

**Practical Guide:
Protecting SAP HANA with IBM
Spectrum Protect and IBM Spectrum
Copy Data Management**

By
Thomas Prause
IBM Germany Research & Development, Boeblingen

Version 1.0
February 13, 2020

© Copyright International Business Machines Corporation 2019, 2020. All rights reserved.
US Government Users Restricted Rights
Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

1 Overview and Objective	4
2 Product Overview	4
2.1 IBM Spectrum Copy Data Management	4
2.2 IBM Spectrum Protect for Enterprise Resource Planning . . .	4
3 Sample Environment	4
4 Requirements and Preparation	5
5 Installation	6
5.1 Installation of IBM Spectrum Copy Data Management	6
5.2 Installation of IBM Spectrum Protect for ERP	6
6 Configuration	7
6.1 Configuration of IBM Spectrum Copy Data Management . . .	7
6.2 Configuration of Data Protection for SAP HANA	8
6.3 Configuration of Expiration Policies	10
6.3.1 Retention Policy in SAP HANA	10
6.3.2 Retention Policy in IBM Spectrum Protect	11
7 Daily Operations	12
8 Recovery Scenarios	13
8.1 Restore Jobs in CDM	13
9 Recover a database to a Specific Point in Time	14
9.1 Snapshot Restore	14
9.2 Recovery of a Restored Snapshot	14
10 Variations of the Recovery Scenario	16
10.1 File System Check During a Restore Operation	16
10.2 Test of a Production Restore Operation	16
Notices and Trademarks	18

Feedback

Feedback is appreciated. Comments or suggestions for changes can be sent to prause@de.ibm.com.

Version Updates

Version 1.0 - initial version

1 Overview and Objective

This document provides guidance for protecting SAP HANA databases with IBM Spectrum Copy Data Management in combination with IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP HANA.

Mission critical databases typically have stricter requirements regarding protection against failures. The combination of the two products provides a comprehensive set of features that can be tailored to the specific needs of the particular environment. This document will showcase some examples in detail.

2 Product Overview

2.1 IBM Spectrum Copy Data Management

IBM Spectrum Copy Data Management (CDM) creates copies of data by using the hardware snapshot functionality of modern storage subsystems like IBM Spectrum Virtualize and makes those copies available for a variety of use cases including backup, recovery, and cloning.

2.2 IBM Spectrum Protect for Enterprise Resource Planning

IBM Spectrum Protect for Enterprise Resource Planning (ERP): Data Protection for SAP HANA supports backup and recovery operations for SAP HANA databases while using all the features of the IBM Spectrum Protect server. IBM Spectrum Protect for ERP can be used to backup both the SAP HANA database and the redo log files.

3 Sample Environment

To better illustrate how the solution can be deployed and used, a real-life system is described, and screen captures and command outputs are included. The environment used for this exercise is an example and not a recommendation to use specific product versions. The sample environment consists of the following components:

Database server

- Database software: SAP HANA 2.0 SPS 04 Revision 40
- Operating system: Red Hat Enterprise Linux (RHEL) Server release 7.7
- Storage: IBM Spectrum Virtualize Version 7.8.0

-
- IBM Spectrum Protect for ERP V8.1.9

CDM server

- IBM Spectrum Copy Data Management 2.2.10.0 (build 139)
- VMware vSphere 6.0
Any VMware version supported by CDM would work. At the time of publication, vSphere 5.1, 5.5, 6.0, 6.5, and 6.7 were supported.

IBM Spectrum Protect server

- IBM Spectrum Protect server V7.1
Any IBM Spectrum Protect server supported by the combination of IBM Spectrum Protect for ERP and the IBM Spectrum Protect application programming interface (API) would work.

To manage a SAP HANA instance, additional utilities are available. These utilities are optional. They do not have to reside on the previously mentioned systems.

- SAP HANA Cockpit Version 2.0.11.11.0
Any version supported by the HANA database is sufficient.
- SAP HANA Studio Version 2.3.43
Any version supported by the HANA database is sufficient.

