WHITE PAPER

SAP HANA BACKUP AND RECOVERY

USING CLOUDSNAP

PURESTORAGE®

TABLE OF CONTENTS

INTRODUCTION

Backing up SAP HANA on Pure Storage® FlashArray™ to cloud is now extremely simple using the Purity CloudSnap feature. Previously, snapshots taken for an SAP HANA system were considered recovery points and not actual backups, since these snapshots resided in FlashArray only. CloudSnap™ enables FlashArray to offload snapshots directly to an S3 bucket, without the need for additional backup software or a cloud gateway. This document will attempt to showcase this feature for a very common use case – backup and recovery of an SAP HANA instance. We will go over the necessary steps to take a CloudSnap of an SAP HANA instance for backup purposes on to a Pure Storage FlashBlade™ system and how to restore the database from the CloudSnap.

CLOUDSNAP BENEFITS

CloudSnap is a self-backup technology built into FlashArray. CloudSnap is natively managed via Pure FlashArray's GUI, CLI, and REST interfaces, and is integrated with the Pure1® Snapshot Catalog. FlashArray connects to AWS via https, and data is encrypted in transit and stored in an encrypted format in an S3 bucket using server side encryption.

Benefits of CloudSnap

- CloudSnap preserves data compression on the wire and in the S3 bucket, saving network bandwidth and increasing storage space efficiency.
- CloudSnap preserves data reduction across snapshots of a volume. After offloading the initial baseline snapshot of a volume, CloudSnap only sends delta changes for subsequent snaps of the same volume. Therefore, there is no back and forth network traffic between FlashArray and the cloud to compute deltas between snapshots, further reducing network congestion and data access costs in the cloud.
- CloudSnap knows which data blocks already exist on FlashArray, so during restores it only pulls back missing data blocks to rebuild the complete snapshot on FlashArray. In addition, CloudSnap uses dedupe preserving restores, so when data is restored from the offload target to FlashArray, it is deduped to save space on FlashArray.

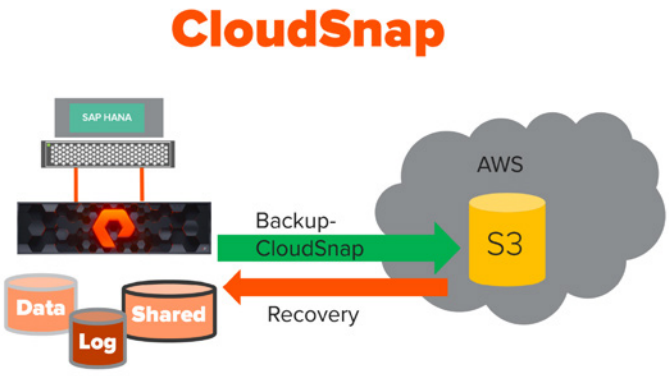

FIGURE 1. CloudSnap backup and recovery

AUDIENCE

The target audience for this document includes storage administrators, SAP system administrators, data center architects, etc., who are interested in deploying and maintaining an SAP environment with SAP HANA as the database layer on a Pure Storage FlashArray. A working knowledge of SAP HANA, server, storage, networks, and data center design is assumed but is not a prerequisite to read this document.

ENVIRONMENT

In order to test this concept, the following environment has been setup in our lab with the following entities:

COMPONENT	TYPE
DATABASE	SAP HANA 1.0 SPS12/SAP HANA 2.0
OPERATING SYSTEM	SUSE Linux 12 SP 2
CLOUD TARGET	AWS (S3)
STORAGE	FlashArray//X

TABLE 1. Lab test parameters

SETTING UP CLOUDSNAP ON FLASHARRAY

The Pure Storage Support team has to enable the offload engine on Purity run before CloudSnap can be used. Enabling the offload engine is a one-time, non-disruptive, data-in-place process, and it involves the following steps:

- Making sure that the version of Purity is version 5.2
- Enabling Purity Run
- Installing and configuring the offload engine on FlashArray

Enabling the offload engine requires:

- A network connection between FlashArray and AWS.
- An AWS account for creating S3 buckets for storing snapshots

For more information, please refer to the CloudSnap Best Practices document.

