

PowerProtect Data Manager: SAP HANA Agent Backup and Recovery

Abstract

This white paper focuses on protecting SAP HANA Database using PowerProtect Data Manager, the next generation data protection platform.

November 2020

Revisions

Date	Description
June 2020	Initial release
November 2020	Revised

Acknowledgment

Author: Sonali Dwivedi

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [11/9/2020] [White paper] [H18578 | PowerProtect Data Manager SAP HANA Database Backup and Recovery]

Table of contents

Revisions	2
Acknowledgment	2
Table of contents	3
Executive summary.....	5
Audience	5
1 Introduction.....	6
2 Architecture	7
2.1 SAP HANA System	7
2.2 BACKINT for SAP HANA Interface	8
2.3 SAP HANA HDBSQL.....	8
2.4 Protection Storage.....	8
2.5 DDBMCON	8
2.6 SQLITE.....	9
2.7 Brownfield Support.....	9
2.8 PowerProtect Data Manager Agent Service.....	9
3 Supported Configurations and Platforms.....	10
4 Installation and Configuration of SAP HANA agent.....	12
4.1 SAP HANA agent Installation.....	12
4.1.1 Integrate the SAP HANA backup program (backint):	12
4.2 Configuration.....	13
4.2.1 Discovery	13
4.2.2 Credentials.....	13
4.2.3 Protection Policy	14
5 SAP HANA Database Backup.....	16
5.1 Centralized Protection Policy Backup Workflow	17
5.2 Self-Service Protection Policy Backup Workflow	18
6 SAP HANA Parallelism	21
7 SAP HANA Backup Discovery Process.....	22
8 SAP HANA Database Recovery.....	23
8.1 Performing a Restore using SAP HANA Studio.....	23
8.2 Performing alternate host recovery	28
8.3 Performing a Restore using the SAP HANA CLI	29
9 Cloud Tiering, Replication, and Backup Deletion.....	30
9.1 Cloud Tiering.....	30

Table of contents

9.2	Replication	30
9.3	Backup Deletion	31
10	Disaster Recovery	32
10.1	Recovering the PowerProtect Data Manager Agent service from a disaster	32
	Conclusion.....	33
	References	33

Executive summary

Like other business critical applications and data structures within a company's IT environment, SAP data is experiencing tremendous growth. This growth is one of the difficulties facing SAP Administrators today. When coupled with shrinking backup windows due to a worldwide user base expecting a 24x7 highly available database and more stringent SLAs, finding the time that is dedicated solely for backup and recovery operations seems nearly impossible.

Traditionally, large SAP environments have not been a good fit for a client-side deduplication backup method because of limited backup windows and matching backup performance required to meet these windows. High CPU requirements on the production host did not provide the performance required without the trade-off of significant negative effects to the production-processing load.

Dell EMC PowerProtect Data Manager featuring our newly introduced application plug-in for SAP HANA now enables SAP and backup administrators effective, flexible, and efficient backup and recovery leveraging deduplication. The SAP HANA agent enables an application administrator to protect and recover the SAP HANA data on the application host. PowerProtect Data Manager integrates with the SAP HANA agent and enables central scheduling for the SAP HANA backups.

In this white paper, we will review the various features contained in this new client; components for SAP HANA backup in PowerProtect Data Manager which is simple, flexible, efficient, and optimized backups for your enterprise-wide SAP HANA deployment.

Audience

This white paper is intended for those DBAs and Backup Administrators looking at deploying efficient backup solutions for enterprise-wide- SAP HANA business critical implementations. Knowledge of SAP HANA and its overall architecture including its integrated backup and recovery tools, terminology, and methods are required.

1 Introduction

PowerProtect Data Manager offers customers the ability to set protection policies, to perform backup and recovery, deduplicate workloads so forth. PowerProtect Data Manager offers centralized oversight of all protected file system copies which makes it simple to track and enforce SLO compliance for backup and recovery, RPOs, and Storage retention lock. It also discovers copies sent to protection storage, catalogs and make protection copies available for compliance measurement to ensure protection compliance and quality of service.

PowerProtect Data Manager UI is a simple wizard for managing entire protection life cycle and designed for creating protection life cycle orchestration, Automated protection policies and Service level agreements (SLA) compliance and actionable insight that optimizes efficiency.

The purpose of this white paper is to show how SAP HANA protection can be done using Dell EMC PowerProtect Data Manager with Dell EMC storage systems as target storage. Use PowerProtect Data Manager with the application agent to perform the following operations:

- Automate the configuration of the application agent backup policy and protection storage settings.
- Create a catalog of backups that are created by the application agent, and then monitor that catalog data to determine if retention policies are being adhered to.
- Manage the life cycle of backups that are run by the application agent—that is, ensure that the backups are marked for garbage collection based on the rules of the retention policy.

PowerProtect Data Manager does not change the way that the application agent works. DBAs create the backups and perform restores. The SAP HANA agent enables an application administrator to protect and recover the SAP HANA data on the application host.

PowerProtect Data Manager integrates with the SAP HANA agent and enables central scheduling for the SAP HANA backups. The SAP HANA agent installation is a command-line process whereby the user installs the required SAP HANA agent and PowerProtect Data Manager software. PowerProtect Data Manager then sets the PowerProtect DD Series Appliance Host, storage unit, user, and password. PowerProtect Data Manager supports the coexistence of the SAP HANA agent and the File System agent on Linux.

2 Architecture

In this section, we will discuss some of the components that will be used in the SAP HANA agent backup solution with PowerProtect Data Manager.

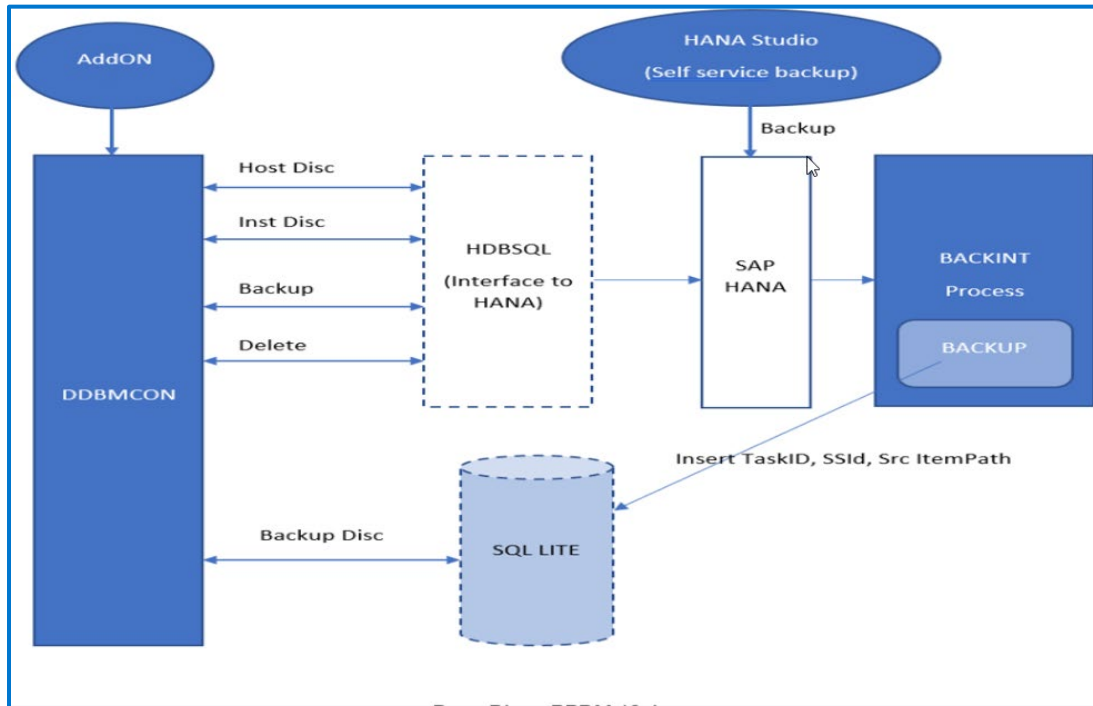


Figure 1 SAP HANA Application Agent Components

2.1 SAP HANA System

SAP HANA architecture is a revolution in columnar database application has been brought about by SAP HANA. It combines the features of in-memory computing and columnar database management system (DBMS), a duo that tremendously enhances speed of operation when performing real-time analysis or running real-time applications. In-memory processing fundamentally uses the system's RAM (main memory) to store data, instead of using disk storage. A columnar DBMS boosts system performance and speed when working with aggregates, by reading only the data that is required to satisfactorily fulfill a query, through effective compression of similar columnar data. This technology effectively deploying this combination thus becomes one of the best suited tools for data warehousing applications.

SAP HANA database consists of multiple servers, for example, name server, index server, preprocessor server, and so on. The databases in an SAP HANA system run different combinations of these servers. The most important server is the index server. It contains the data stores and the engines for processing the data and runs in every tenant database. An SAP HANA system consists of multiple isolated databases and may consist of one host or a cluster of several hosts.

An SAP HANA system is identified by a single system ID (SID) and contains one or more tenant databases and one system database. Databases are identified by a SID and a database name. From the administration perspective, there is a distinction between tasks performed at system level and that performed at database level. Database clients, such as the SAP HANA cockpit, connect to specific databases.

2.2 BACKINT for SAP HANA Interface

You can back up and recover an SAP HANA database using an SAP-certified third-party tool that supports the Backint for SAP HANA interface, which is used to communicate with an SAP HANA database. It has following characteristics:

- Each active host in a distributed SAP HANA system may have one or more volumes to be backed up. When `Backint` for SAP HANA is used to back up a database, several communication processes are started, one for each volume. Backint-based data backups and log backups can be created in parallel.
- A third-party backup tool reads data to be backed up from named pipes and writes the data in accordance with the tool configuration. For a third-party backup tool, the only objects created in the file system are named pipes. Named pipes occupy no space in the file system.

2.3 SAP HANA HDBSQL

SAP HANA `hdbsql` is a command-line tool for running commands on SAP HANA databases. Using SAP HANA `hdbsql`, you can perform SQL statements and database procedures, as well as query information about the database and database objects. SAP HANA `hdbsql` is installed with the SAP HANA software. It accesses databases both on your local computer and on remote computers.

Call SAP HANA `hdbsql` with the command `hdbsql [options]` from the following location `/usr/sap/<SID>/HDB<instance>/exe`. You can perform individual commands interactively or noninteractively. It is also possible to import commands from a file and run them in the background.

2.4 Protection Storage

The first step in SAP HANA integration with PowerProtect Data Manager is to ensure that you have a protection storage target that is configured to store the backup data. You can accomplish that by adding either single or multiple PowerProtect DD Series Appliance. You can add the protection storage from PowerProtect Data Manager UI.

PowerProtect DD Series Appliance operating system version 6.1 or later and the PowerProtect DD Management Console (DDMC). All models of PowerProtect DD Series Appliance are supported.

