PURESTORAGE®

WHITE PAPER

FlashArray™ Protection for SAP HANA®

Technical white paper addressing data protection using storage snapshots, business continuity using Purity ActiveCluster,™ and Multi-Site Disaster recovery solutions.

September 2019

Quick Start

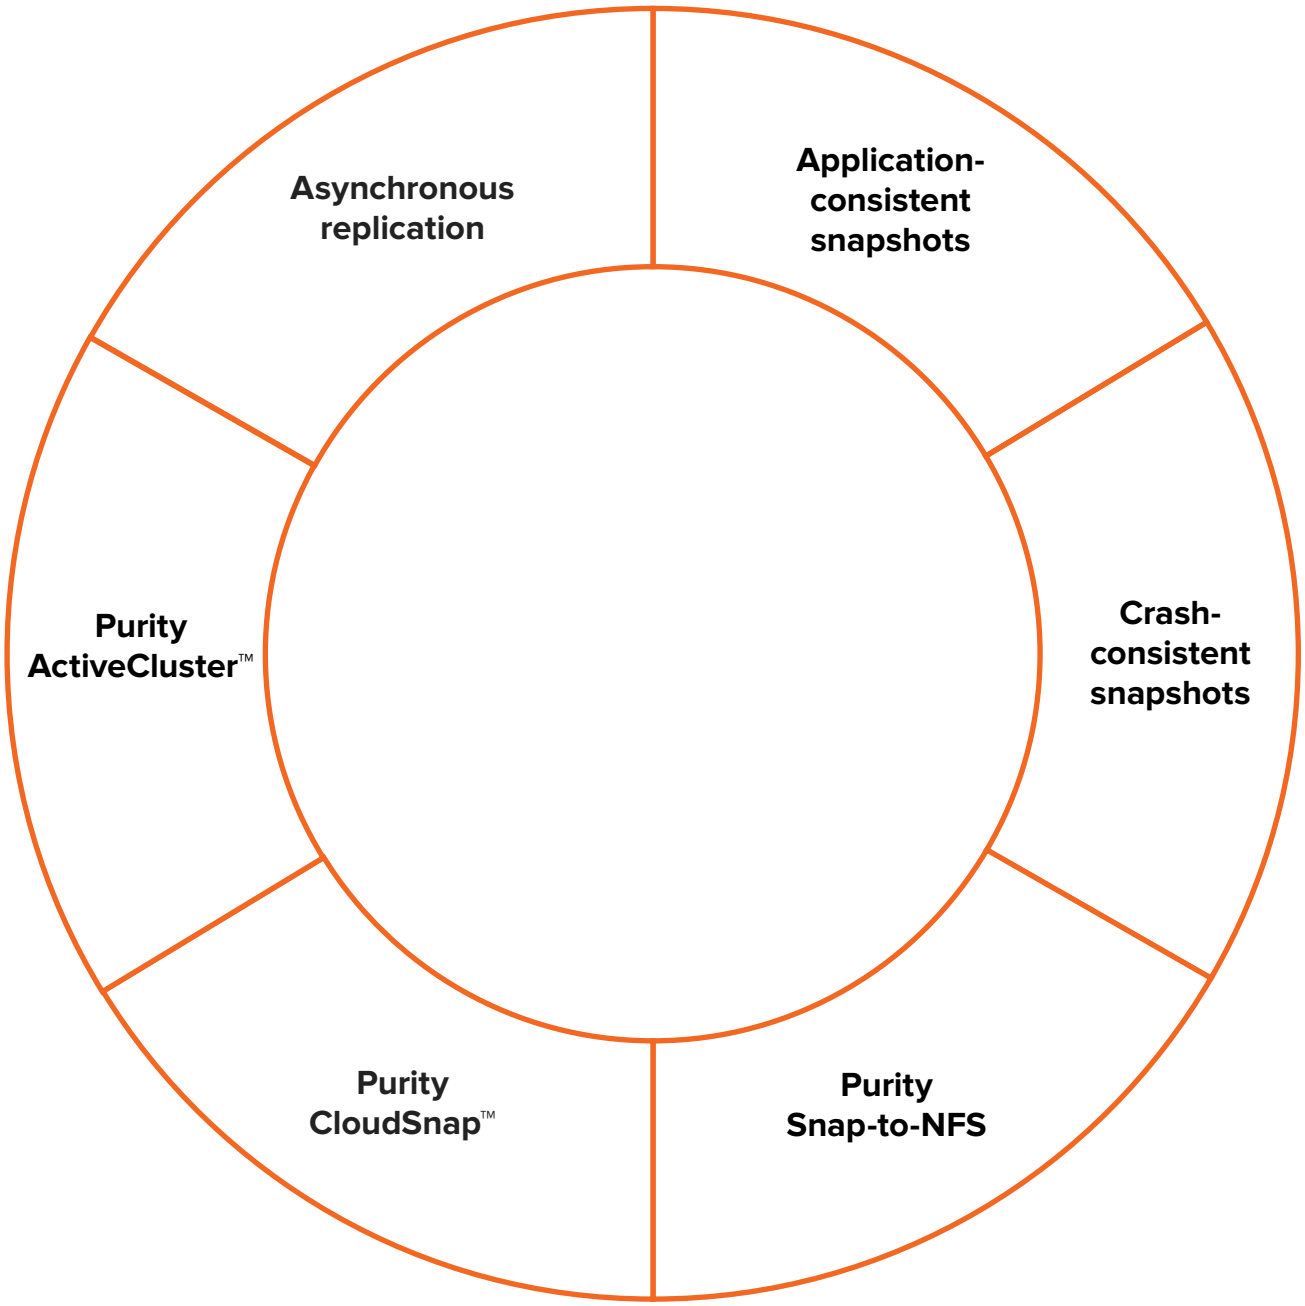
Asynchronous replication
Application-consistent snapshots
Crash-consistent snapshots
Purity Snap-to-NFS
Purity CloudSnap™
Purity ActiveCluster™