For the described solution, SAP HANA Cockpit is the preferred utility because SAP note *2073112 - FAQ: SAP HANA Studio* states “Over time SAP HANA Studio will be replaced by SAP HANA Cockpit ...”

All systems in this sample environment are connected by using 1 Gb Ethernet. In a production environment, the database server and the IBM Spectrum Protect server are typically connected using more advanced technology like 10 GB Ethernet or Fiber Channel (using LAN-free backup operations).

4 Requirements and Preparation

Both IBM Copy Data Management and IBM Spectrum Protect for ERP fully support HANA on RHEL 7 on Linux on Power Systems running in little endian (LE) mode without any additional requirements or preparations.

Before SAP HANA was installed, dedicated volumes for the database persistence files and redo log files were created and assigned to the server (in accordance with the SAP HANA installation instructions). The corresponding file systems were mounted in the following locations: `/hana/data` and `/hana/log`. During installation of SAP HANA, these file systems were

selected for the database and redo log files. The dedicated file system for the redo logs is a requirement from SAP HANA. The log file system is not included in the snapshots and can reside on different storage than the data volume.

5 Installation

5.1 Installation of IBM Spectrum Copy Data Management

IBM Spectrum Copy Data Management is installed by deploying the *.ova file into a VMware vSphere environment. This procedure requires an IP address that will be assigned to the IBM Spectrum Copy Data Management instance.

For instructions, see *Installing IBM Spectrum Copy Data Management as a Virtual Appliance* in the IBM Knowledge Center.

https://www.ibm.com/support/knowledgecenter/en/SS57AN_2.2.7/

5.2 Installation of IBM Spectrum Protect for ERP

IBM Spectrum Protect for ERP: Data Protection for SAP HANA is installed by executing the installer binary on the HANA system. If a graphical user interface is available, the installer will start in GUI mode by default. The console mode can be selected by adding the option `-i console` as an argument.

For a scale-out system, this procedure must be performed on one HANA node only. Required binaries are distributed to additional hosts during configuration.

6 Configuration

6.1 Configuration of IBM Spectrum Copy Data Management

IBM Spectrum Copy Data Management gets configured through the web-interface. This is available at the URL `https://ip-address:8443` where `ip-address` is the IP address that was assigned to the CDM instance during deployment of the *.ova file.

First, the SAP HANA database host has to be registered. CDM requires *system credentials* and *database credentials* to register a HANA database system. See Figure 1.

The *system credentials* are required to run commands like mounting file systems or changing volume groups by using a Secure Shell (SSH) protocol. This user account must be configured to run commands as root through `sudo` without being asked for a password.

The *database credentials* are used to connect to the SAP HANA database. Such connection is required to gather details about the mount points used for persistence and trigger the snapshot in SAP HANA.

The screenshot shows the 'Register Application Server (SAP HANA)' dialog box. It contains the following fields and options:

- Site: Default (dropdown)
- Name: HANA Production (text input)
- Host Address: hanaprd01.syndicate.fun (text input)
- Port: 30013 (text input)
- Run Inventory job after registration
- Type: Virtual, Physical
- System Credential:

Name	Username	Type	
CDM Agent User	cdmagent	System	<input type="button" value="Delete"/>

- Database Credential(s):

Name	Username	Type	Instance	
HANA CDM Operator	hana_cdm_op	SAP HANA	Apply to all instances	<input type="button" value="Delete"/>

- Buttons: OK, Cancel

Figure 1: SAP HANA registration in CDM

A similar procedure must be performed for the storage system, in this case, IBM Spectrum Virtualize. The user must have privileges to create snapshots, map them to a host, and to delete volumes. Figure 2 shows the corresponding configuration pane.