2.5 DDBMCON

PowerProtect Data Manager Agent service interacts with process called `ddbmcon`, and it is responsible for interacting with SAP HANA system using `hdbsql` query.

To initiate Backup using centralized protection policy, `ddbmcon` issues a `hdbsql` command for Backup to SAP HANA which will trigger backup using `backint` process. During the discovery process, `ddbmcon` queries SAP HANA catalog and SQLite to get the list of backups for recovery process and updating the PowerProtect Data Manager catalog.

2.6 SQLITE

SQLite is a C library that provides a lightweight disk-based database that does not require a separate server process and allows accessing the database using a nonstandard variant of the SQL query language. Some applications can use SQLite for internal data storage. It is used in this solution to store information about all types of backups like self-service, centralized and automatic log backup for PowerProtect Data Manager and runs in following sequence:

- Every backup request initiated by sysadmin must go through `backint` process, so `backint` process knows `Saveset_ID` and full path of the file that got written to DD. Backint Process records this information in SQLite on that machine. It is important to note that a single backup job writes several files to DD series so many `Saveset_ID` constitutes single backup. `Backup_ID` is responsible for grouping several `Saveset_ID` in to one backup job. As a result, single backup records several rows in SQL.
- In Multihost system, every machine will have its own `backint` process and SQLite instance. Every host must store the SQLite on shared location (), so that file is accessible from other host.

2.7 Brownfield Support

Brownfield support means if customer was using database application agent and they want to switch to PowerProtect Data Manager SAP HANA agent then they can do it easily by making few changes in configuration files and all backups taken from application direct SAP HANA agent will be valid and will automatically get updated in PowerProtect Data Manager after discovery process is complete.

To enable the PowerProtect Data Manager to discover the SAP HANA backups that was performed by the database application agent, you must update a specific configuration file called `brownfield.utl`.

Note: In the PowerProtect Data Manager UI, ensure that the assets are visible. In the directory `/opt/dpsapps/dbappagent/saphana/brownfield`, rename the `filesap_hana_brownfield_ddbda.utl` as `brownfield.utl`.

2.8 PowerProtect Data Manager Agent Service

- The PowerProtect Data Manager agent service is a REST API-based service that is installed by the application agent on the application host.
- The agent service provides services and APIs for discovery, protection, restore, instant access, and other related operations.
- The PowerProtect Data Manager uses the agent service to provide integrated data protection for the application assets.
- The PowerProtect Data Manager agent service provides important functionality for the application agent operations with the PowerProtect Data Manager.
- The agent service is started during the agent installation by the installer. The agent service runs in the background as a service and you do not interact with it directly.

3 Supported Configurations and Platforms

Following table shows the different topologies supported by PowerProtect Data Manager for SAP HANA Application Agent:

Classification	Features	PowerProtect Data Manager Support
Configuration	Single Database Container (SDC) Single node	Yes
	Single Database Container (SDC) Multi node	Yes
	Multiple Database Container (MDC) Single nodes	Yes
	Multiple Database Container (MDC) Multiple nodes	Yes
	High Isolation Multiple Database Container (MDC)	Yes
	SSL Support for SAP HANA SSL Connections	No
Backup	Full Backup	Yes
	Incremental Backup	Yes
	Differential /Cumulative Backup	Yes
	Snapshot Backup	No
	Dynamic Filters for Automatic Provisioning and Scheduling	Yes
	Data and Log in Separate Storage Units	No
Restore	Full Backup Recovery	Yes
	Full + Differential Backup Recovery	Yes
	Full + Incremental Backup Recovery	Yes
	Full + Differential + Incremental Backup Recovery	Yes
	Centralized Restore using RestAPI	No
	Disaster Recovery	Yes

Classification	Features	PowerProtect Data Manager Support
Database	System Database	Yes
	Tenant Database	Yes
PowerProtect DD Support	Physical PowerProtect DD	Yes
	PowerProtect DD Virtual Edition (DDVE)	Yes
	PowerProtect DD Cloud Tier	Yes
Platform Support	SLES *86_64	Yes
	Red Hat Enterprise Linux *86_64	Yes
Linux Versions	Red Hat Enterprise Linux 6.1 until Red Hat Enterprise Linux 8.1	Yes
	SUSE until SLES 15 SP1	Yes
SAP HANA Version	SAP HANA 1.x until 1.0 SP12	Yes
	SAP HANA 2.X until 2.0 SP4	Yes
Cloning	Clone to DD	Yes
	Clone to Cloud	Yes

4 Installation and Configuration of SAP HANA agent

Verify that the environment meets the following requirements:

- Ensure that all clocks on both the SAP HANA host and PowerProtect Data Manager are time-synced to the local NTP server to ensure discovery of the backups.
- Ensure that the SAP HANA host and the PowerProtect Data Manager network can see and resolve each other.
- Ensure that port 7000 is open on the SAP HANA host.

4.1 SAP HANA agent Installation

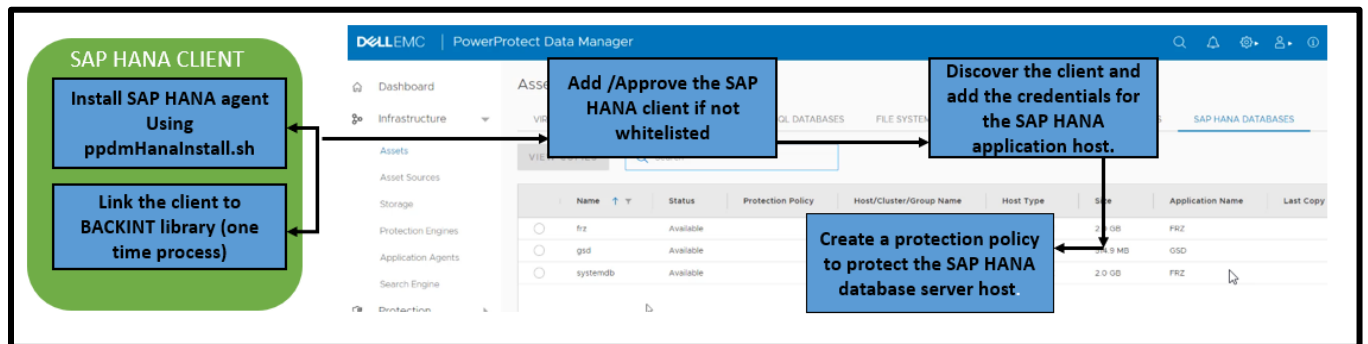
Follow the steps to install the SAP HANA agent:

- Download the SAP HANA agent software package to the SAP host. In the Data Manager UI:
 - Select Agent Downloads from the System Settings menu.
 - Select the SAP HANA agent download package (`dbappagent194_linux_x86_64.tar`).
 - Download the package in the location that you want to install the SAP HANA agent.
- Untar the installer by running `gunzip *` followed by `tar -xvf dbappagent194_linux_x86_64.tar`.
- Run the installation script: `ppdmHanaInstall.sh` as a user root.
- Three installation rpms are installed as part of the installation script.
- Enter the PowerProtect Data Manager server IP address:
`ppdmHanaInstall.sh --server=<PowerProtect_Data_Manager_server_IP_address>`
 - For upgrade `ppdmHanaInstall.sh --upgrade`

4.1.1 Integrate the SAP HANA backup program (`backint`):

- SAP HANA agent installation places the `hdbbackint` program in the `/opt/dpsapps/dbappagent/bin` directory.
- SAP HANA requires the program to be accessible from the `/hana/shared/<SID>/global/hdb/opt` directory. You can either copy the `hdbbackint` file to that directory or create a symbolic link that points from `/hana/shared/<SID>/global/hdb/opt` to the executable file. For example: `ln -s /opt/dpsapps/dbappagent/bin/hdbbackint /hana/shared/S01/global/hdb/opt/hdbbackint`

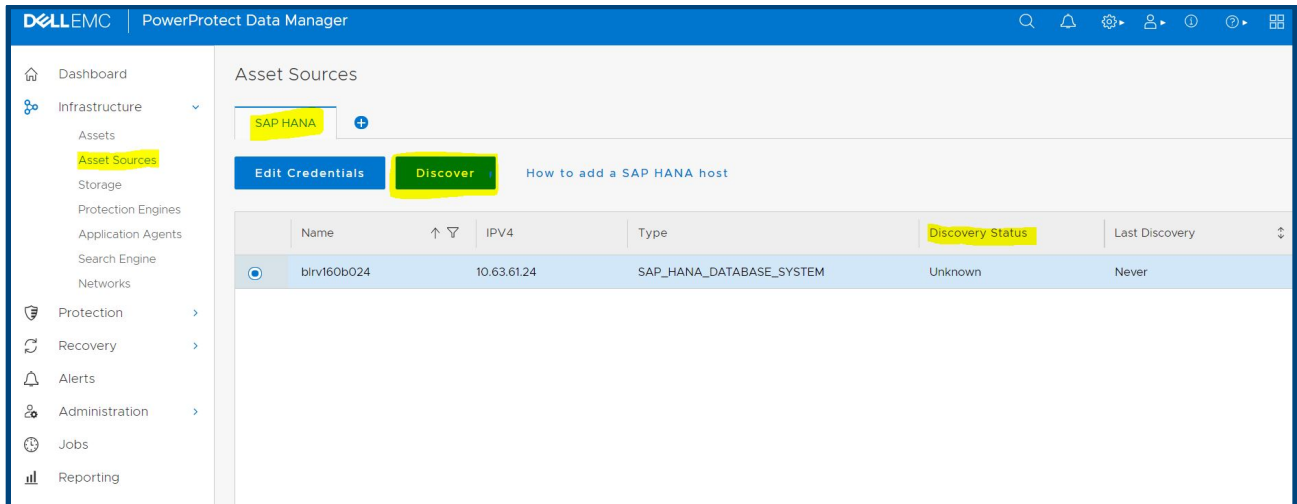
Note: In a multinode environment, install the SAP HANA agent in every node. Then add and discover the application host in PowerProtect Data Manager on each node in the SAP HANA multinode environment.



4.2 Configuration

4.2.1 Discovery

After you register an application host with PowerProtect Data Manager, you can use the Asset Sources window to discover an application host and modify the application host credentials. For application hosts, discovery is required if you want to schedule a backup.