Contents

Quick Start . . . 2

Executive Summary . . . 5

Introduction . . . 6

Architectural Overview . . . 9

Data Protection Solutions . . . 10

SAP HANA Scale Up / Single host system deployment . . . 10

SAP HANA Scale Out/Multiple host (distributed) system deployment . . . 11

SAP HANA Deployment on VMware ESXi™ Hypervisors . . . 13

Business Continuity Solutions (High Availability - ActiveCluster) . . . 14

Availability and live migration within the datacenter . . . 14

Campus or metro availability (multiple datacenter availability) . . . 15

Localised high availability and remote disaster recovery . . . 16

Business Continuity Solutions (Multi-Site Disaster Recovery) . . . 17

One to Many snapshot replication . . . 17

Many to One snapshot and volume replication . . . 18

Data Protection Solutions for SAP HANA . . . 19

Application Consistent Storage Snapshots . . . 20

Overview . . . 20

Configuration and operation . . . 21

SAP HANA Scale Up / Single host system deployment . . . 21

Manual operation . . . 21

Automated operation . . . 30

Recovery Process . . . 34

SAP HANA Scale out / Distributed Host deployment . . . 65

Manual operation . . . 65

Automated operation . . . 78

Recovery Process . . . 83

Virtualised SAP HANA Protection using VMware vSphere snapshots . . . 114

Virtualised SAP HANA Protection using application consistent storage snapshots . . . 118