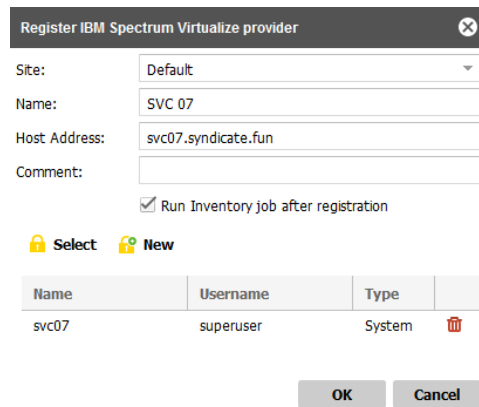


Figure 2: Registering IBM Spectrum Virtualize in CDM

Furthermore, the service level agreement (SLA) policies must be created in CDM. In this example, we use two SLA policies. The first policy creates a regular FlashCopy of the volumes and keep the copies for 5 days. The second one creates an incremental FlashCopy. In contrast to the incremental database backup, this will be a full database backup from the database perspective. But in this case the IBM Spectrum Virtualize will detect which data blocks have been updated since the previous FlashCopy and copy those blocks only. In the sample environment, only one incremental FlashCopy will be retained.

In addition, there are two corresponding jobs in CDM as shown in Figure 3. Each job uses one of the SLAs. The one for the full snapshot is scheduled to run at 3:30 PM each day. The incremental FlashCopy is scheduled to run at 11 AM each day.




Figure 3: CDM job definition

6.2 Configuration of Data Protection for SAP HANA

IBM Spectrum Protect for ERP: Data Protection for SAP HANA is configured by executing the setup script `/opt/tivoli/tsm/tdp_hana/setup.sh`. The script will first prompt for parameters regarding the the HANA instance that will be protected. Then, the script asks which IBM Spectrum

Protect server, node, and management classes to use. Figure 4 shows a screen capture of this part.

```
Pre-installation checks completed successfully. Press enter to start installation.
Do you want to setup the IBM Spectrum Protect configuration now [y/n] y
Please enter the number of full backups that shall be retained (MAX_VERSIONS): 5
Do you want to use automatic password handling (passwordaccess generate) [y/n] n
Please enter the IBM Spectrum Protect server name as defined in dsm.sys: BREMENTS01
Please enter the IBM Spectrum Protect node name: HANA_PLE
Please enter the number of Sessions to the IBM Spectrum Protect server you want to use: 2
Please enter the IBM Spectrum Protect management class(es) for database backups: MDBCONT1
The number of IBM Spectrum Protect management classes specifies the number of
backup copies of redo logs. Please enter the IBM Spectrum Protect management
class(es) for log file backups: MLOGCONT1 MLOGCONT2
Creating/updating keystore entries ...
```

Figure 4: Screen capture of setup.sh

In the sample configuration two management classes are specified for the redo logs. As a result, IBM Spectrum Protect for ERP creates two copies of each redo log on the fly. If at restore time a log is damaged or inaccessible, IBM Spectrum Protect for ERP would automatically switches over to restore data from the other copy without interrupting the recovery process. To create two copies simultaneously, two sessions are required, even though the database backup might be using a single session.

Finally, the script configures SAP HANA to use the BACKINT interface for the backup of database redo log files and the backup catalog.

It is not necessary or possible to schedule the backup of the redo logs and the backup catalog. HANA ensures that the redo logs are backed up once a log file becomes full or a specified length of time has elapsed. By default, this value is set to 15 minutes (HANA configuration parameter `log_backup_timeout_s` set to 900 seconds). The backup catalog itself is backed up after each backup operation, for example, after each full backup and each redo log backup.

Database backups are scheduled by using the IBM Spectrum Protect scheduler. Any other scheduler or the SAP HANA Cockpit could be used as well. Dedicated schedules are available for full and incremental backups of both the tenant database and the system database. In this sample configuration, full and incremental backups are shown in Figure 5, to illustrate this scenario. Technically, it might not be necessary to perform an incremental backup of the rather small HANA system database.

