clone server prerequisites

For setting up a cloning server, the database instances must be installed and configured on the clone server. Ensure that all the prerequisites are in place before you set up the clone server.

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

IBM® Storage Protect Snapshot requires the following settings for cloning:

DB2® and DB2® in an SAP environment

Two database instances must be installed on the clone server. These database instances must be the same version as the database instance on the production server. One of the two database instances is assigned the same name as the production instance (on the production server). The other database instance is assigned the clone instance name. If the clone database name is the same name as the production database name, only one database instance with this common name is required on the clone server.

On the clone server, the database instance with the same name as production database instance must be assigned a user name and group name. The user name and group name must be the same names that are assigned to the database instance on the production server. The same user ID (UID) and group ID (GID) must be used. If the name of the clone database is different from the name of the production database, the user name and UID for the clone instance must be different from the owner name of the production database.

Preparing backup and cloning servers for VMware or KVM

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

Before you begin

Review [“Installation prerequisites for backup and clone servers” on page 37](#) to ensure that all requirements for backup and clone servers are met. These requirements are also required for backup servers on virtual machines.

- 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)

Installing and uninstalling IBM® Storage Protect Snapshot for DB2

Install or uninstall IBM® Storage Protect Snapshot using the graphical installation wizard, or the console wizard in interactive or silent mode

Installing IBM® Storage Protect Snapshot in interactive mode

Install IBM® Storage Protect Snapshot on the production server by using the graphical installation wizard, or the console wizard in interactive or silent mode.

Before you begin

For the most up-to-date requirements, review the *Hardware and Software Requirements* technote that is associated with the IBM® Storage Protect Snapshot release. This technote is available in the *IBM Storage 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 and version, and review the *Pre-Installation Checklist* and *Planning Worksheet*.

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

The files are named like

```
<version>-TIV-TSFCM-<platform>.bin
```

Before you install IBM® Storage Protect Snapshot on AIX® or Linux™, run the Prerequisite Checker to ensure that the prerequisites are met before you proceed with the installation process.

Procedure

To install IBM® Storage Protect Snapshot on the production server, complete these steps.

1. Log on to the production server and use the root user ID. Change to the directory where you downloaded the package file. 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:

```
./<version>-TIV-TSFCM-<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:

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

2. Follow the prompts to install IBM® Storage Protect Snapshot.
3. In the **Pre-Installation Summary**, review your installation settings.
The installation directory for AIX® is `/usr/tivoli/tsfcm/acs_<version>`
The installation directory for Linux™ is `/opt/tivoli/tsfcm/acs_<version>`

If an error occurs during the installation process, correct the errors and restart the installation procedure. Find the `installation.log` file in the installation directory to troubleshoot installation errors.

What to do next

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

Related information

[Prerequisite checker for DB2](#)

[Activating an instance](#)

[Configuring or reconfiguring IBM Storage Protect Snapshot](#)

[Setting up IBM Storage Protect Snapshot on a backup server](#)

Installing in silent mode

To install IBM® Storage Protect Snapshot in silent mode, you must create a properties file.

About this task

You can generate a properties file when you are installing the product in interactive mode. You can use this properties file to install similar setups in silent mode.

Procedure

1. Install IBM® Storage Protect Snapshot in interactive mode and generate a properties file with the following command that is run from the installation directory:

```
./<version>-TIV-TSFCM-<platform>.bin [-i console] -DRECORDFILE=<properties_file>
```

For example,

```
./8.1.0.4-TIV-TSFCM-AIX.bin -DRECORDFILE=/tmp/installation.properties
```

2. Invoke the executable file with the `-i silent` option and the `-f` option to specify the properties file:

```
./<version>-TIV-TSFCM-<platform>.bin -i silent -f <properties_file>
```

The *properties_file* specification must contain a full path.

For example,

```
./8.1.0.4-TIV-TSFCM-AIX.bin -i silent -f /tmp/installation.properties
```

3. Review the `installation.log` file in the installation directory to complete the process.

What to do next

[Activation](#)

Uninstalling the software

Complete the uninstallation procedure to uninstall a version of the product from your system.

Procedure

1. Determine the installation path of the version of the product you want to uninstall. The following paths provide the 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>`.
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>/uninstall/uninstaller.bin`.

- For Linux™, use this command `/opt/tivoli/tsfcm/acs_<version>/uninstall/uninstaller.bin`.

Activating an instance

During the activation process, the necessary files are copied from the installation directory to an instance-specific directory. The installer does not activate the instance. In order to activate an instance, follow the procedure.

Procedure

1. Log in to the production server and use the root user ID. Change to the global installation directory.
2. Run the following command:
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.

```
./setup_db2.sh -a install -d <DB2_instance_owner_$HOME_directory>/sqllib/
```

What to do next

After activating the instance, you must configure the instance to complete the setup.

Related information

[Configuring or reconfiguring IBM Storage Protect Snapshot](#)

[Running the setup script for IBM Storage Protect Snapshot for DB2](#)

[Preparing backup and cloning servers](#)

Setting up IBM® Storage Protect Snapshot separately on backup or clone servers

When you are configuring the production server, it is recommended that you set up a backup or clone server as required. To manage backup or clone servers from the production system with the setup script, Open Secure Shell (OpenSSH) is required. If it is not possible to set up from the production server, you must set up the backup or clone server separately. In this case, all upgrades to IBM® Storage Protect Snapshot must be done separately on all servers.

Before you begin

Ensure that you installed the product, and activated and configured the instance. Take a note of the hostname of the production server and the device classes that the backup or clone server are to use. If the default names were changed, ensure that you know the port of the IBM® Storage Protect management agent (ACSD), and the ACS_DIR directory name. The default ACS_DIR directory is `<instance_owner $HOME>/acs`.

Locate the following items on the production server before you proceed.

- The shared directory in the ACS_DIR directory.
- The `fcmselcert.arm` file in the instance directory. This file is needed when standard CA-signed certificates are not used for server authentication. To find out more about setting up Secure Socket Layer (SSL) and Transport Layer Security (TLS), see [IBM Global Security Kit configuration](#)

Setting up IBM® Storage Protect Snapshot on a backup server

To set up IBM® Storage Protect Snapshot on a backup server separately, complete the following steps on the backup server.

Procedure

Prepare the instance that is to be configured.

1. Install IBM® Storage Protect Snapshot on the backup server.
For more information about how to install the product, see [Installing in interactive mode](#).

2. Activate the instance on the backup server.
For more information about how to activate an instance, see [Activating an instance](#).
3. Log on to the backup system with the instance owner user ID.
4. If standard CA-signed certificates are not used for server authentication, copy the `fcmselcert.arm` file of the production system to the instance directory on the backup system.

```
cd <instance directory>
scp <instance owner>@<production server>:<instance directory>/fcmselcert.arm .
```

5. Check the backup server for the ACS_DIR directory. Use the following command to create the directory:

```
mkdir -p <ACS_DIR>
```

6. Copy the <ACS_DIR>/shared directory on the production system to the ACS_DIR directory on the backup system:

```
cd <ACS_DIR>
scp -r <instance owner>@<production server>:<ACS_DIR>/shared .
```

7. From the instance directory, run the setup script.

```
cd <instance directory>
./setup_db2.sh
```

8. Follow the prompts, ensuring to select **On-Site Backup System configuration**.
9. When you are asked for the **Hostname and port of machine running Management Agent (ACSD)**, enter the hostname of the production server, and the port of the management agent.
If the default port 57328 is used, entering the hostname is sufficient. When asked for the device classes to use for this backup system, enter the backup device classes that the server is to be used for.

Setting up IBM® Storage Protect Snapshot on a clone server

To install IBM® Storage Protect Snapshot on a clone server separately, complete the following steps on the clone server.

Procedure

Prepare the clone instance.

1. Install IBM® Storage Protect Snapshot on the clone server.
For more information about how to install the product, see [Installing in interactive mode](#).
2. Activate the clone instance on the clone server.
For more information about how to activate an instance, see [Activating an instance](#).
3. Log on to the clone system with the clone system owner ID.
4. If standard CA-signed certificates are not used for server authentication, copy the `fcmselcert.arm` file on the production system, to the clone system in the instance directory.

```
cd <clone instance directory>
scp <instance owner>@<production server>:<instance directory>/fcmselcert.arm .
```

5. Check the clone server for the ACS_DIR directory. Create it with this command if it is not there,

```
mkdir -p <ACS_DIR of clone instance>
```

6. Copy the <ACS_DIR>/shared directory on the production system to the ACS_DIR directory of the clone instance:

```
cd <ACS_DIR of clone instance>
scp -r <instance owner>@<production server>:<ACS_DIR>/shared .
```

7. From the clone instance directory, run the setup script.

```
cd <instance directory>
./setup_db2.sh
```

8. Follow the prompts, ensuring to select **On-Site Backup System configuration**.
9. When you are asked for the **Hostname and port of machine running Management Agent {ACSD}**, enter the hostname of the production server, and the port of the management agent.
If the default port 57328 is used, entering the hostname is sufficient. When asked for the device classes to use for this backup system, enter the backup device classes that the server is to be used for.

Upgrading IBM® Storage Protect Snapshot on a backup or clone server

Backup and clone instances must always be kept on the same IBM® Storage Protect Snapshot version and level. When you upgrade the software on the production server, you must upgrade the corresponding backup or clone servers.

Procedure

1. Install IBM® Storage Protect Snapshot.
This procedure is described in [Installing in interactive mode](#).
2. Activate the backup or clone instance.
This procedure is described in [Activating an instance](#).

What to do next

Typically, it is not necessary to run the setup script on a backup or clone server after the initial configuration. Some exceptions to this rule require a reconfiguration of IBM® Storage Protect Snapshot. These exceptions are as follows.

- If any changes were made to the production server or to the port of the management agent, a reconfiguration of IBM® Storage Protect Snapshot on the backup server is required.
- If any changes were made to the device classes.
- If self-signed certificates are used and were changed. All changes to the certificates require a reconfiguration of the backup server. To do this reconfiguration, you must copy the `fcmselfcert.arm` file from the production server before you run the setup script.

In all cases, you must copy the `<ACS_DIR>/shared` directory from the production server to the backup or clone server that is being configured.

Setting up a backup server for multiple production systems in a DB2® HADR environment

To use the same backup system for the primary and secondary nodes in a DB2® HADR environment, you must start the mount agent twice with different IBM® Storage Protect Snapshot profiles, one for each production system. The production systems must use the same certificate.

About this task

Each production server must have a CA-signed certificate or a self-signed certificate file, and its own mount agent. As a result, the backup server needs different directories for each production server. For more information about Protection of DB2® (HADR) standby nodes, see the illustration in [“Preparing the DB2 high-availability disaster recovery server or DB2 standby server”](#) on page 35.

Procedure

1. Configure the backup server for the role of primary server with the IBM® Storage Protect Snapshot for DB2® setup script, `setup_db2.sh`.
2. Go to the `<ACS_DIR>` folder on the backup server, and rename the IBM® Storage Protect Snapshot profile to `profile_<primary_server_hostname>`.
Name the profile with a name of your choice. The `<primary_server_hostname>` is used to match the profile with a corresponding host.

3. In the instance directory, create a new symbolic link with the same name.
This link must point to the renamed profile in the <ACS_DIR>. Name the symbolic link with a name of your choice.
4. Go to the instance directory of the primary server:
<instance_owner_\$HOME>/sql1lib/acs
5. Copy all fcmcert.* and fcmselfcert.arm files, into the instance directory of the standby server.
For example,


```
scp fcmselfcert.arm fcmcert.*
<Instance_owner>@<standby_server_hostname>:<instance_owner_$HOME>/sql1lib/acs
```
6. Switch to the standby server and go to the instance directory. Run the IBM® Storage Protect Snapshot DB2® setup script setup_db2.sh as the instance owner, ensuring to update the backup system with the newly created profile and the self-signed certificates.
7. Switch back to the backup server and go to the <ACS_DIR> folder. Rename the IBM® Storage Protect Snapshot profile.
Use a name similar to profile_<standby_server_hostname>.
8. In the instance directory, create a new symbolic link with the same name to point to the renamed profile in the <ACS_DIR>.
New profile and symbolic link names can be any you choose. The <standby_server_hostname> is used as a convenient way to match the profile with a corresponding server.
9. To finalize the setup, you must manually configure an automatic restart of the two mount agents for the database instance on the system.
Each mount agent must be started with its IBM® Storage Protect Snapshot profile for each production system. The starting method for the daemons depends on the operating system. For more information about configuring to automatically start background daemons, see [Setting up daemons](#)

Configuring or reconfiguring IBM® Storage Protect Snapshot

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

Before you begin

Ensure that you complete the *Installation Planning Worksheet*. See the Planning section, for more information about the *Installation Planning worksheet*.

About this task

Configure IBM® Storage Protect Snapshot by [Running the setup script](#).

When you run the setup script to configure the product, you are asked for instance-specific information as follows.

- The type of instance, whether it is a production, backup, or clone system.
- What functionality is used on the system such as backup or cloning, or for offloading backups to an IBM® Storage Protect server.
- Configuration information for the behavior of IBM® Storage Protect Snapshot.
- Settings for integration with other products such as an IBM® Storage Protect server or IBM® Storage Protect clients.
- Connection information and definitions for corresponding backup and clone instances if they are required.
- Connection information and settings for storage hardware. For more information about storage, see [Configuring storage environments](#).

The *Installation Planning Worksheet* is available here <http://www-01.ibm.com/support/docview.wss?uid=swg21427692>. It contains a list of parameters that are requested when you run the setup script in basic mode. For more information about profile parameters, see [Profile](#). For more information about configuring the

product for specific environments, see [Backing up data](#). For more information about the IBM® Storage Protect Snapshot daemons, see [Setting up daemons](#).

Related information

[Planning](#)

Running the setup script for IBM® Storage Protect Snapshot for DB2®

Run the setup script to configure IBM® Storage Protect Snapshot for DB2.

Before you begin

In most cases, it is sufficient to configure IBM® Storage Protect Snapshot in basic mode rather than advanced mode.

- In basic mode, a subset of parameters is editable and can be modified. For all other parameters, default values are used. The daemons are started automatically.
- In advanced mode, you can configure all parameters. You choose whether IBM® Storage Protect Snapshot daemons are started automatically or not.

For more information about configuration files and different profile parameters, see [Configuration files](#). To read about the profile contents and which parameters are editable in advanced mode, see [Profile](#).

When the setup script is called for the first time for a specific instance, the type of instance is determined. This instance can be the productive instance or an instance for cloning. If this general information on the instance needs to be changed at a later point in time, you must run the setup script again in advanced mode.

Depending on your environment and the functionality to be used, extra configuration files might be required when you are configuring the product. The setup script checks that these files exist, but does not verify the contents. They must be complete before you run any IBM® Storage Protect Snapshot functions.

When you configure offload backups to an IBM® Storage Protect server, the startup script asks if you want to start offloaded tape backups after a snapshot completes. Choose `y` to start the offload immediately after the snapshot backup completes. Alternatively, choose to schedule the offload operation to run at a different time by scheduling backups individually. The backup to the IBM® Storage Protect server can be delayed until the necessary resources in the IBM® Storage Protect server are available. In this case, you must schedule the backup process to start manually. For example, you can add a `crontab` entry. The default value is to run `fcmlcli` as a daemon process on the production server.

You can create multiple entries with different values for some parameters. To create multiple entries, when the script asks if you want to add another instance of a parameter, enter `y`. To delete a parameter entry, when prompted for a parameter value enter `!d`. To display help for a parameter, enter `?` at the prompt for the parameter value. The help is best viewed in a window that is set for at least 130 characters.

Procedure

1. From the production database instance on the production server, log on as the database instance owner, and go to the instance directory.
`<instance_owner_$HOME>/sql/lib/acs/`
2. Run the setup script by entering the following command
`./setup_db2.sh [-advanced]`
3. Follow the setup script instructions.
When you are configuring the production server, select
On-Site Production System configuration with optional remote Backup System configuration
. If OpenSSH is available, it is recommended that you set up backup and clone servers during the configuration process. To do this, select `manage backup systems` or `manage clone instances`.

For more information about configuring a backup or clone server separately, see [Setting up separately on backup or clone servers](#).

Result

After you configured an instance, the following changes are made to the production server:

- The ACS_DIR directory is created. The default location is <instance_owner_\$HOME>/sql/lib/acs.
- The IBM® Storage Protect Snapshot profile is created in the ACS_DIR directory. Parameter values reflect values that you set during the configuration process.
- An <ACS_DIR>/shared directory is created for shared information for the production, backup, and clone servers.
- The IBM® Storage Protect Snapshot repository for the metadata of each backup is created in the ACS_REPOSITORY directory. The ACS_REPOSITORY directory defaults to <ACS_DIR>/acsrepository
- A logs directory for IBM® Storage Protect Snapshot logs and traces is created at <ACS_DIR>/logs.
- For secure communication with IBM® Storage Protect Snapshot daemons, IBM Global Security Kit key database files fcmcert.* are created in the instance directory, with an fcmselfcert.arm file that contains a representation of the production server's self-signed public key. For more information about the secure communication, see [Installing the GSKit](#).
- IBM® Storage Protect Snapshot daemons that are required for the configuration, start automatically unless you stated otherwise when you configured the product in advanced mode.

If a backup server is automatically set up during the configuration process, the following changes are applied to the backup server:

- The ACS_DIR directory is created in the same path as on the production server.
- The necessary IBM® Storage Protect Snapshot are copied to the ACS_DIR directory on the backup server.
- The IBM® Storage Protect Snapshot profile for the backup instance is created in the ACS_DIR directory.
- The <ACS_DIR>/shared directory is copied from the production system.
- A directory for IBM® Storage Protect Snapshot logs and traces is created at <ACS_DIR>/logs.
- For secure communication with IBM® Storage Protect Snapshot daemons, IBM Global Security Kit key database files fcmcert.* are created in the instance directory. Unless standard CA-signed certificates are used for server authentication, the fcmselfcert.arm file of the production system is imported in the keystore.
- The mount agent daemon starts automatically unless you stated otherwise when you configured the product in advanced mode.

If a clone server is automatically set up during configuration, the following changes are applied to the clone server:

- The <ACS_DIR of clone instance> directory is created in the same path relative to the clone instance user \$HOME directory on the production server.
- The necessary IBM® Storage Protect Snapshot files are copied to the ACS_DIR directory on the clone server.
- The IBM® Storage Protect Snapshot profile for the clone instance is created in the <ACS_DIR of clone instance> directory.
- The <ACS_DIR>/shared directory is copied from the production system to the <ACS_DIR of clone instance> directory.
- A directory for IBM® Storage Protect Snapshot logs and traces is created at <ACS_DIR of clone instance>/logs.
- For secure communication with IBM® Storage Protect Snapshot daemons, IBM Global Security Kit key database files fcmcert.* are created in the clone instance directory. Unless standard CA-signed certificates are used for server authentication, the fcmselfcert.arm file of the production system is imported to the keystore.
- The mount agent daemon of the clone instance starts automatically unless you stated otherwise when you configured the product in advanced mode.

Related information

[Profile](#)

[Setting up IBM Storage Protect Snapshot separately on backup or clone servers](#)

Configuring storage environments

IBM® Storage Protect Snapshot requires the connection and configuration information for all storage devices where data is to be protected. For IBM® System Storage® DS8000® environments and IBM® System Storage® SAN Volume Controller with static target allocation, information on the associated volumes is also required.

IBM® Storage Protect Snapshot organizes information on storage hardware in device classes. A device class is a representation of a specific storage device with its connection information and related configuration. There must be a separate **DEVICE_CLASS** defined for each storage device. If the same storage device is used with various functions and configurations, a separate device class must be defined for each configuration. An example of a different function, can be running full and incremental backups on different days of the week.

Device classes are documented as **DEVICE_CLASS** sections in the IBM® Storage Protect Snapshot profile. Parameters in a **DEVICE_CLASS** section describe the characteristics of a storage device. Therefore, they are independent of the protected application, but different parameters are required for the various types of supported storage hardware.

Each **DEVICE_CLASS** can have specific version retention that is controlled through the profile parameter **MAX_VERSIONS**, which is set during configuration. For more information about **DEVICE_CLASS** version control with **MAX_VERSIONS** parameter, see [Device class backup version retention](#).

Not all of the configuration topics are applicable to each storage system. For a list of storage environments and their associated configuration topics, see the following table.

Storage environment	Topics
IBM® System Storage® DS8000®	“Configuring the CIM adapter for SP 800-131A compliant encryption” on page 52 “Select the FLASHCOPY_TYPE” on page 53 “Target set definitions” on page 55 “LVM mirroring environments” on page 57
IBM® System Storage® SAN Volume Controller or IBM® Storwize® family with dynamic target set allocation	“Configuring Storwize family and SAN Volume Controller dynamic target allocation (SVCDTA)” on page 50 “Select the FLASHCOPY_TYPE” on page 53 “LVM mirroring environments” on page 57 “Configuring for remote mirroring” on page 61
IBM® System Storage® SAN Volume Controller or IBM® Storwize® family with static target set allocation	“Configuring the CIM adapter for SP 800-131A compliant encryption” on page 52 “Select the FLASHCOPY_TYPE” on page 53 “Target set definitions” on page 55 “LVM mirroring environments” on page 57 “Configuring for remote mirroring” on page 61
IBM® XIV® Storage System	“LVM mirroring environments” on page 57 “Configuring for remote mirroring” on page 61

For a complete list of parameters for each type of storage hardware, see [“DEVICE_CLASS device” on page 108](#).

Backup version retention

Specify the number backups to be retained by using the **MAX_VERSIONS** parameter when you are configuring IBM® Storage Protect Snapshot.

When you run the setup script, specify how many backup versions to be retained for each device class that is defined. Each device class can have a specific retention rate rather than one defined for all device classes. During configuration, you are asked for the **MAX_VERSIONS** for each **DEVICE_CLASS**.

For example, if two device classes are configured and the first has a **MAX_VERSIONS** set to 2 there are always two versions that are kept for that device class. The oldest version is deleted once a new backup is made. If you configure another device class with **MAX_VERSIONS** set to 3, it always retains three backup versions for that device class. If you do not specify a **USE_FOR** device class option, and you have two device classes in use with a **MAX_VERSIONS** set to 2, there are four versions that are kept in total.

Example of **MAX_VERSIONS** for different device classes profile

In the following example, there are four backup versions to be kept when backups are taken with **DC_XIV1** and four backup versions to be kept when backups are taken with **DC_XIV2**. Three backup versions are kept for backups that are taken with device class **STANDARD**. All other device classes have a **MAX_VERSIONS** value of two, in this case **SOME_OTHER_DC**. That means that 13 backup versions are kept with this profile: two for **SOME_OTHER_DC**, three for **STANDARD**, and four each for **DC_XIV1** and **DC_XIV2**. Thirteen backups are retained for this sample profile.

```
>>> CLIENT
MAX_VERSIONS 2
MAX_VERSIONS 3 USE_FOR STANDARD
MAX_VERSIONS 4 USE_FOR DC_XIV1
MAX_VERSIONS 4 USE_FOR DC_XIV2
DEVICE_CLASS STANDARD DC_XIV1 DC_XIV2 SOME_OTHER_DC
APPLICATION_TYPE GENERIC
TSM_BACKUP NO
NEGATIVE_LIST NO_CHECK
<<<

>>> DEVICE_CLASS STANDARD
...
<<<

>>> DEVICE_CLASS SOME_OTHER_DC
...
<<<

>>> DEVICE_CLASS DC_XIV1
...
<<<

>>> DEVICE_CLASS DC_XIV2
...
<<<
```

In the following example, **MAX_VERSIONS** is set to 2. There are four backup versions retained, two for each device class as there is no **USE_FOR** device class option specified.

```
MAX_VERSIONS 2

>>> DEVICE_CLASS STANDARD
...
<<<
>>> DEVICE_CLASS SOME_OTHER_DC
...
<<<

is the same like:

MAX_VERSIONS 2 USE_FOR STANDARD
MAX_VERSIONS 2 USE_FOR SOME_OTHER_DC

>>> DEVICE_CLASS STANDARD
...
<<<
>>> DEVICE_CLASS SOME_OTHER_DC
...
<<<
```

Configuring Storwize® family 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® Storage 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® Storage Protect Snapshot to authenticate with the private key.

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.

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 SVCDTA 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) <http://www.ibm.com/support/knowledgecenter/STPVGU/welcome?lang=en>.
4. Switch to the IBM® Storage Protect Snapshot instance directory and run the setup script in advanced mode:

```
./setup_db2.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. When prompted to specify FlashCopy® mapping techniques, enter the parameter names based on type of backups, the default state is **DISABLED**.

You can specify the following parameters:

- **IGNORE_FLASHCOPY_MAPPINGS**: The volume group snapshot is created even if the FlashCopy® mappings are present.

Note: If you plan to use the volume group snapshot feature, consider the following key points:

- Any existing backup taken by using the FlashCopy® mapping cannot be restored.
- You cannot switch back to the FlashCopy® mapping.

In case you want to use the FlashCopy® mapping, all volume group snapshots must be deleted.

- **ENFORCE_NO_FLASHCOPY_MAPPINGS**: The volume group snapshot is not created if the FlashCopy® mappings are present.

Volume group snapshots will be generated by the option 'IGNORE_FLASHCOPY_MAPPINGS', but flash copy mapping techniques will be avoided. The older snapshots created using the flash copy mapping technique will not be affected by this. Choose this option if you've recently upgraded from the older Storage Protect Snapshot software version to a newer one.

Option 'ENFORCE_NO_FLASHCOPY_MAPPINGS' will create snapshots, based on volume groups and will not allow snapshots using flash copy mapping based techniques. [DISABLED]

8. Continue configuring IBM® Storage 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.
9. The IBM® Storage Protect Snapshot daemons are automatically restarted by running the setup script either in basic mode or in advanced mode.

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 user is the one specified by the parameter **COPYSERVICES_REMOTE_USERNAME**.

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

Change an existing configuration of IBM® Storage Protect Snapshot for UNIX™ and Linux™ to use SVC dynamic target allocation (DTA) without losing older backups. Change the value of **COPYSERVICES_HARDWARE_TYPE**: **SVC** to **SVCDTA** to update the profile and complete the migration.