Crash Consistent Storage Snapshots . . . 118
Overview . . . 118
Configuration and operation . . . 118
Creating crash consistent snapshots . . . 118
Recovering from crash consistent snapshots . . . 121
Portable Snapshot Technology . . . 127
Configuration for offload functionality . . . 129
Purity Snap-To-NFS . . . 133
Overview . . . 133
Configuration and operation . . . 133
Purity CloudSnap to AWSS3 . . . 137
Overview . . . 137
Configuration and operation . . . 137
Recovering offloaded snapshots . . . 144
Business Continuity Solutions for SAP HANA . . . 147
High Availability – ActiveCluster . . . 147
Overview . . . 147
Pods . . . 149
Transparent failover . . . 149
The Pure1 Cloud Mediator . . . 150
On-Premises Failover Mediator . . . 150
Multipathing . . . 151
Uniform host access . . . 151
Symmetric Active/Active host access . . . 152
Configuration and operation . . . 154
ActiveCluster Glossary of Terms . . . 154
Removing ActiveCluster Configuration . . . 162
Multi-Site Disaster Recovery . . . 163
Overview . . . 163
Configuration and operation . . . 164
Adding additional replication target(s) . . . 174
Adding an asynchronous snapshot replication target to an ActiveCluster configuration . . . 176
Third site Disaster recovery for Active Cluster . . . 176
References . . . 183

Executive Summary

Pure Storage® offers a range of solutions to solve data protection and business continuity problems posed to organizations using the SAP HANA in-memory data platform. The FlashArray™ product family is recommended as persistence storage for data and transaction logs in any SAP HANA deployment as it offers organizations a range of benefits and incredible flexibility in solving multiple business problems in a single storage device.

Some of the problems which organizations are faced with in any SAP HANA deployment are how to ensure the continued availability of the solution when various components fail, how to meet short recovery point objectives and how to ensure business rules are adhered to with regards to recovering to a point in time or in the event of a complete data loss.

Pure Storage FlashArray and the Purity operating environment are capable of meeting these needs with an easy to use and built in management interface, a consistent understanding of any SAP HANA deployment type and without any additional license costs.

In order to meet the various business needs for data protection and business continuity the following solutions are provided by FlashArray:

- **Data protection** – Space efficient volume snapshots combined with SAP HANA storage snapshots to create application consistent recovery points in order to protect business data over time. Using offload features such as Snap-to-NFS and Snap-to-Amazon® S3 any storage snapshot can be transported from the storage array to inexpensive network file system storage or an Amazon S3 bucket, ensuring that any snapshot can be used during a full recovery in the event of rectifying system or data loss.
- **Business Continuity for device loss** – A single FlashArray offers a range of redundancies in a single chassis such as redundant controllers, port redundancy and protection against the loss of a single solid-state storage device. To further guarantee availability of the solution in the event of an entire FlashArray failure, ActiveCluster can be used as a synchronous replication solution providing an Active/Active architectural configuration where data can be read and written to any of the two arrays.
- **Business continuity for site loss** – using FlashRecover replication one or many FlashArray devices can be used as the target of an asynchronous relationship where volume snapshots can be copied to multiple sites and locations. These snapshots can then be used as recovery points in the event of a site loss.

Any of these solutions can be combined with one another to reduce business risk and create multiple resilient recovery domains ensuring organizations can focus on core business operations without the fear of disruption.

The purpose of this white paper is to provide a detailed insight into how storage snapshots, ActiveCluster and FlashRecover replication can be implemented to meet the various business needs SAP HANA has for availability and recoverability. Various reference architectures and deployment types are discussed herein attempting to highlight the differences in approach required for each scenario.

Introduction

Pure Storage® FlashArray™ is set a software defined, all flash block storage products catering to multiple business needs and use cases. The FlashArray product line is offered in 3 distinct classes:

— **FlashArray//M** - Offers exceptional block storage value for general purpose consolidation of on-premises workloads.

— **FlashArray//X** - The first all-flash, 100% NVMe® storage solution designed for a range of solutions deployed on-premises.

— **Pure Cloud Block Store™ on Amazon Web Services** – Block storage delivered natively in the cloud, powered by Purity software.