4.2.2 Credentials

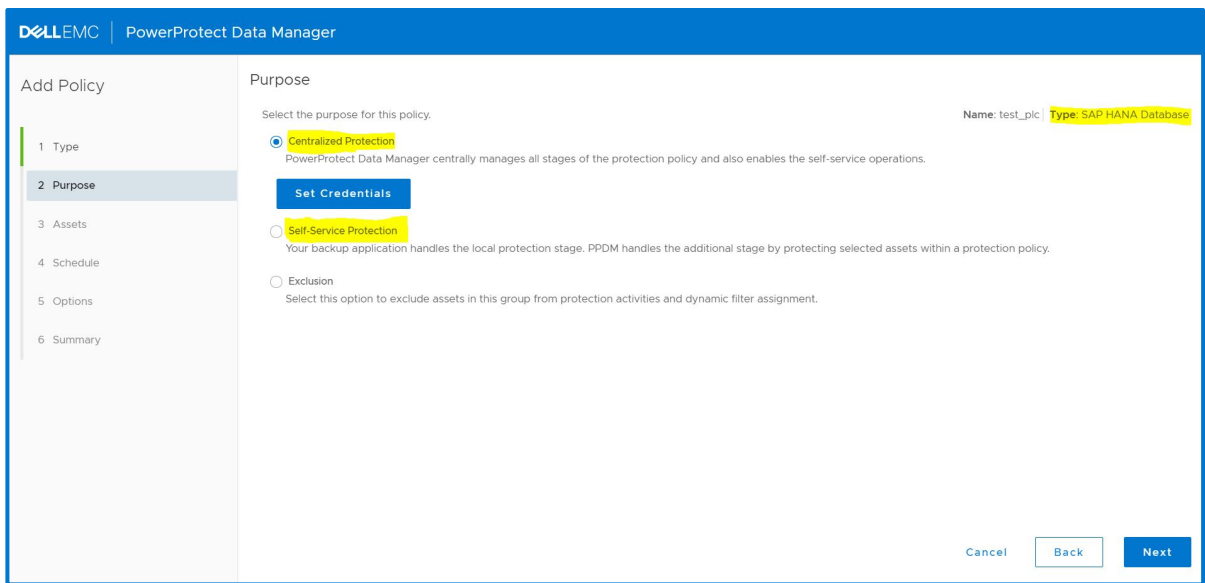
You must add credentials to the SAP HANA database so that PowerProtect Data Manager can access the database to create backups. Following points explain the requirements for credentials:

- Credentials that you set at the host level supersede the credentials that you set at the protection policy level.
- An SAP HANA instance can have one system database `SystemDB` and multiple tenant databases. The `SystemDB` database and each tenant database can have its own credentials. A tenant database can be backed up with either its own credentials or the `SystemDB` credentials. Both the `SystemDB` database and tenant databases support two types of credentials:
 - **SAP HANA Database** credentials, for which you must specify the database username and password.
 - **SAP HANA User key** credentials, for which you must specify the user key (created with `hdbuserstore` command) and operating system username of the SAP HANA instance.
- To create credentials, click Add to open the Add Credentials dialog box, where you can set the authentication type as either Username/Password or User Key and then specify the credential values. For a tenant database, you can specify the credentials of either the `SystemDB` database or the tenant database itself. When you specify the `SystemDB` credentials for a tenant database, select the `IsSystemDB` check box; otherwise, clear the `IsSystemDB` check box.

4.2.3 Protection Policy

Use the PowerProtect Data Manager UI to add a protection policy to protect the SAP HANA databases using following instructions:

- For an SAP HANA database, you can select from three types:
 - To use PowerProtect Data Manager to manage all protection centrally, select Centralized Protection
 - To use SAP HANA to create local backup protection, select Self-Service Protection. PowerProtect Data Manager creates a protection policy and manages extra stages.
 - If there are SAP HANA assets within the protection policy that you plan to exclude from data protection operations, select Exclusion.
- When you are creating protection policies for SAP HANA multinode databases, ensure that all nodes in the multinode environment are powered on. Register each node, one after the other, with an interval of a few minutes between each node. Otherwise the protection might fail.
- You cannot add assets from a single SAP HANA host in two or more protection policies with the same PowerProtect DD Series Appliance. You must add additional protection policies for a different PowerProtect DD.
- You can add a cloud tier schedule to a protection policy for SAP HANA centralized and self-service backups to cloud tier. Both SAP HANA centralized and self-service protection policies support cloud tiering. You can create the cloud tier schedule from both primary and replication stages. Schedules have a weekly or monthly recurrence and a retention time of 14 days or greater.

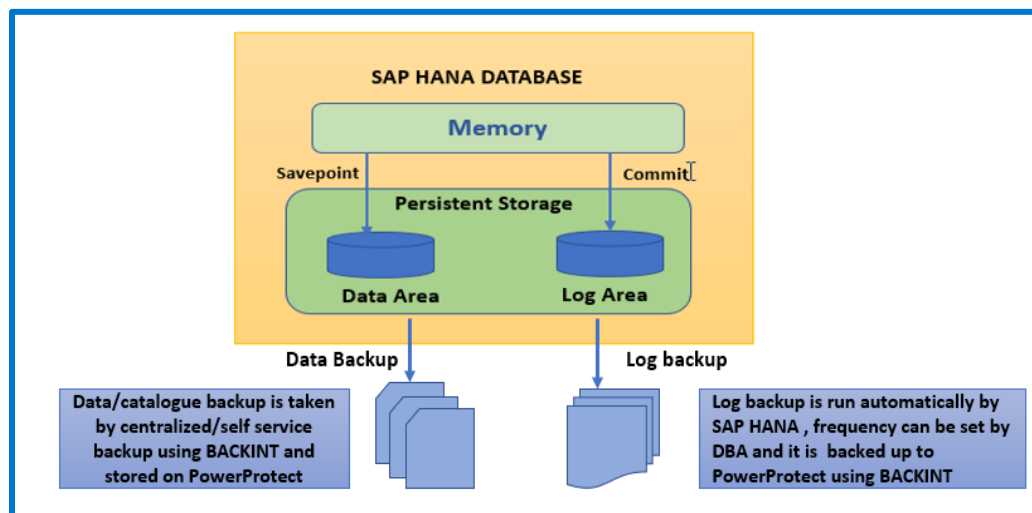


Note: For detailed configuration steps, refer to [PowerProtect Data Manager for SAP HANA Agent User Guide](#)

5 SAP HANA Database Backup

PowerProtect Data Manager can manage and monitor data protection and replication for SAP HANA assets through integration with the SAP HANA agent. In SAP HANA environment, there are three backups that needs to be run. Each of this gets backed up a separate backup and it will have a `backup_ID` (stored in `M_BACKUP_CATALOG_FILES`). HANA maintains different catalog `SystemDB`; every `TenantDB` will have a separate catalog file. Items to backup are as follows:

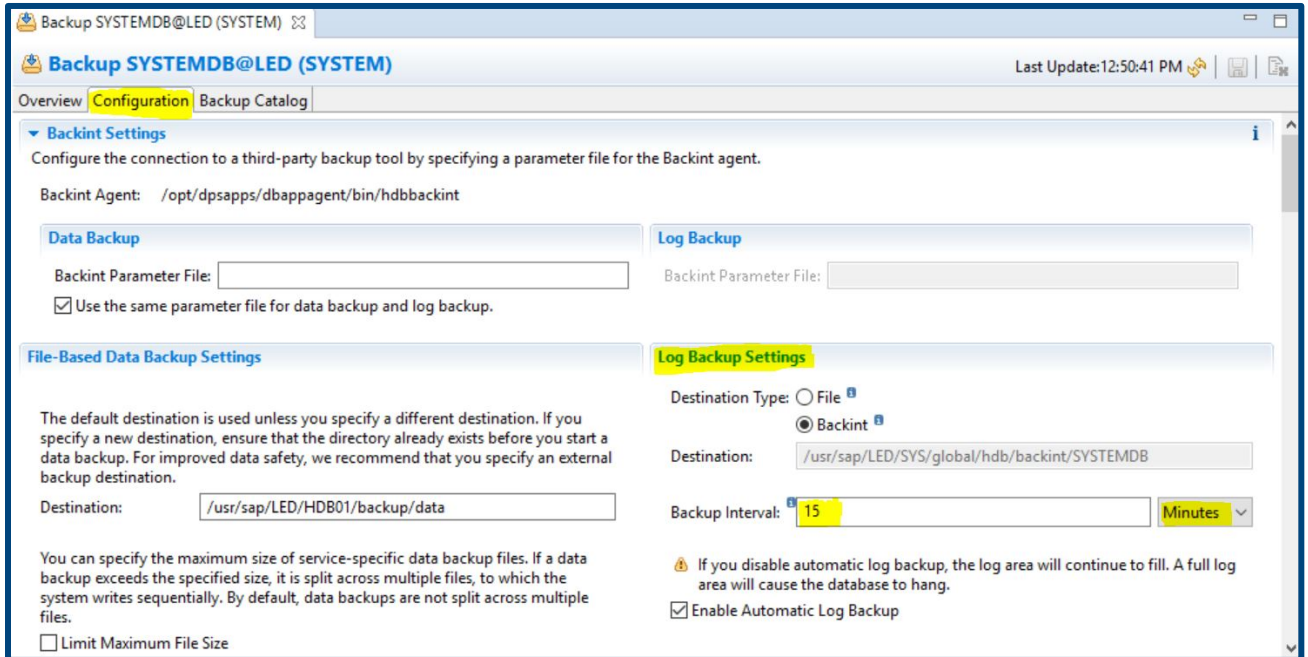
- **Data backup:** A data backup includes all the data (System Database, Tenant Database) that is required to recover the database to a consistent state. With a data backup, only the data is backed up; unused space in the database is not backed up. The data area is backed up in parallel for each of the SAP HANA services. If SAP HANA is running on multiple hosts, a data backup includes all the service-specific backup parts for all the hosts.
- **Log files:** Log backup is triggered by HANA automatically and sent directly to Protection Storage. User can control the frequency of Log backup. Frequency can be modified using HANA Studio. This configuration is per database.
- **Catalog:** HANA launches two backups in Full backup command they get performed right after each other. Full backup of data and once that is complete, Catalog backup.



- PowerProtect Data Manager creates SAP HANA backups and manages remote replication copies based on the Protection Policy (PLC). PowerProtect Data Manager performs the backup and replication operations based on the protection policy and governed by the SLA. SAP HANA can be backed up through:
 - Automatic backup by the PowerProtect Data Manager Centralized Protection policy.
 - Manual backup by the SAP HANA Studio/SAP HANA command line and governed by the PowerProtect Data Manager Self-Service Protection policy.
- Types of backup available with PowerProtect Data Manager SAP HANA Agent: Full data backup, Incremental backup and Differential backup. The SAP HANA database server automatically performs periodic backups of the redo logs for a database. You do not run the log backups with any backup commands as is the case for database backups.
- SAP HANA 2.0 SPS 00 supports the user-configurable interval mode for log backups. The Interval mode enables the creation of log backups after the service-specific timeout is reached,

instead of when the log segment becomes full. A log backup can include multiple log segments.

- The interval mode setting limits the number of log backups in high-transaction databases, which reduces the impact on the databases. You can configure the interval mode by setting the following parameter in the SAP HANA global.ini file:
`log_backup_interval_mode = service`
- SAP HANA 2.0 requires that the backup catalog backups and transaction log backups be configured separately. In previous SAP HANA releases, the backup catalog was automatically saved as part of the transaction log backups. Ensure that both the backup catalog backups and transaction log backups are configured to use the database application agent through the Backint interface.