About this task

Use the following information to modify an existing IBM® Storage Protect Snapshot profile and reconfigure the product to SVCDTA.

Procedure

1. To start the configuration process, run the setup script.
2. Choose (m) to modify the profile.
3. Change the value of **COPYSERVICES_HARDWARE_TYPE** from SVC to SVCDTA.
4. If **MAX_VERSIONS** is set to ADAPTIVE, you must return to the CLIENT 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** <TCP/IP host name>.
6. Enter the user name for the primary storage device **COPYSERVICES_USERNAME**.
The default value is superuser.
7. Enter the path and the file name of the private SSH key file in parameter **SVC_SSHKEY_FULLPATH**.
For example,

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

8. Accept the defaults for the remaining parameters, or change the values where required.

Note: Do not change the FlashCopy® type.

Result

The updated profile is saved, and if required you can specify a backup system or quit the configuration. IBM® Storage Protect Snapshot for UNIX™ and Linux™ is configured to use the SAN Volume Controller storage adapter with dynamic target allocation.

After reconfiguring the profile for SVCDTA, all backups are managed by the new dynamic target allocation adapter, including all backups that were taken with the CIM adapter. Continue to mount, restore, delete, and expire backups as usual. When **FLASHCOPY_TYPE** is set to INCR, you can also refresh backups as required.

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® Storage 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 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 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.
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.

Select the FLASHCOPY_TYPE

DS8000®, SAN Volume Controller, and Storwize® family 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® Storage 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® Storage 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® family dynamic target allocation you do not have to delete any old backups.

Table 5: Selecting the FLASHCOPY_TYPE for DS8000®, SAN Volume Controller, and Storwize® family		
FLASHCOPY_TYPE	DS8000®	SAN Volume ControllerStorwize® family
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 COPYactivities 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 COPYFlashCopy® but with fewerCOPYactivities in the background. For more information, see Note 1 and Note 2 in this table.

FLASHCOPY_TY PE	DS8000®	SAN Volume ControllerStorwize® family
NOCOPY	Can be mounted remotely, but cannot be restored.	<p>Can be mounted remotely and can be restored.</p> <p>Can be used to create a FlashCopy® to a space-efficient target, but does not offer protection from physical failures to the source volume.</p> <p>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 theAUTOEXPANDoption. In this scenario, the target is allocated more physical storage to prevent it going offline.</p>
<p>Note 1: If space-efficient source volumes are used in combination with space-efficient target volumes, IBM® Storage Protect Snapshot can be configured to use FLASHCOPY_TYPECOPY, INCR, or NOCOPY. If fully allocated source volumes are used in combination with space-efficient target volumes, then IBM® Storage Protect Snapshot can be configured to use FLASHCOPY_TYPECOPY, 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_TYPENOCOPY is possible.</p> <p>Note 2: The information in Note 1 only applies if you use SAN Volume Controller and Storwize® family static target allocation. If you use SAN Volume Controller and Storwize® family dynamic target allocation, the profile parameter ALLOW_ALL_FLASHCOPY_TYPES is not available.</p>		

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

Table 6: 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 ControllerIBM® 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 ControllerIBM® Storwize® family with dynamic target allocation	Yes	Yes	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. See parameter USE_WRITABLE_SNAPSHOTS in XIV System DEVICE_CLASS in the Reference section.

Target set definitions

IBM® Storage Protect Snapshot requires target sets to be defined for SAN Volume Controller, Storwize® family, and DS8000®. A target set contains volumes that are allocated in the same storage subsystem, and are of the same size as the volumes that contain the data to be protected. The number of target sets determines the number of backup versions that can be kept at a time.

Define targets by using target set definition files, as follows:

- SAN Volume Controller
- Storwize® family
- DS8000®

Alternatively, use a naming convention for SAN Volume Controller and Storwize® family 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® family 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® Storage 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 files” on page 132](#).

If IBM® Storage 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® Storage 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® family storage solutions, IBM® Storage 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 # IBM Storage Protect Snapshot 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
                        # definition (source name following target name)
TARGET_VOLUME 40C13158 40613158
TARGET_VOLUME 40D13158 40713158
TARGET_VOLUME 40E13158 40813158
<<<
```

Multi-partition DB2® target set definition file

Multi-partition DB2® database target set definitions must be specified for each partition.

As a result, the contents of the target set definition file is separated into multiple sections. One section is used for each partition, as shown in this example:

```
>>> TARGET_SET SET_1
                        # IBM Storage Protect Snapshot determines a target for every source
>>> PARTITION NODE0000
TARGET_VOLUME 40913158
```

```

TARGET_VOLUME 40A13158
<<<
>>> PARTITION NODE0001
TARGET_VOLUME 40B13158
TARGET_VOLUME 50913158
TARGET_VOLUME 50A13158
TARGET_VOLUME 50B13158
TARGET_VOLUME 51713158
<<<
>>> PARTITION NODE0002
TARGET_VOLUME 51813158
TARGET_VOLUME 52113158
TARGET_VOLUME 52313158
<<<
<<<

```

A single backup server environment (for a multi-partition DB2® database) can be distributed across multiple servers. In this situation, make sure that the target volumes in the target set definition file are assigned to the correct host. For best results, assign all target volumes in the target set definition file (and that are associated with the same partition) to the same host. In environments where multiple partitions are on the same volume group (volume sharing), target set definition files need to specify certain values. For more information, see [“Manage target volumes files for your storage system” on page 133](#).

Referring to target set definitions from the profile

The following example is a section from an IBM® Storage 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® family.

IBM® Storage Protect Snapshot supports using a naming convention, instead of a definition file, for target set definitions on SAN Volume Controller and Storwize® family storage systems. IBM® Storage 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® Storage 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
```

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
<<<
```

Configuring a HyperSwap environment

Run the setup script to modify HyperSwap parameters.

```
./setup_db2.sh -advanced
```

Obtain the relevant pool name and IO group settings from your storage administrator. At least one device class is required and the pool name and IO group of the HyperSwap master site must be added to this device class.

Modify the SVC_POOLNAME and SVC_IOGROUP settings when prompted by the setup script.

Specify HyperSwap pool names and IO groups with the SVC_POOLNAME and SVC_IOGROUP parameters when you are configuring IBM® Storage Protect Snapshot in your HyperSwap environment.

The SVC_POOLNAME parameter indicates the name of the IBM Storwize family or IBM System Storage SAN Volume Controller storage pool that is used to create target volumes for IBM® Storage Protect Snapshot backups:

- If a value is not specified, then an IBM® Storage Protect Snapshot backup will fail.
- If one value is specified, the target disk will be located in the storage pool that equals this value.
- If two values are specified, the target disk will be located in the storage pool equal to the first value, and the vDisk mirror will be created in the storage pool equal to the second value.

The SVC_IOGROUP parameter specifies the input and output (IO) group with which to associate the target volumes of the Snapshot relationships.

The HyperSwap parameters are maintained in the DEVICE_CLASS section in the profile shown in this example.

```
>>> DEVICE_CLASS YOUR_CLASS_NAME
COPYSERVICES_HARDWARE_TYPE GENERIC
COPYSERVICES_ADAPTERNAME svc/SvcAdapter.jar
COPYSERVICES_SERVERNAME svc05
COPYSERVICES_USERNAME superuser
BACKUP_HOST_NAME AIX2
SVC_POOLNAME standard_pool
SVC_IOGROUP io_grp0
...
```

In this example, change the SVC_POOLNAME setting to the pool name supplied by the HyperSwap administrator. Allocate the IO group name SVC_IOGROUP to the group where this pool belongs.

```
>>> DEVICE_CLASS YOUR_CLASS_NAME
COPYSERVICES_HARDWARE_TYPE GENERIC
COPYSERVICES_ADAPTERNAME svc/SvcAdapter.jar
COPYSERVICES_SERVERNAME svc05
COPYSERVICES_USERNAME superuser
BACKUP_HOST_NAME AIX2
SVC_POOLNAME hyperswap_pool
SVC_IOGROUP io_grp1
...
```

Backup and clone server assignment

With IBM® Storage 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® Storage 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_clone** or **refresh_clone** command from the command-line interface.
- By issuing a backup command in environments where a forced mount is required during a backup operation.
- When IBM® Storage Protect Snapshot is used with IBM® Storage Protect and you offload backups to IBM® Storage Protect.

The information that you enter during the installation and configuration of IBM® Storage 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® family

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® family

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. This setting ensures that targets associated with a single device class are mounted from the same backup or clone server.

For all IBM® Storage 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.

db2 backup use snapshot

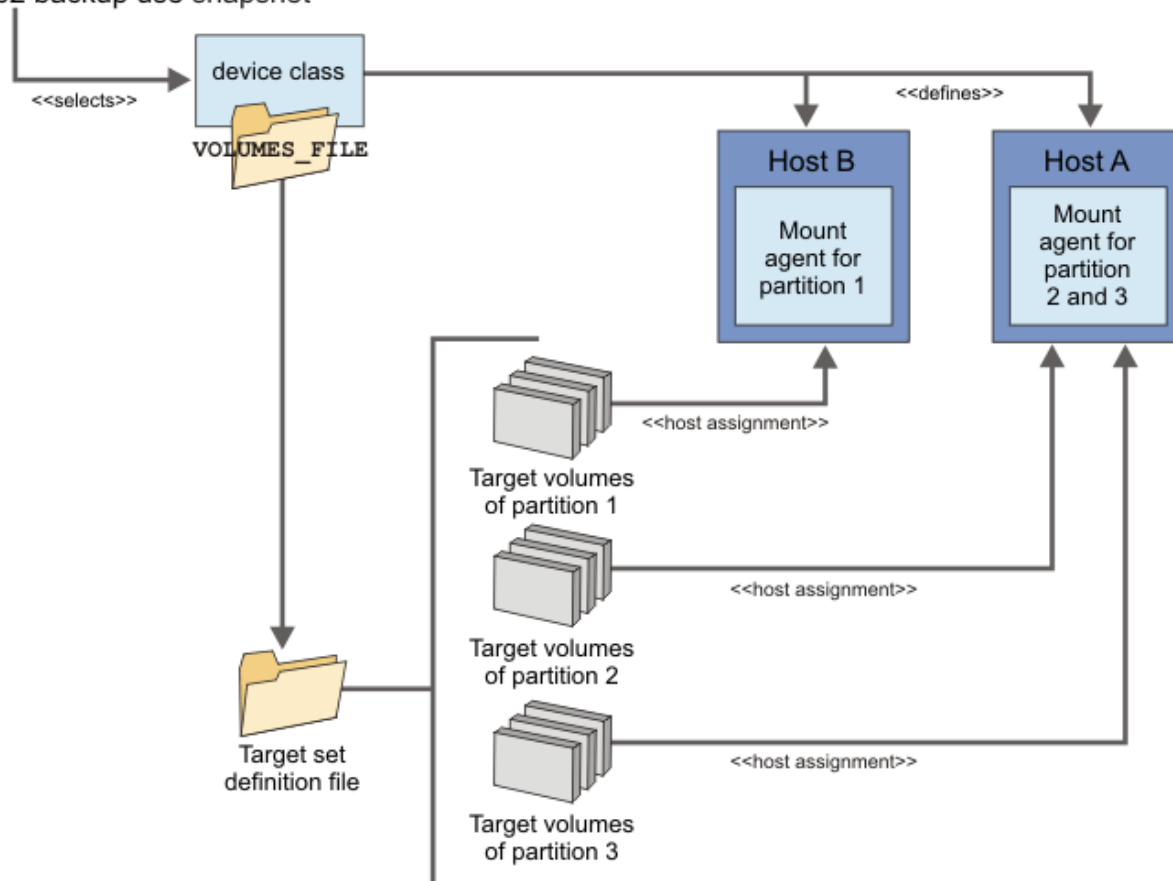


Figure 7: IBM® Storage Protect Snapshot host assignments for DB2®.

Managing backups and clones with the **DEVICE_CLASS** parameter

Use the **DEVICE_CLASS** parameter in the **CLIENT** section of the IBM® Storage 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® Storage 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 storage 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® family
- The type of snapshot backup to be completed (for example, incremental or copy)
- The mount location of the backup
- Whether a snapshot backup should be offloaded to an IBM® Storage Protect server.

The **DEVICE_CLASS** parameter is specified in the client section of IBM® Storage 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 DB2® backup command line

```
OPTIONS DEVICE_CLASS=device class
```

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® Storage 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 offloaded to the IBM® Storage Protect server:

```
>>> CLIENT
TSM_BACKUP LATEST USE_FOR MIRROR_2
DEVICE_CLASS MIRROR_1 MIRROR_2
[...]
```

This example creates backups of a partitioned DB2® database with five partitions. While partition 0 uses **DEVICE_CLASS PARTITION0**, the partitions 1-4 use **DEVICE_CLASS PARTITIONX**:

```
>>> CLIENT
DEVICE_CLASS PARTITION0 ON_DBPARTITIONNUMS 0
DEVICE_CLASS PARTITIONX ON_DBPARTITIONNUMS 1 2 3 4
[...]
```

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 the IBM® Storage Protect server:

```
>>> 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 the IBM® Storage Protect server. 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
[...]
```

```

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® Storage 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® Storage 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® family Global Mirror and Metro Mirror
- IBM® XIV® Storage System Synchronous Remote Mirroring and Asynchronous Remote Mirroring

Configuring for remote mirroring on SVC

When you configure IBM® Storage 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.

About this task

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

Procedure

1. On the SAN Volume Controller or Storwize® family 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® family, 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?
```

Configuring for remote mirroring on XIV

When you configure IBM® Storage 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.

About this task

To configure IBM® Storage 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® Storage 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 daemons

By default, IBM® Storage Protect Snapshot daemons are started and restarted automatically. If you need to stop and start the daemon processes for any reason, you can do this as required.

About this task

Important: If you want to configure the background daemons in a clustered environment, add the commands that are listed here to the High Availability (HA) startup scripts on the production system.

- `<instance directory>/acsd`
- `<instance directory>/acsgen -D`
- `<instance directory>/fcmcli -D` (optional)

The following daemons are started for the productive instance:

- The management agent, `acsd`

- The generic device agent, *acsgen -D*
- The offload agent, *fccli -D*, if offloaded backups to an IBM® Storage Protect server are to be started after a snapshot backup.

For each backup or clone instance, a mount agent, *acsgen -D -M*, must be started.

- `<instance directory>/acsgen -D -M -s <deviceclass>[,<deviceclass>...]`
- `<clone instance directory>/acsgen -D -M -s <clone_deviceclass>[,<clone_deviceclass>...]`

When you run the setup script in advanced mode, you can request that the daemons are not to be started automatically by the operating system. In this case, daemons must be started and set up to be restarted by other means.

Result

For more information about daemons, see [“Administrative commands” on page 147](#).

Postinstallation and post-configuration tasks

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

Use the setup script to update the profile and configure IBM® Storage 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.

In a physically partitioned DB2® environment, installation and configuration are only required on one node of the production server: the master production server.

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

If OpenSSH is not available, follow the instructions for [“Setting up IBM Storage Protect Snapshot separately on backup or clone servers” on page 42](#) 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® Storage Protect Snapshot on the production server. For details about how to stop an activated IBM® Storage Protect Snapshot instance, see IBM® Storage 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® Storage Protect Snapshot on the backup server.
- Changes to the scheduling policy for offloaded IBM® Storage 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 `fcmselfcert.arm` file to the backup server before the setup script is run to configure the backup server again.

In these cases, stop IBM® Storage Protect Snapshot on the production server before reconfiguring the backup server. Otherwise, you are prompted to stop IBM® Storage Protect Snapshot on the production server.

Upgrading

To upgrade to a newer version of IBM® Storage Protect Snapshot, you must install that newer version. The updates are effective for instances only after you activate them with the new version of the product.

Before you begin

When you are upgrading an instance, make sure to update its backup and clone systems.

Procedure

1. Install the new version of the product as described here, [“Installing IBM Storage Protect Snapshot in interactive mode” on page 40](#)

2. After the product is installed successfully, the application-specific instances must be activated with the new version.
The activation copies the installation binary files from the global installation directory into the application-specific installation directory. Follow these steps to activate the instances, [“Activating an instance” on page 42](#)
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.
This action updates the profile with new parameters and removes deprecated parameters, or renames them if required. Select the upgrade from the wizard and choose the option to **update IBM® Storage Protect Snapshot installation**. Follow the instructions to run the setup script as described here, [“Running the setup script for IBM Storage Protect Snapshot for DB2” on page 46](#).

If the backup or clone system was installed, activated, and configured separately, follow the instructions in [Setting up backup or clone servers](#).
4. Uninstall the older version of the product.
Follow the instructions described here, [“Uninstalling the software” on page 41](#)

Upgrading from IBM® Tivoli® Storage FlashCopy® Manager version 3.1

The use of consistency groups is required for IBM® Storage Protect Snapshot. The profile parameter **USE_CONSISTENCY_GROUPS** is no longer supported and is removed automatically from an IBM® Tivoli® Storage FlashCopy® Manager v3.1 profile when the upgraded instance is configured.

About this task

Use the setup script to migrate a profile that was created for IBM Tivoli Storage FlashCopy Manager V3.1 to IBM® Storage Protect Snapshot.

Procedure

1. Log in to the production server with the instance owner ID, go to the instance directory, and run the setup script with the following command:

```
./setup_db2.sh
```

2. Follow the setup script instructions.

What to do next

For each IBM® Tivoli® Storage FlashCopy® Manager Version 3.1 profile that has the **USE_CONSISTENCY_GROUPS** parameter, repeat the steps to automatically remove the parameter.

Protecting your data with IBM® Storage Protect Snapshot

Backing up data

Create snapshot backups of your databases and applications with IBM® Storage Protect Snapshot. Integrated the product with IBM® Storage Protect clients to offload backups to an IBM® Storage Protect server.

DB2® backup commands

When you are backing up DB2® databases, use specific commands for your database configuration and backup type.

The following table summarizes the commands that are needed for backing up DB2® databases for different database configurations and corresponding backup types.

Summary of backup commands for DB2® databases

Table 7: Summary of backup commands for DB2® databases				
Database configuration	Snapshot backup (disk only)	Backing up to an IBM® Storage Protect server		
		From production database (tape only)	Integrated with snapshot	From existing snapshot
DB2®	<code>db2 backup ... use snapshot ...</code>	<code>db2 backup ... use tsm</code>	<code>db2 backup ... use snapshot ...¹</code>	<code>fcmlcli -f tape_backup^{2, 3}</code>
DB2® in an SAP environment	<code>db2 backup ... use snapshot ...</code>	<code>db2 backup ... load <library> or backom</code>	<code>db2 backup ... use snapshot ...¹</code>	<code>fcmlcli -f tape_backup^{2, 3}</code>
Note:				
1. The profile parameter TSM_BACKUP is set to YES, MANDATE, or LATEST, and the offload agent (fcmlcli) is running in daemon mode on the production server.				
2. The profile parameter TSM_BACKUP is set to YES, MANDATE, or LATEST, and the offload agent (fcmlcli) is not running in daemon mode.				
3. The <code>fcmlcli -f tape_backup</code> operation must be issued from the production system.				

The **db2 backup database** command with the `use snapshot` option is described in the *DB2® Command Reference* guide.

When you restore a backup image, the whole backup image is restored. Therefore, any data on the file system that is updated or created after the backup is taken is overwritten. To ensure that nothing is overwritten unintentionally, do not store any other data in the data and log volume groups. Any other data that is stored on the volume groups is processed and included in the backup image.

Extra files that are not associated with the database backup operation but are stored in the volume groups can cause the backup to fail. Use the **NEGATIVE_LIST** parameter in the profile to control the processing of these files. For information about this parameter, see the [“CLIENT” on page 101](#) profile section for details.

Backing up DB2® databases

To protect data against the possibility of loss, ensure that you back up the DB2® database. Back up single partition, and multi-partition databases with DB2® backup commands.

About this task

For volume group snapshot, refer [“Configuring Storwize family and SAN Volume Controller dynamic target allocation \(SVCDTA\)” on page 50](#).

Procedure

1. Use the **db2 backup** commands when you are backing up DB2® databases with IBM® Storage Protect Snapshot.
2. Log on to one of the production hosts and use the DB2® instance owner ID.
3. Depending on the database you are backing up, type in one of the following commands:
 - Single partition database: **db2 backup db dbname use snapshot**
 - Multi-partition database: **db2 backup db dbname on all dbpartitionnums use snapshot**
4. (Optional) To specify a different profile when you back up the database, enter the following command:
db2 backup db dbname on all dbpartitionnums use snapshot options "PROFILE=path to new profile or name of new profile"

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.

Serial and parallel backup modes for multipartition DB2® database

DB2® backs up database partitioning feature (DPF) partitions in either serial mode or parallel mode.

In some later versions of DB2®, the backup mode is manipulated by the value of the DB2® registry variable DB2®_PARALLEL_ACS. In versions of DB2® that do not support this registry variable, the backup mode is determined by the DB2® environment.

Serial mode

Backups are done in serial mode when the DB2® registry variable DB2®_PARALLEL_ACS is set to NO. It is the default if DB2®_WORKLOAD is not set or if it is set to a value different from SAP. A restore operation is always performed on a single partition. In serial mode, the partitions are processed sequentially. The sequence of events is as follows: each partition is suspended. Then the snapshot is created, and the partition is resumed before the next partition is processed. IBM® Storage Protect Snapshot returns an error during a backup if multiple partitions share a physical volume.

Parallel mode

This is the default setting when DB2®_WORKLOAD is set to SAP. IBM® Storage Protect Snapshot allows multiple logical partitions to share a physical volume when these logical partitions are grouped in a **PARTITION_GROUP**. In this situation, **fcmcli** must be used to perform the restore. The restore proceeds in parallel for all the logical partitions that are specified in **PARTITION_GROUP**. In parallel mode that is the default mode for an SAP DB2® database, all partitions are suspended before DB2® issues snapshot requests. The requests are then performed in parallel on all partitions. Backups are done in parallel mode if the DB2® registry variable DB2®_PARALLEL_ACS is set to YES.

Related information

[Commands that use DB2 utilities](#)

[The fcmcli offload agent](#)

Backing up DB2® to an IBM® Storage Protect server

IBM® Storage Protect Snapshot relies on standard DB2® mechanisms to back up a snapshot image to the IBM® Storage Protect server.

IBM® Storage Protect Snapshot uses the following products to offload backups to an IBM® Storage Protect server:

- IBM® DB2® Storage Protect client for DB2® environments.
- IBM® Storage Protect for Enterprise Resource Planning, with the option for DB2® in an SAP environment.
- IBM® Storage Protect backup-archive client, which is used to offload backups from the DB2® standby server.

Note: In an SAP environment, IBM® Storage Protect Snapshot can be configured to use the DB2® IBM® Storage Protect agent. See the IBM® Storage Protect Snapshot **VENDOR_LIB** profile parameter in the offload section for details.

Use one of the following methods to offload a snapshot backup to an IBM® Storage Protect server:

- When the offload agent (fcmcli) is running, it queries periodically for snapshot backups where the **TSM_BACKUP** configuration parameter is set to YES, MANDATE, or LATEST, in the IBM® Storage Protect Snapshot profile file. The offload agent initiates a tape backup from the snapshot target set when the snapshot is created. The offload agent must be running in daemon mode on the production server.
- The **TSM_BACKUP** parameter is set to YES, MANDATE, or LATEST but the offload agent is not running then you must enter the following command **fcmcli -f tape_backup**. You must enter the command as the database instance owner on the production server. This command offloads a previously generated snapshot backup to an IBM® Storage Protect server.

IBM® Storage Protect Snapshot, IBM® Storage Protect for Enterprise Resource Planning, IBM® DB2® Storage Protect client for DB2® environments, and IBM® Storage Protect backup-archive client use their own configuration files.

To offload a backup to an IBM® Storage Protect server, an OFFLOAD section must be present in the IBM® Storage Protect Snapshot profile file. This OFFLOAD section is created during the installation when the option to offload backups is selected.

When a profile is created or changed, the IBM® Storage Protect Snapshot configuration setup wizard, prompts you to specify whether offloaded backups are run. When YES is specified, the OFFLOAD section is added to the profile and the **TSM_BACKUP** parameter is added to the CLIENT section of the profile.

Backing up DB2® (HADR) standby servers

You can use IBM® Storage Protect Snapshot to create snapshot backups of a DB2® High-Availability Disaster Recovery (HADR) standby server.

Although DB2® (HADR) standby server cannot be directly backed up to a IBM® Storage Protect server, snapshot backups and subsequent backups of the snapshot files are possible. To ensure application consistency of the DB2® (HADR) standby database, the database must be deactivated or stopped before the snapshot is created. Then, after the snapshot is created the database must be reactivated or started again. Sample scripts that show you how to shut down and restart the DB2® (HADR) standby server are in the instance directory. You can customize these sample scripts. The scripts are available when IBM® Storage Protect Snapshot is configured for DB2® (HADR) standby server support. The sample scripts can be used as parameter values for the IBM® Storage Protect Snapshot profile **PRE_FLASH_CMD** and **POST_FLASH_CMD** parameters in the DB2®STANDBY section of the profile file.

Enter the following commands to create snapshot backups of a DB2® (HADR) standby servers:

1. Log on to the DB2® (HADR) standby server and use the database instance owner user ID.
2. Change to the instance directory. For example, *<DB2 instance owner \$HOME>/sql1lib/acs/*
3. Enter the following command:

```
./fcmcli -f backup_db2standby
```

Before IBM® Storage Protect Snapshot starts the backup process, IBM® Storage Protect Snapshot communicates with DB2® on the primary site to determine the list of database files that must be protected.

After a DB2® takeover, the DB2® (HADR) standby server becomes the DB2® (HADR) primary server. After this takeover, use the DB2® snapshot backup command for backup. Do not use the **backup_db2standby** command. IBM® Storage Protect Snapshot supports sharing the profile for both types of backups. You need only

to use a different command to start your database backup that is dependent on the current DB2® HADR database role.

Use the same device classes for the DB2® snapshot backups, and for the DB2® (HADR) standby server backups.

Note: For DS8000® and SAN Volume Controller a DB2® snapshot backup can overwrite the backup from a DB2® (HADR) standby server backup. This situation occurs if there are different device classes for two backup types, but the same target volumes are used. The **INQUIRE** command incorrectly indicates that both backups are valid.

Because IBM® Storage Protect Snapshot software cannot directly communicate with the standby server, IBM® Storage Protect Snapshot queries the DB2® instance that is running on the primary server by using DB2® remote connections. The software assumes that the file names for the production server database and the standby server database are identical. IBM® Storage Protect Snapshot can back up in situations where the DB2® automatic storage managed table spaces are in use, which can cause table space container IDs on the HADR primary and standby server to be mismatched.

You cannot use DB2® commands to create a DB2® backup of the standby databases and offload directly to an IBM® Storage Protect server. You can use IBM® Storage Protect Snapshot to start a file-based IBM® Storage Protect backup from the snapshot using the IBM® Storage Protect backup-archive client.

You cannot use DB2® for a database restore from the file-based IBM® Storage Protect backups. Use the IBM® Storage Protect backup-archive client to restore the files that are backed up to the IBM® Storage Protect server. Recover the database by using the DB2® **RECOVER** command.

Snapshot backup of individual mirrors

IBM® Storage Protect Snapshot supports mirroring.

Mirroring by using the AIX® logical volume manager (LVM mirroring)

IBM® Storage Protect Snapshot provides LVM mirroring support for the following storage devices:

- DS8000®
- IBM® XIV® Storage System
- Storwize® family
- SAN Volume Controller

For those devices, IBM® Storage Protect Snapshot creates a snapshot 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® Storage 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 snapshot per source volume. However, in LVM environments, each source volume is mirrored. Therefore, IBM® Storage Protect Snapshot can create two INCREMENTAL snapshot backups for DS8000®.

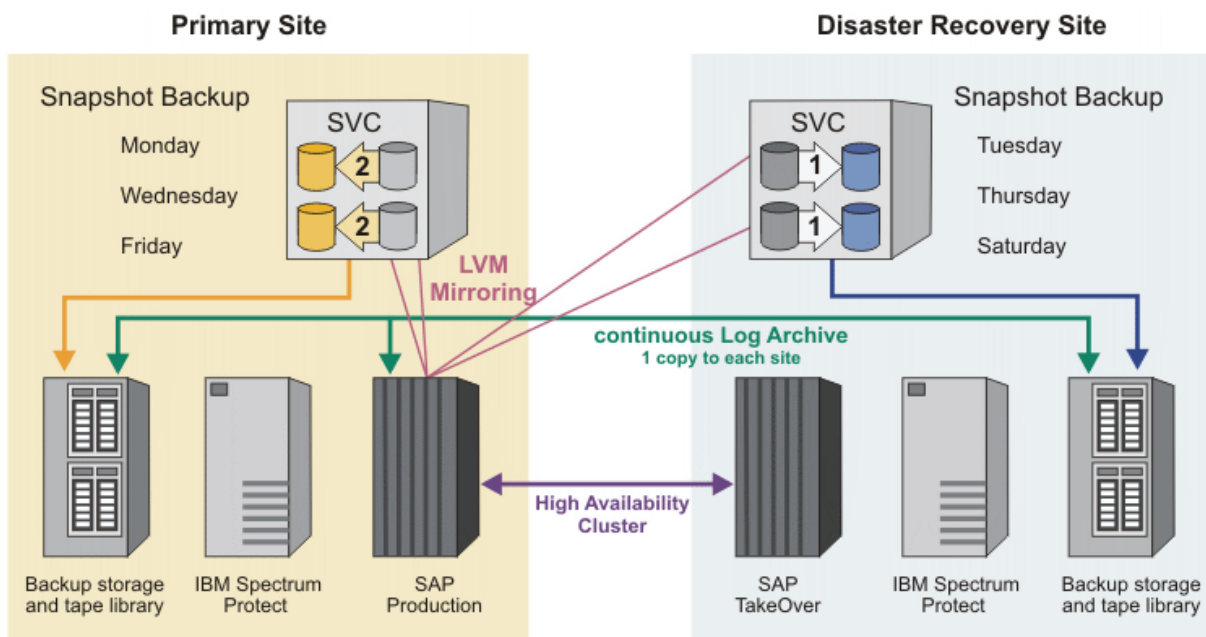


Figure 8: Cross-site mirrored SAP database that is protected with IBM® Storage Protect Snapshot and an IBM® Storage Protect server.

AIX® LVM mirroring advantages

- Only one of the two LVM mirrors are used for the snapshot. Using one mirror saves the number of needed target volumes and reduces the time that is needed for creating the snapshot.
- Avoids unnecessary performance degradation within the storage system.
- All LVM mirrors on the production system remain synchronized when the snapshot is created.
- Online or offline snapshot backups can be created in both LVM mirrored and non-LVM mirrored environments. The backup and restore procedures as provided in the applicable documentation remain unchanged.
- The snapshot 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 snapshot backup request.
- IBM® Storage Protect Snapshot provides information about asymmetrical LVM mirror setups when they are encountered. The snapshot backup can fail in such environments, indicating a general deficiency of the high-availability setup.

IBM® Storage Protect Snapshot requires that the LVM mirroring sets are in different storage subsystems. For example, different SAN Volume Controller clusters. Complete mirrors must be stored on both storage clusters. If this setting is not possible, IBM® Storage Protect Snapshot continues processing for those clusters where a complete image of the application can be found.

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

During a restore operation, IBM® Storage 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.

Support of AIX® enhanced concurrent capable volume groups

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

Heterogeneous device mirroring

IBM® Storage 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

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® 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® Storage Protect Snapshot is installed in a supported environment.
You must have a 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 to configure IBM® Storage Protect Snapshot for remote mirroring.
When you configure for remote mirroring, the following parameters are set in the `DEVICE_CLASS` section:
 - **COPYSERVICES_REMOTEYES**
 - **COPYSERVICES_REMOTE_SERVERNAME** <SERVER_NAME>
 - **COPYSERVICES_REMOTE_USERNAME** <USER_NAME>
 - **TAKEOVER_HOST_NAME** <HOST_NAME>

3. At the end of the setup script, 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 that was 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` that is in use is selected.
5. From the production host, start the remote snapshot backup by typing in the following command:
DB2® database, remote snapshot backup, single partition database

```
db2 backup db dbname use snapshot
```

or, DB2® database, remote snapshot backup, multi-partition database

```
db2 backup db dbname on all dbpartitionnums use snapshot
```

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. A snapshot consistency group is 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 available space on the target storage pool so that it can create the snapshot. Increase the capacity on the requested storage pool, or free up items that are using existing capacity, if there is not enough space. 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:

```
fcmccli -f inquire_detail
```

What to do next

When you complete 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: **fcmccli -f mount**
- Unmount the backup, from a command prompt window, by entering the following command: **fcmccli -f unmount**

Related information

[Mounting and unmounting snapshots on a secondary system](#)

[Updating DEVICE_CLASS device for mirroring](#)

Logical Volume Manager support (AIX® only)

You can use IBM® Storage 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 snapshot backup is created such that only one mirror is being copied.

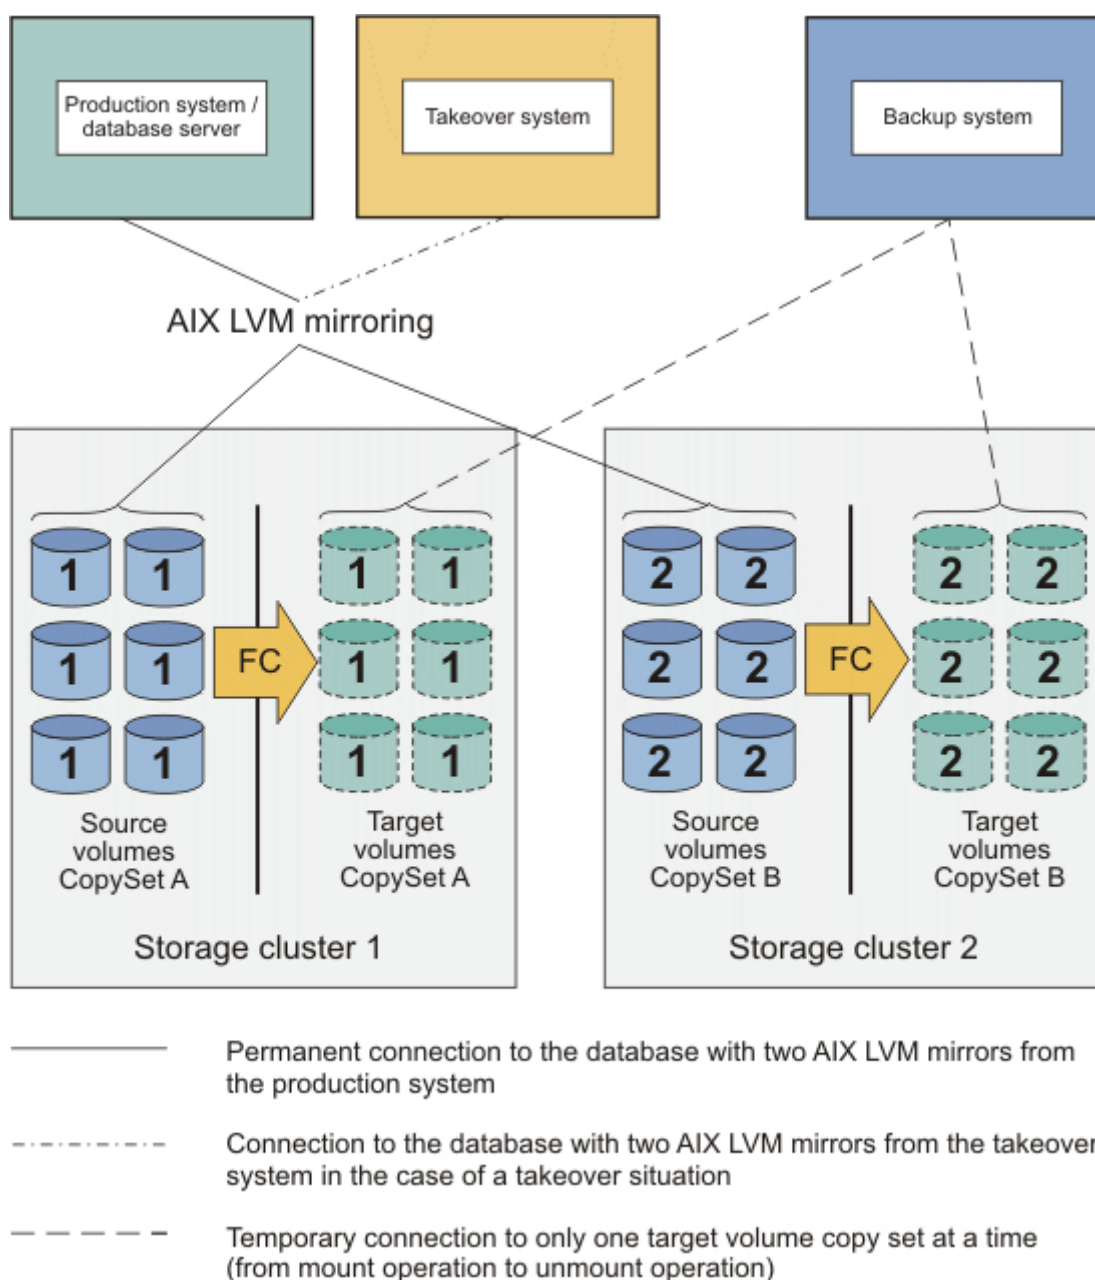


Figure 9: IBM® Storage Protect Snapshot in an LVM environment

Using Snapshot backup in a HyperSwap environment

You can avail of the ultimate level of data availability by backing up storage with IBM® Storage Protect Snapshot support for IBM HyperSwap technology.

Backing up data to a HyperSwap pool works the same way as backing up to a non-HyperSwap pool. Backing up is carried out using the master site of the HyperSwap pairing. The auxiliary pool is used only to implement HyperSwap high availability.

You can use the relevant commands described in the topic [Backing up DB2 databases](#).

Using Volume Group Snapshot in IBM FlashSystems

IBM® Storage Protect Snapshot supports volume group snapshot.

IBM® Storage Protect Snapshot supports backup, in-place restore, delete, and offload operations from single site and Policy-Based High Availability (HA) topologies for the IBM FlashSystem volume group snapshots.

To enable the volume group snapshot feature, see Step “7” on page 51 in the “Configuring Storwize family and SAN Volume Controller dynamic target allocation (SVCDTA)” on page 50 topic.

On IBM FlashSystem, volume group snapshots are created with the SPS_<Snapshot - ID> prefix.

Note: If you plan to use the volume group snapshot feature, consider the following key points:

- Any existing backup taken by using the FlashCopy® mapping cannot be restored.
- You cannot switch back to the FlashCopy® mapping.

In case you want to use the FlashCopy® mapping, all volume group snapshots must be deleted.

Restoring data

Restore databases with IBM® Storage Protect Snapshot by restoring from a snapshot on the storage subsystem, or restoring data from IBM® Storage Protect.

Restoring DB2® databases

To restore a DB2® database, open a command-line window.

Prerequisite for restoring a DB2® database in PBHA topology

Confirm with IBM FlashSystem storage administrator that the host location refers to the backup site.

1. Execute the **lshost** command.

```
IBM_FlashSystem:SPS_CL_30:superuser>lshost
id name port_count iogrp_count status site_id site_name host_cluster_id host_cluster_name protocol owner_id owner_name portset_id portset_name
partition_id partition_name draft_partition_id draft_partition_name ungrouped_volume_mapping location_system_name
0 AIXora_183 1 4 online no sps-130 scsi 64 portset64
0 p1_30
```

Figure 10: Confirming site location

The output indicates that the host location (location_system) refers to the backup site.

2. If the location of the host is different than the backup site, issue the following commands from any site:

```
# chhost -location <site-2-name> <hostname>
```

Refer IBM FlashSystem documentation for more details.

3. If you want to restore the snapshot backup from site2, issue the following command from any site, before restore operation:

```
# chhost -location <site-2-name> <hostname>
```

Note: Only supported on the AIX platform.

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

Table 8: Summary of Restore Commands for DB2®		
Database configuration	Snapshot restore	Restore from IBM® Storage Protect
DB2®	db2 restore ... use snapshot ...	db2 restore ... or db2 recover ...
DB2® in an SAP environment	db2 restore ... use snapshot ...	db2 restore ... or db2 recover ... or backom

Depending on the options that are specified in the **db2 backup ... use snapshot ...** command when the snapshot is created, both backup types (snapshot and IBM® Storage Protect) can be restored. The background copy must be completed before you restore a snapshot backup even if the snapshot backup request completed successfully. When you restore from snapshots, all objects on target volumes that are created in the backup snapshot process are restored.

The default behavior when you restore data from a snapshot image is a full database offline restore of all paths that make up the database. This restore includes all containers, the local volume directory, and the database path but does not include the database log files. To restore a snapshot, issue the following command:

```
db2 restore db DBNAME use snapshot taken at TIMESTAMP
```

Where *TIMESTAMP* is a string that represents the time stamp of the backup image and is optional if there is only one backup image in the source specified. The log files are excluded by default from the restore unless you specify the `logtarget include` option. To restore log files, table spaces, and directories, issue the following command:

```
db2 restore db DBNAME use snapshot taken at TIMESTAMP logtarget include
```

Note: Using some storage systems, the snapshot restore requires a certain amount of 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.

In a more complex environment, use the following commands to restore a DB2® multi-partition snapshot:

```
(catalog node first)
db2_all "<<+0< db2 restore db H80 use snapshot without prompting"
```

```
(remaining nodes)
db2_all "<<-0< db2 restore db H80 use snapshot without prompting"
```

If the **PARTITION_GROUP** parameter was specified when the snapshot was created, you cannot use the **db2 restore** command to restore the snapshot. For this scenario, use the IBM® Storage Protect Snapshot **fcmccli** command with the following syntax:

```
fcmccli -f restore -d <dbname> -B <backupID> -P <partition group name>
```

If more than one partition group is used during the backup, the restore operation must be completed for all partition groups.

Restoring DB2® HADR and DB2® standby server environments

Use the `fcmccli -f restore_db2standby` command to restore a IBM® Storage Protect Snapshot snapshot backup of a DB2® HADR standby server. Before you restore, ensure that you stop the DB2® HADR function and deactivate the DB2® database. IBM® Storage Protect Snapshot stops the DB2® manager before the restore operation. If the DB2® HADR is still running or the database is still active, IBM® Storage Protect Snapshot cannot stop the DB2® manager and the restore is not successful.

A restore operation of an offloaded tape backup from IBM® Storage Protect can be done by running IBM® Storage Protect backup-archive client commands on the DB2® HADR standby server. A proxy node setup can be used.

IBM® Storage Protect Snapshot maintains a backup ID that can uniquely identify any snapshot backup. Backup-archive client backups are not associated with backup IDs. Therefore, there is no correlation between a snapshot backup and its corresponding IBM® Storage Protect backup, other than the time when either backup was created. In configurations where the offloaded IBM® Storage Protect backup is scheduled at a different time from the snapshot backup, a time correlation can be difficult to achieve.

Tip: If you want to maintain that relationship explicitly in the offload section of the IBM® Storage Protect Snapshot profile, set the **MODE** parameter to **ARCHIVE**. This way, IBM® Storage Protect Snapshot can correlate both backup methods by using the snapshot backup ID as the archive description for the offloaded backup.

DB2® backup history file overview

DB2® provides its own history file that stores all information about backup, restore, and changes in the database (such as adding containers to a tablespace).

Issue one of these commands to list information from the backup history file:

```
db2 list history backup all for <SID>
```

or

```
db2 list history rollforward all for <SID>
```

For more information about the **db2 list history** command, see *IBM® DB2® Command Reference*.

To restore a backup that was performed on the local production system, you can find the timestamp of the backup with the **db2 list history** command.

Restoring data with remote mirroring

Restore data on a remote site with IBM® Storage 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® Storage 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 DB2® 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 **ACSD** 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® Storage Protect Snapshot snapshot local repository after the remote backup.

Procedure

1. Start the DB2® database manager with the following command: **db2start**
2. Update the IBM® Storage Protect Snapshot configuration parameters with the setup script wizard. Specifically, set the **ACSD** parameter to use the *acsd* on the takeover host. Do not use the *acsd* of the production host.
3. Start the IBM® Storage Protect Snapshot *acsd* daemon on the takeover host.
4. From the backups that are displayed, select the remote backup to use for the restore. The backups are displayed when you enter the query command on the takeover host.
For example, **db2acsutil query show details**
5. Start the restore by entering the following command: `db2 restore db <db_name> use snapshot taken at <backup_timestamp>`

Result

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.

Using Snapshot restore in a HyperSwap environment

You can avail of the ultimate level of data availability by restoring storage with IBM® Storage Protect Snapshot support for IBM HyperSwap technology.

You can use the relevant commands described in the topic [Restoring DB2 databases](#).

HyperSwap restore limitations

Snapshot restore operations on master and auxiliary volumes are not possible without interrupting HyperSwap relations. During the interruption, high availability is lost. This is because current HyperSwap implementations do not support the creation of a FlashCopy mapping when the FlashCopy target volume is a master or auxiliary volume in a HyperSwap or remote copy relationship. Furthermore, current HyperSwap implementations do not support pausing or stopping a HyperSwap relationship. To create a FlashCopy mapping, the HyperSwap volume copy on the auxiliary site must be removed. At this point, the FlashCopy mappings for the change volumes and dual IO routing are also removed.

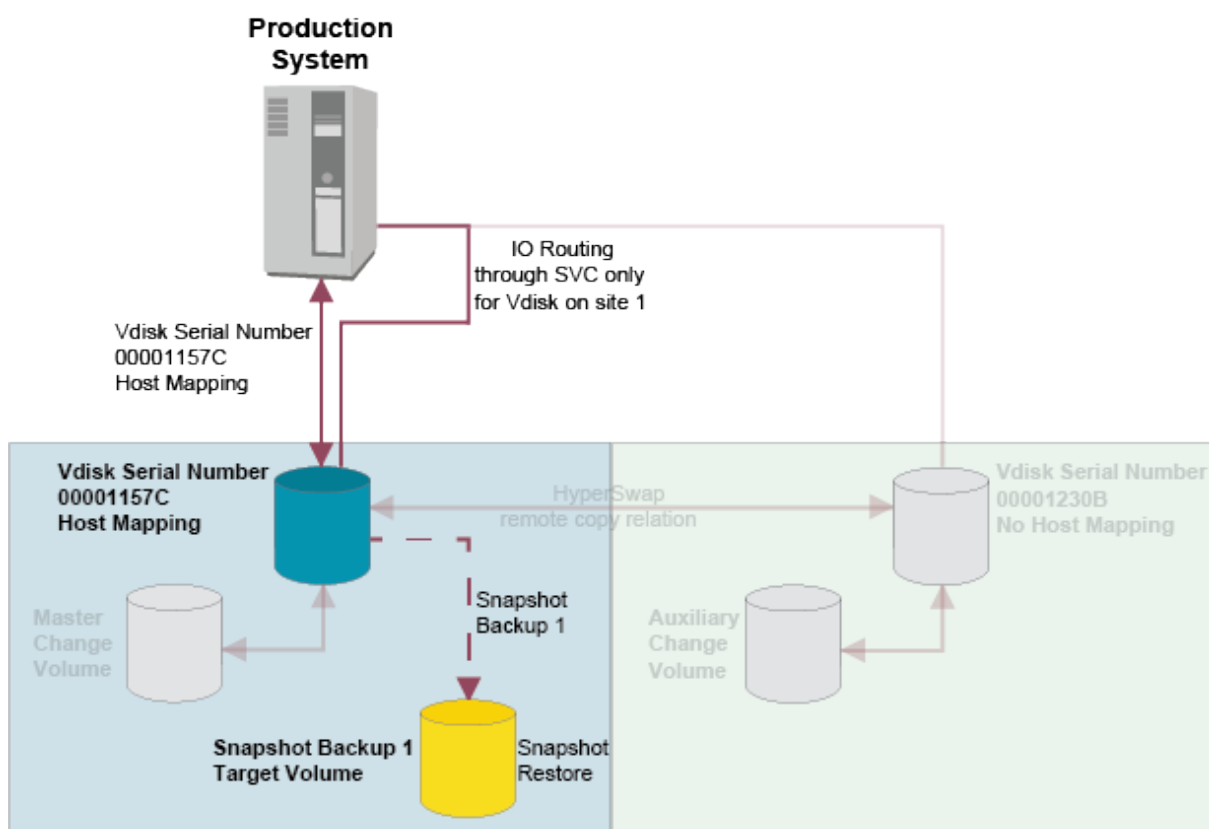


Figure 11: IBM® Storage Protect Snapshot after a HyperSwap restore.

After a Snapshot HyperSwap restore, all HyperSwap settings must be reconfigured to a consistent and synchronized state. Current HyperSwap implementations do not support the creation of a remote copy relationship on master or auxiliary volumes that are the target of a FlashCopy mapping. When synchronization is complete, you can use a dynamically created shell script to automatically reconfigure the required relationships and IO settings. The reconfiguration script is named `<timestamp>_hyperrel.sh`, where `<timestamp>` is the current date and time. Run the reconfiguration script from the `acs` directory that is created when Snapshot is installed.

If successful, the reconfiguration script terminates silently. If for example the volume has copies in two sites using HyperSwap already or if other errors from the storage system are reported, the script will display those.

You can automate the execution of the reconfiguration script after a timeout interval defined by the `HYPERSWAP_RESTORE_TIMEOUT` parameter in the Snapshot configuration file. However, this approach runs the risk of the failure of the reconfiguration operation if the Snapshot restore is late finishing. If the time required to perform a resynch is unpredictable, run the reconfiguration script manually.

Resynchronizing the HyperSwap pools and restoring FlashCopy relationships can take from minutes to hours, depending on the size of the data and the available bandwidth on storage network connections.




Figure 12: IBM® Storage Protect Snapshot after running the reconfiguration script.

Cloning databases

IBM® Storage 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® Storage 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® Storage Protect Snapshot, a cloning process can be started with an online or offline source database. For online IBM® Storage 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® Storage Protect Snapshot cloning process. Each clone target requires a unique database name. IBM® Storage 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.

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® Storage Protect Snapshot

Create a database clone with IBM® Storage 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® Storage 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.

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® Storage Protect Snapshot cloning.

IBM® Storage 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® Storage 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 files for DB2® is: `/db2/P01/acs/preprocessing.ini`. When adding processing configuration files, place each script on a separate line as shown in the following example:

```
/db2/P01/acs/scripts/PreProcessing_stopsap.sh
/db2/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® Storage 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 `db2c01`. This addition requires that the owner of the script to be identical to the user who is 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 db2c01 -C C01 -X /db2/C01/acs/preprocessing.ini
./fcmcli -f postproc_clone -u db2c01 -C C01 -Y /db2/C01/acs/postprocessing.ini
```

If a DB2® SQL script needs extra command-line options, use a shell script as a wrapper with the DB2® SQL commands. IBM® Storage Protect Snapshot requires that the DB2® SQL script use a semicolon delimiter: ;. If another delimiter character is specified, use another shell script wrapper that calls the DB2® SQL script with the correct delimiter character in the command-line entry. This example shows a shell script wrapper that specifies the correct delimiter character:

```
your delimiter character/bin/ksh
DELIMITER=your delimiter character
db2 -td${DELIMITER} -svf PostProcessing_script1.sql
```

Usability states of snapshot backup operations

To view the usability states of a DB2® snapshot backup, enter the following command: `./fcmcli -f inquire_detail`. The output from this command shows a comma-separated list of the current usability states.