The sample configuration uses a generic name, `/hana/shared/backup_DB_full.sh`, for the backup script. This name is useful because a script with the same name can exist on different HANA systems. If one of those systems is associated with this schedule, the backup would run there as well. The different database names and users may be

```

Domain          *      Schedule Name      Action      Start Date/Time
-----
TST1_TDP_V3    HANA_FULL          CMD         10/21/2019 22:00:00
TST1_TDP_V3    HANA_INCR          CMD         10/21/2019 13:00:00
TST1_TDP_V3    HANA_SYSDB_FULL   CMD         10/21/2019 05:00:00
TST1_TDP_V3    HANA_SYSDB_INCR   CMD         10/21/2019 16:00:00

Protect: HAMBURGTSM02>q sched * hana_full f=d

      Policy Domain Name: TST1_TDP_V3
      Schedule Name: HANA_FULL
      Description:
      Action: Command
      Subaction:
      Options:
      Objects: /hana/shared/backup_DB_full.sh
      Priority: 5
      Start Date/Time: 10/21/2019 22:00:00

```

Figure 5: IBM Spectrum Protect schedule definitions

specified within the script (see Figure 6). So the scripts might be unique for each HANA system. But the schedule can still be used for multiple systems. The scheduler runs as root user. So the command in the script must switch the user before backing up the data to get the required environment set.

```

pleadm>cat /hana/shared/backup_DB_full.sh
#!/bin/bash

# full tenant db backup
/bin/su - pleadm -c "hdbsql -i 0 -U TSM -d ple \"backup data using backint ('full_$(date +%y%m%d)')\""
pleadm>

```

Figure 6: Sample backup script

6.3 Configuration of Expiration Policies

Two options are available to delete obsolete backup generations. SAP HANA can expire backups by using the retention policy settings in the SAP HANA Cockpit. The other option is to use the data expiration function of IBM Spectrum Protect for ERP. Do not confuse this function with the expiration functions in IBM Spectrum Protect copy groups. IBM Spectrum Protect for ERP provides its own mechanism to delete obsolete backup generations.

6.3.1 Retention Policy in SAP HANA

In the SAP HANA Backup Configuration (within SAP HANA Cockpit), a retention policy can be defined (see Figure 7). After the policy is specified, SAP HANA periodically cleans up the backup catalog and optionally deletes

expired backups in IBM Spectrum Protect using the BACKINT interface. The CDM snapshots are triggered outside of SAP HANA. Although the snapshots are listed in the SAP HANA backup catalog, they cannot be deleted by HANA. The SLAs defined in CDM have their own retention settings. The preferred method is to align the settings in CDM with the settings in SAP HANA so that SAP HANA does not delete backups that are still available in CDM. Keeping backup entries in SAP HANA that are no longer available in either CDM or IBM Spectrum Protect does not do any harm. But if a backup entry is deleted in SAP HANA, the backup can no longer be restored.