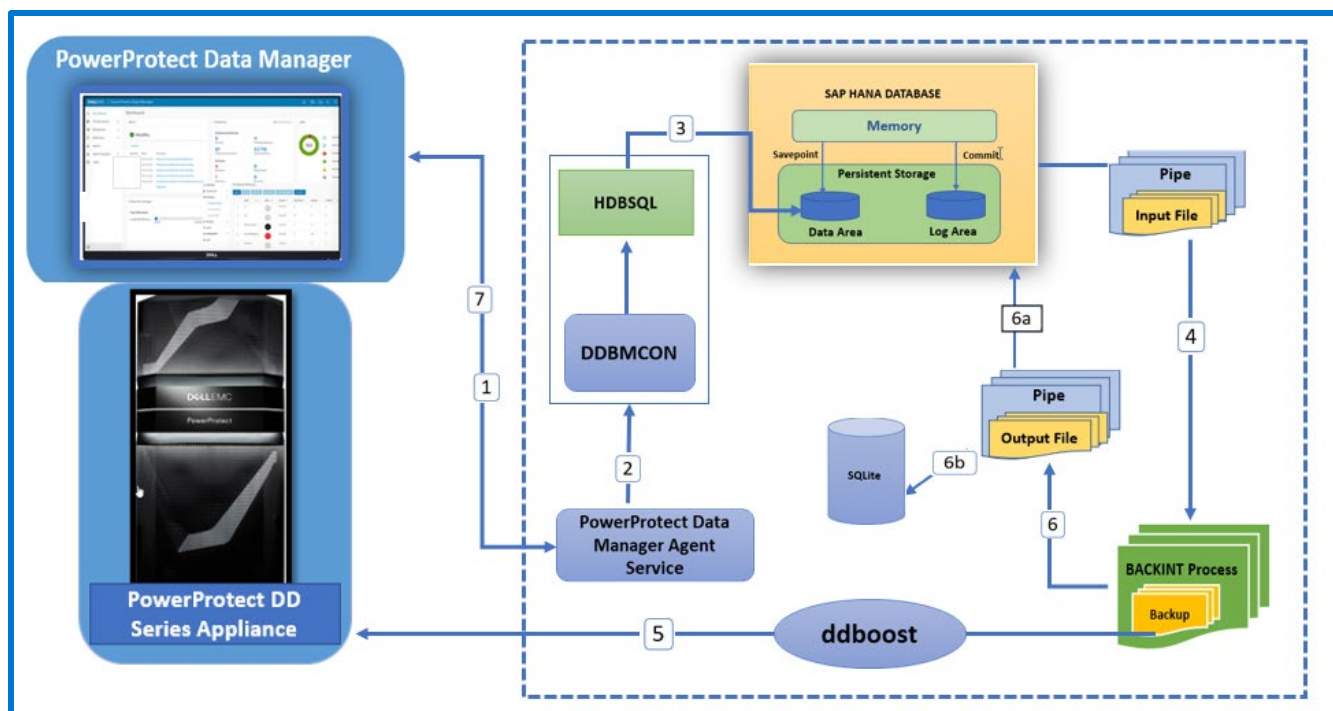


- PowerProtect Data Manager automatically creates the two configuration files as follows, where <SID> is the system ID of the SAP HANA system. These configuration file pathnames are predetermined and cannot be changed:
 - User configuration file:
`/usr/sap/<SID>/SYS/global/ppdmconfig/<SID>_<database_name>_data_user.utl`
 - Auto configuration file:
`/usr/sap/<SID>/SYS/global/ppdmconfig/<SID>_<database_name>_data_auto.utl`

5.1 Centralized Protection Policy Backup Workflow

When PowerProtect Data Manager Administrator creates a protection policy for SAP HANA databases, the Centralized Protection option enables the PowerProtect Data Manager to centrally manage the entire life cycle of data protection operations for the SAP HANA database except the Log backup that is done automatically by SAP HANA.

The data protection attributes are specified when the Centralized Protection policy is created: Type, Purpose, Assets, Schedule, Retention and SLA. After the Protection Policy creation is completed the Lockbox is automatically created. The backup workflow is as follows:



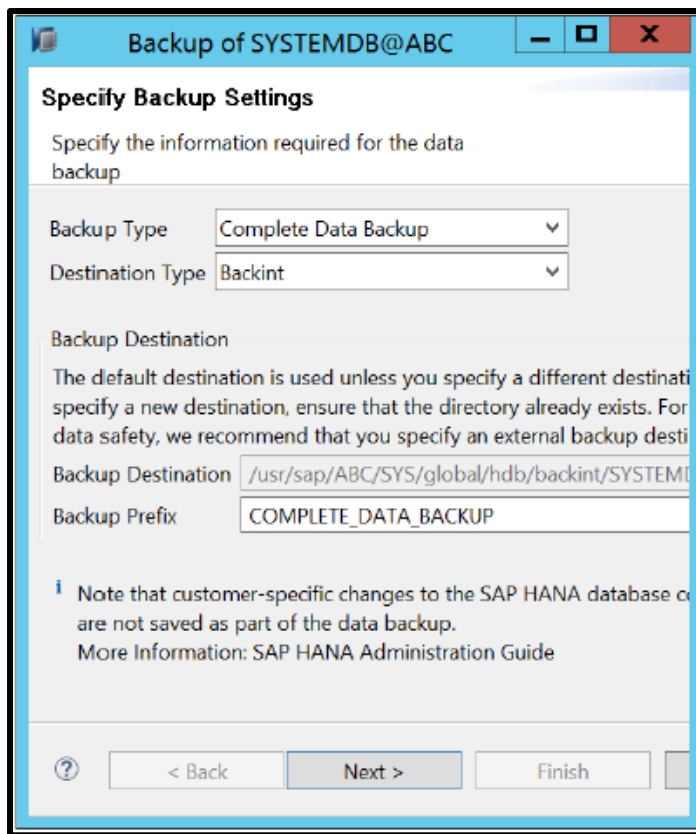
1. Backup schedule starts and triggers the Protection Policy. PowerProtect Data Manager sends a REST API request to PowerProtect Data Manager Agent Service.
2. PowerProtect Data Manager Agent service triggers the `ddbmcon` interface and requests a backup.
3. `Ddbmcon` generates the `hdbsql` command-line shell to connect to SAP HANA Data Area.
4. SAP HANA writes the data to `pipe` and provides the input file to `backint` process.
5. `Backint` process reads the data from this `pipe` and sends it to `ddboost` which transfers the data to PowerProtect DD Series Appliance.
6. `Backint` gives output file to HANA along with `SaveSetID`.
 - 6.a HANA writes `SaveSetID`, `Source File Name`, `BackupID` to `M_BACKUP_CATALOG_FILES` tables.
 - 6.b While writing the stream to Protection Storage `hdbbackint insert SavesetPath, BackupID and TaskID` to `SQLite` which will be used during the backup discovery process.
7. When the discovery process runs PowerProtect Data Manager Agent Service will inform PowerProtect Data Manager ADM agent about the backup status and it will become visible for recovery in PowerProtect Data Manager UI under assets SAP HANA agent.

5.2 Self-Service Protection Policy Backup Workflow

To enable self-service protection, when you create the SAP HANA protection policy, select Self-Service Protection. To perform a self-service or manual backup of an SAP HANA database, you can run either of the supported SAP HANA tools on the SAP HANA database server:

a. SAP HANA Studio user interface

In SAP HANA Studio, you must specify the database for backup and enable the backup to use the `Backint` interface. For Backup Type, you can select one of the three supported backup types: Complete Data Backup, Differential Backup, Incremental Backup.

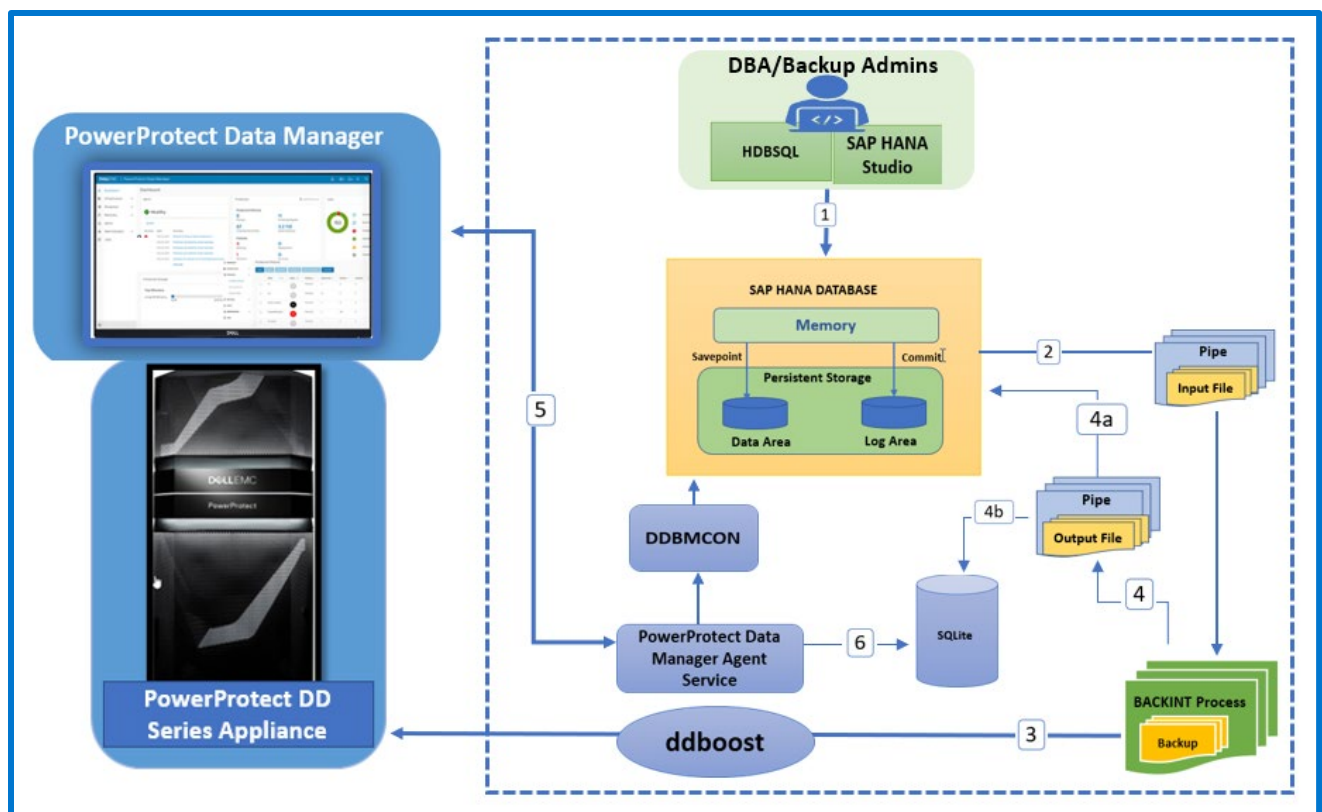


b. SAP HANA command-line interface (CLI) with the hdbsql command

To perform a database backup, you can run the appropriate `hdbsql` command. For example:

```
hdbsql -U <user_key> "backup data using backint
('/usr/sap/<SID>/SYS/global/hdb/backint/<prefix>')"
```

- `<user_key>` is the user store key created with the SAP HANA `hdbuserstore` tool.
- `<SID>` is the system ID of the SAP HANA system.
`/usr/sap/<SID>/SYS/global/hdb/backint/` is optional in the `hdbsql` command.
- `<prefix>` is an optional prefix for the backup name. If you specify `<prefix>` in the `hdbsql` command, then the backup name will include the prefix, `<prefix>_databackup_#_#`.