After the offload engine is enabled by the Pure Support team, customers can start using it by connecting an S3 target to FlashArray and selecting it as a replication target in **Protection Groups**. Once it is configured, you will see **Offload Targets enabled** under **Storage tab → Array section** in GUI as shown below

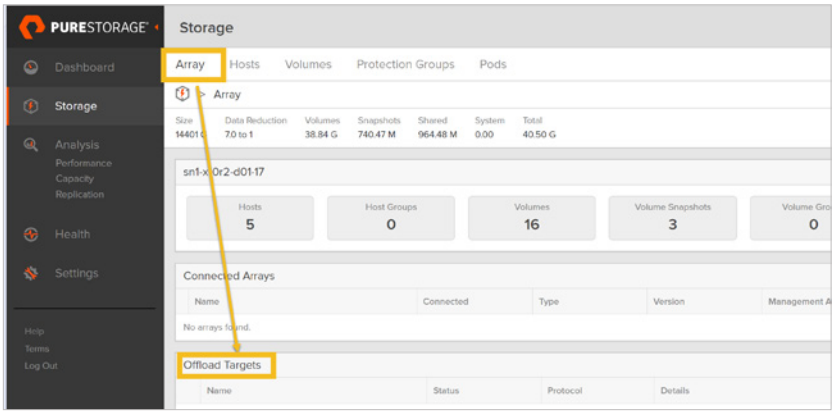

FIGURE 2. Storage tab, Array pane

In the **Offload Targets** section, create a new S3 Offload target by **clicking the 3 dots**.

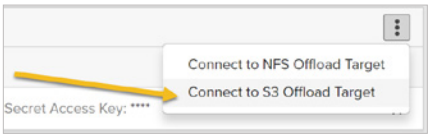

FIGURE 3. Connecting to an S3 target

Then add a new **Offload Target**, in this case the AWS S3 bucket information. The following information needs to be provided:

1. Name for the S3 target on FlashArray
2. Access Key ID and the Secret Access Key for the AWS IAM user
3. S3 bucket name and initialize it

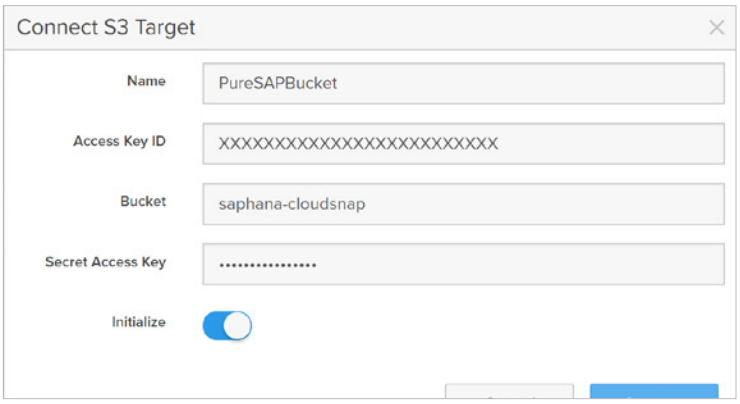

FIGURE 4. AWS S3 bucket information

Once it is created, you can see the offload target added, as shown below:

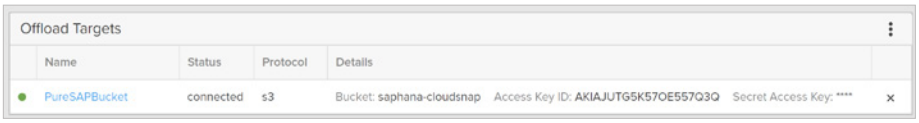

FIGURE 5. New offload target added

The SAP HANA system has the following volumes. We need to add the data volume to the protection group.

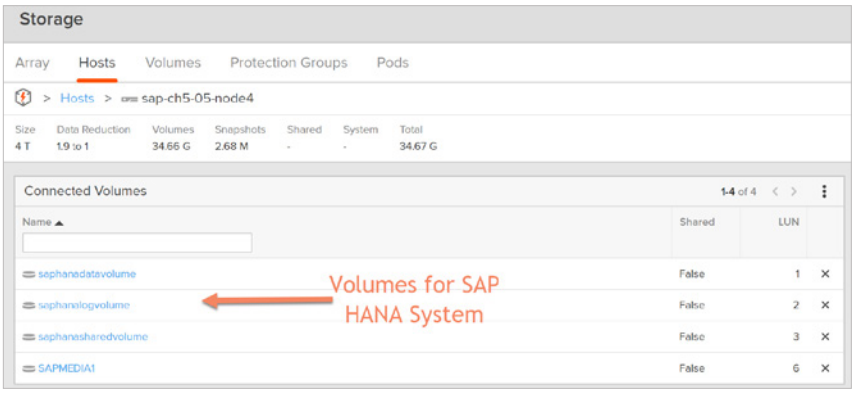

FIGURE 6. SAP HANA volumes

Create a **protection group** and add the **SAP HANA data volume** to this protection group. Then select **Add**, shown below, to add the Offload Target. This means the protection group is ready to replicate data to the S3 bucket on AWS.