```
...
Type      Partition  Backup-ID      TSM Backup-ID  State      DevClass  TargetSet  BackgroundCopy
BytesToBeFlashCopied

#BACKUP NODE0000  20130827214721  -            IN_PROGRESS MIRROR1   1          3.000 GB of 3.000 GB
3.000 GB
UsabilityStates: REMOTELY_MOUNTABLE,REPETITIVELY_RESTOREABLE,SWAP_RESTOREABLE,PHYSICAL_PROTECTION,FULL_COPY,
TAPE_BACKUP_PENDING,TAPE_BACKUP_FAILED
...
```

A snapshot backup generation can have one or more of the following states to indicate how it can or must be used by other IBM® Storage Protect Snapshot functions.

Table 9: Usability states	
Usability state value	Meaning
REMOTELY_MOUNTABLE	Backup data can be mounted from a remote system.
REPETITIVELY_RESTOREABLE	Backup data can be restored. The image can be used multiple times.
DESTRUCTIVELY_RESTOREABLE	Data can be restored. After the restore, other backups and possibly the backup to be restored can potentially be deleted.
SWAP_RESTOREABLE	Restore is possible by using the backup volumes directly rather than copying the data back to the source volumes.
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_RESTOREABLE .
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.

Usability state value	Meaning
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. After it finishes, the usability state is changed to FULL_COPY .
TAPE_BACKUP_PENDING	Indicates that a requested backup to an IBM® Storage Protect server 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, the usability state is changed to TAPE_BACKUP_COMPLETE . 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® Storage Protect Snapshot offload agent. If the backup fails, this state is reset. In a DB2® DPF environment, this state is used to indicate that some partitions are successfully backed up and others are still pending backup to tape.
TAPE_BACKUP_COMPLETE	Indicates that the requested tape backup is finished by the IBM® Storage Protect Snapshot offload agent. In a DB2® DPF environment, the TAPE_BACKUP_COMPLETE state is set when all the partitions of the database are successfully backed up to tape.
TAPE_BACKUP_FAILED	Indicates that the tape backup of the IBM® Storage Protect Snapshot offload agent was not successful and the operation is not retried. In a DB2® DPF environment, the TAPE_BACKUP_FAILED state is set when at least one partition of the database failed during tape backup.
CLONE_DATABASE	Indicates that an IBM® Storage Protect Snapshot cloning operation was started.
RESTORING	Indicates that an IBM® Storage Protect Snapshot restore operation was started.

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® Storage Protect Snapshot. The black arrows indicate IBM® Storage 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® Storage Protect Snapshot backup operation. Depending on the storage system that is used some states differ.

For example, on XIV®, the snapshot backup is immediately restorable and the restore can be repeated multiple times. After successfully processing on DS8000®, and SAN Volume Controller and Storwize family devices, the **BACKGROUND_MONITOR_PENDING** usability state is changed to **FULL_COPY** and **PHYSICAL_PROTECTION** by a monitoring daemon (**acsngen -D**) when aCOPYsnapshot backup was requested.

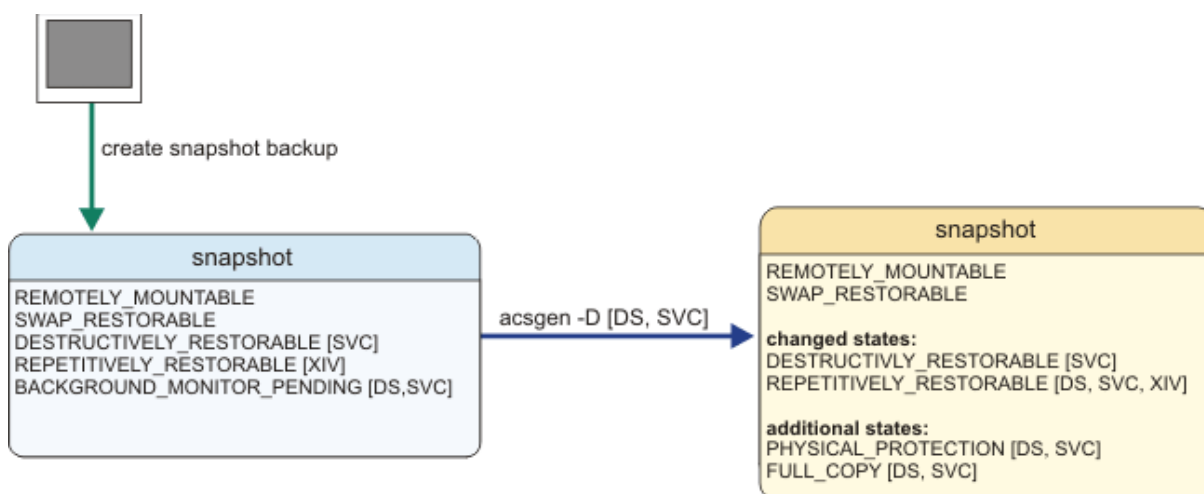


Figure 13: Usability States during snapshot backup

Snapshot restore

The second state diagram shows the usability state changes during an IBM® Storage Protect 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 in a **RESTORING** state. The background monitor daemon (**acsgen -D**) resets both states when the copy process in the storage system finishes.

For XIV® there is no usability state change during restore processing.

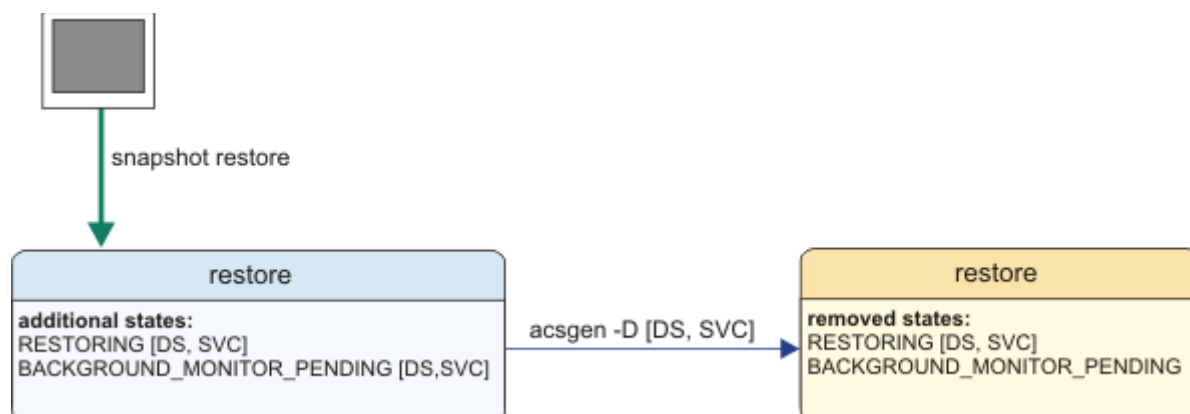


Figure 14: Usability states during snapshot restore

Snapshot delete

The next state diagram shows the usability state changes during an IBM® Storage Protect Snapshot delete operation. There are two types of delete operations: delete and delete with force option. For both types, the snapshot backup is set to the **DELETING** state. After processing completes, the background monitor daemon (**acsgen -D**), switches the states to **DELETED**.

Snapshots on XIV® systems are deleted, and the snapshot backup is removed from the IBM® Storage Protect Snapshot repository by the background monitor daemon.

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

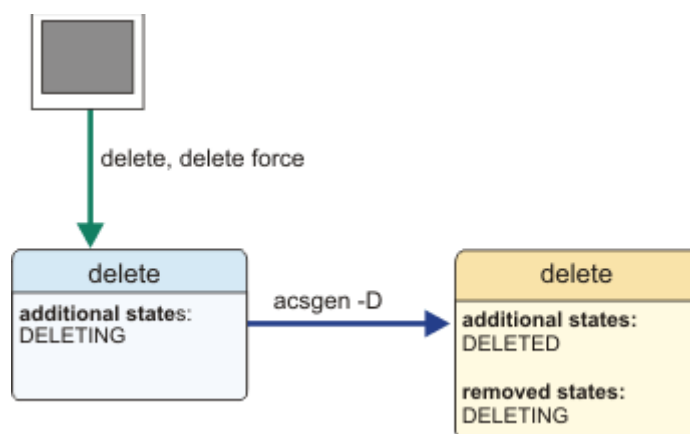


Figure 15: Usability states during snapshot delete

Snapshot mount

The next state diagram shows the usability state changes during an IBM® Storage Protect 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® Storage 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 16: Usability states during snapshot mount

Snapshot offload

The last state diagram shows the usability state change during an IBM® Storage Protect 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 -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® Storage 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.

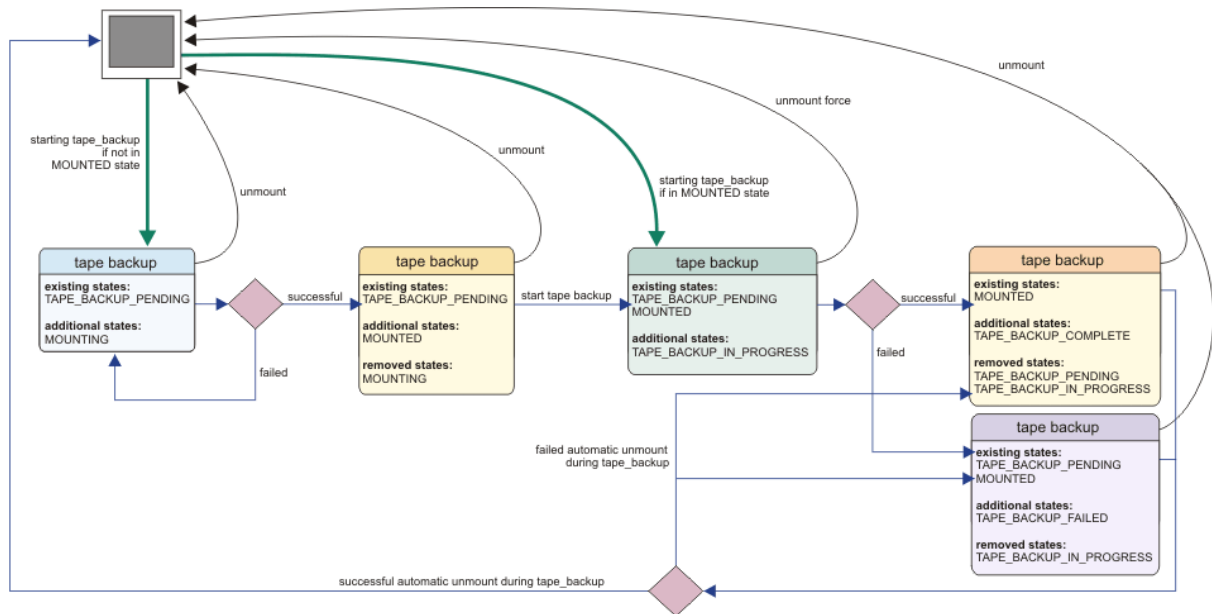


Figure 17: Usability states during snapshot offload

The usability state **TAPE_BACKUP_PENDING** can be removed by using the IBM® Storage Protect Snapshot function **fcmlcli -f 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® Storage Protect Snapshot repository.

Troubleshooting

There are multiple resources for support.

The following list identifies the various ways that you can find information online:

- [IBM® Storage Protect Snapshot wiki](#) on the developerWorks® site.
- [Service Management Connect](#)
- [IBM® Storage 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® Storage 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® Storage 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® Storage Protect Snapshot operations.

Log and trace files are written during backup and restore processing by these products:

- DB2®
- IBM® Storage Protect Snapshot
- Storage system
- CIM
- IBM® Storage Protect for Enterprise Resource Planning
- Operating system

Log files and trace files

Refer to these examples of the log and trace files that are maintained by IBM® Storage Protect Snapshot.

IBM® Storage 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® Storage 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® Storage Protect Snapshot. These files are in ACS_DIR/logs.

The following table identifies the log files that are created for IBM® Storage Protect Snapshot.

Table 11: IBM® Storage 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

Note:

- *timestamp* is the date (yyyymmdd)
- *longtimestamp* is the date and time (yyyymmddHHMMSS)
- *function* is a value of backup, restore, inquire, delete, mount, unmount, tsm, or clone

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.

The following table identifies the trace files that are created for IBM® Storage Protect Snapshot.

Table 12: IBM® Storage Protect Snapshot trace files	
Component	File
Management Agent (acsd)	acsd.id.trace
Application client (for DB2®, the Snapshot Backup Library)	client.instance.db name.node.id.trace
Generic Device Agent (acsgen)	acsgen.hostname.device class.node num.id.trace acsgen.hostname.function.id.trace acsgend.hostname.id.trace
Device Agent for IBM® XIV® Storage System Devices	xivadapter_id_function.trace
Device Agent for CIM Devices (DS8000®, SAN Volume Controller, Storwize® family)	fmcima.hostname.function.id.trace fmcima.hostname.device class.node num.id.trace
Offload Agent (tsm4acs)	tsm4acs.host.id.trace
fcmcli	fcmcli.host.id.trace
RMAN (when started by IBM® Storage Protect Snapshot)	rman.SID.id.log

Notes:

- Names ending in -dare daemon processes (started with -Doption).
- id* 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.
- device class* 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.
- instance* and *db hostname* can be *undef* for query and delete requests that are started with db2acsutil.
- node num* 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.
- function* is backup, delete, restore, mount, unmount, or reconcile.

The following table identifies the return codes that are used by IBM® Storage Protect Snapshot.

Table 13: IBM® Storage Protect Snapshot return codes		
Reason code	Explanation	User response
0	Operation is successful	None
1	Operation terminated successfully with warnings	The IBM® Storage Protect Snapshot operation was successful but warning messages were reported. Check the IBM® Storage Protect Snapshot summary log file and the therein referenced detail log files for more information.
2	Operation terminated with error	The IBM® Storage Protect Snapshot operation failed. Check the IBM® Storage Protect Snapshot summary log file and the therein referenced detail log files for more information.

The following table identifies the installer exit codes that are used by IBM® Storage Protect Snapshot.

Table 14: IBM® Storage 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® Storage 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® Storage Protect Snapshot installer log file installation.log in the installation directory for details.
>=1000	<p>The operation terminated with error</p> <div> <p>Note: There more error codes with numbers greater than or equal to 1000 which all mean that some kind of error occurred.</p> </div>	One or more of the actions from the installation sequence caused a unrecoverable error. See the IBM® Storage Protect Snapshot installer log file installation.log in the installation directory for details.

The following table identifies the reason codes that are used by the DB2® database.

Table 15: DB2® vendor reason codes		
Reason Code	Explanation	User Response
0	The operation is successful.	None
2	Communication error with device	The IBM® Storage Protect Snapshot operation failed. Check the db2diag.log and the IBM® Storage Protect Snapshot summary log file for details.
3	The DB2® and vendor products are incompatible	The IBM® Storage Protect Snapshot operation failed during initialization of the IBM® Storage Protect Snapshot vendor library. The DB2® API version does not match the IBM® Storage Protect Snapshot vendor library version. Check the db2diag.log for details.
6	Object specified cannot be found	The IBM® Storage Protect Snapshot operation failed because the requested object cannot be found in the IBM® Storage Protect Snapshot repository. Check the db2diag.log and the IBM® Storage Protect Snapshot summary log file for details.
8	Invalid user ID specified	The IBM® Storage Protect Snapshot operation failed because an invalid user ID was specified on the db2 command line. Check the db2diag.log.
9	Invalid password provided	The IBM® Storage Protect Snapshot operation failed because an invalid password was specified on the db2 command line. Check the db2diag.log.

Reason Code	Explanation	User Response
10	Invalid options specified	The IBM® Storage Protect Snapshot operation failed because an invalid db2 command-line option was specified. Check the db2diag.log.
11	Initialization failed	The IBM® Storage Protect Snapshot operation failed because the IBM® Storage Protect Snapshot vendor library cannot be initialized. Check the db2diag.log and the IBM® Storage Protect Snapshot summary log file for details.
14	End of data reached	Not an error condition.
18	Device error	The IBM® Storage Protect Snapshot operation failed. Check the IBM® Storage Protect Snapshot summary log file for details.
19	Warning	The IBM® Storage Protect Snapshot operation is successful with warning messages. Check the IBM® Storage Protect Snapshot summary log file for details.
21	More data to come	Not an error condition.
26	Delete object fails	The IBM® Storage Protect Snapshot delete operation failed. Check the IBM® Storage Protect Snapshot summary log file for details.
29	Abort request failed	The IBM® Storage Protect Snapshot abort request failed. Check the IBM® Storage Protect Snapshot summary log file for details.
30	Unexpected Error	The IBM® Storage Protect Snapshot operation failed. Check the IBM® Storage 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® Storage Protect Snapshot operation failed because the object specified for a restore or query is not under the control of IBM® Storage Protect Snapshot. It might be under control of IBM® Storage Protect for Enterprise Resource Planning, for example. Check the db2diag.log and the IBM® Storage Protect Snapshot summary log file for details.
34	Another database or application is using the same storage groups	The IBM® Storage Protect Snapshot snapshot backup operation failed because another database or application is using the same storage group. Check the db2diag.log and the IBM® Storage Protect Snapshot summary log file for details.

Storage system log and trace files

Storage system log and trace files are updated during IBM® Storage 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® Storage 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® family master console produce log and trace output.

GPFS™ log files

IBM® General Parallel File System (GPFS™) log files are updated during IBM® Storage Protect Snapshot, or IBM® Storage Protect Snapshot Custom Applications operations.

The GPFS™ log files are in `/var/adm/ras` directory on each GPFS™ node, and start with the prefix `mmfs.log`. The most current GPFS™ log file can be found by using the symbolic link `/var/adm/ras/mmfs.log.latest`. See the information about GPFS™ log files and troubleshooting procedures in the *IBM® General Parallel File System (GPFS™) for Clusters: Problem Determination Guide (GA76-0415-08)*.

IBM® Storage Protect for Enterprise Resource Planning log and trace files

IBM® Storage Protect for Enterprise Resource Planning 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® Storage Protect for Enterprise Resource Planning *Installation and User's Guide* for details concerning logs and traces within IBM® Storage Protect for Enterprise Resource Planning.

Important: A trace file can be requested by specifying the `TRACEFILE` parameter in the IBM® Storage Protect for Enterprise Resource Planning 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® Storage 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  
starttrconsistgrp consist group id>
```

Question

(XIV® only) One of the following issues exists.

- The remote mirroring is not operational.
- For XIV® synchronous mirroring, the state of the consistency group is not consistent synchronized.
- For XIV® asynchronous mirroring, the state of the consistency group is notRPO_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® Storage 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:

```
# lslv 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 DB2® setup problems

When you are setting up IBM® Storage Protect Snapshot, certain setup scenarios might require extra steps. The solutions outline the steps to follow.

Setting up a backup server for multiple production systems in a DB2® HADR environment

If you use one backup system for multiple production systems, you must use different profiles for each of those production systems. The self-signed certificates must be copied from the primary server to each production system. This setup enables the mount agent to connect to each production system. Every production system in the setup must use the same self-signed certificate.

Problem

The backup system has a DB2® instance with the same name as the DB2® instance on the production systems, including the same instance home directory, and user ID. When you set up a IBM® Storage Protect Snapshot backup system with the DB2® setup script on the secondary DB2® server, the backup system configuration for the primary DB2® server is overwritten. The profile on the backup server points to the `acsd` on the primary DB2® server. For this reason, it cannot connect because the IBM® Storage Protect Snapshot primary DB2® server uses different self-signed certificates than the secondary DB2® server.

Solution

To use the same backup system for the primary and secondary node in a DB2® HADR environment, you must start the mount agent twice with different IBM® Storage Protect Snapshot profiles, one for every production system. Each production system must use the same self-signed certificate. For the steps required, see [“Setting up a backup server for multiple production systems in a DB2 HADR environment” on page 44](#).

Internet Protocol Version 6 (IPv6) support

The IBM® Storage 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® Storage 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 configuring the product, 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.

Configuration files

When you complete the setup script, the information you enter is used to configure IBM® Storage Protect Snapshot.

IBM® Storage Protect Snapshot uses the following configuration files:

- Profile
- Target volumes file
- Password file

Profile

When you complete the setup script, the information you enter is used to create the profile. Do not edit the profile manually, run the setup script to make any amendments to the profile. Each section of the profile includes parameters and options that determine how the IBM® Storage Protect Snapshot backs up and restores data in your environment. The following information explains the various parameters and options.

When you use IBM® Storage Protect Snapshot executable commands, the profile is identified by the value that is specified for option -p.

For DB2® databases, the profile is identified by the `PROFILE` vendor option. The elements of the profile are not case-sensitive. Section and parameter names are indicated with uppercase letters.

The profile is divided into the following sections:

- GLOBAL
- ACSD
- CLIENT
- `DEVICE_CLASS device`
- OFFLOAD
- OFFLOAD_DB2STANDBY
- CLONING
- DB2STANDBY

In some cases, there are 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 all backup and clone systems.

To overwrite IBM® Storage Protect Snapshot profile parameters for DB2®, use vendor options.

GLOBAL

The GLOBAL section contains information that is required and used by all IBM® Storage Protect Snapshot components. The components reference the information in the GLOBAL section during the startup process. Changes to the Global section require a restart of IBM® Storage Protect Snapshot.

IBM® Storage Protect Snapshot can be installed on multiple systems within a database 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® Storage Protect server. When IBM® Storage Protect Snapshot is installed on multiple systems within a database environment, there is only one active management agent. The location of this management agent is specified in GLOBAL section by using the **ACSD** parameter.

Other parameters in the GLOBAL section specify the location for logging, tracing, and password files. On the backup and clone servers, 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*. The ACSD 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® Storage Protect Snapshot data. If the repository is lost, any previously created backup cannot be restored.

CLIENT

The CLIENT section contains all parameters that relate to back up operations, including parameters for database applications, whether an IBM® Storage 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 device. A DEVICE_CLASS section describes the characteristics of a storage device 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. At least one DEVICE_CLASS section is required for the configuration of the management agent.

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. The same DEVICE_CLASS value cannot be specified in the CLIENT and CLONING sections at the same time.

During backup, the value of the DEVICE_CLASS parameter that is used is recorded in the IBM® Storage Protect Snapshot repository. The same DEVICE_CLASS must be used when you are restoring the backup. 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. This password is used by IBM® Storage Protect Snapshot to authenticate to the management interface of the storage device that is represented by the DEVICE_CLASS section. You can specify the password during configuration with the setup script, or by using the following **fcmcli** command:

```
fcmcli -f password
```

OFFLOAD

The parameters and options in the OFFLOAD section determine how a snapshot is transferred to IBM® Storage Protect server. The information is sent to the offload agent, (*fcmcli -D*).

To configure IBM® Storage Protect Snapshot for protection of both active and standby nodes of a DB2® high-availability disaster recovery (HADR) environment, have two OFFLOAD sections in the profile: OFFLOAD and OFFLOAD_DB2STANDBY.

When the offload agent is started, it connects to the management agent and queries for snapshot backups that are backed up with the **TSM_BACKUP** parameter that is set to YES. If this parameter and option is found, the offload agent mounts the snapshot on a backup server and initiates an IBM® Storage Protect backup by using one of the following applications:

- IBM®DB2® Storage Protect client for DB2® environments
- IBM® Storage Protect for Enterprise Resource Planning

OFFLOAD_DB2STANDBY

The OFFLOAD_DB2STANDBY section uses the same semantics as the OFFLOAD section, except that the section is evaluated when an offloaded backup from a DB2® HADR secondary occurs. In most environments, the profile also includes an OFFLOAD section to protect databases when the DB2® HADR primary and secondary nodes switch.

CLONING

The CLONING section contains the parameters that are used for cloning operations. The section is ignored for all other operations.

DB2STANDBY

The DB2STANDBY section is used when you are backing up the standby node of a DB2® HADR environment. Although IBM® Storage Protect Snapshot also evaluates all parameters that are specified in the CLIENT section during HADR operations, some parameters that are specified in the DB2STANDBY section override the corresponding parameters in the CLIENT section.

Example

The following example is a sample profile for DB2® including cloning and backup of a standby node of a DB2® HADR environment.