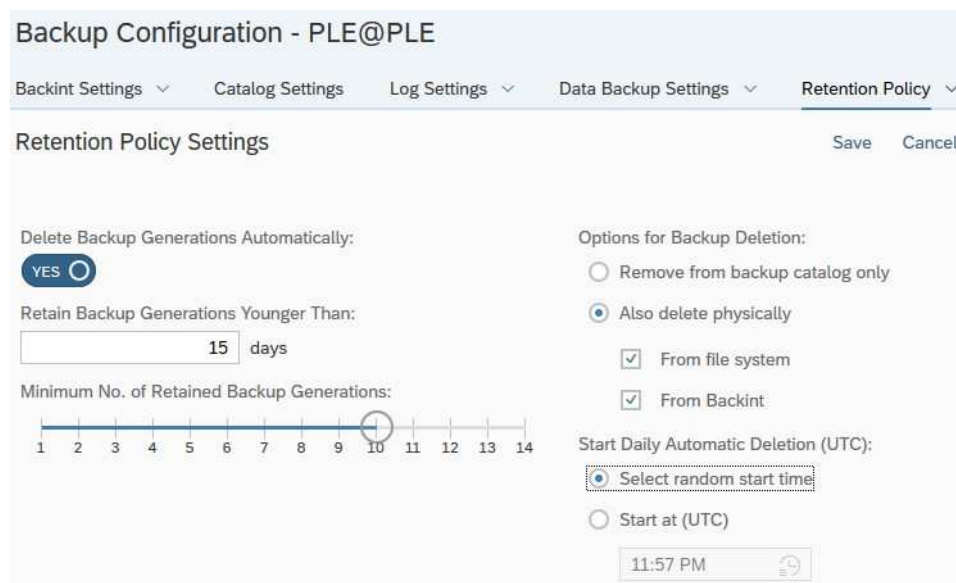


Figure 7: HANA retention policy

6.3.2 Retention Policy in IBM Spectrum Protect

During configuration of IBM Spectrum Protect for ERP, the number of backup generations to keep can be specified. This is set by the script as a value of the profile parameter `MAX_VERSIONS`. A backup generation spans a full database backup and all subsequent incremental or differential backups as well as all redo log backups until the next full database backup. The archive copy group of the management classes in the IBM Spectrum Protect server should be configured to not expire the data. In this way, IBM Spectrum Protect for ERP can control the expiration process.

If SAP HANA Cockpit is not used or it is not configured to expire and clean up old entries in the backup catalog, it is useful to periodically delete

obsolete entries from the backup catalog manually. Otherwise, the backup catalog continues to grow and this will impact the performance of backup and recovery operations significantly. All backup entries that are older than the oldest backup that is still available for restore (through CDM or IBM Spectrum Protect for ERP) are obsolete. All backups that expire can be removed from the backup catalog. In SAP HANA Studio select the oldest backup of destination type BACKINT that must be retained in the SAP HANA backup catalog. Right mouse click opens the context menu which has an option *Delete Older Backups ...* This in turn opens an wizard where it is possible to specify if backups shall be removed from the catalog only or also from the backup location. Since backups are expired and removed due to expiration it is sufficient to just remove them from the catalog. The same functionality is available from the backup catalog in SAP HANA Cockpit.

Use only one of the retention features. Setting retention policies in both SAP HANA and IBM Spectrum Protect for ERP might produce unexpected results.

7 Daily Operations

After both products are configured, the amount of required monitoring and maintenance is minimal. Backups are created according to the defined schedule. However, monitoring should be established to detect failed backups. SAP HANA will issue alerts in the cockpit when a backup fails or when no backup was made for some period of time. By default, a warning is issued if there is no database backup or the last full database backup is older than 7 days. This level of monitoring might be insufficient for production systems. Depending on business requirements, it might be useful to monitor whether the schedule is executed correctly in order to detect problems with the scheduler itself.

All backups will be listed in the SAP HANA backup catalog. Figure 8 shows a screen capture of the backup catalog in the SAP HANA Cockpit with snapshots from CDM as well as full and incremental backups from IBM Spectrum Protect for ERP.

SAP HANA can handle a situation in which backups recorded in the catalog no longer exist, regardless of whether the snapshots were created through the BACKINT interface with IBM Spectrum Protect or if they are snapshots in the storage system. The actions described in Chapter 6.3 are intended to keep the backup catalog in sync with the existing backups. But in case of restore operation, HANA checks whether the backup subject to restore is still available before attempting the restore operation.