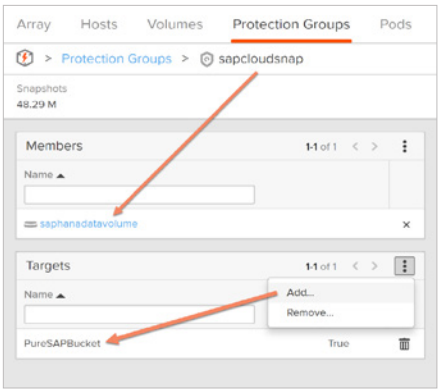

FIGURE 7. New protection group added

PROCEDURE TO USE CLOUDSNAP BACKUP FOR SAP HANA

SAP HANA best practice dictates using the **PREPARE** statement prior to taking a backup copy. This step creates a database internal snapshot based on a system-wide **savepoint** executed during the **PREPARE**. During this period, no other backup should run, but log backup can continue as configured.

Step 1: Prepare Statement

From SAP HANA Studio, perform the following steps:

In Backup node ➔ Context menu ➔ Manage Storage Snapshot ➔ Select **Prepare**

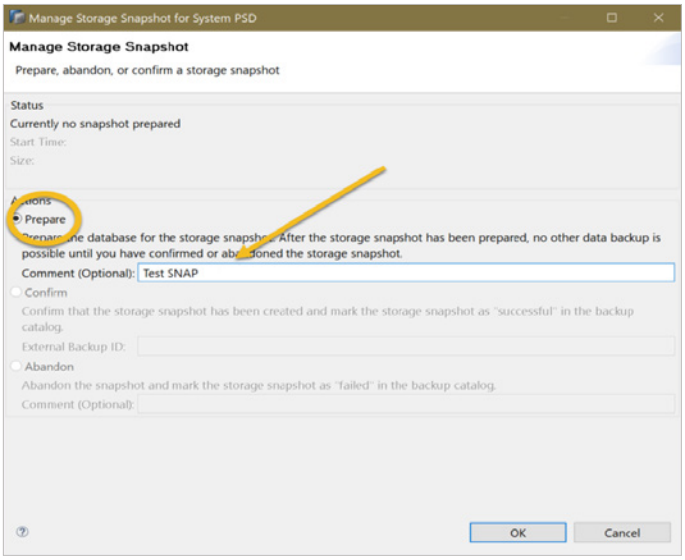

FIGURE 8. Manage Storage Snapshot pane

Alternatively, this step can also be performed with this SQL statement on the index server (the database user as **BACKUP OPERATOR** authorization):

```
BACKUP DATA CREATE SNAPSHOT BACKUP OPERATOR
```

In SAP HANA 2.0, you need to do the following to put SAP HANA in create snapshot mode:

```
BACKUP DATA FOR FULL SYSTEM CREATE SNAPSHOT
```

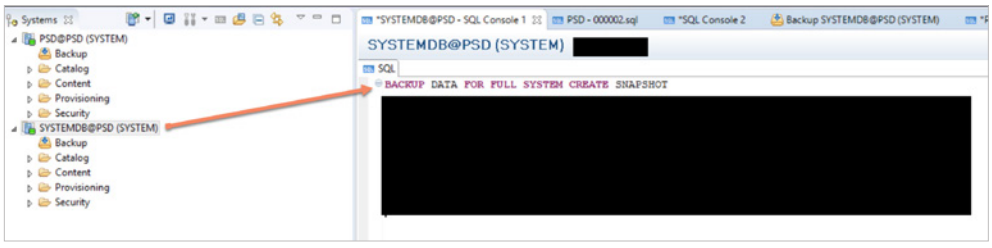

FIGURE 9. Backup Operator authorization

Now the database is prepared for the storage snapshot. Verify the backup details, either through the backup catalog tab in SAP HANA studio (make sure you capture important details, like ID) or by issuing the following SQL statement to get the backup details using the view **M_BACKUP_CATALOG**:

```
select * from M_BACKUP_CATALOG where "ENTRY_TYPE_NAME" ='data snapshot'.
```