```
>>> GLOBAL
ACS_DIR /db2/AS1/acs
ACSD auntmathilda 54331
# TRACE NO
<<<

>>> ACSD
ACS_REPOSITORY /db2/AS1/acs/acsrepository
# SYNCHRONOUS_RECONCILE RESTORE_AND_DELETE
# REPOSITORY_LABEL TSM
<<<

>>> CLIENT
# BACKUPIDPREFIX DB2_---
APPLICATION_TYPE SAP
# PARTITION_GROUP
TSM_BACKUP YES
# MAX_VERSIONS ADAPTIVE
# LVM_FREEZE_THAW AUTO
# NEGATIVE_LIST no_check
# TIMEOUT_FLASH 120
# GLOBAL_SYSTEM_IDENTIFIER
# DEVICE_CLASS STANDARD
<<<

>>> OFFLOAD
BACKUP_METHOD DB2
OPTIONS @/db2/AS1/dbs/vendor.env.%DB2NODE
# PARALLELISM AUTO
# NUM_SESSIONS 1
# NUM_BUFFERS AUTO
# BUFFER_SIZE AUTO
<<<

>>> DB2STANDBY
DB2_PRODUCTIVE_SERVER akuma DB2_db2as1_svc
DB2_ALIAS AS1
PRE_FLASH_CMD /db2/AS1/acs/db2_pre_flash_script.cmd
POST_FLASH_CMD /db2/AS1/acs/db2_post_flash_script.cmd
DB2_USERNAME db2as1
DB2_AUTH_TYPE SERVER
<<<

>>> OFFLOAD_DB2STANDBY
BACKUP_METHOD TSM_CLIENT
# MODE FULL
ASNODENAME AS1TEST
# DSM_DIR
# DSM_CONFIG
# VIRTUALFSNAME fcm
<<<

>>> CLONING
DEVICE_CLASS CLONE1 USE_FOR_CLONING AS2
<<<

>>> DEVICE_CLASS STANDARD
COPYSERVICES_HARDWARE_TYPE DS8000
COPYSERVICES_PRIMARY_SERVERNAME uncletitus
# COPYSERVICES_SECONDARY_SERVERNAME
COPYSERVICES_USERNAME tsmuser
# COPYSERVICES_COMMPROTOCOL HTTPS
# COPYSERVICES_CERTIFICATEFILE NO_CERTIFICATE
# COPYSERVICES_SERVERPORT 5989
FLASHCOPY_TYPE NOCOPY
# COPYSERVICES_TIMEOUT 6
# RESTORE_FORCE NO
# LVM_MIRRORING NO
TARGET_SETS VOLUMES_FILE
VOLUMES_FILE /db2/AS1/acs/acsvolumes/fct.file
<<<

>>> DEVICE_CLASS CLONE1
COPYSERVICES_HARDWARE_TYPE DS8000
```

```

COPYSERVICES_PRIMARY_SERVERNAME uncletitus
# COPYSERVICES_SECONDARY_SERVERNAME
COPYSERVICES_USERNAME tsmuser
# COPYSERVICES_COMMPROTOCOL HTTPS
# COPYSERVICES_CERTIFICATEFILE NO_CERTIFICATE
# COPYSERVICES_SERVERPORT 5989
# FLASHCOPY_TYPE COPY
# COPYSERVICES_TIMEOUT 6
# RESTORE_FORCE NO
# LVM_MIRRORING NO
TARGET_SETS VOLUMES_FILE
VOLUMES_FILE /db2/AS1/acs/acsvolumes/fct.file.cloning
CLONE_DATABASE YES
<<<

```

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® Storage Protect Snapshot and the location of log files. The parameters are independent of the storage solution, and database 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® Storage 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® Storage Protect Snapshot.

shared

The subdirectory contains information that is shared among all IBM® Storage 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 the setup script, the information in these directories is promoted to all systems that belong to the instance where IBM® Storage Protect Snapshot is configured. When you run configuration tasks separately, you must promote the information manually.

Default

`<instance_owner_$HOME>/acs`

Advanced mode only

Yes

ACSD

The host name of the system where the management agent is running and the port it is listening to. The following format is used for **ACSD**: `<hostname> <port>`

This parameter must be identical on all systems where IBM® Storage 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

`localhost 57328`

Advanced mode only

Yes

ENFORCE_TLS12

IBM® Storage 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.

Only use this parameter when advised to do so by IBM® Support.

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 denotes the directory of the IBM® Storage Protect Snapshot repository. This directory is used for all operations, and must be in a secure location. If the repository is lost, all backups become unavailable to IBM® Storage Protect Snapshot even if they remain in the storage device.

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® Storage Protect Snapshot reports a failure. The IBM® Storage Protect Snapshot repository cannot be in the main IBM® Storage 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® Storage 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, an error is reported.

Default

<ACS_DIR>/acsrepository.

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® Storage Protect Snapshot to synchronously reconcile and delete snapshot backups. The following options are possible for this parameter.

NO

Specify this option when you do not want 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. This process is useful for storage systems that delete snapshot backups during an IBM® Storage Protect Snapshot backup or cloning operation. SAN Volume Controller and Storwize® family storage systems can delete backups during a restore operation.

RESTORE_AND_DELETE

Use the **RESTORE_AND_DELETE** option to start a synchronous delete and reconcile process as part of a restore and delete operation. This option is useful for storage systems that can delete snapshot backups during an IBM® Storage Protect Snapshot restore process. For example, the Storwize® family and SAN Volume Controller storage systems can delete backups during a restore and delete operation. The **RESTORE_AND_DELETE** option is also useful if you manually delete snapshot backups and use the force option (**-f**) on DS8000®, SAN Volume Controller, or Storwize® family storage systems.

Default

The default for this parameter is **RESTORE_AND_DELETE**.

Advanced mode only

YES

CLIENT

IBM® Storage Protect Snapshot uses DB2® specific parameters to configure DB2® backup and restore operations. These parameters are defined in the client section of the IBM® Storage Protect Snapshot profile configuration file.

The ability to create an IBM® Storage Protect backup from a snapshot requires an IBM® Storage Protect Snapshot license.

APPLICATION_TYPE

This parameter specifies the environment. The following list identifies the possible options:

DB2®

Use as a DB2® database.

SAP

Use as a DB2® in an SAP environment database.

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>]  
(partitioned DB2 databases) [ON_DBPARTITIONNUMS <list of partitions>]
```

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® Storage 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® Storage 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 snapshot 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® Storage Protect Snapshot processes all symbolic links of files or directories.

NO

Use this option to ensure that IBM® Storage Protect Snapshot does not process symbolic links of files or directories.

Default

YES

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® Storage 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 set to **YES**, then this parameter is set with the following option: **LVM_FREEZE_THAWYES**. If the file system does not support freeze actions, the **AUTO** value sets the parameter to **LVM_FREEZE_THAWNO**.

For more information, see [“Interdependency of LVM_FREEZE_THAW and TARGET_DATABASE_SUSPEND” on page 132](#).

Default

AUTO

Advanced mode only

Yes

MAX_VERSIONS

This parameter is specified for each device class and specifies the number of snapshot versions to store for each device class. The following list identifies the possible options:

```
MAX_VERSIONS ADAPTIVE
MAX_VERSIONS 2
MAX_VERSIONS 3 USE_FOR DC_TEST
```

ADAPTIVE

The maximum number varies depending on the available space. IBM® Storage 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 per each device class. 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.

Optional

When you use `valuen` you can add a device class after it `USE_FOR <device class>` to specify that **MAX_VERSIONS_n** is valid for that device class only.

Default

None

Advanced mode only

No

When you add another **MAX_VERSIONS** parameter, specify values based on the following criteria:

- **MAX_VERSIONS** with a specific value for a specific device class.
- **MAX_VERSIONS** with a default value for all device classes that have no **MAX_VERSIONS** already specified.
- **MAX_VERSIONS** with an `adaptive` value. This value must be used only when there are no other values set for any **MAX_VERSIONS** for any device classes.

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 snapshot 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 snapshot 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® Storage 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

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 snapshot 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 snapshot restore:

- **PREPARERESTORE**
- **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® Storage Protect backup from a snapshot. IBM® Storage 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® Storage Protect backup is created from the snapshot. The following list identifies the possible options:

YES

This option creates an IBM® Storage Protect backup from a snapshot. If the IBM® Storage Protect backup operation does not successfully complete, the target set can be reused.

MANDATE

This option creates an IBM® Storage Protect backup from a snapshot. However, the target set cannot be reused until the IBM® Storage Protect backup successfully completes.

LATEST

This option removes a backup request to IBM® Storage Protect from a previous backup. When a new snapshot with **TSM_BACKUP** set to **LATEST**, **YES**, or **MANDATE** is created, IBM® Storage 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® Storage 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® Storage 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® Storage 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

CLONING

The **CLONING** section of the IBM® Storage 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 DB2® databases:

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> [conditions]
```

There is an optional *conditions* statement that can be used. The following code sample includes an example of how to use the *conditions* statement.

The condition statement is optional. When you use the condition statement, use the following syntax:

```
[USE_AT <days of week>] [FROM <time> TO <time>]  
[ON_DBPARTITIONNUMS <list of partitions>]
```

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.

CLONE_DATABASE_MEMORY

Use this parameter to change the amount of initial shared memory that is allocated by the database on the clone server. The value that is specified is used to modify the DB2 database shared memory size configuration parameter (DATABASE_MEMORY) on the clone server. The value must be a number, and documented limits for the database configuration parameter DATABASE_MEMORY apply. If the CLONE_DATABASE_MEMORY parameter is set to 0, the DATABASE_MEMORY value set for the productive database is used for the clone database.

Default

The same amount of initial shared memory that is allocated by DB2 on the production server.

Advanced mode only

Yes

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. IBM® Storage 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 snapshot 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® Storage Protect Snapshot profile file. The following list identifies the possible options:

YES

Use this option to ensure that IBM® Storage 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® Storage 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 snapshot operation.

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® Storage Protect Snapshot includes the complete volume groups where the specified files or directories are located in the snapshot.

In DB2® Data Partitioning Feature environments that contain multiple partitions, a different file name for each partition can be specified. In this scenario, because different file names are specified for different partitions, specify the **ON DBPARTITIONNUM** parameter. Each partition number is separated by a space.

The default value is an empty list. This value prevents extra volumes from participating in the snapshot 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 DB2®, the **DATABASE_SCHEMA** parameter is required. The default database schema is determined by the `db6_schema` environment variable. These environment variables are set on the production database instance owner environment. If the environment variable is not set, the default database schema value is `SAPR3`. 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 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 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 snapshot 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 snapshot 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 snapshot 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.

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 cloning operation ends and an error message is displayed. The default value is 3600 seconds.

You can specify one of these phase values for a snapshot backup. For example: **TIMEOUT_PREPARE**

- **PARTITION**
- **PREPARE**
- **VERIFY**
- **CLOSE**

Advanced mode only

Yes.

DEVICE_CLASS *device*

The IBM® Storage Protect Snapshot profile can contain one or more **DEVICE_CLASS** sections. The device class section configures IBM® Storage Protect Snapshot for use with a particular storage or file system solution. The parameters do not depend on the database 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.

For more information about setting different device class **MAX_VERSIONS**, see [Backup version retention](#).

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 values 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.

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_USERNAME** parameter cannot be used.

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 IBM® XIV® Storage System Storage System parameters

The parameters that are defined in the device class section of the IBM® Storage Protect Snapshot profile, configure IBM® Storage 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 IBM® XIV® 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 indicates whether the device class is used for cloning. 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.

The following example shows the **CLONE_DATABASE** parameter that is specified in the **DEVICE_CLASS** section for **DEVICE_CLASS STANDARD**.

```
>>> 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

NO

Advanced mode only

No.

COPYSERVICES_HARDWARE_TYPE

This parameter is required. Only one device can be specified.

XIV

Specify theXIVoption, when the database is stored on the IBM® XIV® 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 IBM® XIV® 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 to log on to the IBM® XIV® Storage System.

Default

superuser

Advanced mode only

No.

LVM_MIRRORING

Set this parameter toYESif 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 between0and24inclusive.

Default

12

Advanced mode only

Yes.

PATH_TO_XCLI

This parameter specifies the path where the IBM XIV® command-line interface, *XCLI*, is installed. There is no default value. This parameter is only valid when **COPYSERVICES_HARDWARE_TYPE** specifiesXIV.

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

The parameters that are defined in the device class section of the profile file, configure IBM® Storage Protect Snapshot for UNIX and Linux for use with the IBM® Storwize® family or IBM® System Storage® SAN Volume Controller storage systems.

When you configure, you have a choice of Storwize® family and SAN Volume Controller device types. Depending on which device type you select, the parameter values that are required vary.

Specify **SVC** or **DTA** when prompted by the setup script with a choice of storage system types (**COPYSERVICES_HARDWARE_TYPE**). You can select one of the following device types:

Storwize® family and SAN Volume Controller dynamic target allocation (SVC DTA)

IBM® Storage Protect Snapshot for UNIX and Linux dynamically allocates target volumes on the storage system during the backup process.

Storwize® family and SAN Volume Controller static target allocation (SVC)

Before you start the backup process, you must manually create target volumes on the storage system. Also, predefined volumes must be defined in an IBM® Storage Protect Snapshot configuration file or must match a specific naming pattern.

DEVICE_CLASS parameters for static target allocation

The device class parameters for static target allocation are defined in the **DEVICE_CLASS** section of the IBM® Storage Protect Snapshot profile. These parameters configure IBM® Storage Protect Snapshot to use static target allocation with the IBM® Storwize® family or IBM® System Storage® SAN Volume Controller storage systems.

CLONE_DATABASE

This parameter indicates whether the device class is used for cloning. 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.

Default

No

Advanced mode only

No

COPYSERVICES_HARDWARE_TYPE

This parameter is required. Only one device can be specified.

SVC

Specify the *SVCoption*, when the database is stored on either the SAN Volume Controller or the Storwize® family 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® family, use the *Storwize V7000 user* name that you use to log on to the Storwize® family.

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 IBM® Storwize® family or IBM® System Storage® SAN Volume Controller requires communication in secure mode. For this scenario, clients such as IBM® Storage 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® Storage Protect Snapshot profile parameter **COPYSERVICES_COMMPROTOCOL** must specify HTTPS. For this scenario, the **COPYSERVICES_CERTIFICATEFILE** parameter can define a certificate file name, and IBM® Storage 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 system, backup or clone 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. This host manages the SAN Volume Controller master console or the embedded CIM Agent in the Storwize® family 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® family, the **COPYSERVICES_PRIMARY_SERVERNAME** parameter must point to the Storwize® family cluster.

Default

localhost

Advanced mode only

No

COPYSERVICES_SERVERPORT

This parameter identifies the server port number on the CIM Agent. 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® family 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® Storage 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 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.

Default

NO

Advanced mode only

Yes

TARGET_SETS

This parameter indicates how target volumes are specified. The following options are available:

VOLUMES_FILE

This parameter specifies that a file is used to specify the target volumes. The name of the file must be specified in the **VOLUMES_FILE** parameter.

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® Storage 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

Specify **VOLUMES_FILE** if the target sets are passed in a target volumes file (.fct). Its actual name must be given in parameter **VOLUMES_FILE**. Specify the fully qualified file name.

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® Storage Protect Snapshot clone on space-efficient targets. For this device class, use space-efficient targets and set **FLASHCOPY_TYPE** toNOCOPY. FlashCopy® backups cannot be stored on the same source volumes.

NO

Do not create an IBM® Storage 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 toNO.

Default

NO

Advanced mode only

Yes

ALLOW_ALL_FLASHCOPY_TYPES

Use this parameter when IBM® Storage Protect Snapshot is configured with **FLASHCOPY_TYPE**COPY, 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® Storage Protect Snapshot to be configured to use **FLASHCOPY_TYPE**COPY, or **FLASHCOPY_TYPE**INCRwhen 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** toNOCOPYonly.

Default

NO

Advanced mode only

Yes

SVC_CLEAN_RATE

This parameter specifies the cleaning rate for the FlashCopy® mapping. A value from 0 to 150 can be entered.

Default

None

Advanced mode only

Yes

SVC_COPY_RATE

This parameter specifies the priority that the SAN Volume Controller or Storwize® family gives to the FlashCopy® background process for the current backup or restore. A value from 0 to 150 can be entered.

A value of 150 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® family. 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 device class parameters for dynamic target allocation are defined in the **DEVICE_CLASS** section of the IBM® Storage Protect Snapshot profile. These parameters configure IBM® Storage Protect Snapshot to use dynamic target allocation with IBM® Storwize® family or IBM® System Storage® SAN Volume Controller storage systems.

CLONE_DATABASE

This parameter indicates whether the device class is used for cloning. 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.

Default

No

Advanced mode only

No

COPYSERVICES_HARDWARE_TYPE

This parameter is required. Only one device can be specified.

GENERIC

Specify the **GENERIC** option when the storage system is SAN Volume Controller or Storwize® family 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® family, specify the Storwize® family 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 that is specified for the **COPYSERVICES_USERNAME** parameter.

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 that is 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 0 to 150.

The **SVC_COPY_RATE** parameter applies only for full copy backups (FLASHCOPY_TYPE COPY). For space-efficient backups (FLASHCOPY_TYPE NOCOPY), the copy rate is implicitly set to 0.

A value of 150 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

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.

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® Storage 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.

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® family. 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 SVCDTA 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_REMOTEYES, 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 that is used to create target volumes.

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 storage system that is used to create target volumes.

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 applies only 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® Storage Protect Snapshot profile, configure IBM® Storage 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 indicates whether the device class is used for cloning. 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.

Default

No

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 any supported IBM DS8000 storage device.

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® Storage 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® Storage Protect Snapshot profile parameter **COPYSERVICES_COMMPROTOCOL** must specify HTTPS. For this scenario, the **COPYSERVICES_CERTIFICATEFILE** parameter can define a certificate file name, and IBM® Storage 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 should not be used unless 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 theCOPYOption but theINCRoption 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® Storage Protect backups are taken from disk copies. This type of backup creates less burden on the storage system than for theCOPYOption.
- 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® Storage Protect server is possible even if the parameter is set toNOCOPY.

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 toYESin the profile 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** toYESfor a specific restore, you can create a temporary profile.

Default

NO

Advanced mode only

Yes

TARGET_SETS

This parameter indicates how target volumes are specified. The following list identifies the possible options:

VOLUMES_FILE

Specify **VOLUMES_FILE** if the **BACKUP_HOSTNAME** is set toPREASSIGNED_VOLUMES. The actual file name of the target volumes file (.fct) is specified in parameter **VOLUMES_FILE**.

Default

None.

Advanced mode only

No.

VOLUMES_FILE

This parameter specifies the name of the target volumes file (.fct). Specify the fully qualified file name.

Default

None.

Advanced mode only

No.

OFFLOAD

The OFFLOAD section of the profile configuration contains information that is related to IBM® Storage 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 DB2®, and DB2® 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 DB2®, and DB2® 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.

- DB2® and DB2® in an SAP environment -DB2
- DB2® standby server -TSM_CLIENT

Default

Preset by the setup script, according to the environment.

Advanced mode only

Yes.

OPTIONS

A file specification must be a fully qualified file name. If IBM® Storage Protect for Enterprise Resource Planning is being used, the IBM® Storage Protect for Enterprise Resource Planning DB2® vendor options file, `vendor.env`, must be specified.

To be able to set up individual partitions in a DB2® environment in a different manner, the placeholder string `%DB2NODE` can be embedded in the options string. At run time, the string is replaced with the appropriate partition number for which the backup was run. This placeholder can be part of the vendor options file entry, and depending on the partition, allows different configuration files. For example, if there are two partitions as follows:

```
OPTIONS @/db2/T01/tdpr3/vendor_%DB2NODE.env
```

The following two files are referred:

```
/db2/T01/tdpr3/vendor_0.env  
/db2/T01/tdpr3/vendor_1.env
```

The first file is used for partition 0, the second for partition 1. Specifying this parameter overrides the value that is specified by the **VENDOROPT** database configuration parameter.

options string

Specifies options to be used for this IBM® Storage Protect backup operation. The string is passed directly to the backup utility.

@filename

Specifies that the options to be used for the IBM® Storage Protect backup operation are contained in a file that is on the backup server. The string is passed directly to the backup utility.

Default

The default value is an empty string.

Advanced mode only

No.

DB2_OPTIONS

This parameter allows certain extra DB2® backup command options to be set during an offload backup to IBM® Storage Protect server. The options string that is specified by the **DB2_OPTIONS** parameter is passed directly to the DB2® backup command.

Note: If you use **DB2_OPTIONS** in combination with **NUM_BUFFERS**, **BUFFER_SIZE** or **PARALLELISM**, errors that relate to incorrect order of options are sent for the DB2® backup operation. Resolve the errors by not using the **NUM_BUFFERS**, **BUFFER_SIZE**, and **PARALLELISM** parameters with the **DEDUP_DEVICE** option in the **DB2_OPTIONS** string. If this cannot be avoided, then specify the corresponding DB2® options **WITH num-buffers BUFFERS, BUFFER buffer-size**, and **PARALLELISM n** with the **DB2_OPTIONS** parameter and the **DEDUP_DEVICE** option. Ensure to use the correct order of options that is defined in the DB2® documentation for the DB2® BACKUP command.

If multiple options must be specified in **DB2_OPTIONS** parameter string, you must enclose the complete options string in single or double quotation marks. The string (without the quotation marks) is passed to the DB2® backup command. The following example shows you how to define more than one option:

```
DB2_OPTIONS "DEDUP_DEVICE UTIL_IMPACT_PRIORITY 50"
```

<options string>

Specifies options to be used for this IBM® Storage Protect backup operation. The string is passed directly to the DB2® backup command.

Default

The default value is an empty string.

Advanced mode only

Yes.

DBPARTITIONNUM

This parameter can be overridden by the **tsm4acs** command option -N.

ALL

All partitions

partition list

Comma-separated list of partition numbers

Default

ALL

Advanced mode only

Yes.

PARALLELISM

Note: If you use **DB2_OPTIONS** in combination with **NUM_BUFFERS**, **BUFFER_SIZE** or **PARALLELISM**, errors that relate to incorrect order of options are sent for the DB2® backup operation. Resolve the errors by not using the **NUM_BUFFERS**, **BUFFER_SIZE**, and **PARALLELISM** parameters with the **DEDUP_DEVICE** option in the **DB2_OPTIONS** string. If this cannot be avoided, then specify the corresponding DB2® options **WITH num-buffers BUFFERS, BUFFER buffer-size**, and **PARALLELISM n** with the **DB2_OPTIONS** parameter and the **DEDUP_DEVICE** option. Ensure to use the correct order of options that is defined in the DB2® documentation for the DB2® BACKUP command.

The following list identifies the options for this parameter:

n

Number of table spaces that can be read in parallel by the backup utility.

AUTO

DB2® calculates an optimum value.

Default

AUTO

Advanced mode only

Yes.

NUM_SESSIONS

The following list identifies the options for this parameter:

n

Number of I/O sessions to be created between DB2® and the IBM® Storage Protect server.

Default

1

Advanced mode only

No.

NUM_BUFFERS

The following list identifies the options for this parameter:

Note: If you use **DB2_OPTIONS** in combination with **NUM_BUFFERS**, **BUFFER_SIZE** or **PARALLELISM**, errors that relate to incorrect order of options are sent for the DB2® backup operation. Resolve the errors by not using the **NUM_BUFFERS**, **BUFFER_SIZE**, and **PARALLELISM** parameters with the **DEDUP_DEVICE** option in the **DB2_OPTIONS** string. If this cannot be avoided, then specify the corresponding DB2® options **WITH num-buffers BUFFERS, BUFFER buffer-size**, and **PARALLELISM n** with the **DB2_OPTIONS** parameter and the **DEDUP_DEVICE** option. Ensure to use the correct order of options that is defined in the DB2® documentation for the DB2® BACKUP command.

n

The number of buffers to be used by DB2®

AUTO

The DB2® software calculates the optimum value for this parameter.

Default

AUTO

Advanced mode only

Yes.

BUFFER_SIZE

The following list identifies the options for this parameter:

Note: If you use **DB2_OPTIONS** in combination with **NUM_BUFFERS**, **BUFFER_SIZE** or **PARALLELISM**, errors that relate to incorrect order of options are sent for the DB2® backup operation. Resolve the errors by not using the **NUM_BUFFERS**, **BUFFER_SIZE**, and **PARALLELISM** parameters with the **DEDUP_DEVICE** option in the **DB2_OPTIONS** string. If this cannot be avoided, then specify the corresponding DB2® options **WITH num-buffers BUFFERS, BUFFER buffer-size**, and **PARALLELISM n** with the **DB2_OPTIONS** parameter and the **DEDUP_DEVICE** option. Ensure to use the correct order of options that is defined in the DB2® documentation for the DB2® BACKUP command.

n

The value of this parameter specifies the size, in 4 KB pages, of the buffer that is used by the DB2® software when the backup image is built. The minimum value is 8 pages.

AUTO

If the backup is automatically started, the DB2® software calculates the optimum value.

Default

AUTO

Advanced mode only

Yes.

PARALLEL_BACKUP

The following list identifies the options for this parameter:

YES

The IBM® Storage Protect backup of all participating partitions is run in parallel.

NO

The IBM® Storage Protect backups of all participating partitions are sequentially run. If **NUMBER_BACKUPS_IN_PARALLEL** is set to a value greater than 0, then the specified number of partitions is backed up in parallel. For more information about parallel backups, see the description of **NUMBER_BACKUPS_IN_PARALLEL**.

Default

No

Advanced mode only

Yes.

DATABASE_MEMORY

For IBM® Storage Protect Snapshot software to correctly protect DB2® database data, the DB2® database must be started on the backup system during the offload of this database to the IBM® Storage Protect server. To start this database on the backup system, the DB2® software needs the database memory size that is specified on the DB2® database on the production system.

This parameter specifies the size of DB2® database shared memory on the backup system. By specifying an empty string or 0, the DB2® memory size that is specified in the DB2® database configuration on the production system is used.

Default

0

Advanced mode only

Yes.

NUMBER_BACKUPS_IN_PARALLEL

In DB2® DPF environments with more than one DB2® partition, start the offloaded IBM® Storage Protect backup in parallel for multiple DB2® partitions. In large DB2® DPF environments with more than one backup system, the performance of the overall offloaded IBM® Storage Protect backup can be increased if the IBM® Storage Protect backups of each DB2® partition can be started in parallel on each of the backup systems. This profile parameter specifies the degree of parallelism to use during the offloaded backup. If you specify a positive integer value for this parameter, the specified number of IBM® Storage Protect backups start in parallel on each of the backup systems. For example, if a value of 4 is specified, four backups are started in parallel on each backup system. If the parameter value is 0, no parallelism is used. If you want to use a value greater than 0 for **NUMBER_BACKUPS_IN_PARALLEL**, set the **PARALLEL_BACKUP** value to NO.

Default

0

Advanced mode only

Yes.

VENDOR_LIB

DEFAULT

Use the library or agent corresponding to the database environment.

DB2_TSM_AGENT

Use the DB2® IBM® Storage Protect agent.

TSM4ERP

Use IBM® Storage Protect for Enterprise Resource Planning.

library

Use the fully qualified name for the custom library.

Default

DEFAULT

Advanced mode only

Yes.

DB2STANDBY

The DB2STANDBY section is the same as the CLIENT section, except when it is configured for a DB2® database that acts as a DB2® HADR secondary. After takeover, when the database is active, the CLIENT section is used. The parameters do not depend on the storage device.

The following list provides the applicable parameters, a description of each parameter, and default values:

PRE_FLASH_CMD

This parameter identifies the command script, or executable file, used to quiesce the DB2® standby server or DB2® High-Availability Disaster Recovery (HADR) secondary server. This quiescing occurs immediately before the snapshot operation begins.

When the value of this parameter includes command arguments, place the values for the command arguments between quotation marks. For example: "

This parameter is required for the profile, or it needs to be specified from the command-line interface. If specified from the command-line interface, the parameter overrides the corresponding parameter in the profile.

Default

There is no default value.

Advanced mode only

No.

POST_FLASH_CMD

This parameter identifies the command script, or executable file, used to resume the DB2® standby server or DB2® High-Availability Disaster Recovery (HADR) secondary server immediately after the snapshot is created.

When the value of this parameter includes command arguments, place the values for the command arguments between quotation marks. For example: "

This parameter is required for the profile, or it needs to be specified from the command-line interface. If specified from the command-line interface, the parameter overrides the corresponding parameter in the profile.

Default

There is no default value.

Advanced mode only

No.

DB2_PRODUCTION_SERVER

This parameter contains the following two values:

host name or TCP/IP name

Specify either the host name or TCP/IP name of the DB2® server where the HADR primary server, the production system, is running.

TCP/IP port

The TCP/IP port where the DB2® production database instance is listening for remote connections (DB2® database manager configuration parameter **SVCENAME**).