The backup workflow in self-service backup is as follows:

1. Backup is started by SAP HANA admin at scheduled timed using either `hdbsql` CLI or SAP HANA Studio and request is sent to SAP HANA.
2. SAP HANA writes the data to pipe and give the `input file` to `backint` process.
3. The `backint` Process reads the data from this `pipe` and sends it to `ddboost` which transfers the data to PowerProtect DD Series Appliance.
4. The `backint` gives output file to HANA along with `SaveSetID`.
 - 4a. HANA writes `SaveSetID`, `Source File Name`, `BackupID` to `M_BACKUP_CATALOG_FILES` tables.
 - 4b. While writing the stream to Protection Storage `hdbbackint insert SavesetPath`, `BackupID` and `TaskID` to `SQLite` which will be used during the backup discovery process.
5. PowerProtect Data Manager sends a request hourly to PowerProtect Data Manager Agent service to check the backup catalog called discovery process then updates the catalog information and sets the retention for application agent backup.

6 SAP HANA Parallelism

For SAP HANA, the PowerProtect Data Manager SAP HANA agent requires one protection storage stream for each backed-up pipe. To enable multi-stream SAP HANA backups for a centralized protection policy, you can set the parallelism value in the PowerProtect Data Manager UI. The parallelism value controls the number of parallel streams that the SAP HANA agent uses to write the backup data to Data Domain.

For example, if an SAP HANA scale-out system has 12 running services, then 12 streams are required to back up the data. Following are some specifications of parallelism :

- Starting with SAP HANA SPS 09, each service can also back up multiple logs for each backup, as controlled by the database parameter `max_log_backup_group_size`.
- For a multistream backup starting with SAP HANA SPS 11, the SAP HANA agent can use multiple SAP HANA channels to write the backup data for each service. The SAP HANA agent uses a separate SAP HANA channel to write each stream of data to the Protection Storage. To specify the number of channels to use for the backup, up to a maximum of 32 channels, you can set the SAP HANA parameter `parallel_data_backup_backint_channels`. SAP HANA opens the corresponding number of pipe files for the backup, and the SAP HANA agent saves each stream as a separate save set.
- To specify the maximum number of concurrent backup or restore streams, you set the parallelism parameter in the `hdbbackint` configuration file. For example, if the `parallel_data_backup_backint_channels` parameter is set to 12 on the SAP HANA server, then 12 streams are used for the backup, which produces 12 save sets. If the `hdbbackint` parallelism parameter is set to 6, then a maximum of 6 streams are backed up concurrently, and new streams start as other streams complete.
- Prior to SAP HANA SPS 11, an SAP HANA backup always has a parallelism of 1 per `hdbbackint` process. If the SAP HANA parameter `parallel_data_backup_backint_channels` is set in SAP HANA SPS 11 or later, the default value of this parallelism parameter is 8.

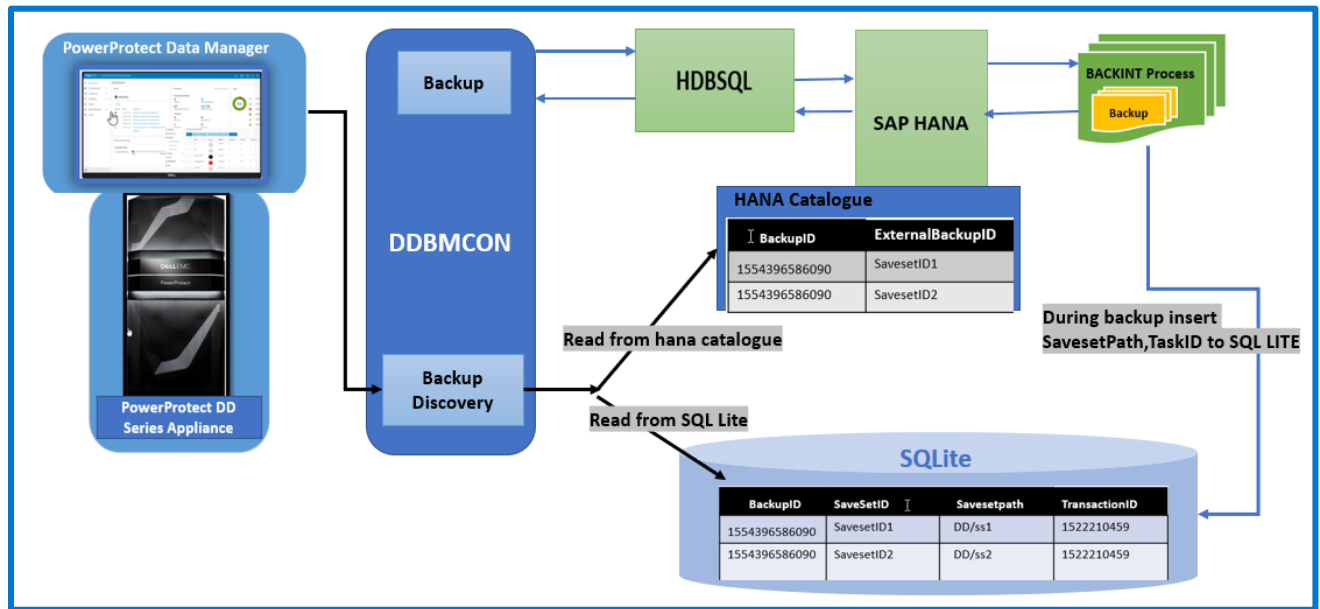
NOTE: A restore uses the same number of streams as the backup and ignores the `parallel_data_backup_backint_channels` parameter setting.

- The SAP HANA storage unit typically uses the following number of streams during a backup and restore:
 - If parallelism is set in the `hdbbackint` configuration file:
 - Number of services x parallelism
 - If parallelism is not set:
 - Number of services x `max_log_backup_group_size`
- Use the following guidelines to set the parallelism value in the PowerProtect Data Manager UI:
 - Set the parallelism to a value that is equal to or very close to the `parallel_data_backup_backint_channels` parameter setting. When the two settings are equal or almost equal, the SAP HANA agent can take full advantage of the parallelism provided by SAP HANA.
 - Set the parallelism to a value that is not more than the maximum number of connections provided by Data Domain.

7 SAP HANA Backup Discovery Process

Backup discovery fetches some data from SQLite and some from SAP HANA catalog and uses both for constructing Backup Disco output. Following points explain the process:

- SAP HANA catalog has all the info for creating Backup Disco output, except the Saveset path. In order to capture the Saveset path, `hdbbackint` process writes DD path to SQLite. Row in SQLite and Row in HANA Catalog can be connected using Saveset timestamp.
- Both rows will have `saveset` timestamp. SQLite is kept on local disk. In a scale-out environment, database could be spanned across multiple nodes and during backup, each node will backup few streams only. `hdbbackint` process writes DD path to local SQLite only.
- During backup discovery agent must send Backup discovery to every node for knowing all the stream that makes up one backup. `BackupID` created by HANA will identify the backup that may have several streams.
- Each stream is identified using `Saveset` timestamp. Agent distinguishes between Self-service and Centralized backup by looking at the `TaskID` stored in SQLite table.
- Backup discovery command needs database credentials for reading HANA catalog. PowerProtect Data Manager sends these credentials as part of Self-service Configuration. PowerProtect Data Manager Agent would have stored these credentials in lockbox.
- Backup Discovery should fail if there is any error when discovering records for that database. This is because, if Agent returns partial data, then retention will get affected.



8 SAP HANA Database Recovery

You can perform database restores directly to the SAP HANA application host by using the SAP HANA agent.

To perform a restore and recovery of an SAP HANA database, you can run either of the supported SAP HANA tools on the SAP HANA database server:

- SAP HANA Studio user interface (recommended method)
- SAP HANA command-line interface (CLI) using an SQL command

The following topics provide details about the restore and recovery procedures with these supported SAP HANA tools. The [SAP HANA documentation](#) provides more information about the SAP HANA restore and recovery tools.

Note: When you restore an SAP HANA backup that was performed by the database application agent (brownfield recovery) the restore must use the same UTL file that was originally used for the backup. Before you start the restore, run the SAP HANA Studio user interface and `set data_backup_parameter_file`, `log_backup_parameter_file`, and `catalog_backup_parameter_file` to all point to the UTL file that was used for the backup.

For Greenfield recovery - In SAP Hana client, go to path `/usr/sap/sid/sys/global/ppdmconfig` to set the Boolean variable `PPDM_Managed=True` inside file `sid_dbname_data_user.utl`

To recover SAP HANA database, the database needs to be shut down. During recovery, the database therefore cannot be accessed by end users or applications. It may be necessary to recover the SAP HANA database in the following situations:

- A disk in the data area is unusable.
- A disk in the log area is unusable.
- Because of a logical error, the database needs to be reset to its state at a point in time.
- You want to create a copy of the database.