Check the data volumes mount point and see if there is a file created as you see below **snapshot_databackup_0_1**

```
saphananode3:/hana/data/PSD/mnt00001/hdb00001 # ls
__DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__  datavolume_0000.dat  landscape.id  snapshot_databackup_0_1
saphananode3:/hana/data/PSD/mnt00001/hdb00001 #
```

FIGURE 10. File created

But be sure to freeze your file system (XFS) using **xfs_freeze** after you perform the above operation:

```
xfs_freeze -f  <Data Volume mount point>
```

And unfreeze it after taking the snapshot

```
xfs_freeze -u  <Data volume mount point>
```

Step 2: Create Protection Group Snapshot

To take a snapshot using the Pure Storage GUI, go to the **Storage** tab and select the protection group.

Click **Add** on the right and select **Create Snapshot**. This creates a near-instantaneous copy regardless of the size of the volume.

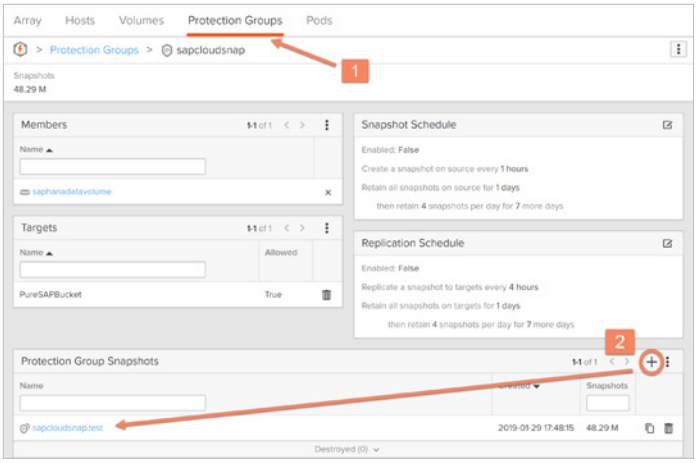

FIGURE 11. Pure Storage GUI, Storage tab

Make sure you select **Replicate Now**, as this will replicate the snapshot to the S3 target on AWS.

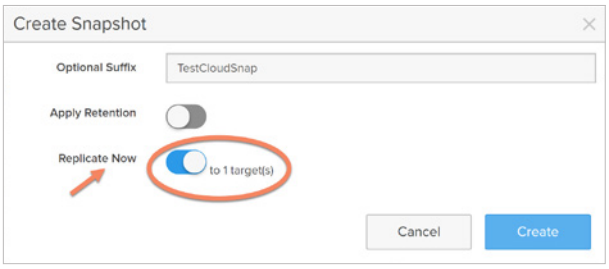

FIGURE 12. Create Snapshot pane

Check if the snapshot is replicated completely to the S3 target, as shown below:

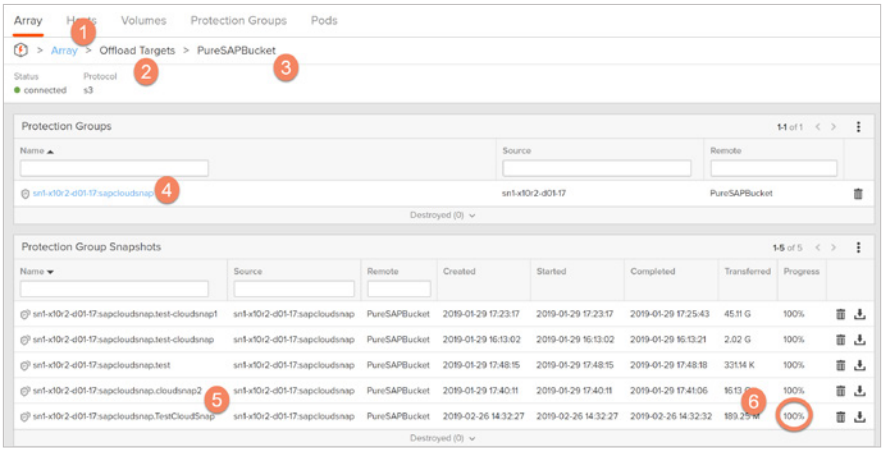

FIGURE 13. Viewing the S3 target

Step 3: Confirm or Abandon the Snapshot in SAP HANA

After creating the snapshot, SAP HANA has to be made aware that it has been created in order to confirm or abandon the snapshot. All storage snapshots are recorded in SAP HANA in the backup catalog with a status of successful or failed. This is very important, as SAP HANA has to clear its internal state once it is confirmed. This action can either be done in SAP HANA studio or via SQL statements.