Both value are separated by a space and both values are required to be specified.

Default

There is no default value for this required parameter.

Advanced mode only

No.

DB2_ALIAS

Specify the alias name of the DB2® database running on the HADR primary server.

Default

There is no default value for this required parameter.

Advanced mode only

No.

DB2_USERNAME

This parameter specifies the DB2® user that is used to connect from the HADR standby server to the DB2® database running on the HADR primary server.

Default

The user name of the user who starts the *fcmbli* operation.

Advanced mode only

Yes.

DB2_AUTH_TYPE

This optional parameter is used to specify the value of the DB2® instance **AUTHENTICATION** parameters on the DB2® HADR primary server. The following list identifies valid values:

SERVER

Authentication of the user name and password takes place at the server.

CLIENT

Authentication of the user name and password takes place at the client.

SERVER_ENCRYPT

Specifies that authentication takes place on the node containing the target database, and that the authentication password is to be encrypted.

DATA_ENCRYPT

Specifies that authentication takes place on the node containing the target database, and that connections must use data encryption.

GSSPLUGIN

Specifies that authentication takes place using an external GSS API-based plug-in security mechanism.

Default

SERVER_ENCRYPT

Advanced mode only

No.

TSM_BACKUP

To create an IBM® Storage Protect server backup from a snapshot, install IBM® Storage Protect Snapshot on a backup server. The offload agent runs to trigger an IBM® Storage Protect server backup from any snapshot created with **TSM_BACKUP** set to YES, MANDATE, or LATEST.

YES

Create an IBM® Storage Protect server backup from this snapshot. If the IBM® Storage Protect server backup operation does not successfully complete, reuse the target set.

MANDATE

In contrast to YES, do not reuse the target set until the IBM® Storage Protect server backup completes.

LATEST

When a snapshot backup was performed with **TSM_BACKUP LATEST** and the offloaded backup to an IBM® Storage Protect server has either not started or has failed, any new snapshot backup with option **TSM_BACKUP** set to LATEST, YES, or MANDATE, removes the backup request to the IBM® Storage Protect server from the previous backup. This removal prevents backup requests to the IBM® Storage Protect server from queuing if the requests are not completed in time.

NO

Keep the snapshot backup and do not use the backup as a source for a subsequent tape backup operation.

TSM_ONLY

The backup is automatically marked for deletion during the unmount operation after the IBM® Storage Protect server backup has completed. This deletion occurs whether or not the backup was successful.

USE_FOR *list of device classes*

This attribute can be combined with any of these options to limit its application to snapshots performed with particular device classes as specified in the profile. Any number of device classes that are listed must be separated by spaces.

Note: (DB2®) The ability to create an IBM® Storage Protect server backup from a snapshot requires an IBM® Storage Protect Snapshot license.

Default

There is no default value.

Advanced mode only

Yes.

MAX_VERSIONS

When the DB2® system acts as the DB2® standby server or as a high-availability disaster recovery secondary server, this parameter is evaluated, instead of the parameter specified in the CLIENT section. This parameter accepts the following options:

ADAPTIVE

The maximum number varies depending on the available space. IBM® Storage Protect Snapshot reuses the oldest target set as the target for the current backup.

n

The value specified as *n* represents the maximum number of snapshot versions to maintain. When this limit is reached, the oldest version is deleted.

Default

ADAPTIVE

Advanced mode only

Yes.

DEVICE_CLASS

When the DB2® system acts as the DB2® standby server or as a high-availability disaster recovery secondary server, this parameter is evaluated, instead of the parameter specified in the CLIENT section. When specifying the options for this parameter, use the following syntax:

<list of device classes> [<conditions>]

During backup, IBM® Storage Protect Snapshot uses one of the device classes listed in the *<list of device classes>* of the **DEVICE_CLASS** statement for which the *<conditions>* evaluates to *true*. If multiple *<conditions>* statements evaluate to true the operation fails. For the device classes listed in the *<list of device classes>*, separated by spaces, IBM® Storage Protect Snapshot use the device class that follows the device class which was used most recently for the next backup operation. If the last device class in the list was used during the most recent backup, or if no device class in the list was used for a backup, IBM® Storage Protect Snapshot uses the first device class in the list. The value of the **DEVICE_CLASS** parameter uses the following syntax:

```
[USE_AT <days of week>] [FROM <time> TO <time>]
(partitioned DB2 databases) [ON_DBPARTITIONNUMS] <list of partitions>
```

Multiple sections representing different devices are possible. Any such section can be selected using the **DEVICE_CLASS** profile parameter or vendor option. At restore time, IBM® Storage Protect Snapshot uses the same **DEVICE_CLASS** value that was used during the backup.

Additional **DEVICE_CLASS** sections are added to the IBM® Storage Protect Snapshot profile automatically by the setup script when you add additional instances of the **DEVICE_CLASS** parameter to the **CLIENT** section.

Default

STANDARD

Advanced mode only

Yes.

OFFLOAD_DB2STANDBY

The OFFLOAD_DB2STANDBY section has the same semantics as the OFFLOAD section, except that it is evaluated whenever an offloaded backup from a DB2® HADR secondary occurs instead of the DB2® HADR primary.

The parameters in the OFFLOAD_DB2STANDBY section are the same as the OFFLOAD section. The parameters do not depend on the storage device. The parameter **BACKUP_METHOD** is set to TSM_CLIENT for this standby variant of the offload section.

Changing profile parameters

Except for the GLOBAL and ACSD sections, changes to the profile take effect immediately and do not require restarting IBM® Storage Protect Snapshot. Updates to the GLOBAL and ACSD sections require a restart of IBM® Storage Protect Snapshot.

About this task

To change the GLOBAL, ACSD, or any other sections of the profile, complete the following steps:

Procedure

1. Start the setup script by entering the following command:

```
cd <instance directory>
./setup_db2.sh
```

To run the setup script in advanced mode, use the `-advanced` option with the setup script command. If you run the setup script in advanced mode, you can change all parameters and override default values.

2. Follow the setup script instructions that are displayed.

- When you run the setup script, select **manage backup systems** or **manage clone instances** as required. This step is required when changes were made to the GLOBAL or ACSD sections so that all backup and clone systems are updated.

Interdependency of LVM_FREEZE_THAW and TARGET_DATABASE_SUSPEND

The **LVM_FREEZE_THAW** and **TARGET_DATABASE_SUSPEND** parameters are interdependent.

These two IBM® Storage 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.

The following table summarizes the actions taken depending on the values of the two parameters: This table identifies how **LVM_FREEZE_THAW** and **TARGET_DATABASE_SUSPEND** parameters affect the actions that can be completed.

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 files

Snapshot backups on DS8000®, SAN Volume Controller, and Storwize® family with static target allocation, 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 snapshot backups. The target volumes file, with extension .fct, identifies the target volumes to be used for an IBM® Storage 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 name 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.

Related information

[Examples](#)

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® family storage systems, need the **TARGET_SETS** parameter to specify the target volumes file, **VOLUMES_FILE**. The details are shown in the following table:

Table 17: Managing target volume LUNs by storage system	
DS8000®	SAN Volume Controller and Storwize® family
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.

For IBM® Storage 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® Storage Protect Snapshot 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.

VOLUMES_FILE for partitioned DB2

On a partitioned DB2® environment, use the **VOLUMES_FILE** parameter to create target set definitions for specific **PARTITION** sections. This setting is required when two partitions are accessing the same **TARGET_SET** during a single backup operation.

IBM® Storage Protect Snapshot supports volume sharing on DB2 partitions when the **TARGET_SETS** parameter is set to **VOLUMES_FILE** and the **PARTITION_GROUP** parameter is specified in the **DEVICE_CLASS**. The following example shows that configuration:

```
<ACS_DIR>/profile:  
...  
>>> DEVICE_CLASS STANDARD  
...  
PARTITION_GROUP GROUP_A 0 1  
PARTITION_GROUP GROUP_B 2 3  
...  
TARGET_SETS VOLUMES_FILE  
VOLUMES_FILE <ACS_DIR>/acsvolumes/volumes_file.fct  
<<<  
  
<ACS_DIR>/acsvolumes/volumes_file.fct:  
>>> TARGET_SET 1  
>>> PARTITION GROUP_A  
TARGET_VOLUME J01acs_td_0  
TARGET_VOLUME J01acs_tl_0  
TARGET_VOLUME J01acs_td_1  
TARGET_VOLUME J01acs_tl_1  
...  
<<< PARTITION GROUP_A  
  
>>> PARTITION GROUP_B  
TARGET_VOLUME J01acs_td_2  
TARGET_VOLUME J01acs_tl_2  
TARGET_VOLUME J01acs_td_3  
TARGET_VOLUME J01acs_tl_3  
...  
<<< PARTITION GROUP_B  
<<< TARGET_SET 1  
  
>>> TARGET_SET 2  
  
...  
<<< TARGET_SET 2
```

Related information

[DS8000 target volume parameter settings](#)

[SAN Volume Controller and Storwize family target volume parameter settings](#)

[Target set handling for cloning](#)

Changing target set definitions system

You can extend an existing target set definition file by either adding a target set or by adding another target volume to an existing target set.

If you want to remove a volume from an existing target set, ensure that backups on the affected target set are deleted first. All FlashCopy relations of volumes in the target set must also be withdrawn. If you want to remove a target set from a target set definition file, ensure that backups on the affected target set are deleted first. All FlashCopy relations of volumes in the target set must be withdrawn.

Related information

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® Storage Protect Snapshot attempts to match source volumes to suitable targets within a target set during backup.

Table 18: TARGET_VOLUME parameters	
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.

Related information

[Target volumes file examples](#)

[Target set and target volumes files](#)

SAN Volume Controller and Storwize® family 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® Storage Protect Snapshot software attempts to match source volumes to suitable targets within a target set.

Table 19: TARGET_VOLUME parameters (SAN Volume Controller and Storwize® family)	
Parameter Name	Value
<code>TARGET_VOLUME</code> <code><target volume virtual disk name></code> <code><source volume virtual disk name></code> <code><source volume size></code>	<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 relationship 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 <code>svdfsrc4</code> must be used in a FlashCopy® with the target volume with virtual disk name <code>svdftgt4</code>.</p> <pre>TARGET_VOLUME svdftgt4 svdfsrc4 Size=2.0_GB</pre> <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:</p> <pre>TARGET_VOLUME svdftgt4 - -</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 or greater than the size of the source volume.• The source and target volumes that are 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 files” on page 132](#).

Changing or deleting target volume or target set definitions

The following actions are possible with target sets:

- 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 146](#) with the `force` option.

For SAN Volume Controller 6.1 or later and Storwize® family, with IBM® Storage 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.

Example

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 information

[SAN Volume Controller and Storwize family target volumes file example](#)

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 snapshot 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® Storage 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® family, 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 snapshot backup of the same database, must not be on the same cluster. If you are cloning databases in an AIX® Logical Volume Mirroring (LVM) environment, use FlashCopy® cloning on one of the clusters, and FlashCopy® backup on the other 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 multi-partition DB2®:

```
>>> 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 or a list of clone database names
>>> PARTITION name of partition
          # e.g. NODE0000 for partition 0 or NODE0001 for partition 1, ...
          # or the name of a PARTITION_GROUP
      TARGET_VOLUME target [source] [size]
[...]
```

The target volumes file (specified by the VOLUMES_FILE parameter) must have the following syntax for single partition DB2®:

```
>>> 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® Storage Protect Snapshot password file

To access the storage system where the database volumes are stored, IBM® Storage 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® Storage Protect Snapshot agents are running in a distributed environment across multiple servers or in DB2® DPF partitions, 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. When backup and clone servers are set up with the setup script, the passwords are automatically made available on all servers.

A password file is created during the IBM® Storage 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 master 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.

Commands and scripts

A list of various commands and scripts that are used with IBM® Storage 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® Storage Protect Snapshot can be issued from the command line.

- 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 the database name as shown in the command **db2 list database directory**
`-i` is the instance name as returned from the **db2ilist** command

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® Storage Protect Snapshot repository. You can create and manage database clones from the command-line interface.

Commands that use DB2® utilities

IBM® Storage Protect Snapshot fully integrates with DB2® backup utilities like `db2 backup` and `db2 restore` commands. In addition, IBM® Storage Protect Snapshot provides the **fcmlcli** command-line interface, which you can use for extra tasks for example, to restore backups from a DB2® partition groups.

For detailed information about how to use DB2® to create snapshot backups, see the *DB2® Command Reference*.

You can use the options in the following table as parameters in the `<option string>` option to be specified with the following commands:

- `db2 backup db <dbname> [...] use snapshot options <option string>`
- `db2 restore db <dbname> [...] use snapshot options <option string>`
- `db2acsutil [...] options <option string>`

Where `<option string>` has the form `<parameter>[=<value>] [<parameter>[=<value>] ...]`.

PROFILE

Absolute path and file name of profile. The default value is `<ACS_DIR>/profile`.

TSM_BACKUP

Options include YES, MANDATE, LATEST, and NO. Applies to only **db2 backup**.

DELETE_FORCE

This option applies to only **db2acsutil**. With **db2acsutil delete**, withdraws any FlashCopy® relations for the target set represented by the backup (applicable to CIM devices only). A manual withdraw of FlashCopy® relations is needed for a restore when multiple target sets are in use. At least one target set other than the one to be restored is in a **NOCOPY** or **INCR** FlashCopy® relation.

With **db2acsutil query**, also lists backups that are deleted without the `DELETE_FORCE` option.

For DS8000® and SAN Volume Controller, incremental FlashCopy® relations are not withdrawn.

DEVICE_CLASS

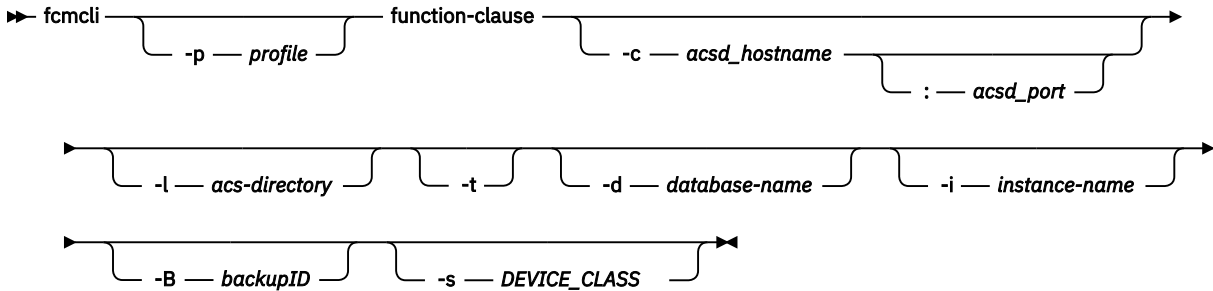
In the profile, see *device_section_name*. Applies only to **db2 backup** command.

fcmcli -f *_db2standby

Information for the support of DB2® standby server and DB2® HADR environments is provided.

With IBM® Storage Protect Snapshot for Custom Applications you can create snapshot backups of a DB2® standby server and DB2® HADR environments. You can back up to IBM® Storage Protect from a snapshot by using the IBM® Storage Protect backup-archive client.

Figure 18: *fcmcli* command



Where:

-p profile

Specifies the full profile name. The default value: *<INSTANCE_DIR>/profile*

-c acsd-hostname

Specifies the name of the server where the management agent (*acsd*) is running. The default value: *localhost*

acsd-port

Specifies the TCP/IP port number or service name on which the management agent (*acsd*) is listening. The default value: *57328*

-l acs-directory

Specified the directory where the logs and shared directories are located. The default value: *<ACS_DIR>*

-t

Specifies to start with the trace facility turned on. By default, the trace facility is turned off.

-d database-name

Specifies the database name. Required for -F option. There is no default value.

-i instance-name

Specifies the instance name that applies to the command. Required for -F option. There is no default value.

-B backup ID

Specifies the Backup ID as displayed by **fcmcli -f inquire [_detail]** or **db2acsutil**. There is no default value.

-s DEVICE_CLASS


Specifies the name of the *DEVICE_CLASS* section in the profile that is used for the backup operation. The default value is specified in the profile.

The values for the *function-clause* parameter are described in the following sections.

Function-clause: FlashCopy® operations of DB2® HADR databases

The following functions are supported by the **fcmcli** command option -f *function* for FlashCopy® backups of DB2® HADR databases:

Figure 19: *fcmcli* command functions



Where:

-E preflash command

This option overrides the value of the **PRE_FLASH_CMD** parameter that is specified in the **CLIENT** section of the profile. The **preflash** command is run on the production server and can be a script. The return code of the **preflash** command is evaluated. The following return codes are used:

0

Indicates a successful operation. The IBM® Storage Protect Snapshot backup operation continues.

Any value other than 0

Indicates an unsuccessful operation. The IBM® Storage Protect Snapshot backup operation ends.

There is no default value.

-Q postflash command

This option overrides the value of the **POST_FLASH_CMD** parameter that is specified in the **CLIENT** section of the profile. The **postflash** command is run on the production server and can be a script. The return code of the **postflash** command is evaluated. The following return codes are used:

0

Indicates a successful operation. The IBM® Storage Protect Snapshot backup operation continues.

Any value other than 0

Indicates an unsuccessful operation. The IBM® Storage Protect Snapshot backup operation ends.

There is no default value.

The **-f backup_db2standby** function operates like the **fmccli -f backup** function. However, instead of requesting the user to specify a list of files for backup, this command retrieves this information from the active DB2® instance. It is important that the names of the files and file systems from the active DB2® database (HADR primary) match the names of files and file systems of the standby server (HADR secondary). The editable sample **preflash** and **postflash** command files assist you to shut down and restart the DB2® standby server (HADR secondary). This action creates a consistent backup image.

The following figure illustrates the workflow that is executed from this command.

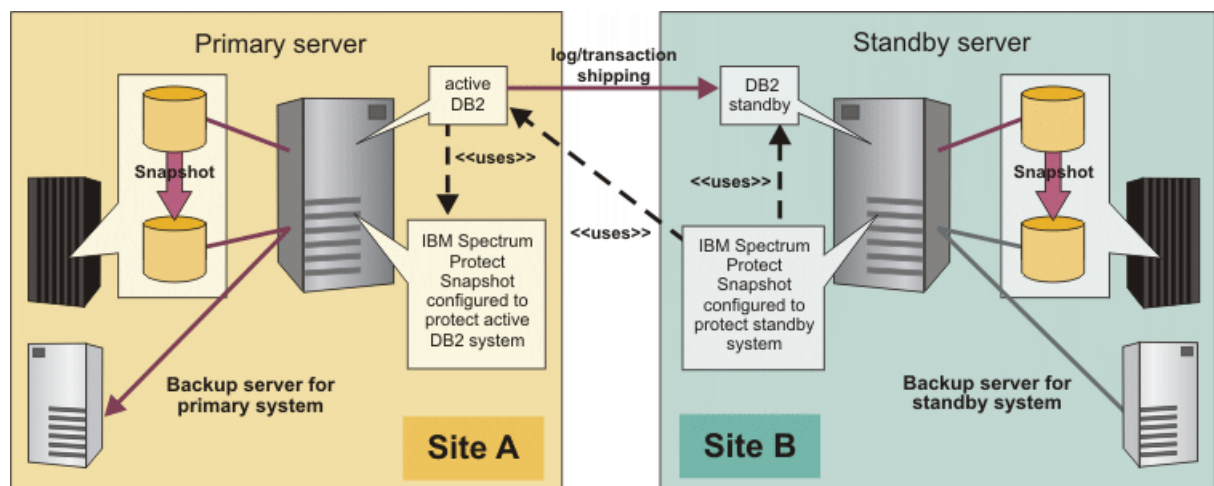


Figure 20: Protection of DB2® HADR standby nodes with IBM® Storage Protect Snapshot for Custom Applications

1. It starts with connecting to the active DB2® database (for example site A) to retrieve a list of files from the active database.
2. This list of files is then translated into a list of volumes and LUNs in site B, using the assumption that the names of DB2® files and file systems in site A are identical to the names of files and file systems in site B.
3. The preflash script is invoked. This script typically stops the DB2® standby system, in order to enable IBM® Storage Protect Snapshot to create a consistent snapshot backup. The return code of the **preflash** command is evaluated as follows:

0

Successful. The IBM® Storage Protect Snapshot backup operation continues.

Any value other than 0

Unsuccessful. The IBM® Storage Protect Snapshot backup operation terminates.

4. Once the snapshot has been taken, the **postflash** command is invoked, which typically starts the DB2® standby system again. The return code of the **postflash** command is evaluated as follows:

0


Successful. The IBM® Storage Protect Snapshot backup operation continues.

Any value other than 0

Unsuccessful. The IBM® Storage Protect Snapshot backup operation terminates.

The following functions are supported by the **fcmccli** command option -ffunction for FlashCopy® restore, inquire, and delete of DB2® HADR databases:

Figure 21: fcmcli command



The following list identifies the options to use when you use the **fcmccli** command in DB2® standby server and DB2® HADR environments:

-f restore_db2standby

Use to restore a DB2® HADR snapshot backup. There is no default value.

-f inquire_db2standby

Use to query the backup repository and list all available backups. There is no default value.

-f inquire_detail_db2standby

Use to query the backup repository and list all available backups in detail. There is no default value.

-f delete_db2standby

Use to unmount and delete a snapshot of a DB2® standby server. There is no default value.

-i instance-name

The instance name that applies to the command. This option is required for -F option. There is no default value.

-F

The DELETE_FORCE flag applies to the following commands:

- **inquire_db2standby**
- **inquire_detail_db2standby**
- **delete_db2standby**

When you run the **delete_db2standby** command, for only DS8000® and SAN Volume Controller storage systems, any FlashCopy® relations that are set for the target set represented by the backup are withdrawn. When you run the **inquire_detail_db2standby** command, backups that are deleted with the DELETE_FORCE flag are deleted.

There is no default value.

-t

Start with the trace facility turned on. By default, the trace facility is turned off.

-v

Displays the version. There is no default value.

-h

Displays the help information. There is no default value.

-B

The backup ID as displayed by the following commands: **fcmccli -f inquire [_detail]** or **db2acsutil**. There is no default value.

-s DEVICE_CLASS


The name of the `DEVICE_CLASS` section in the profile that is used for the backup operation. The default value is specified in the profile.

If the **fcmccli** command finishes the request without an error or if there were no candidates for the request, the return code is 0. If one or more non-critical issues are reported, the return code is 1. Resolve these issues to prevent more issues from occurring. If an error is reported, the return code is 2.

The following sections describe the details of the various functions that are specified with the **-f** option of the **fcmccli** command.


Cloning commands

You can use the IBM® Storage Protect Snapshot command-line interface, *fcmccli*, to create and manage clones of component databases.



Syntax for obtaining version or help information:

Figure 22: *fcmccli help*



Where:

-p profile

Full profile name. Default value: *<instance directory>/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.

-d database-name.

The name of the database.

-v

Show version.

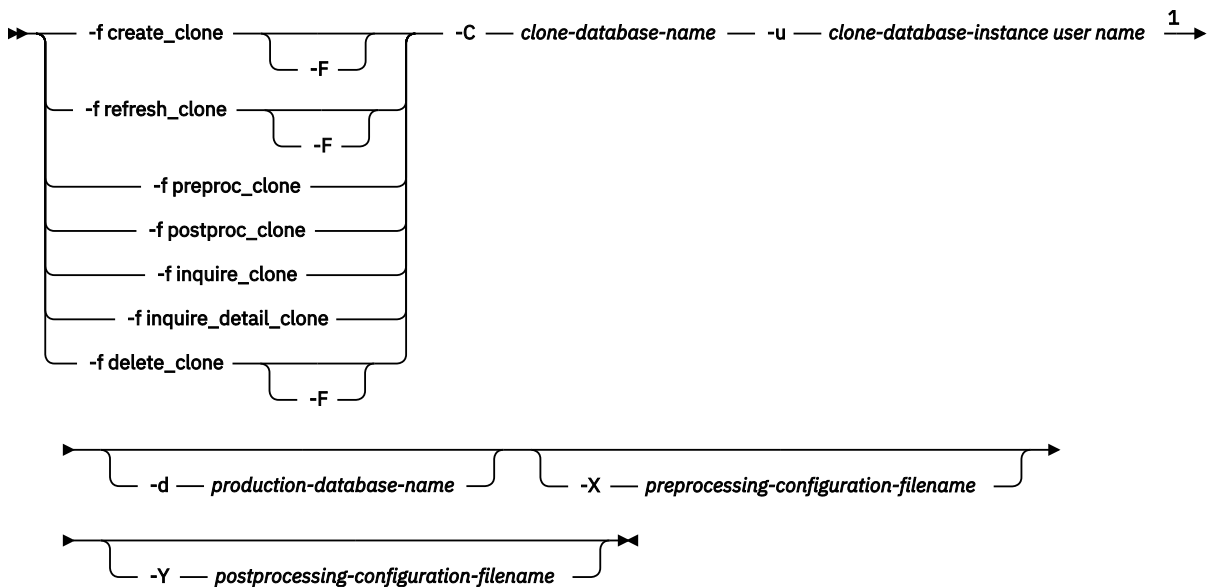
-h

Show help text.

The values for the function-clause parameter are described in the following sections.

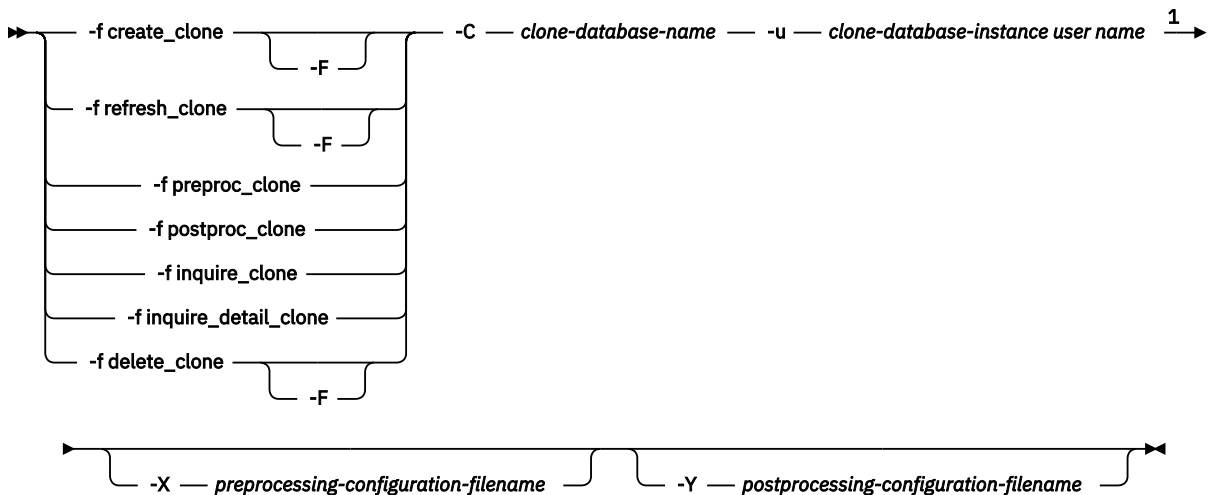
FlashCopy® cloning function-clauses

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.



Notes:

¹ This option is not required for -f *inquire_clone* or -f *inquire_detail_clone* commands.

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 directory where the IBM® Storage Protect Snapshot product files are located. The *fcmcli* command identifies the name of the production database.

- For DB2® databases, the DB2® database directory is used. If only one entry exists in the database directory, this entry is used to identify the production database name.

- If the `-d production database name` option is specified, this value overwrites the value that is identified in the previous step. Also, this value is used to identify the production database name. Its use with DB2® depends on the number of database entries in the database directory. If more than one entry is in the DB2® database directory, the `-d` option is required.

Example

For a DB2® database the cloned database is not moved to the clone target instance, when the target database name is identical to the production database name. If required to move the cloned database, rename the database by providing an alternative name. For example, two database instances are created on the clone server, the production instance that is named *db2prodins* and the clone instance named *db2cloneins*. To clone the production database *PROD* and move it to the clone target instance, issue the following command:

```
fcmcli -f create_clone -C CLONE -u db2cloneins -d PROD
```

Deleting snapshot backups

IBM® Storage 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® Storage 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.