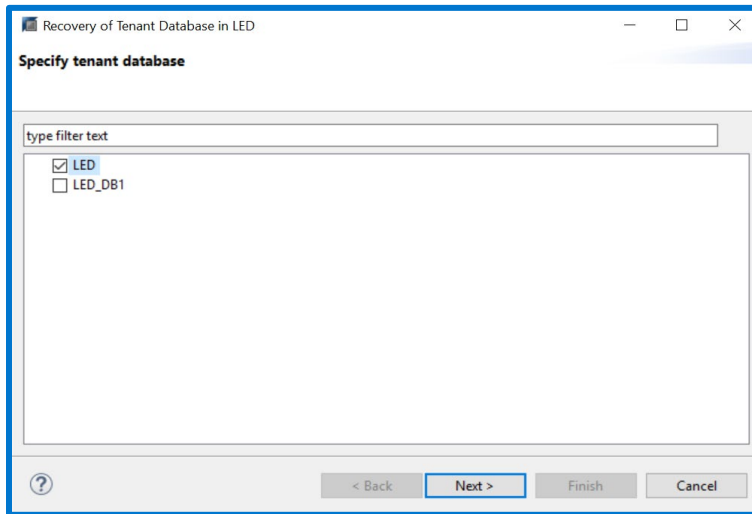
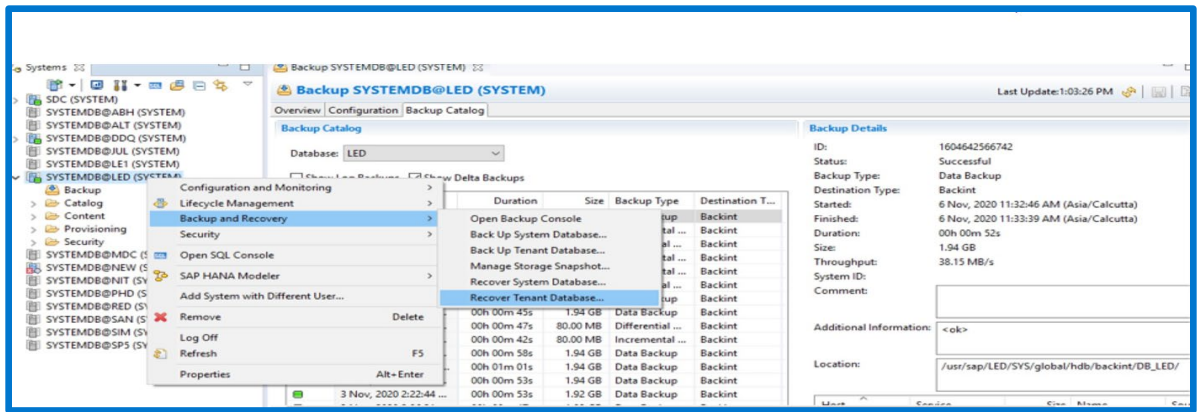
There are a few different types of recovery:

- **Most Recent State:** Used for recovering the database to the time as close as possible to the current time. For this recovery, the data backup and log backup must be available since last data backup and log area are required to perform the above type recovery
- **Point in Time:** Used for recovering the database to the specific point in time. For this recovery, the data backup and log backup must be available since last data backup and log area are required to perform the above type recovery
- **Specific Data Backup:** Used for recovering the database to a specified data backup. Specific data backup is required for the above type of recovery option.
- **Specific Log Position:** This recovery type is an advanced option that can be used in exceptional cases where a previous recovery failed.

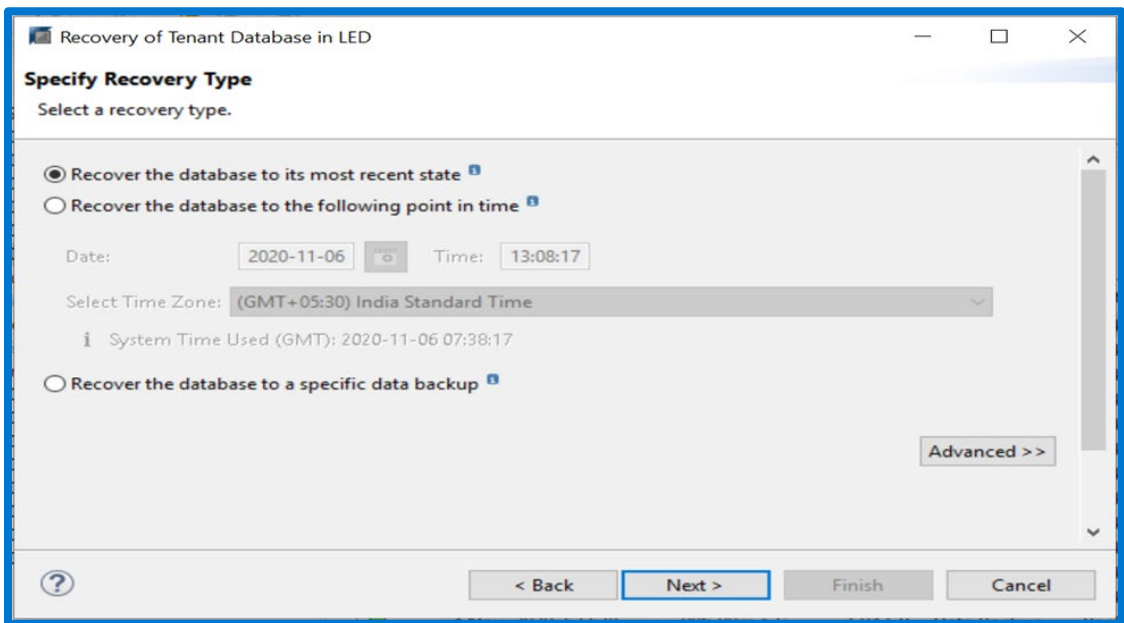
8.1 Performing a Restore using SAP HANA Studio

Follow the steps below for restore:

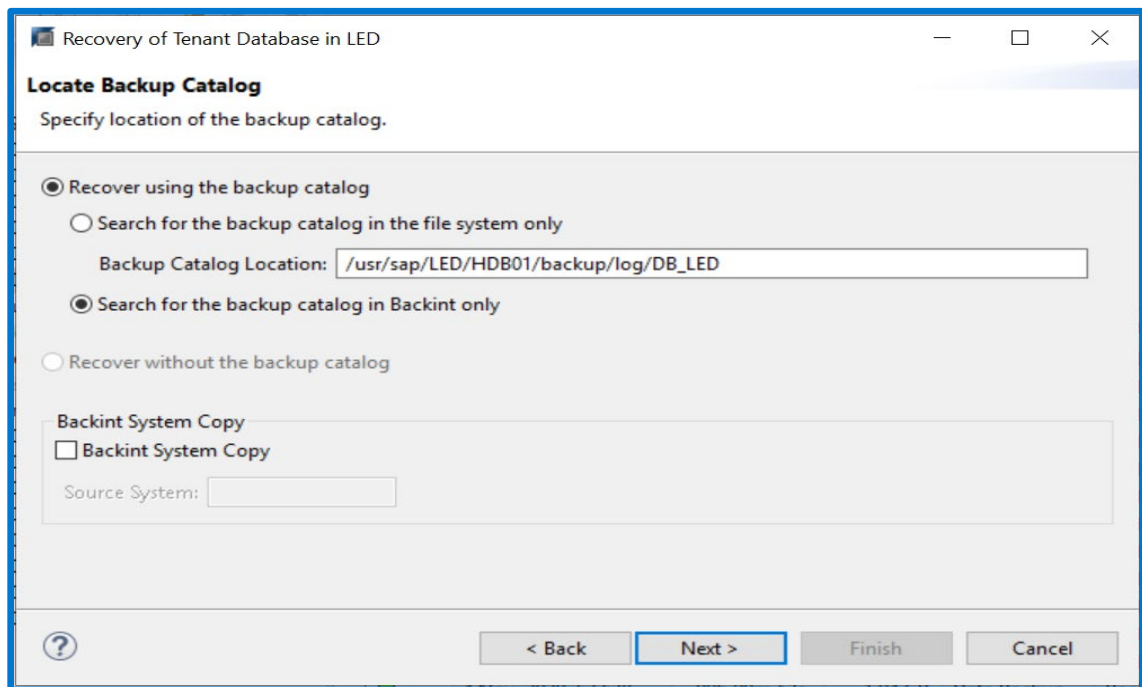
1. Start SAP HANA Studio.



2. In the recovery user interface, go to the **Specify Recovery Type** page

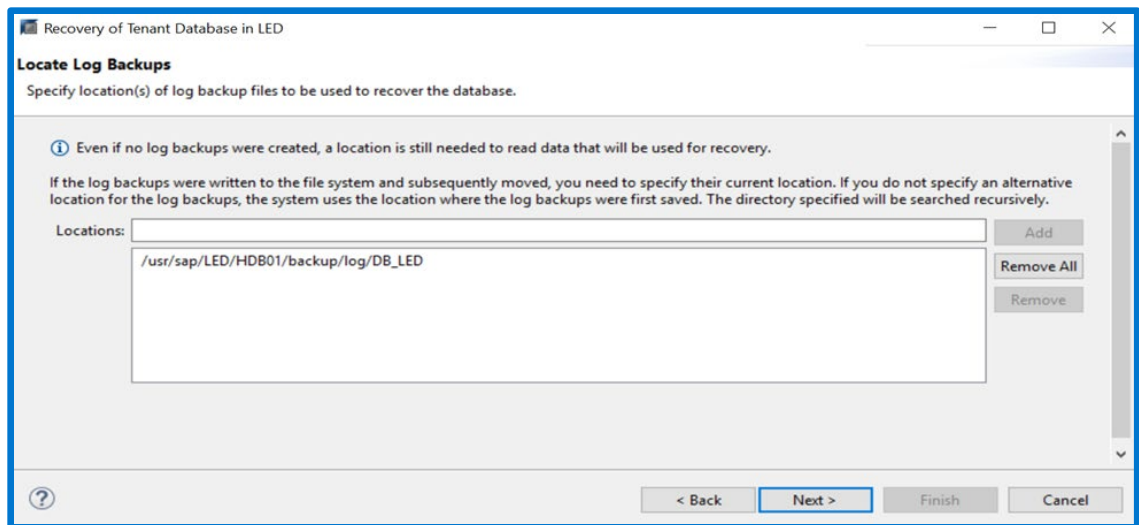


- a. Select one of the following types of recovery for the required SID:
 - **Recover the database to its most recent state**—Recovers the database as close as possible to the current time.
 - **Recover the database to the following point in time**—Recovers the database to a specific point-in-time.
 - **Recover the database to a specific data backup**—Restores only the specified data backup.
 - b. In an exceptional case when a previous recovery has failed, click Advanced >>, and then select Recover the database to the following log position.
 - c. Click Next >.
3. On the Locate Log Backups page, specify the locations of the log backup files to use to perform recovery.