USING SAP HANA STUDIO

Go to **Manage Storage Snapshot** and this time you will see only two options: **Confirm** or **Abandon**. Confirm the storage snapshot by using the External Backup ID generated. This is the same ID that was captured in Step 1.

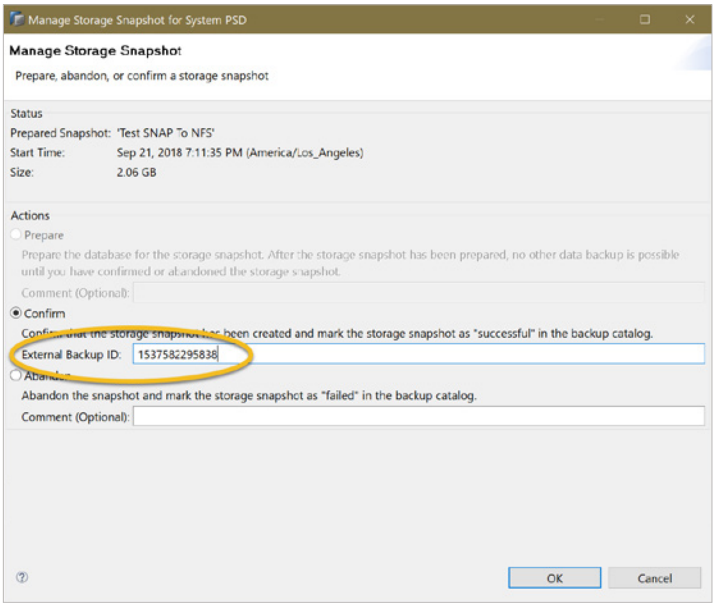

FIGURE 14. Manage Storage Snapshot pane

USING SQL STATEMENTS

You can use the following SQL statement to perform the same action as above:

```
BACKUP DATA CLOSE SNAPSHOT BACKUP_ID <BACKUP_ID> SUCCESSFUL
```

If it is SAP HANA 2.0, you need to do the following:

Go to the SQL console for the system database and run the following. This will update the backup catalog with this successful backup id.

```
BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID 1516141675109 SUCCESSFUL 'use this'
```

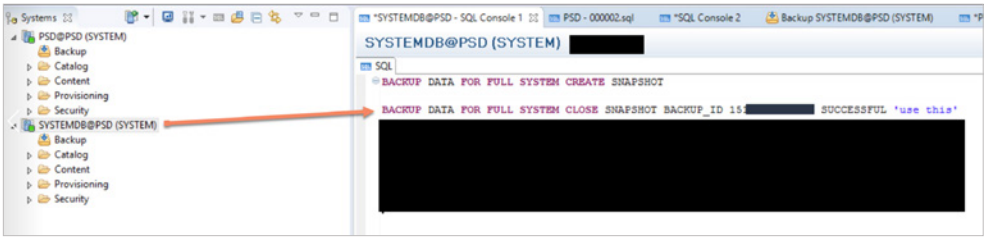

FIGURE 15. SQL console for the system database

RECOVERING SAP HANA FROM A SNAPSHOT

After a storage snapshot is created, it can be used to recover the SAP HANA instance. There are multiple options for recovering an instance using a storage snapshot.

Recover the database to a specific storage snapshot

- Uses a specific storage snapshot for recovery
- All logs will be lost; it is like beginning a new database lifecycle

Start the recovery process in SAP HANA studio:

System ➔ Backup and Recovery ➔ Recover System

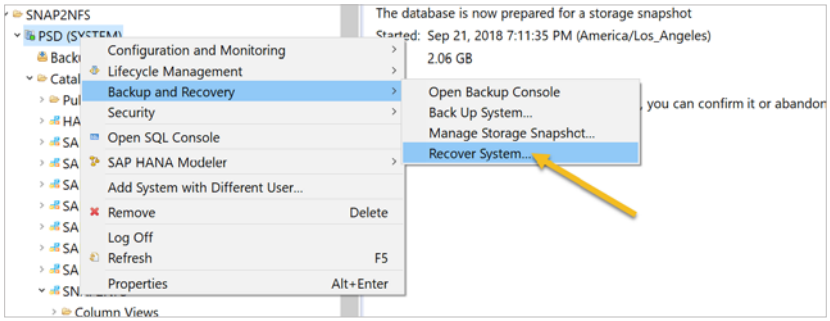

FIGURE 16. SAP HANA studio

In SAP HANA 2.0, go to the system database of this SAP HANA system and in the context menu select **Backup and Recovery ➔ Recover System Database** which is SYSTEMDB@PSD. This will shut down the system.