```
fcmcli -f unmount [-B <backupID>]
```
2. Based on the use of this target set, any existing source, and target snapshot relationships (such as INCR or NOCOPY) must be withdrawn. Run the following command:

```
db2acsutil delete
```

Result

Note: For IBM® XIV® Storage System, these commands delete the snapshot backup in the IBM® Storage 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® Storage 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.

```
fcmcli -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:
`db2acsutil delete options "DELETE_FORCE"`

Result

The withdrawal of the source and target FlashCopy® relationship is done by the IBM® Storage 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® Storage Protect Snapshot repository.

To check the status of snapshot backups in the IBM® Storage Protect Snapshot repository, use one of the following commands:

```
For DB2®, fcmcli -f inquire[_detail]
or, db2acsutil query status
```

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_RESTOREABLE, SWAP-RESTORABLE, PHYSICAL_PROTECTION,
FULL_COPY, TAPE_BACKUP_PENDING
```

Tip: The `db2acsutil query status` command does not show all information that is shown by the `inquire_detail` command. That is because `db2acsutil` is a tool delivered by DB2® which only knows a subset of all possible states that can be set by IBM® Storage Protect Snapshot.

Administrative commands

You can use commands to administer IBM® Storage Protect Snapshot.

Administrative commands are available for you to do the following tasks:

- Start, stop, or configure IBM® Storage Protect Snapshot.
- Mount or unmount a snapshot backup on a secondary system.
- Create a backup to IBM® Storage Protect from a snapshot if you have IBM® Storage Protect configured in your environment

To use the commands to automate operations for IBM® Storage Protect Snapshot, add entries to the *cron* table (*crontab*) file. Because there are so many ways to implement IBM® Storage 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® Storage 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® Storage Protect Snapshot installation program. The setup script can also be used to manually set up IBM® Storage 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_db2.sh

You can use the setup script for the following purposes:

- Upgrade of IBM® Storage 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 installation directory.

- Initial configuration and reconfiguration:

```
setup_type.sh
```

The setup script must be run as the database instance owner.

- 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.

- Starting an activated instance:

```
setup_type.sh -a start -d <Instance owner $HOME directory>
```

The command must be run as the database instance owner.

- 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 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® Storage Protect Snapshot, run the command on the production system before the backup system.

The setup script can be used to install IBM® Storage 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® Storage Protect Snapshot.

The script must be run from the database instance-specific installation directory:

- (DB2®) *Instance owner \$HOME directory/sqlllib/acs/*

The default action, *setup*, is run 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® Storage Protect Snapshot cannot be stopped, stop IBM® Storage 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_db2.sh` to configure IBM® Storage Protect Snapshot for DB2®.
The following values are available for action. The instance directory name-option, is required for all explicit actions.

disable

This call can be issued as the root or instance owner. It stops IBM® Storage Protect Snapshot and removes all daemon jobs. To reactivate IBM® Storage Protect Snapshot, call the script without parameters.

If IBM® Storage Protect Snapshot cannot be stopped, stop IBM® Storage Protect Snapshot on the production system before running `setup_type.sh -a install -d <Instance owner $HOME directory>`.

For DB2® databases, change `<INSTANCE owner $HOME directory>` to `<INSTANCE owner $HOME directory>/sqllib`.

install

This call needs to be issued with the root user ID. When issued, the following actions are completed:

1. Stops IBM® Storage 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® Storage Protect Snapshot installation directory to the instance-specific installation directory (`<instance directory>`)
3. Sets the appropriate access rights for the binary files.
4. Restarts IBM® Storage Protect Snapshot (`setup_type.sh -a start -d <Instance owner $HOME directory>`).

For DB2® databases, change `<INSTANCE owner $HOME directory>` to `<INSTANCE owner $HOME directory>/sqllib`

The steps to start and stop IBM® Storage Protect Snapshot are skipped if it is not configured.

If IBM® Storage Protect Snapshot cannot be stopped, stop IBM® Storage Protect Snapshot on the production system before running `setup_type.sh -a install -d <Instance owner $HOME directory>`.

For DB2® databases, change `<INSTANCE owner $HOME directory>` to `<INSTANCE owner $HOME directory>/sqllib`.

start

This call can be issued as the root or instance owner. The call starts a previously installed and configured version of IBM® Storage 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® Storage Protect Snapshot that is currently running. This call updates the configured daemon jobs and checks that IBM® Storage 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® Storage Protect Snapshot is running on the production system. To successfully stop IBM® Storage Protect Snapshot in those environments, stop IBM® Storage Protect Snapshot on the production system.

This option is not required for the default setup function.

Setting or changing passwords with the Configuration Wizard

You can set or change passwords with the **Configuration Wizard**.

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 [“Installing and setting up IBM Storage Protect Snapshot” on page 32](#).

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 Storage Protect Snapshot installation
2) start IBM Storage Protect Snapshot services
3) stop IBM Storage Protect Snapshot
4) uninstall IBM Storage 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_<type>.sh` script or the `fcmcli -f password` command to change the IBM® Storage Protect Snapshot passwords.

The `fcmcli -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 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.

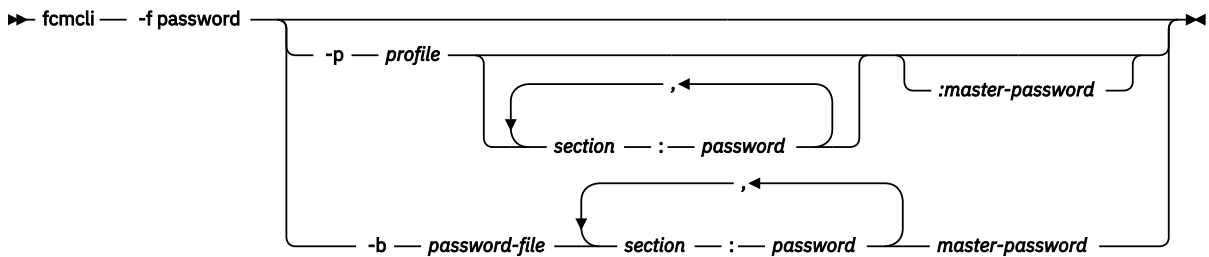
The interactive mode is the preferred method for setting passwords. Using this method, the passwords are verified by testing the connections to the corresponding storage devices, management agent, or database. For the non-interactive mode, the command syntax is verified but no additional validations are performed.

Note: 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.

Tip: To ensure that backup servers are also updated by SSH if applicable, use the setup scripts to modify any passwords.

Use the following syntax to change the passwords for intercommunication between IBM® Storage Protect Snapshot components, and communication to DB2® databases and to storage devices.

Figure 23: *fcmcli* command: **-f password**



Parameters

-p profile

Specify the full path and name of the profile that is used. If the path is not specified, the profile file in the current working path is used.

In interactive mode, the command searches the profile for the DB2STANDBY, and DEVICE_CLASS sections and then requests you to enter the relevant passwords.

-b password-file

Specify the password file to be created or updated. By default, the `shared/acsd.pwd` password file is in the directory that is specified by the **ACS_DIR** parameter. This parameter is included in the GLOBAL section of the profile file. This information is read from one of the following profiles:

- When the `-p` option is not specified, the profile file in the current working directory is used.
- When the `-p` option is specified, the profile file that is specified by this option is used.

-b password-file

Specify the password file to be created or updated. By default, the `shared/acsd.pwd` password file is in the directory that is specified by the **ACS_DIR** parameter. This parameter is included in the GLOBAL section of the profile file. This information is read from one of the following profiles:

- When the `-p` option is not specified, the profile file in the current working directory is used.
- When the `-p` option is specified, the profile file that is specified by this option is used.

sectionname:password

Specify the password for the user account that is referenced by the DB2STANDBY, and DEVICE_CLASS sections of the profile. To specify the password for the DEVICE_CLASS section, replace the *sectionname* variable with the `DEVICE_CLASS:device class name` variable for example, `DEVICE_CLASS:STANDARD`. Use this syntax when you specify the password: `DEVICE_CLASS:device class name:password`. No spaces are allowed between the *sectionname:password* syntax.

:masterpassword

Specify the master password that is used to authenticate a library or agent to the `acsd` management agent. Alternatively, enter the value *auto* to enable IBM® Storage 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® Storage 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® Storage Protect Snapshot. When the IBM® Storage Protect Snapshot setup script is run on the production server, it automatically creates the `fcmselfcert.arm` file. It is stored on the production server in the default installation path. The `fcmselfcert.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® Storage 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® Storage 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 fcmcert.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® Storage 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® Storage 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 fcmcert.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 fcmcert.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 fcmcert.kdb -stashed -label "FCM server  
certificate"
```

To rename the CA signed certificate issue the following command:

```
gsk8capicmd_64 -cert -rename -db fcmcert.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 fcmcert.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® Storage 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

Query managed capacity

Use the **managed_capacity** command to display information about IBM® Storage 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® Storage 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 **fcmcli -f managed_capacity** syntax is as follows:

```
fcmcli -f managed_capacity [-p profile] [-c] [-o <output_directory>]
```

-p

Specify the name of the IBM® Storage 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 managed capacity for the profile in */db2/DAB/sqllib/acs*:

```
fcmcli -f managed_capacity -p /db2/DAB/sqllib/acs/profile
```

Output:

```
FMM0461I Created tracefile '/db2/DAB/sqllib/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 */db2/DAB/sqllib/acs* as comma-separated values:

```
fcmcli -f managed_capacity -p /db2/DAB/sqllib/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 http://public.dhe.ibm.com/storage/tivoli-storage-management/front_end_capacity_measurement_tools/

Background daemons

For IBM® Storage 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 script commands.

To support high availability environments where the `/etc/inittab` cannot be used, you can instruct the `setup_db2.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.

If you must deviate from the standard installation, the management agent offers the following command options for customization:

Figure 24: *acsd* management agent



Syntax for obtaining version or help information:

Figure 25: *acsd* management agent help




Table 20: 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 <i>acsd</i> sections of the configuration profile.	<i><instance directory>/profile</i>	
-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

Option	Description	Default	Overrides profile parameter
-d acs-directory	Name of IBM® Storage Protect Snapshot directory	ACS_DIR	
-b password-file	File in which the IBM® Storage Protect Snapshot management agent password is stored (in encrypted form). See notes.	ACS_DIR/shared/pwd.acsd	No corresponding profile parameter.
-v	Display version and help information	None	N/A
-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.

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.

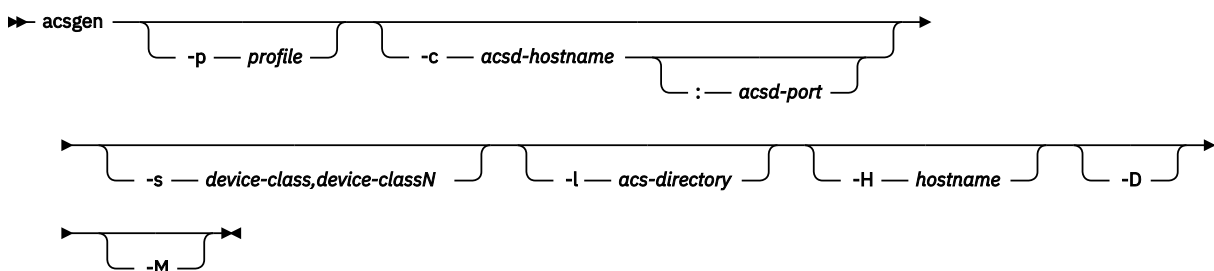
Generic device agent: *acsgen*

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:

Figure 26: *acsgen* generic device agent



Syntax for obtaining version or help information:

Figure 27: *acsgen* generic device agent help



Description of *acsgen* options with default values if applicable.

Table 21: Options for starting the generic device agent, <i>acsgen</i>		
Option	Description	Default
-p profile	Full profile name.	<instance_directory>/profile
-c acsd-hostname	Name of the server where the management agent, <i>acsd</i> , is running.	localhost
acsd-port	TCP/IP port number or service name on which the management agent, <i>acsd</i> , is listening.	57328
-s device-class	Section in the profile that pertains to the device class. Specify multiple device classes by separating each device class by a space.	STANDARD
-l acs-directory	Directory where the logs and shared directories can be found.	<ACS_DIR>
-D	Start as daemon. The -a option defines the usability states that the device agent responds to. Valid only when started from the following path: /etc/inittab	Run and end.
-H hostname	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 hostname command.
-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® Storage Protect is requested • Database files on JFS file systems • Database files on AIX® LVM mirrored volumes • The database is not suspended • Cloning databases 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

Description of *acsgen* options with default values if applicable.

Table 22: Options for starting the generic device agent, <i>acsgen</i>		
Option	Description	Default
-p profile	Full profile name.	<instance_directory>/profile
-c acsd-hostname	Name of the server where the management agent, <i>acsd</i> , is running.	localhost
acsd-port	TCP/IP port number or service name on which the management agent, <i>acsd</i> , is listening.	57328
-s device-class	Section in the profile that pertains to the device class. Specify multiple device classes by separating each device class by a space.	STANDARD
-l acs-directory	Directory where the logs and shared directories can be found.	<ACS_DIR>
-D	Start as daemon. The -aoption defines the usability states that the device agent responds to. Valid only when started from the following path: /etc/inittab	Run and end.
-H hostname	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 hostname command.
-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® Storage 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

Description of *acsgen* options with default values if applicable.


Table 23: Options for starting the generic device agent, <i>acsgen</i>		
Option	Description	Default
-p profile	Full profile name.	<instance_directory>/profile

Option	Description	Default
-U	Additional client or cloning sections must be named.	<use_client_section_name>, <use_cloning_section_name>
-c acsd-hostname	Name of the server where the management agent, <i>acsd</i> , is running.	<i>localhost</i>
acsd-port	TCP/IP port number or service name on which the management agent, <i>acsd</i> , is listening.	<i>57328</i>
-s device-class	Section in the profile that pertains to the device class. Specify multiple device classes by separating each device class by a space.	<i>STANDARD</i>
-l acs-directory	Directory where the logs and shared directories can be found.	<ACS_DIR>
-D	Start as daemon. The -aoption defines the usability states that the device agent responds to. Valid only when started from the following path: /etc/inittab	Run and end.
-H hostname	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 hostname command.
-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® Storage 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® Storage Protect Snapshot commands are available to mount or unmount a snapshot backup on a secondary system.

Figure 28: *fcmcli* command



Where:

-p *profile*

Full profile name. Default value: *<instance directory>/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

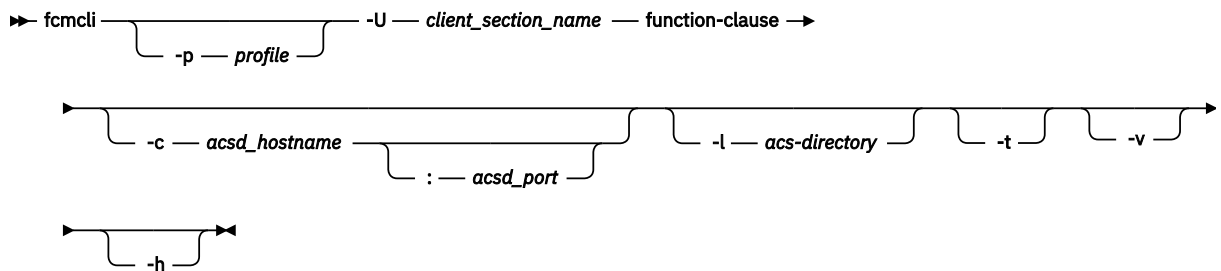
Show version.

-h

Show help text.

The return code of the **fcmcli** 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.

Figure 29: *fcmcli* command



Where:

-p *profile*

Full profile name. Default value: *<instance directory>/profile*

-U *client_section_name*

If there is one CLIENT and one CLONING section this can be omitted. For profiles with multiple CLIENT and CLONING sections, -U must be defined.

-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 **fcmccli** 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 **fcmccli** command option **-f**'function' for mount and unmount:

Figure 30: **-f** mountfunction-clauses




Figure 31: **-f** mountfunction-clauses



Where:

-f mount

Mount 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 **fcmccli -f inquire [_detail]** command.

The following functions are supported by the **fcmccli** command option **-f**'function' for forced unmount:

Figure 32: **-f** unmountfunction-clause with force option



The following functions are supported by the **fcmccli** command option **-f**function for forced unmount:

Figure 33: **-f** unmountfunction-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 **fcmccli -f inquire [_detail]** command.

Where:

Unmount snapshot target set.

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.


`client_section_name`If there is one CLIENT and one CLONING section this value can be omitted. For profiles with multiple CLIENT and CLONING sections, -U must be defined.

The Backup ID as displayed by `fcmcli -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 **fcmcli** command option -f'function' for mount and unmount:

Figure 34: `-f mountand-f unmountfunction-clauses`



-f mount

Mount snapshot target set.

-f unmount
Unmount snapshot target set.

Database name.


-i *instance-name*

Instance name to apply to the command. There are no limitations.

The Backup ID as displayed by `fcmcli -f inquire [_detail]` or `db2acsutil` command.

The following functions are supported by the **fcmlcli** command option `-f'function'` for forced unmount:

Figure 35: `-f unmountfunction-clause` with `force` option



-f unmount

Unmount snapshot target set.

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

Database name.

Instance name to apply to the command. There are no limitations.

-B backup ID

The Backup ID as displayed by `fcmcli -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 108 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 more 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 **-B** parameter is mandatory. Here are two examples. The first one is generic:

```
fcmcli -f mount -U <client_section_name> -B <backup-id>
```

```
fcmcli -f mount -d <database-name> -i <instance-name> -B <backup-id>.
```

Here is a specific example:

```
fcmcli -f mount -d DB1 -i db2db1 -B 20151020131727.
```

Note: If the option **-B** is omitted, the oldest backup still in state `tape_backup_pending` is selected implicitly.

(DB2®) In a DPF environment with multiple partitions, IBM® Storage Protect Snapshot always mounts all partitions that are associated with a snapshot backup operation.

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® Storage 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. Use **-B** to identify the specific backup as follows:

```
-B <backup-id>
```

(DB2®) In a DPF environment with multiple partitions, IBM® Storage Protect Snapshot unmounts all partitions that are associated with a snapshot backup.

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 **fcmcli 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® Storage Protect Snapshot for unmounting.

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 **fcmcli unmount** function is used 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 -B option is specified to uniquely identify the backup that is involved.

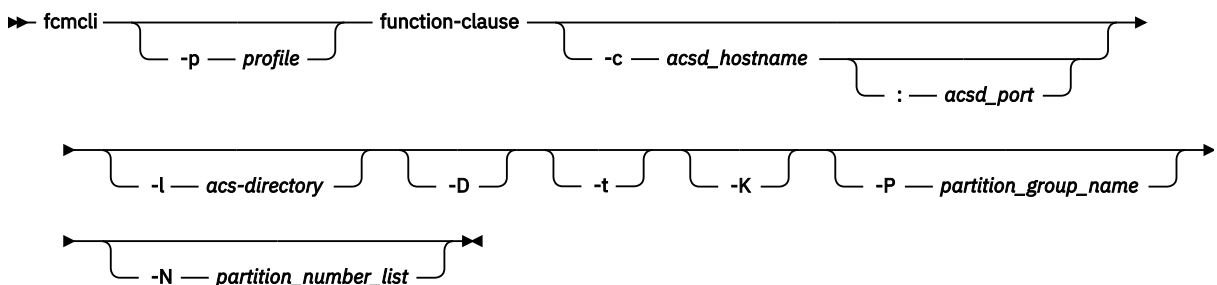
Integration with IBM® Storage Protect

If IBM® Storage Protect is set up and configured in your environment, you can create a backup to IBM® Storage Protect from a snapshot.

The **fcmcli** offload agent

The offload agent is a daemon process that manages offloaded backups to IBM® Storage Protect. The agent also provides a command line interface offering functions for managing IBM® Storage Protect backups.

Figure 36: **fcmcli** command



Where:

-p profile

Full profile name. Default value: *<instance directory>/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® Storage 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® Storage 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.

FlashCopy® offload operations

In a DB2® environment, the installation of IBM® Storage Protect Snapshot enables the offload function that otherwise is not available.

The following functions are supported by the *fcmdi* command option **-f** *function*:

Figure 37: -f tape_backupfunction-clause



Where:

-f tape_backup

Backs up a snapshot target set to IBM® Storage Protect.

-d *database-name*

The name of the database.

-i *instance name*

The name of the instance to apply the action against. This parameter is required for the `-f` option. There are no limitations.

-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.


-K

In a multi-partition environment, the partitions remain mounted when all participating partitions are successfully offloaded to IBM® Storage 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® Storage Protect backup operation.

The functions `mount`, `unmount`, or `tape_backup` cannot run in parallel on the same backup server.

The function `update_status` is supported by the **fcmcli** command:

Figure 38: -f update_statusfunction-clause



Where:

-f update status

Update the usability state after a snapshot backup operation completes to one of the following states:

- Offload a snapshot backup to IBM® Storage Protect (**TSM_BACKUP=yes**).

- Do not offload a snapshot backup to IBM® Storage Protect (**TSM_BACKUP=no**) if it was scheduled to offload.

-d database-name

The name of the database.

-i instance-name

The name of the instance to apply the action against. This parameter is required for the -f option. There are no limitations.

-B backup ID

The Backup ID as displayed by **fcmccli -f inquire [_detail]** or **db2acsutil** command.

-S TSM_BACKUP=yes/no

Use this option to inform IBM® Storage Protect Snapshot that a IBM® Storage Protect backup is no longer required or that a IBM® Storage Protect backup is requested. This option is only valid with the **update_status** function.

The *fcmccli* process connects to the *acsd* management agent and then runs the function that is specified with the -f option. After the operation is finished, the *fcmccli* agent passes notification to the *acsd* agent and the snapshot backup repository is backed up. When started as a daemon (-D option), as is the case for the standard */etc/inittab* entry, *fcmccli* runs offloaded tape backup operations. This action results in a synchronous tape backup of all snapshot backups. When a new snapshot is started with **TSM_BACKUP YES**, the offload agent starts the backup to tape when it becomes available for mounting on a backup system (**REMOTELY_MOUNTABLE**).

If the offload agent completes the request without an error and there are no candidates for the request, the return code is 0. If there is a warning, the return code is 1. If there is an error, the return code is 2.

-f tape_backup

This offload agent command backs up data to tape storage.

Note: IBM® Storage Protect for Enterprise Resource Planning must be installed on the production and backup server if you use IBM® Storage Protect Snapshot in an SAP environment with DB2®.

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. This setting applies to all backups. The management agent updates the usability state with **TAPE_BACKUP_PENDING**. The IBM® Storage Protect Snapshot offload agent then picks up all snapshot backups in the state **TAPE_BACKUP_PENDING** and backs them up to tape. The *fcmccli -f backup* operation must be issued from the production system.

To start the offload backup to tape, enter the command:

```
fcmccli -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® Storage Protect.

(DB2®) Offloaded tape backups are not be registered in the DB2® history on the production system. They are reflected in the DB2® history on the offload system when the assigned volumes have not been overwritten. For more information, see [“DB2 backup history file overview” on page 76](#).

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:

```
fcmccli -f tape_backup filter_arguments
```

If multiple snapshot backups of a database are candidates for offloading to tape, the IBM® Storage 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® Storage 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® Storage 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® Storage 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 *fcmccli* 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® Storage Protect fails, the IBM® Storage Protect Snapshot software can try the backup operation again.

In DPF environments, offloaded backups can be performed only when the snapshot was created on all partitions. If the Offload Agent retries a backup operation, it only backs up those partitions that have not been backed up successfully.

In DPF environments with more than one partition it can be useful to perform the offloaded IBM® Storage Protect backup in parallel on multiple partitions. This is useful when the DB2® DPF database is spread across more than one backup system. The profile parameter **NUMBER_BACKUPS_IN_PARALLEL** in the **OFFLOAD** section can be used to specify the degree of parallelism. For more information, see [“Configuration files” on page 96](#).