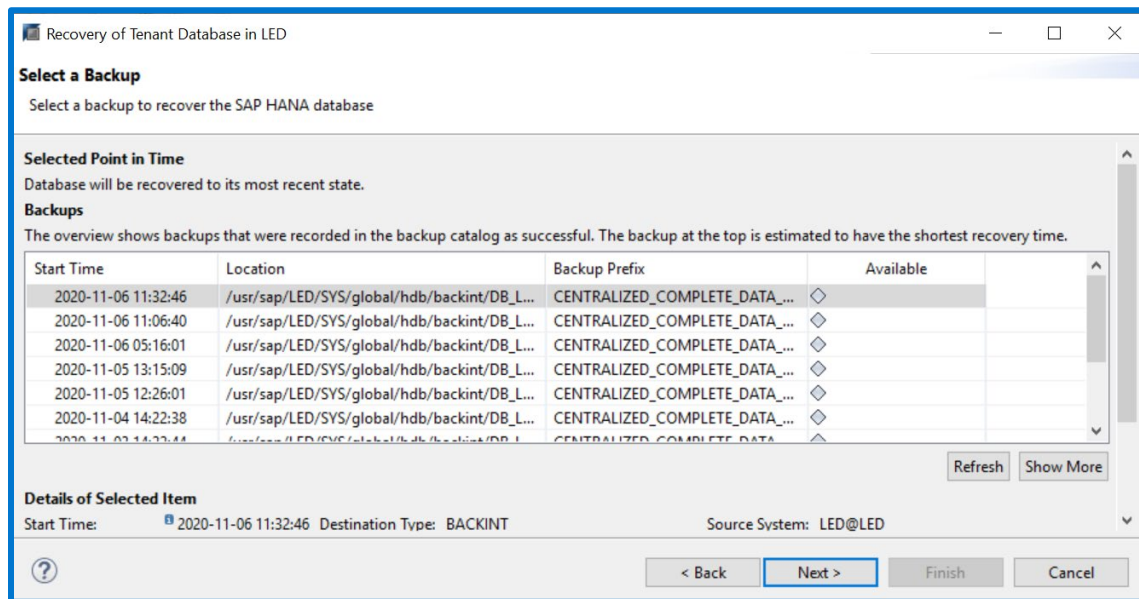


The backup locations for the single database containers and the multitenant database containers are as follows:

- **Single database container**—`/usr/sap/<SID>/SYS/global/hdb/backint/`
- **Multitenant database container** —`/usr/sap/<SID>/SYS/global/hdb/backint/SYSTEMDB`
- **Multitenant database container (Tenant database)** —
`/usr/sap/<SID>/SYS/global/hdb/backint/DB_<tenantDB>`

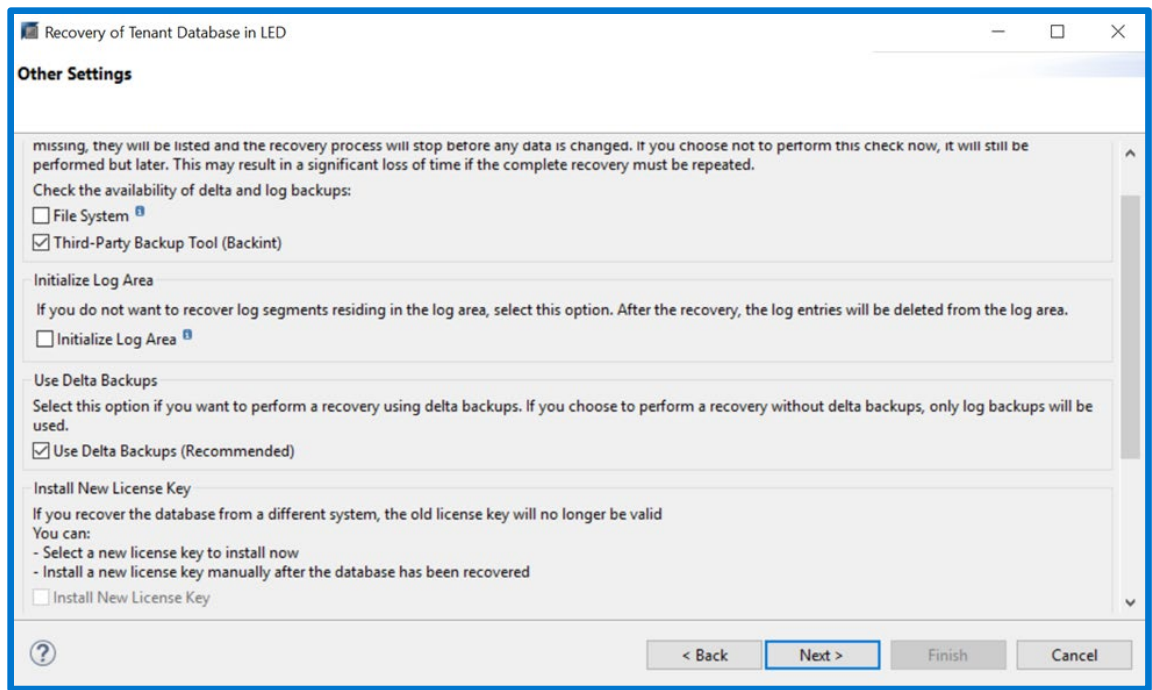


4. On the **Select Data Backup** page, select the database backup for either restore or recovery.

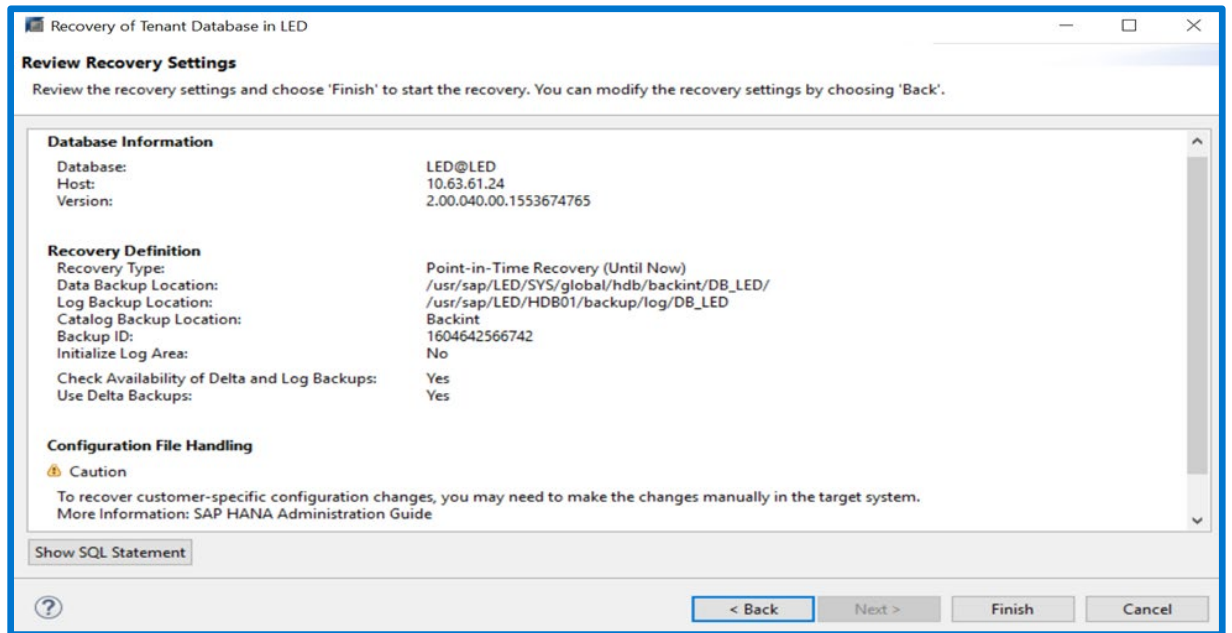


To check the availability of a backup, select the backup in the table, and then click Check Availability. The Available column in the table displays either a green icon if the backup is available or a red icon if the backup is not available.

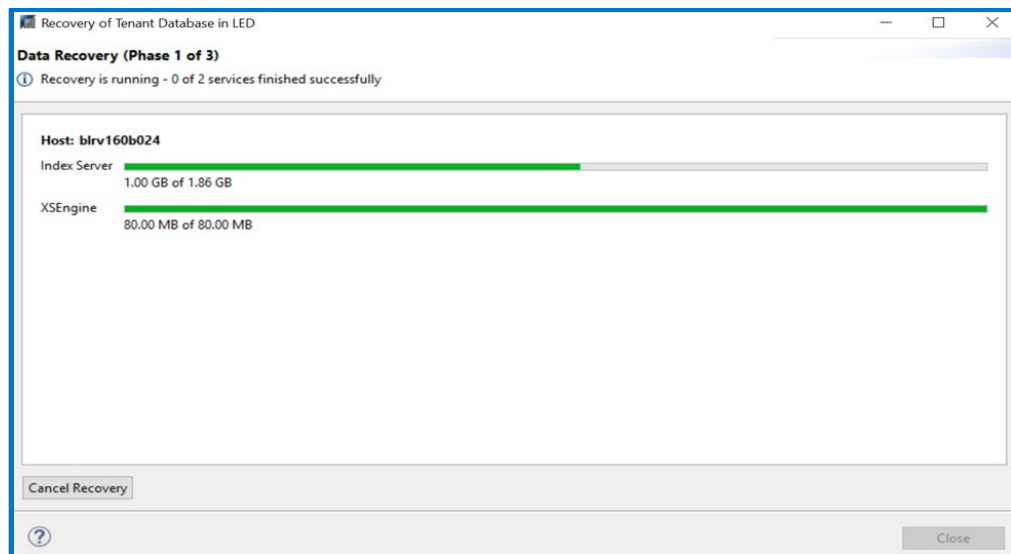
5. On the Other Settings page, select the required options, and then click Next >.



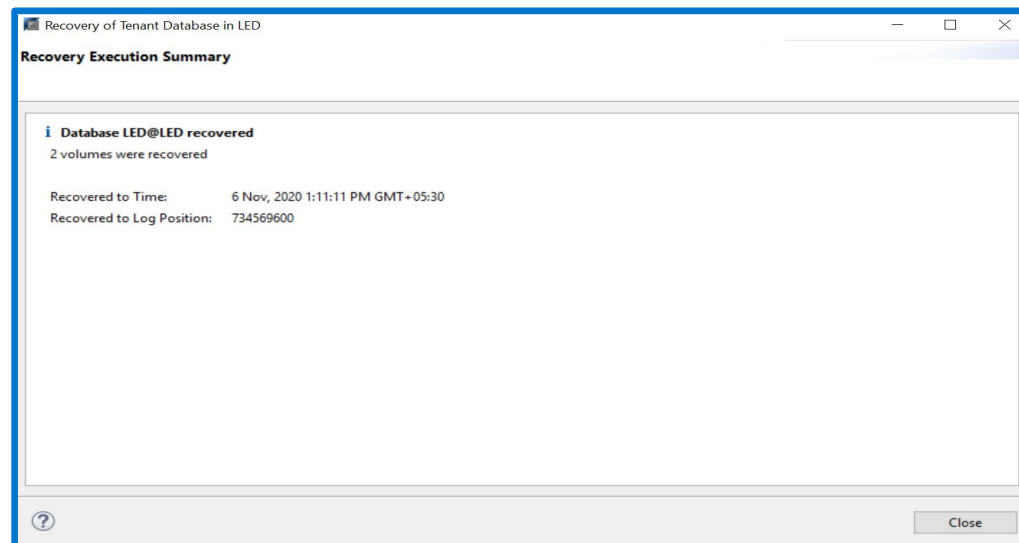
6. On the Review Recovery Settings page, review the information, and then click Finish.



7. The recovery progress appears for each service that includes the name server, the index server, and the statistics server.