Key differentiators of the FlashArray product line are that the storage offers an effortless experience, behaves in an efficient manner by offering deduplication and compression without a reduction in performance and offers an Evergreen™ product model to increase capacity and performance without the need to keep buying new storage products. With Cloud Block Store on AWS Pure Storage further extends its ability to provide hybrid solutions to seamlessly and effortlessly move between cloud and on-prem while maintaining the benefits of both.

Fig. 1 Pure Storage FlashArray//M product models, capacity and Evergreen upgrade options.

Fig. 2 Pure Storage FlashArray//X product models, capacity and Evergreen upgrade options.

Each FlashArray is operated by the Purity operating environment, the software defined solution for flash management, basic and advanced software defined data services and storage API's. Purity offers a range of mechanisms in which the effective data protection of a business solution can be achieved at different levels, these features are collectively known as Purity Protect.

Purity Protect encompasses the following features and functionality:

— Full business continuity with Purity ActiveCluster™

— Multi-site replication

— Space-efficient local and remote snapshots

Purity ActiveCluster is a business continuity solution where Active/Active synchronous replication is enabled between two FlashArray devices at no extra cost for transparent failover mediated by an instance of the Pure1® management software. Utilizing this solution organizations can achieve zero recovery point (RPO) and zero recovery time (RTO) objectives for tier 1 applications such as hypervisors, databases and file storage services. As both FlashArray devices in the cluster are Active simultaneously the second array is not wasted and can be used for further read-only operations in the solution.

Fig. 3 Pure Storage Multi-Site Active/Active Stretch Cluster design topology.

Multi-site replication offers the flexibility to replicate volumes and snapshots in a 1: Many, Many:1 or Many: Many manners between FlashArray devices. This is typically appropriate for data sharing, centralised data protection strategies and disaster recovery.

SAP HANA is an in-memory data platform engineered, marketed and sold by SAP® SE. The data platform offers and in-memory, column oriented relational database management system primarily focused on data processing and analysis in a performant manner. Even though SAP HANA is an in-memory data platform, it requires a high performing persistence layer which is based storage area network (SAN) or network attached storage (NAS). Typical deployments require that a data and log area be provisioned for the persistence layer where both are required to meet a range of key performance indicators (KPI) for production instances. One of the characteristics of the log area is that it is required to achieve as low a latency as possible, this is due to the architecture of SAP HANA where transactions/operations are entered simultaneously in memory and serialized as a write operation to the log volume. As soon as both complete the transaction/operation is considered completed. The low latency requirement becomes more apparent when dealing with systems which perform enormous amount of transaction-processing style operations as when the log area comes under load, the entire data platform must not slow down and provide a high performant service to all users and applications. SAP HANA is offered in two distinct deployment scenarios:

— **Scale up** – a single compute and memory domain scaled by adding more resources directly to the domain. i.e. increasing memory or compute capacity.

— **Scale out** – multiple compute and memory domains scaled by adding additional similar domains alongside existing deployments. Also provides for hardware failures through the provisioning of a standby node for the instance i.e. adding another server to an existing server pool and extended the SAP HANA instance to utilise it.

Pure Storage FlashArray devices offer SAP HANA deployments a distinct advantage due to the following reasons:

— The FlashArray//X product line is a 100% NVMe storage solution providing low latency for both the log and data areas in any SAP HANA Tailored Datacenter Integration (TDI) deployment.

— Like other vendors, Pure Storage offers turn-key appliances for SAP HANA covering networking, compute and storage. . But Pure Storage pursues an SAP TDI strategy. So both FlashArray and our FlashStack SAP appliance are SAP TDI-certified to provide organizations with the flexibility to choose the best, cost-effective and appropriate solution that meets their needs. And either deployment option scales from on-prem into the cloud. .

— The Evergreen product model allows organizations to increase performance, scalability and capacity over time without the need to purchase entirely new storage.

— Pure Storage FlashArray includes a range of data services aimed at enabling customers to realize the full potential of their SAP HANA deployment, namely ActiveCluster and Multi-Site replication.