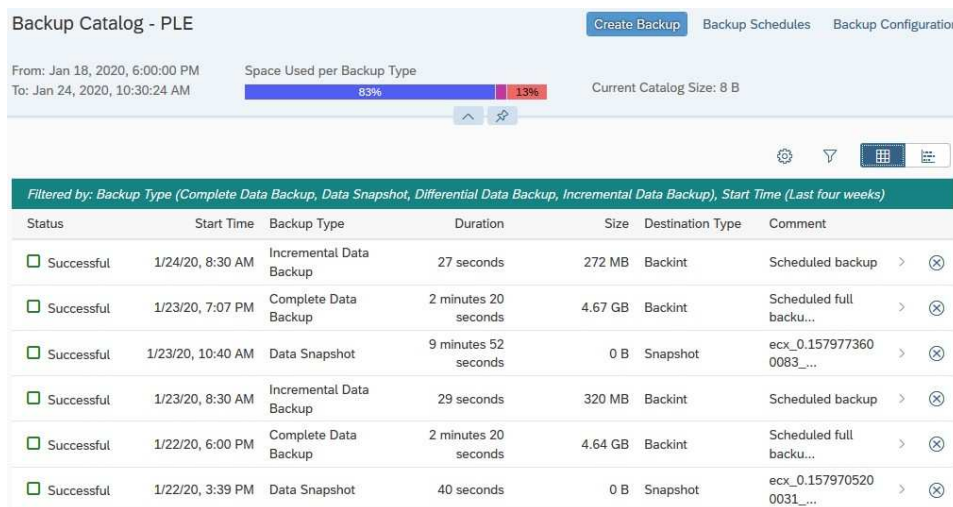


Figure 8: HANA backup catalog

8 Recovery Scenarios

This document focuses on the combination of IBM Spectrum Copy Data Management and IBM Spectrum Protect for ERP capabilities. For this reason, the document does not cover the restore of full database backups from IBM Spectrum Protect nor the recovery of only a snapshot backup. Furthermore, it is assumed that the hardware used for the SAP HANA database is still intact and all required database and redo log backups in both CDM and IBM Spectrum Protect are available.

If only a specific tenant database must be recovered, this can be accomplished by using the IBM Spectrum Protect for ERP backups. With SAP HANA SPS 04, snapshot backup and recovery of an individual tenant are not supported because using a snapshot backup for recovery will affect the system database and all tenant databases.

8.1 Restore Jobs in CDM

IBM Spectrum Copy Data Management supports two types of restore jobs: *instant disk restore* and *instant database restore*. As the name implies, an instant database restore operation restores an immediately usable database. This means that the database will be recovered as part of the restore operation. The database will contain the data as of the time when the snapshot was created. The advantage is that the database can be used right after the restore is completed. But no further logs can be applied. This type of restore is useful for scenarios like database cloning and will not be discussed further in this document.

The *instant disk restore* operation restores a snapshot without running a recovery operation. The restored volumes can then be used to run a HANA database recovery that restores and applies redo logs in order to recover the database to a specific point in time.

9 Recover a database to a Specific Point in Time

To recover a database to a specific point in time, both a full database backup and subsequent redo logs are required. With IBM Spectrum Protect for ERP, this restore and recovery can be completed in a single step using the SAP HANA Studio or SAP HANA Cockpit. But a full database restore operation might take several hours. By restoring a snapshot with CDM, the recovery time can be shortened dramatically to some minutes.

The restore operation occurs in two steps. First, the snapshot must be restored by using CDM. Second, the redo logs are restored from the IBM Spectrum Protect server by using the SAP HANA Cockpit.

9.1 Snapshot Restore

To restore a snapshot created with CDM, a job of type *restore* must be created. The *Instant Disk Restore* template must be selected for the job. The attributes of this job define which system is restored to which target. To restore a snapshot in a production environment, the source and target must be identical.

In the *Advanced Options* pane further options can be configured. If the *Revert* option is set to *User Selection*, CDM will first attach the volumes to the host and mount them under a new path. Later you can decide to revert the content of the snapshot volumes back to the original volumes. This may be useful when you first want to confirm that the copied volumes are intact before you remove the original volumes. If the *Revert* option is set to *Enabled*, CDM will immediately restore the backup to the source volumes by performing a reverse Flash-Copy operation which is considered a production restore. The operation is shown in Figure 9.