8. A confirmation message appears when the recovery completes.



8.2 Performing alternate host recovery

To enable an alternate host recovery, copy the required user configuration file and auto configuration file from the source machine to the destination machine. Modify configuration files

<SID>_<database_name>_data_user.utl and <SID>_<database_name>_data_auto.utl.

You must replicate the directory path of the configuration files on the destination machine, including the same permissions, and then copy the user and auto configuration files from the source machine to the directory path. For example, if the source instance is SID, then create the following path on the destination machine:

```
/usr/sap/SID/SYS/global/ppdmconfig
```

When the configuration files are in place, the steps for performing the alternate host recovery are the same as the steps for restoring the database to the same host.

8.3 Performing a Restore using the SAP HANA CLI

To perform a recovery of an SAP HANA backup, starting with SAP HANA SPS 07, you can use an SQL command.

- To run the recovery SQL command, you must be logged in to the SAP HANA system as the operating system user <SID>adm. In a scale-out multinode environment, you must run the command from the active node.
- The recovery command must first include the `HDBSettings.sh` and `recoverSys.py` scripts in that order. The `HDBSettings.sh` script sets the environment for recovery. The `recoverSys.py` script processes the recovery SQL command and performs the database recovery. The following example command performs a recovery of an SAP HANA backup:

```
./HDBSettings.sh recoverSys.py --command="recover database until timestamp  
'2014-02-28 00:00:00' using data path('/usr/sap/AAA/SYS/global/hdb/backint/')  
using log path('/usr/sap/AAA/SYS/global/hdb/backint/') using BACKUP_ID  
1593520382350" --wait
```

- You must include the `--wait` option for proper completion of the `recoverSys.py` script.
- The [SAP HANA User Guide](#) provides details about the SQL command syntax of the recovery command.
- To recover an SAP HANA MDC database with the SAP HANA CLI, you must log in to the SAP HANA SystemDB database and run the required recover command. For example, run the following command to recover the database application agent backup of an SAP HANA MDC database:

```
RECOVER DATABASE FOR DB0 UNTIL TIMESTAMP '2015-08-29 01:25:25' USING DATA  
PATH('/usr/sap/AAA/SYS/global/hdb/backint/DB_DB0/') USING LOG PATH  
( '/usr/sap/AAA/HDB00/backup/log/DB_DB0' ) USING BACKUP_ID 1440811502099 CHECK  
ACCESS USING BACKINT
```

Restrictions for Recovery:

- SAP HANA database cannot be recovered to SAP HANA database with a lower software version. The SAP HANA database software version used for the recovery must always be the same version or higher than the SAP HANA database used to create the backup.
- There can be any number of hosts in the target system, if the number and type of services is identical in both the source and the target system.
- With SAP HANA multitenant database containers, recovery using snapshots is not supported yet. Tenant databases can be recovered by the relevant system database only.

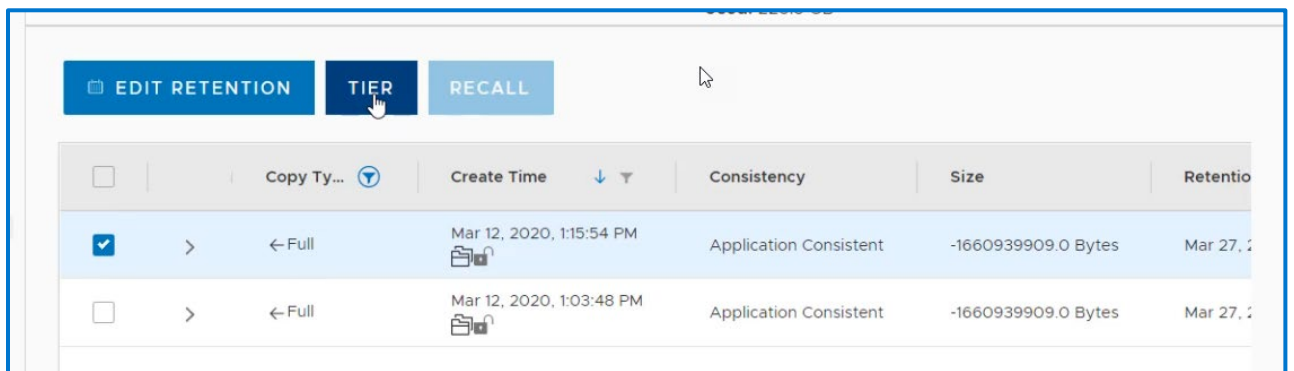
Note: For more details on recovery check [PowerProtect Data Manager SAP HANA agent guide](#)

9 Cloud Tiering, Replication, and Backup Deletion

9.1 Cloud Tiering

The PowerProtect Data Manager cloud tier feature works in tandem with the PowerProtect DD Cloud Tier feature to move PowerProtect Data Manager backups from PowerProtect DD systems to the cloud. This provides long-term storage of PowerProtect Data Manager backups by seamlessly and securely tiering data to the cloud.

- From the PowerProtect Data Manager UI, you configure cloud tier to move PowerProtect Data Manager backups from PowerProtect DD to the cloud, and you can perform seamless recovery of these backups.
- PowerProtect DD cloud storage units must be preconfigured on the PowerProtect DD system before they are configured for cloud tier in the PowerProtect Data Manager UI. The [PowerProtect DD Operating System Administration Guide](#) provides further information.
- Both SAP HANA centralized and self-service protection policies support cloud tiering. You can create the cloud tier schedule from both primary and replication stages. Schedules have a weekly or monthly recurrence and a retention time of 14 days or greater.
- The restore operations of backups that have been tiered to the cloud are identical to normal restore operations. The PowerProtect Data Manager recalls a copy of the backup from the cloud to the local (active) tier of the PowerProtect DD Series Appliance, and then performs a restore of the backup from the active tier to the client. The status appears as Cloud. After the restore, the backup copy is removed from cloud tier, and is stored on the active tier of the PowerProtect DD Series Appliance for 10 days, after which it is deleted.



9.2 Replication

PowerProtect Data Manager can manage and monitor data protection and replication for SAP HANA assets through integration with the SAP HANA agent. Replication can be configured while creating the protection policy for self-service or centralized backup and retention can be set for replicated data. To enable replication, ensure that you add a protection storage as the replication location.

9.3 Backup Deletion

In addition to deleting backups upon expiration of the retention period, PowerProtect Data Manager enables you to manually delete backup copies from the PowerProtect DD Series Appliance.

Starting with PowerProtect Data Manager version 19.6, you can perform a backup copy deletion that deletes only a specified part of a backup copy chain, without impacting the ability to restore other backup copies in the chain. When you select a specific backup copy for deletion, only that backup copy and the backup copies that depend on the selected backup copy are deleted:

- When you select to delete a full backup copy, any other backup copies in the chain that depend on the full backup copy are also deleted.
- When you select to delete an incremental backup copy, any incremental backup copies that depend on the selected incremental backup copy are also deleted. The whole backup chain is not deleted.
- When you select to delete a differential backup copy, any incremental backup copies that depend on the selected differential backup copy are also deleted. The whole backup chain is not deleted.
- When you select to delete a log backup copy, any other log backup copies that depend on the selected log backup copy are also deleted. The whole backup chain is not deleted.

Note:

You can delete a full backup copy only if an archive log exists and the previous full backup copy is available. If you delete a full backup of an incomplete chain, you cannot restore the whole chain until you perform the next full backup.

Deletion of the last remaining full backup from Data Domain does not delete the entry for that backup in the SAP HANA catalog. This behavior applies to both scheduled deletion and manual deletion operations. Deletion of log backups might affect a point-in-time restore that depends on those log backups.

- Regarding expired copy deletion for a backup chain, the full backup expires last. The other incremental, differential, and log backups can be expired and deleted earlier.

10 Disaster Recovery

For a comprehensive disaster recovery plan, you must ensure that you can reconstruct the computing environment and all the SAP HANA server files associated with maintaining data on the application host. To prepare for a disaster recovery of the SAP HANA server host, use the following guidelines:

Maintain accurate and complete records of the network and system configurations. Keep all the original software media and the following items in a safe location:

- Original operating system media and patches
- Device drivers and device names
- File system configuration
- IP addresses and hostnames

To ensure the recovery of customer-specific settings of the database, back up the SAP HANA configuration files by following the [SAP HANA Administration Guide](#). These files are not backed up during database backups. Ensure that you have a current full backup of the database and all the redo logs as required for a recovery of the database to the current time.

Perform the following tasks for a disaster recovery.Procedure:

1. Set up the SAP HANA database system to the same configuration as when the data was backed up.
2. Set up the same SAP HANA configuration file for the recovery, as was used to perform backups.

Set the `CLIENT` parameter in the SAP HANA configuration file to the hostname of the host, to which the data was backed up.

3. Re-create the lockbox on the database host.
4. Ensure that all the required database and log backup images are available.
5. To recover the database and redo logs, run SAP HANA Studio.

Note: To get the step-by-step details, please refer to [PowerProtect Data Manager for SAP HANA Agent User Guide](#).

10.1 Recovering the PowerProtect Data Manager Agent service from a disaster

You can perform self-service restores of application assets by using a file system or application agent, regardless of the state of the agent service or PowerProtect Data Manager. The information in this section describes how to bring the agent service to an operational state to continue if a disaster occurs and the agent service datastore is lost.

The agent service periodically creates a backup of its datastore in the `<agent_service_installation_location>/dbs/v1/backups` repository. If all these backups are lost, the agent service can still start. The agent service discovers all the application systems, assets, and backup copies on the system again, and notifies PowerProtect Data Manager. Depending on when the failure occurred, the agent service might not be able to find older backup copies for some asset types. As a result, the centralized

deletion operations might fail when cleaning up the database vendor catalog or removing older backups that are taken before the asset is added to PowerProtect Data Manager.

By default, the agent service backs up consistent copies of its datastore files to the local disk every hour and keeps the copies for 7 days. Each time the agent service backs up the contents of the datastore, it creates a subdirectory under the `<agent_service_installation_location>/dbs/v1/backups` repository. The subdirectories are named after the time the operation occurred, in the format `YYYY-MM-DD_HHMM-SS_epochTime`.

By default, the datastore repository is on the local disk. To ensure that the agent service datastore and its local backups are not lost, it is recommended that you back up the datastore through file system backups. You can also change the datastore backup location to a different location that is not local to the system. To change the datastore backup location, update the values in the `config.yml` file

Conclusion

Dell EMC PowerProtect Data Manager delivers complete control of SAP HANA backup and disaster recovery to SAP HANA database administrators. The advanced integration between PowerProtect Data Manager and SAP HANA provides a faster, more efficient database backup and restore solution.

References

[PowerProtect Data Manager SAP HANA Agent Guide](#)

[PowerProtect Data Manager Administration and User Guide](#)

[SAP HANA Backup and Recovery Wiki](#)

[PowerProtect DD Administration Guide](#)