SAP HANA offers its own internal business continuity approach for data protection and replication solutions each of which is integrated into the core product. Business continuity can be achieved by utilising a scale out instance with one or more failover nodes, backups to NFS or a backint certified storage target, storage replication and system replication.

The purpose of this white paper is to provide a detailed overview of the different business protection and continuity areas provided by Purity Protect for SAP HANA deployed on a FlashArray. Different sections explored are how to configure each solution area or an appropriate business case for each.

Architectural Overview

Architectures of Pure Storage solutions with SAP HANA covered within this white paper are data protection, high availability and multi-site disaster recovery.

All of the outlined architectures rely on the Purity operating environment's ability to create snapshots of volumes on an on-premise FlashArray solution or hybrid Cloud Block Store deployment. Data protection architectural solutions will encompass the ability of volume snapshots to be transported to a third platform for retention and recovery purposes. High availability architectural solutions focus on eliminating business downtime risk by creating a storage failure domain both inside a single datacenter and across a campus or metro deployment area. Business continuity architectural solutions are similar to the outline of data protection solutions but provide for the ability to replicate snapshots to other FlashArray or Cloud Block Storage devices.

The limit of snapshots created by the Purity operating environment are not isolated to data protection or business continuity solutions. Once the volume snapshot is created it can be used to enable data mobility architectural solutions. Data mobility in the instance of SAP HANA is defined to be the architectural federation of 2 or more FlashArray storage devices where volumes and snapshots can be transported from a resource constrained system to another with more resource availability. This process enables organizations to continue to function in an efficient manner without an impact on responsiveness. This is particularly useful for SAP HANA deployments where resource constraints are a critical issue as business needs scale over time.

SAP HANA Important information and terminology

SAP HANA releases are based around a major revision, minor version and a minor database revision for patches and patch rollups (e.g. SAP HANA 2.0 SPS 03 Database Revision 034). Currently there are only two major revisions (SAP HANA 1.0 and SAP HANA 2.0), 12 minor versions for SAP HANA 1.0 and 4 minor versions for SAP HANA 2.0. The minor versions are identified as "Support Package Stacks" or "SPS" abbreviated where the most recent release is SAP HANA 2.0 SPS04.

SAP HANA is available for Intel®-®-based hardware platforms (x86_64) and IBM® Power® Systems (PowerPC® architecture). It is also only supported to run on SUSE® Enterprise Linux® (SLES) and Red Hat® Enterprise Linux (RHEL), further information can be found at https://launchpad.support.sap.com/#/notes/2235581 (SAP Support login required). It is recommended to use the "SUSE Linux Enterprise Server for SAP Applications" and "Red Hat Enterprise Linux for SAP Solutions".

Pure Storage offers a flexible SAP HANA solution aimed at tailored datacenter integration scenarios. Organizations have the opportunity to choose the most suitable server, networking and storage platform separately to meet both efficiency and cost business requirements.

With the release of SAP HANA 1.0 SPS 09 a new feature called multitenant database containers (MDC) was introduced. A multitenant database container is a single system with a system database which stores and maintains a system-wide landscape to allow configuration and monitoring of the overall system while allowing for multiple other databases to be deployed in parallel to one another. Each system can be made up of multiple tenant databases which are isolated from one another in terms of application data and user management. For SAP HANA 1.0 SPS 09 onwards MDC is an optional feature until SAP HANA 2 SPS 01 where it becomes the only operation mode.

With the release of SAP HANA 2.0 SPS03, the use of Intel® Optane™ DC Persistent memory is supported and provides a streamlined data tiering perspective and lower TCO through reduced downtime. Even with the existence of persistent memory SAP HANA requires performant storage with data services as a persistence layer in order to meet many of the business requirements set out by organizations.