When the job definition is saved, it can be run immediately to perform a restore. But ensure that the database instance subject to restore (system and all tenant databases) is shut down. The job takes a few minutes to complete. Then, the recovery of the database can start.

9.2 Recovery of a Restored Snapshot

After the volumes are successfully restored, the system database and all tenant databases must be recovered. In contrast to the CDM *instant database*



Figure 9: Restore job definition

restore job, an *instant disk restore* job does not recover the database. The instant disk restore job makes it possible to restore the database to any point in time after the creation of the backup or snapshot. This is achieved by applying redo logs. Even different recovery points for the system and tenant databases would be possible. Keep in mind that for the system database a recovery to a specific point in time requires the use of Structured Query Language (SQL) command. The SAP HANA Cockpit can be used to recover the system database to the most recent state only.

This type of recovery is started in the SAP HANA Cockpit by selecting the *Recover Database* function in the *System Overview*. In the first pane, you can specify whether the recovery should proceed until the latest point in time available (end of logs) or to a specific point in time. When advancing further through the Recovery Wizard pages, a list of available backups as shown in Figure 10 is presented.

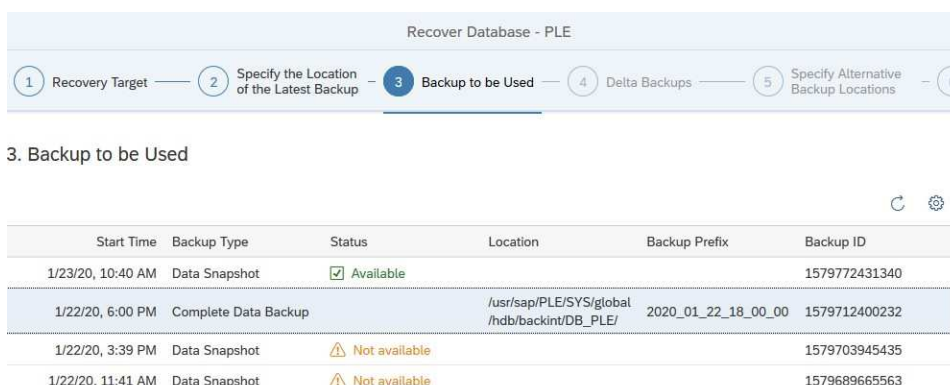


Figure 10: List of backups available for recovery

in this case, HANA detected that the snapshot is available in the file system. When this entry is selected, there will be no further restore action. HANA will use the snapshot available in the file system and apply redo logs to recover the database to the selected point in time. If instead the backup on IBM Spectrum Protect (*Complete Data Backup*) would be selected, HANA would restore this backup by using IBM Spectrum Protect for ERP first.

After the system database is successfully recovered, tenant databases must be recovered as well. With SAP HANA SPS 04, it is not possible to recover individual tenants out of a snapshot. The recovery of the tenants is performed in the same way as the recovery of the system database as described previously.