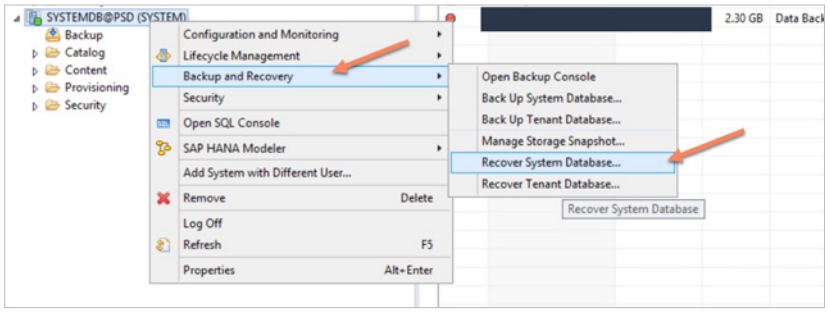

FIGURE 17. SAP HANA 2.0

Select the following option: **Recover the database to a specific data backup or storage snapshot**.

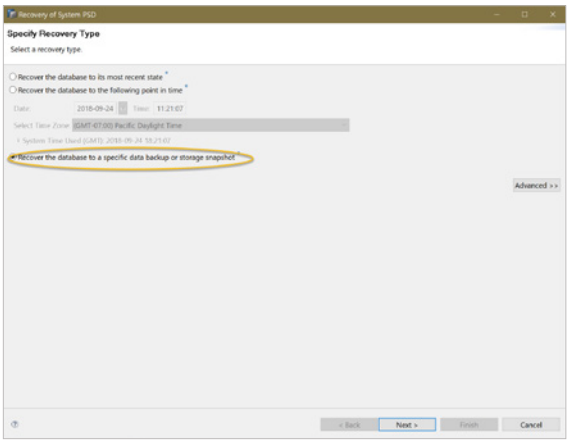

FIGURE 18. Sepecify Recovery Type pane

Then select **Specify backup without catalog**.

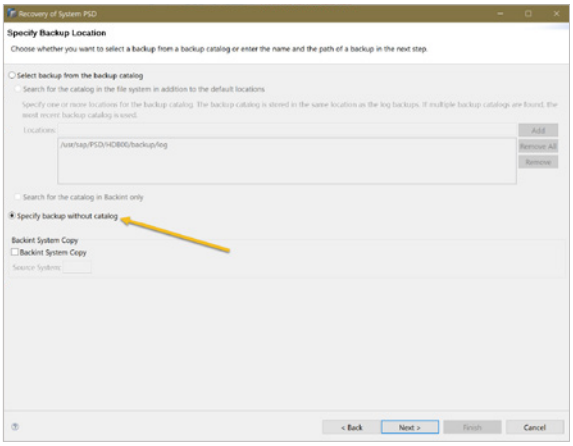

FIGURE 19. Sepecify Backup Location pane

Select **Snapshot** as Destination Type.

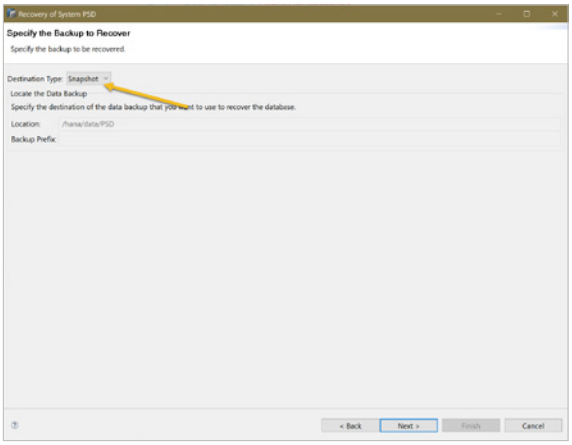

FIGURE 20. Sepecify the Backup to Recover pane

Initialize Log Area.