Data Protection Solutions

SAP HANA Scale Up / Single host system deployment

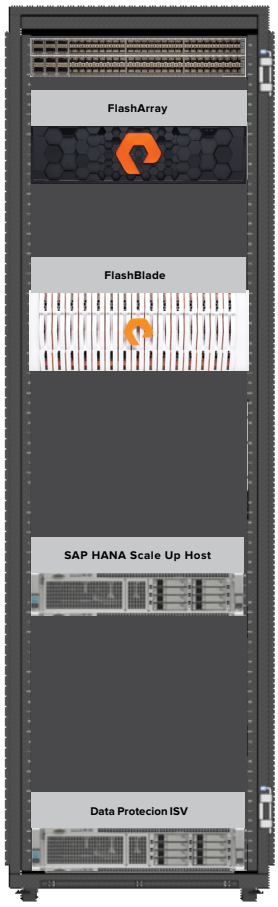

SAP HANA Scale up deployments commonly utilise a set of volumes mounted over direct attached storage, Fibre channel or network file system (NFS) connectivity.

These volumes are then formatted with a file system (typically XFS) and are then mounted at various locations within the Linux operating system.

In this document the volumes for SAP HANA are mounted at the locations set out below:

— Install/base/binaries directory - **/hana/shared**

— Data persistence directory - **/hana/data**

— Transaction log directory - **/hana/log**

— Backup directory - /**hana/backup**

Included in the scale up deployment is a Pure Storage FlashBlade™ and Data Protection independent software vendor to manage the protection and retention of transaction logs. When using FlashBlade as a target for block volume snapshot offload this is considered a rapid recovery scenario, utilising the performance capabilities of this device to shorten recovery operations.

Transaction log backups can be sent directly to a local volume formatted with a file system or Network File System (NFS) target; however, this approach does not include retention management of protected transaction logs.

SAP HANA Scale Out/Multiple host (distributed) system deployment

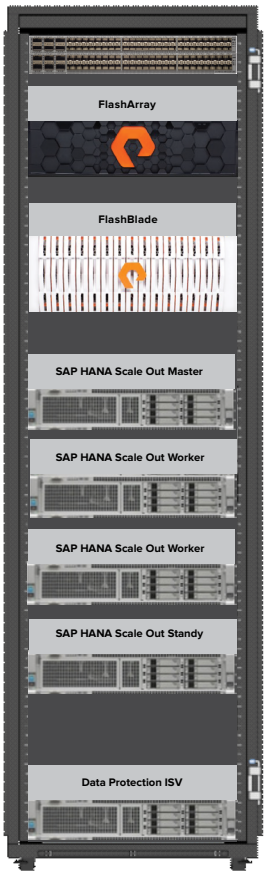

SAP HANA Scale out deployments are made up of multiple hosts where each can be configured as an active worker or idle standby host. The minimum required hosts for a highly available scale out system is a single master worker, a worker and a standby host (2 + 1 topology). While having standby nodes for high availability is not necessary in a scale out deployment, it is strongly advised.

When deploying a scale out system it is mandatory to utilise a network file system (NFS) for the install/base/binaries directory and this directory must be shared across all of the hosts (both worker and standby) which will be included in the deployment. In a scale out deployment each host will have its own data and log volumes, and all volumes for each worker must be presented to each standby host.

In this document the volumes for the SAP HANA scale out deployment (3 worker and one standby host) are set out below:

— SAP HANA scale out master
/hana/shared, /hana/data, /hana/log, /hana/backup

— SAP HANA scale out worker 1
/hana/shared, /hana/data, /hana/log, /hana/backup

— SAP HANA scale out worker 2
/hana/shared, /hana/data, /hana/log, /hana/backup

— SAP HANA scale out standby
/hana/shared
Master, worker 1 and worker 2 Data Volume
Master, worker 1 and worker 2 Log Volume

Included in the scale up deployment is a Pure Storage FlashBlade and Data Protection independent software vendor to manage the protection and retention of transaction logs. The FlashBlade will be used as the shared NFS storage provider for the **/hana/shared** directory on each host.