10 Variations of the Recovery Scenario

10.1 File System Check During a Restore Operation

An *instant disk restore* job creates copies of the volumes that are subject to restore and those copies can be used to test the restore operation. For this test, a dedicated job should be created that has the *Revert* option set to *Disabled*. The job should not use the original mount point, but should append a time stamp. CDM first creates new copies of the volumes subject to restore. These new volumes are then added to a new volume group on the target host and will be mounted under the selected mount point as shown in Figure 11.

```
[root@ulmhana5 ~]# ls -l /hana/data*/PLE/mnt00001/h*/d*
-rw-r--r-- 1 pleadm sapsys 1883504640 Jan 23 10:40 /hana/data_1579774348863/PLE/mnt00001/hdb00001/datavolume_0000.dat
-rw-r--r-- 1 pleadm sapsys 5352718336 Jan 23 10:40 /hana/data_1579774348863/PLE/mnt00001/hdb00002.00003/datavolume_0000.dat
-rw-r--r-- 1 pleadm sapsys 201359360 Jan 23 10:40 /hana/data_1579774348863/PLE/mnt00001/hdb00003.00003/datavolume_0000.dat
-rw-r--r-- 1 pleadm sapsys 1883504640 Jan 23 11:15 /hana/data/PLE/mnt00001/hdb00001/datavolume_0000.dat
-rw-r--r-- 1 pleadm sapsys 5352718336 Jan 23 11:08 /hana/data/PLE/mnt00001/hdb00002.00003/datavolume_0000.dat
-rw-r--r-- 1 pleadm sapsys 218103808 Jan 23 11:09 /hana/data/PLE/mnt00001/hdb00003.00003/datavolume_0000.dat
[root@ulmhana5 ~]# df /hana/data*
Filesystem                1K-blocks    Used Available Use% Mounted on
/dev/mapper/PLEdata_vg-PLEdata_lv 26201600 7103468 19098132 28% /hana/data
/dev/mapper/PLEdata_vgSam-PLEdata_lv 26201600 7087424 19114176 28% /hana/data_1579774348863
[root@ulmhana5 ~]#
```

Figure 11: File system attached

At this stage, it is possible to confirm that the file system of the copied volumes is intact and contains the expected content before destroying the original file system. To complete an actual restore operation, CDM can be instructed to revert the copies to the original volumes, as shown in Figure 12.

10.2 Test of a Production Restore Operation

Another potential scenario is to run a production-like recovery job for testing purposes only.

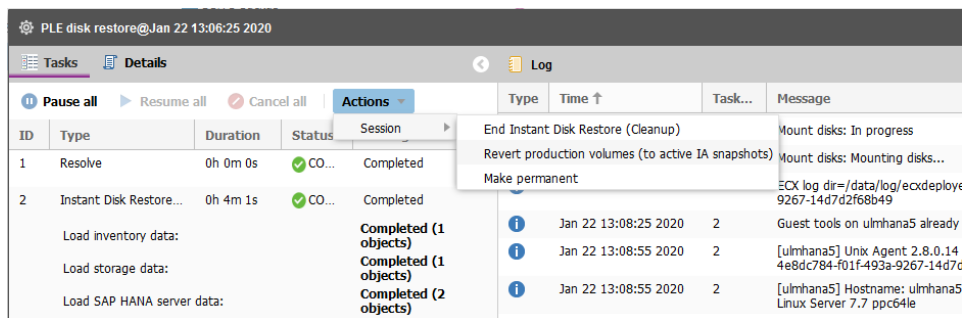


Figure 12: Revert volumes

In this case, the *instant disk restore* job must have the *Revert* option set to *Disabled*, and the file system must be mounted to the original mount point. In order to run the job the database instance (including all tenants) must be shut down and the file system for the data files must be unmounted. CDM will mount the copied volumes under the original path. In this way HANA can recover the data residing on these volumes as for a production recovery job. This has the advantage that a complete recovery could be tested during downtime of the production database without impacting the original volumes of the database. This test can be used to estimate the amount of time required to recover the database in a production situation. The required time depends not only on the restore speed of the redo logs. The time also depends on the speed of the database server and storage system while applying the logs and other factors. For this reason, you cannot precisely calculate the required time. However, you can obtain a realistic estimate of the time by running the described test.

After the test is finished, you must select the *End Instant Disk Restore (Cleanup)* option to remove the temporary database volumes. When the cleanup is completed, the original volume can be mounted and the database can be restarted to resume normal operations.

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 provide 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 client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

The information herein is subject to change before the products described become available.

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 actual people or business enterprises 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.

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.
- Red Hat and Red Hat Enterprise Linux are registered trademarks of Red Hat, Inc. in the U.S. and other countries.
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- VMware vSphere is a registered trademark of VMware, Inc. in the United States and certain other countries.
- All other trademarks are the property of their respective owners.