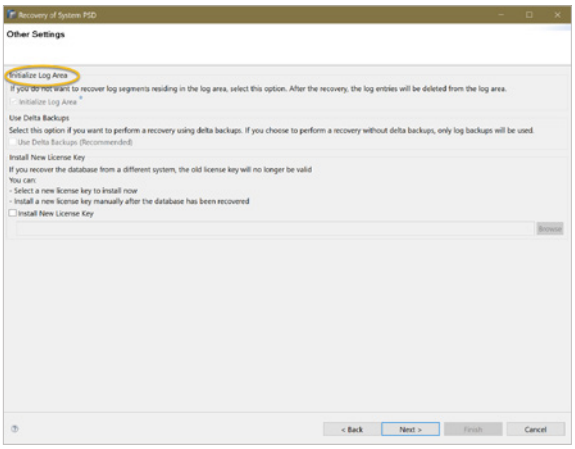

FIGURE 21. Initializa Log Area pane

Before you press **Next**, make sure the snapshot is recovered from the S3 target.

Recovering SAP HANA Snapshot from S3 Target

Even if the snapshot for this protection group is deleted, we can recover the snapshot from the S3 target with the following steps:

Go to the **Storage** tab of the Pure Storage GUI. Select the offload target.

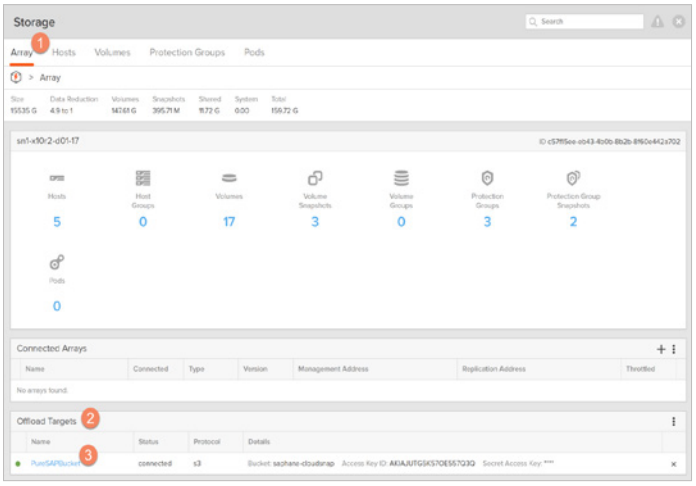

FIGURE 22. Selecting an offload target

Select the **download snapshot** option shown in the GUI.

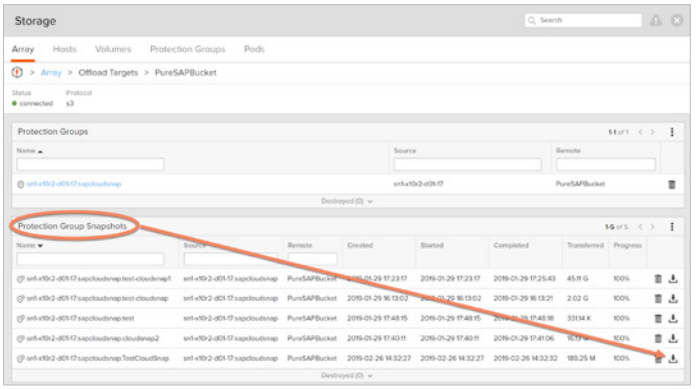

FIGURE 23. Selecting the download snapshot option

This will show the list of existing snapshots for this protection group. Select the snapshot and it will download from the S3 target.

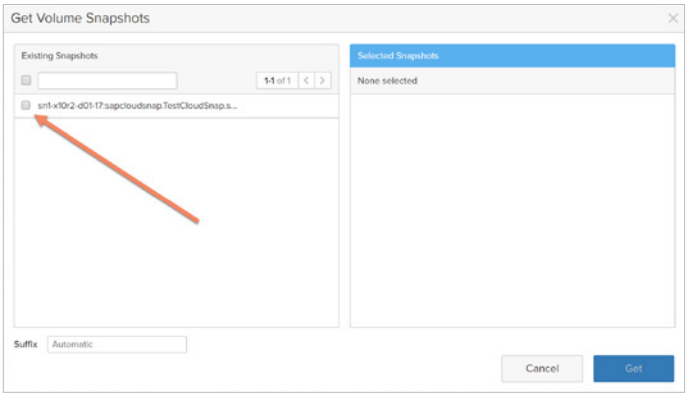

FIGURE 24. Get Volume Snapshots pane

Once the snapshot is downloaded, you can go the **Volumes** section in **Storage** tab (Pure Storage GUI). You will see this snapshot:

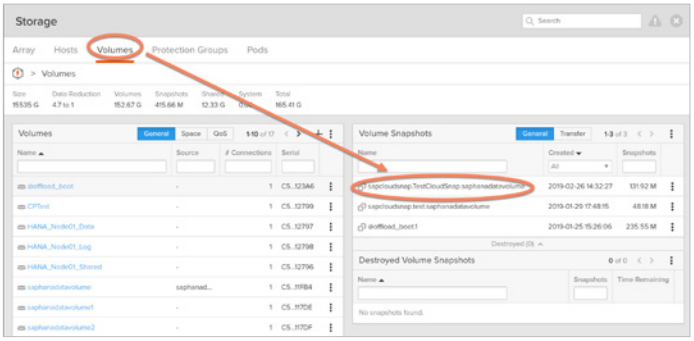

FIGURE 25. Volumes tab

You can restore this volume from this snapshot. Please unmount it before restoring.

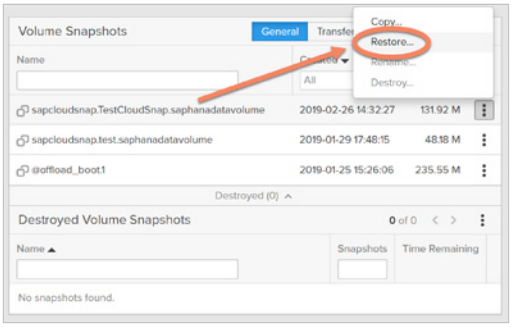

FIGURE 26. Restoring a volume

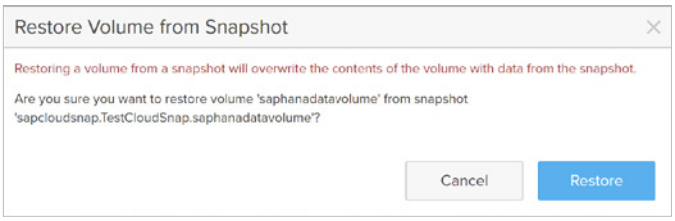

FIGURE 27. Restore Volume from Snapshot pane

Once it is restored continue with the **Recovery** process. The SAP HANA system will be recovered from this storage snapshot.

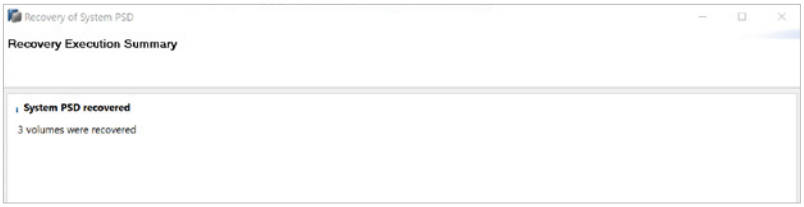

FIGURE 28. Recovery Execution Summary

In SAP 2.0, until now it has only recovered the System database. Now we need to continue with tenant recovery, as shown below:

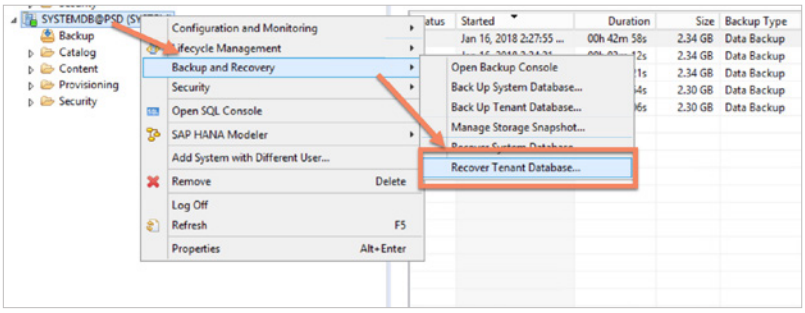

FIGURE 29. Initiating tenant recovery

Select the tenant which needs to be recovered, in this case a PSD.

FIGURE 30. Specifying the tenant database

Repeat the same steps we did for the system database. This completes the recovery for the SAP HANA 2.0 MDC system as well.

REFERENCES

1. Pure Storage FlashArray Data Sheet
2. Pure Storage Flash Storage

ps_wp17p_sap-hana-backup-and-recovery-using-cloudsnap_01