-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® Storage Protect (**TSM_BACKUP=yes**), or to not offload a snapshot (**TSM_BACKUP=no**). It is possible to offload a snapshot backup to IBM® Storage 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
```

IBM® Global Security Kit configuration

IBM® Storage 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 security standards as defined in the Special Publications 800131 (SP 800-131). GSKit is automatically installed by IBM® Storage Protect Snapshot.

This security standard requires longer key lengths, stronger cryptographic algorithms, and incorporates TLS Protocol version 1.2.

During the installation, IBM® Storage 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 by the setup script through the existing SSH remote deployment mechanisms.

If you do not use the SSH remote deployment capabilities of IBM® Storage Protect Snapshot, you must complete the following steps:

1. Globally install GSKit on each server by activating the instance. The required installation files are available in the `gskit_install` subdirectory of the IBM® Storage Protect Snapshot instance directory.
2. Manually copy the `fcmselcert.arm` file to the backup and cloning servers. The manually copied certificate is imported automatically when the setup script is run on the backup or cloning servers.

To install or reinstall GSKit separately, enter the command, `./setup_db2.sh -a install_gskit -d <instance directory>`

Alternatively, use a CA-signed certificate. The signed certificate can be from an internal or external certificate authority (CA). When SP800-131 encryption is enforced by setting the **ENFORCE_TLS12** profile parameter to YES in the IBM® Storage Protect Snapshot profile, 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 servers.

The following GSKit files are installed by IBM® Storage Protect Snapshot:

- A key database file, `fcmcert.kdb`, is in the instance 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, `fcmcert.rdb`, is in the instance 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® Storage Protect Snapshot creates a key database file.
- An encrypted stash file, `fcmcert.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-signed 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 `fcmselcert.arm` file to the backup and cloning servers.

- A certificate revocation list file, `fcmcert.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.

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® Storage Protect Snapshot, you must set the **ENFORCE_TLS12** parameter to YES in the IBM® Storage 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® Storage Protect Snapshot must be deleted to allow IBM® Storage Protect Snapshot to create new self-signed certificates. To remove any existing self-signed certificates, go to the IBM® Storage Protect Snapshot instance directory and enter the following command:

```
rm fmcert.*
```

Note: Do not delete certificates signed by certificate authority (CA). 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. When you uninstall GSKit, you remove the entire global GSKit installation from the system.

If required, you can uninstall GSKit with the following steps.

1. Log in with the root user ID.
2. Change to the IBM® Storage Protect Snapshot instance directory.
3. Run the setup script to uninstall GSKit, as follows.

```
./setup_db2.sh -a uninstall_gskit -d <instance_directory>
```

Examples

Refer to these IBM® Storage Protect Snapshot examples when you are configuring, updating, or following product tasks.

DB2® overall disk layout example

Refer to this example when you are configuring the disk layout for a DB2® environment.

The following figure shows file systems that are involved in an example disk layout.




Figure 39: Example overall disk layout for a DB2® 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

```
/db2/D01
/db2/D01/db2dump
/db2/D01/db2event
/db2/D01/sqllib
/sapmnt/D01
/usr/sap/D01
/usr/sap/trans
/opt/IBM/db2/V10.5
/db2/D01/acs (ACS_DIR)
```

2. Source volume disks on the production system (db_disk category) for the file systems

```
/db2/D01/sapdata1
/db2/D01/sapdata2
/db2/D01/sapdata3
/db2/D01/sapdata4
/db2/D01/sapdata5
/db2/D01/sapdata6
/db2/D01/sapdata7
/db2/D01/sapdata8
/db2/D01/sapdata9
/db2/D01/sapdata10
/db2/D01/sapdata11
/db2/D01/sapdata12
/db2/D01/sapdata13
/db2/D01/sapdata14
/db2/D01/sapdata15
/db2/D01/sapdata16
/db2/D01/sapdata17
/db2/D01/sapdata18
/db2/D01/sapdata19
/db2/D01/sapdata20
/db2/D01/sapdata21
/db2/D01/sapdata22
/db2/D01/sapdata23
/db2/D01/sapdata24
/db2/D01/sapdata25
/db2/D01/sapdata26
/db2/D01/sapdata27
/db2/D01/sapdata28
/db2/D01/sapdata29
/db2/D01/sapdata30
/db2/D01/sapdata31
/db2/D01/sapdata32
/db2/D01/sapdata33
/db2/D01/sapdata34
/db2/D01/sapdata35
/db2/D01/sapdata36
/db2/D01/sapdata37
/db2/D01/sapdata38
/db2/D01/sapdata39
/db2/D01/sapdata40
/db2/D01/sapdata41
/db2/D01/sapdata42
/db2/D01/sapdata43
/db2/D01/sapdata44
/db2/D01/sapdata45
/db2/D01/sapdata46
/db2/D01/sapdata47
/db2/D01/sapdata48
/db2/D01/sapdata49
/db2/D01/sapdata50
/db2/D01/sapdata51
/db2/D01/sapdata52
/db2/D01/sapdata53
/db2/D01/sapdata54
/db2/D01/sapdata55
/db2/D01/sapdata56
/db2/D01/sapdata57
/db2/D01/sapdata58
/db2/D01/sapdata59
/db2/D01/sapdata60
/db2/D01/sapdata61
/db2/D01/sapdata62
/db2/D01/sapdata63
/db2/D01/sapdata64
/db2/D01/sapdata65
/db2/D01/sapdata66
/db2/D01/sapdata67
/db2/D01/sapdata68
/db2/D01/sapdata69
/db2/D01/sapdata70
/db2/D01/sapdata71
/db2/D01/sapdata72
/db2/D01/sapdata73
/db2/D01/sapdata74
/db2/D01/sapdata75
/db2/D01/sapdata76
/db2/D01/sapdata77
/db2/D01/sapdata78
/db2/D01/sapdata79
/db2/D01/sapdata80
/db2/D01/sapdata81
/db2/D01/sapdata82
/db2/D01/sapdata83
/db2/D01/sapdata84
/db2/D01/sapdata85
/db2/D01/sapdata86
/db2/D01/sapdata87
/db2/D01/sapdata88
/db2/D01/sapdata89
/db2/D01/sapdata90
/db2/D01/sapdata91
/db2/D01/sapdata92
/db2/D01/sapdata93
/db2/D01/sapdata94
/db2/D01/sapdata95
/db2/D01/sapdata96
/db2/D01/sapdata97
/db2/D01/sapdata98
/db2/D01/sapdata99
/db2/D01/sapdata100
```

Source volume disks on the production system (db_log category) for the file system

```
/db2/D01/log_dir
```

3. Local disks on the production system (p_db_disk category) for the file systems

```
/db2/D01/log_archive
/db2/D01/log_retrieve
```

4. Local disks on the backup system (b_disk category) for the file systems

```
/db2/D01
/opt/IBM/db2/V9.5
/db2/D01/acs (ACS_DIR)
```

5. Disks for the IBM® Storage Protect server (TSM_disk category) for the file systems

```
/tsmdb
```

DB2® profile example

The profile file provides parameters that customize how IBM® Storage Protect Snapshot works within a particular environment. Use this example to verify that the configuration of the profile is correct for your DB2® data partitioning environment.

The following sample profile is used for a DB2® DPF environment. The profile specifies that snapshots are completed from 06am to 8pm and snapshots with offloaded backups for the remaining time. These offloaded backups use two different SAN Volume Controller clusters, one from 00am to 6am and the other cluster from 8pm to 12pm. Each cluster is backed up by a separate backup system, even though that fact is not obvious in the example.

```
>>> GLOBAL
ACS_DIR /db2/D01/acs
ACSD acsprod5 57328
TRACE NO
<<<
>>> ACSD
ACS_REPOSITORY /db2/D01/acs/acsrepository
# ADMIN_ASSISTANT NO
# REPOSITORY_LABEL TSM
<<<
>>> OFFLOAD
BACKUP_METHOD DB2
OPTIONS @/db2/D01/tdp_r3/vendor.env
# PARALLELISM AUTO
# NUM_SESSIONS 1
# NUM_BUFFERS AUTO
# BUFFER_SIZE AUTO
<<<
>>> CLIENT
# BACKUPIDPREFIX DB2_---
APPLICATION_TYPE DB2
TSM_BACKUP YES
TSM_BACKUP NO USE_FOR SVC_NO_TSM
MAX_VERSIONS ADAPTIVE
# LVM_FREEZE_THAW AUTO
# NEGATIVE_LIST NO_CHECK
# TIMEOUT_FLASH 120
# GLOBAL_SYSTEM_IDENTIFIER
DEVICE_CLASS SVC2 FROM 00:00 TO 05:59
DEVICE_CLASS SVC_NO_TSM FROM 06:00 TO 20:00
DEVICE_CLASS SVC1 FROM 20:01
<<<
```

```
>>> DEVICE_CLASS SVC_NO_TSM
COPYSERVICES_HARDWARE_TYPE SVC
COPYSERVICES_PRIMARY_SERVERNAME 192.168.1.104
# COPYSERVICES_USERNAME superuser
# SVC_COPY_RATE 80
# SVC_CLEAN_RATE 50
# COPYSERVICES_COMMPROTOCOL HTTPS
# COPYSERVICES_CERTIFICATEFILE NO_CERTIFICATE
COPYSERVICES_SERVERPORT 5989
FLASHCOPY_TYPE COPY
# COPYSERVICES_TIMEOUT 6
# RESTORE_FORCE NO
LVM_MIRRORING YES
TARGET_SETS VOLUMES_FILE
VOLUMES_FILE /db2/D01/acs/acsvolumes/fct.file
BACKUP_HOST_NAME PREASSIGNED_VOLUMES
<<<
```

```

>>> DEVICE_CLASS SVC1
COPYSERVICES_HARDWARE_TYPE SVC
COPYSERVICES_PRIMARY_SERVERNAME 192.168.1.101
# COPYSERVICES_USERNAME superuser
# SVC_COPY_RATE 80
# SVC_CLEAN_RATE 50
# COPYSERVICES_COMMPROTOCOL HTTPS
# COPYSERVICES_CERTIFICATEFILE NO_CERTIFICATE
# COPYSERVICES_SERVERPORT 5989
FLASHCOPY_TYPE INCR
# COPYSERVICES_TIMEOUT 6
# RESTORE_FORCE NO
LVM_MIRRORING YES
TARGET_SETS VOLUMES_FILE
VOLUMES_FILE /db2/D01/acs/acsvolumes/fct.file
BACKUP_HOST_NAME PREASSIGNED_VOLUMES
<<<
>>> DEVICE_CLASS SVC2
COPYSERVICES_HARDWARE_TYPE SVC
COPYSERVICES_PRIMARY_SERVERNAME 192.168.1.104
# COPYSERVICES_USERNAME superuser
# SVC_COPY_RATE 80
# SVC_CLEAN_RATE 50
# COPYSERVICES_COMMPROTOCOL HTTPS
# COPYSERVICES_CERTIFICATEFILE NO_CERTIFICATE
COPYSERVICES_SERVERPORT 5989
FLASHCOPY_TYPE INCR
# COPYSERVICES_TIMEOUT 6
# RESTORE_FORCE NO
LVM_MIRRORING YES
TARGET_SETS VOLUMES_file
VOLUMES_FILE /db2/D01/acs/acsvolumes/fct.file
BACKUP_HOST_NAME PREASSIGNED_VOLUMES
<<<

```

Target volumes file examples

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 S97p5l1_t1 - -
    TARGET_VOLUME S97p5l2_t1 - -
  <<<
  >>> PARTITION NODE0003
    TARGET_VOLUME S97p5l3_t1 - -
    TARGET_VOLUME S97p5l4_t1 - -
  <<<
<<<

#=====#

```

SAN Volume Controller and Storwize® family target volumes file example

Refer to this example when you are editing the target volumes file for an SAN Volume Controller or Storwize® family storage system configuration.

```

#***** First sample *****#
#

#=====#

>>> TARGET_SET VOLUMES_SET_1
TARGET_VOLUME svdftgt1 svdfssrc2 -
TARGET_VOLUME svdftgt2 svdfssrc3 -
TARGET_VOLUME svdftgt3 svdfssrc4 -
TARGET_VOLUME svdftgt4 svdfssrc5 -
TARGET_VOLUME svdftgt5 svdfssrc6 -
<<<

#=====#

```

Example

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® Storage 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® Storage 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® Storage 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
...
<<<
```

SAP in a multi-partition DB2® target volume file example

The following sample shows a configuration of SAP in a multi-partition DB2® environment.

ACS_DIR/profile:

```
>>> DEVICE_CLASS STANDARD
...
TARGET_SETS VOLUMES_FILE
VOLUMES_FILE <ACS_DIR>/acsvolumes/volumes_file.fct
<<<
```

ACS_DIR/acsvolumes/volumes_file.fct:

```

>>> TARGET_SET 1
>>> PARTITION NODE0000
TARGET_VOLUME J01acs_td_0
TARGET_VOLUME J01acs_tl_0
<<<

>>> PARTITION NODE0001
TARGET_VOLUME J01acs_td_1
TARGET_VOLUME J01acs_tl_1
<<<

>>> PARTITION NODE0002
TARGET_VOLUME J01acs_td_2
TARGET_VOLUME J01acs_tl_2
<<<

>>> PARTITION NODE0003
TARGET_VOLUME J01acs_td_3
TARGET_VOLUME J01acs_tl_3
<<<
<<<

>>> TARGET_SET 2
>>> PARTITION NODE0000
TARGET_VOLUME J01acs_td_2_0
TARGET_VOLUME J01acs_tl_2_0
<<<

>>> PARTITION NODE0001
TARGET_VOLUME J01acs_td_2_1
TARGET_VOLUME J01acs_tl_2_1
<<<

>>> PARTITION NODE0002
TARGET_VOLUME J01acs_td_2_2
TARGET_VOLUME J01acs_tl_2_2
<<<

>>> PARTITION NODE0003
TARGET_VOLUME J01acs_td_2_3
TARGET_VOLUME J01acs_tl_2_3
<<<
<<<

```

SAN Volume Controller and Storwize® family target volumes file example

Refer to this example when you are editing the target volumes file for an SAN Volume Controller or Storwize® family storage system configuration.

```

#***** First sample *****#
#
#=====#

>>> TARGET_SET VOLUMES_SET_1
TARGET_VOLUME svdftgt1 svdfsrc2 -
TARGET_VOLUME svdftgt2 svdfsrc3 -
TARGET_VOLUME svdftgt3 svdfsrc4 -
TARGET_VOLUME svdftgt4 svdfsrc5 -
TARGET_VOLUME svdftgt5 svdfsrc6 -
<<<

#=====#

```

Example

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® Storage 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® Storage 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® Storage 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
...
<<<
```

DB2® HADR standby server profile example

This example contains a sample profile that shows configuration parameters for a DB2® HADR standby server.

```
>>> GLOBAL
ACS_DIR /db2/CET/acs
ACSD ehonda 63000
# TRACE YES
<<<

>>> ACSD
ACS_REPOSITORY /db2/CET/acs/acsrepository
REPOSITORY_LABEL CET
<<<

>>> CLIENT
BACKUPIDPREFIX CET___
APPLICATION_TYPE DB2
# MAX_VERSIONS ADAPTIVE
# LVM_FREEZE_THAW AUTO
# NEGATIVE_LIST NO_CHECK
# TIMEOUT_FLASH 120
# GLOBAL_SYSTEM_IDENTIFIER
# DEVICE_CLASS STANDARD
<<<

>>> DB2STANDBY
DB2_PRODUCTIVE_SERVER akuma 60050
DB2_ALIAS CET
PRE_FLASH_CMD /db2/CET/acs/scripts/pre_flash_cmd
POST_FLASH_CMD /db2/CET/acs/scripts/post_flash_cmd
# DB2_USERNAME db2cet
DB2_AUTH_TYPE SERVER
<<<

>>> DEVICE_CLASS STANDARD
COPYSERVICES_HARDWARE_TYPE SVC
COPYSERVICES_PRIMARY_SERVERNAME 192.168.1.104
# COPYSERVICES_USERNAME superuser
# CLONE_DATABASE NO
SVC_COPY_RATE 95
# SVC_CLEAN_RATE 50
# COPYSERVICES_COMMPROTOCOL HTTPS
# COPYSERVICES_CERTIFICATEFILE NO_CERTIFICATE
# COPYSERVICES_SERVERPORT 5989
FLASHCOPY_TYPE INCR
# COPYSERVICES_TIMEOUT 6
# RESTORE_FORCE NO
# LVM_MIRRORING NO
# RECON_INTERVAL 12
BACKUP_HOST_NAME sagat
TARGET_SETS TS1 TS2
TARGET_NAMING %SOURCE_%TARGETSET
<<<
```

Accessibility features for the IBM Storage® 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 Storage® Protect family of products includes the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Storage® Protect family of products uses the latest W3C Standard, [WAI-ARIA 1.0 \(www.w3.org/TR/wai-aria/\)](http://www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 and [Web Content Accessibility Guidelines \(WCAG\) 2.0 \(www.w3.org/TR/WCAG20/\)](http://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® Documentation is enabled for 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 Storage® 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 \(www.ibm.com/able\)](http://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.

The registered trademark Linux® is used pursuant to a sublicense from the Linux® Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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.

Red Hat®, OpenShift®, Ansible®, and Ceph® are trademarks or registered trademarks of Red Hat®, Inc. or its subsidiaries in the United States and other countries.

UNIX® is a registered trademark of The Open Group in the United States and other countries.

VMware, VMware vCenter Server™, and VMware vSphere™ are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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 Storage® Protect family of products.

See the [IBM Storage® Protect glossary](#).

Index

51, 12, 156, 155

backup to 67

examples

target volumes file 172, 173, 173, 175, 175

target volumes file mirror setup 174

integration 164

log files 92

planning 24, 20, 20

target volumes file (.fct)

parameter settings 135, 136, 136

trace files 92

-

-f mount 163

-f unmount 163

A

accessibility features 178

ACS_DIR 99

ACS_REPOSITORY 100

ACSD 99, 100

activating instance 42

ADMIN_ASSISTANT 100

administrative commands 147

APAR 87

application agents

14, 14, 14

APPLICATION_TYPE 101

Asynchronous Remote Mirroring 26

audience 9

B

background daemons

description 155

backing up DB2 66

backup 73, 73

backup

to 67, 164

backup commands 66, 140

backup environment

backup servers

assignment 58

prerequisites 38

clone servers

assignment 58

installation

non-remote 42

preparing

backup servers 37, 37

clone servers 37, 37

upgrading 44

backup procedure

DB2 66

DB2

serial mode 67

to 67

DB2 database

parallel mode 67

DB2 HADR server 68

DB2 standby server 68

backup server

virtual machine 39

block device mapping 31

C

CA certificate 151

capacity planning 16

changing passwords 150

CIM

156

generic device agent 156

log files 91

trace files 91

CLIENT 101

Clone database

parameters 109

clone devices

24, 20, 20

clone server

virtual machine 39

cloning

commands 144

configuration files 80

target volumes file (.fct)

parameter settings 137

CLONING

parameters 105

commands 140

commands

administrative 147

backup 140

cloning 144

delete 140

inquire 140

restore 140

communication ports 18

Configuration 45

configuration files

configuration file 96

password file 138

configuring 46

configuring

backup server 43, 42, 42, 44

cloning server 43, 42

remote mirroring 61, 61, 62

Configuring SVC 50

consistency groups 28

control files

cloning files

- description 80
- profile
 - structure 96
- target volumes file (.fct)
 - cloning parameter settings 137
 - description 132
 - parameter settings 135, 136, 136

COPYSERVICES

- CERTIFICATEFILE 111
- COMMPROTOCOL 111
- HARDWARE_TYPE 111
- PRIMARY_SERVERNAME 111
- SERVERNAME 111
- SERVERPORT 111
- TIMEOUT 111
- USERNAME 111

D

database cloning

- commands 80
- description 13
- postprocessing 80
- preprocessing 80
- process 79

Db2 9

DB2 46, 76

DB2

- backup history file 76
- commands 140
- environment 35
- examples
 - disk layout 170
 - profile 171
- LIST HISTORY command 76
- vendor options 140

DB2 HADR standby server

- examples
 - profile 177

DB2STANDBY 128

defining target sets

- DB2
 - multi-partition 55
 - naming convention 56
 - single partition 55

delete commands 140

deleting

- snapshot backup 146

developerWorks wiki 87

device agents

- 14, 14
- CIM adapter 14
- command-line interface 14
- offload agent 14
- query capacity 14
- volume group takeover script 14

device class

- parameters 111, 120

DEVICE_CLASS 59, 101, 111

DEVICE_CLASS

- MAX_VERSIONS 48
- parameters 116, 108

devices

- 23, 23, 24, 20, 20

- preparing 18

storage

- log files 91
- setting up 48
- trace files 91

disability 178

disk layout

- examples
 - DB2 170

DPF

- backup modes 67

Dynamic Target Allocation 50

Dynamic target allocation 22, 21

E

encryption

- SP 800-131A 52

ENHANCED_PARTITIONING 101

environment

- backup servers 37, 37

backup servers

- assignment 58
- prerequisites 38

- clone servers 37, 37

clone servers

- assignment 58

- DB2 35

examples 170

examples

DB2

- disk layout 170
- profile 171

DB2 HADR standby server

- profile 177

target volumes file

- 172, 173, 173, 175, 175
- mirror setup 174

F

FlashCopy backup

- reconciliation 25

FLASHCOPY_TYPE 111

fmquery

- description 154

G

generic device agent 156

GLOBAL 99

Global Mirror 26

GSKit

- configuring 168

- FIPS 168

- installing 168

H

HyperSwap [73, 29, 73, 77](#)
HyperSwap configuration [57](#)

I

IBM Documentation [9](#)
incremental backups
 MAX_VERSIONS [21](#)
individual mirrors [69](#)
inquire commands [140](#)
installation
 backup server
 non-remote [42](#)
 clone server [43, 42](#)
 description [32](#)
 prerequisites [32, 64](#)
installation planning [45](#)
instance ID [42](#)
Internet Protocol Version 6 [94](#)
introduction [12](#)

K

keyboard [178](#)
KVM [31](#)

L

log files
 [88, 92](#)
 CIM [91](#)
 GPFS [92](#)
 storage subsystems [91](#)
 summary [88](#)
LVM mirroring [57](#)
LVM_FREEZE_THAW [101](#)
LVM_FREEZE_THAW
 TARGET_DATABASE_SUSPEND [132](#)

M

management agent [155](#)
MAX_VERSIONS [101](#)
Metro Mirror [26](#)
Migration [64](#)
mirroring
 examples
 target volumes file [174](#)
 individual [69](#)
 LVM [69, 72](#)
 overview [72](#)
multi-partition
 target set definitions [55](#)
multipath
 KVM check [31](#)

N

naming convention [56](#)
NEGATIVE_LIST [101](#)

new [9](#)

O

OFFLOAD [124](#)
offload agent
 fcmcli [164](#)
 tape_backup [166](#)
 update_status [167](#)
OFFLOAD_DB2STANDBY [131](#)
overview [12](#)

P

PARTITION_GROUP [101](#)
password administration [150](#)
password file [138](#)
passwords
 changing [150](#)
 setup script [150](#)
planning [16, 29, 33, 77](#)
planning
 capacity [16](#)
 checklist [16](#)
preinstallation checklist [45](#)
preinstallation planning [16](#)
prerequisite checker [16, 33](#)
prerequisite checker
 installing [33](#)
Prerequisite checker
 results [34](#)
Prerequisite Checker tool
 Uninstall [34](#)
prerequisites
 backup servers [38](#)
 hardware [32, 64](#)
 software [32, 64](#)
problem determination
 general procedure [87](#)
product support [87](#)
profile
 examples
 DB2 [171](#)
 DB2 HADR standby server [177](#)
 structure [96](#)
 target set definitions
 naming convention [56](#)
publications [9](#)

Q

query managed capacity (fmquery)
 description [154](#)

R

RECON_INTERVAL [111](#)
reconciliation
 FlashCopy backup [25](#)
remote mirror [76](#)
remote mirroring [28](#)

- repository
 - snapshot backup
 - status [147](#)
- REPOSITORY_LABEL [100](#)
- restore commands [140](#)
- restore procedure
 - DB2 [74](#), [74](#)
 - DB2 HADR server [74](#)
 - DB2 standby server [74](#)
 - GPFS [74](#)
- restoring [76](#)

S

- Service Management Console [87](#)
- setup script [46](#)
- setup script
 - description [148](#)
 - values [149](#)
- snapshot [12](#)
- snapshot
 - mounting [159](#)
 - unmounting [159](#)
- snapshot backup [82](#)
- snapshot backup
 - deleting [146](#)
 - modes with DPF partitions [67](#)
 - repository
 - status [147](#)
- snapshot devices
 - [24](#), [20](#), [20](#)
- SSH key file [50](#)
- status
 - repository [147](#)
- storage solutions
 - preparing [18](#)
- storage subsystems
 - log files [91](#)
 - preparing
 - [23](#), [23](#), [18](#)
 - setting up [48](#)
 - trace files [91](#)
- SVC [28](#), [50](#), [26](#)
- SVC_COPY_RATE [111](#)
- SVCDTA [22](#), [21](#), [26](#)
- SVCDTA SVC Migrating to new adapter [51](#)
- Synchronous Remote Mirroring [26](#)

T

- target set definitions [55](#)
- target set definitions
 - DB2
 - multi-partition [55](#)
 - files [55](#)
 - naming convention [55](#), [56](#)
- target volumes
 - storage systems [133](#), [135](#)
- target volumes file
 - examples
 - [172](#), [173](#), [173](#), [175](#), [175](#)
 - mirror setup [174](#)
- target volumes file (.fct)
 - description [132](#)
 - parameter settings
 - [135](#), [136](#), [136](#)
 - cloning [137](#)
- TARGET_NAMING [111](#)
- TARGET_SETS [111](#)
- TIMEOUT_PHASE [101](#)
- TRACE [99](#)
- trace files
 - [88](#), [92](#)
 - CIM [91](#)
 - storage subsystems [91](#)
 - summary [88](#)
- troubleshooting [93](#)
- troubleshooting
 - general procedure [87](#)
- TSM_BACKUP [101](#)

U

- uninstalling [41](#)
- Upgrade
 - production server [64](#)
- upgrading
 - backup server [44](#)
 - process [32](#)
- usability states [82](#)

V

- version 8.1.11 [9](#)
- VOLUMES_FILE [111](#)

© Copyright International Business Machines Corporation 2001, 2025

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