SAP HANA Deployment on VMware ESXi™ Hypervisors

SAP HANA is supported to be deployed on VMware vSphere® as a virtual machine or set of virtual machines. Both Scale up and scale out deployment scenarios can be configured, but it is advised that when deploying production instances, a scale up configuration be used. Some of the benefits of a virtualised SAP HANA instance are as follows

— Provisioning instances of SAP HANA in virtual machines is significantly faster

— Live migrations of SAP HANA instances can be performed using VMware vSphere vMotion®

— Standardised high availability using VMware vSphere High Availability (HA)

This document will only look at SAP HANA scale up deployment scenarios when utilising VMware vSphere as a platform.

Business Continuity Solutions (High Availability - ActiveCluster)

Availability and live migration within the datacenter

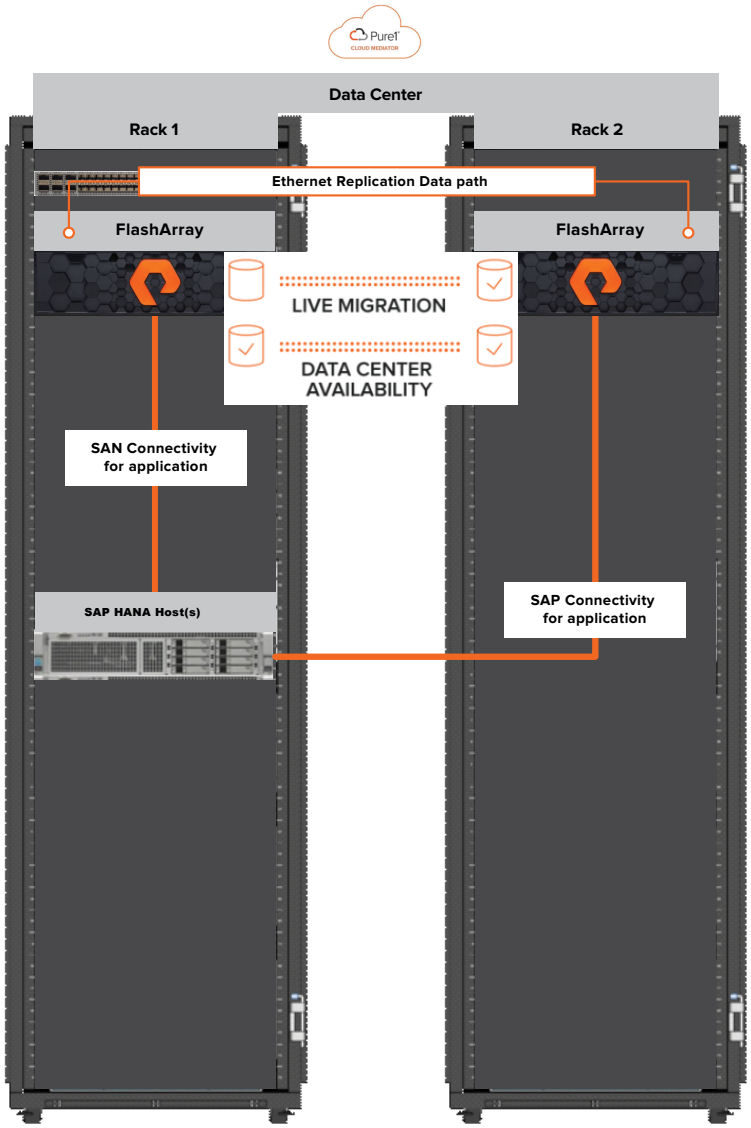

Fig. 4 ActiveCluster inside a Datacenter deployment scenario.

For business continuity within a single datacenter, ActiveCluster is used to ensure the availability of storage for a single SAP HANA deployment (both scale up and scale out) within a single datacenter. This is a synchronous clustering of a separate set of the data volumes used for the SAP HANA deployment. Block storage data and operations are performed over ethernet between the two FlashArray storage appliances. Pure1 (Data storage, management and support) is used to mediate the availability of the storage in an ActiveCluster deployment.

ActiveCluster can also be used to migrate live data between two sets of systems without the need for any downtime or performance degradation.

Campus or metro availability (multiple datacenter availability)

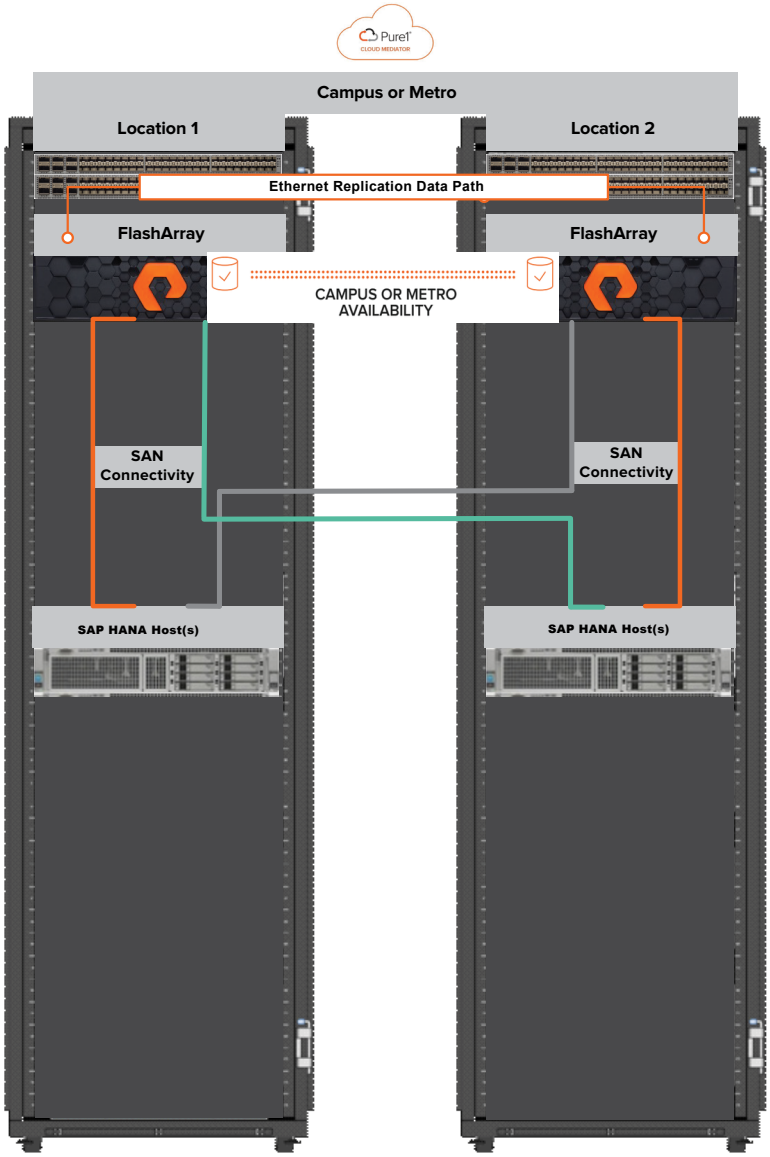

Fig. 5 ActiveCluster across campus or metro availability zones with secondary compute at the redundant site.

For business continuity across multiple datacenters, ActiveCluster can be used as both a high availability and disaster recovery solution based on the replication of data and log disks to remote storage attached to a secondary SAP HANA system (both scale up and scale out). This solution can be configured to operate both synchronous and asynchronously over ethernet depending on the customer need and business rules. When using this deployment, the secondary system is not wasted and can be used for test, quality assurance or read only operations on the production data. Pure1 is used to mediate the availability of the storage in this deployment scenario and can provide information on which system should be utilised for business operations.

Localised high availability and remote disaster recovery

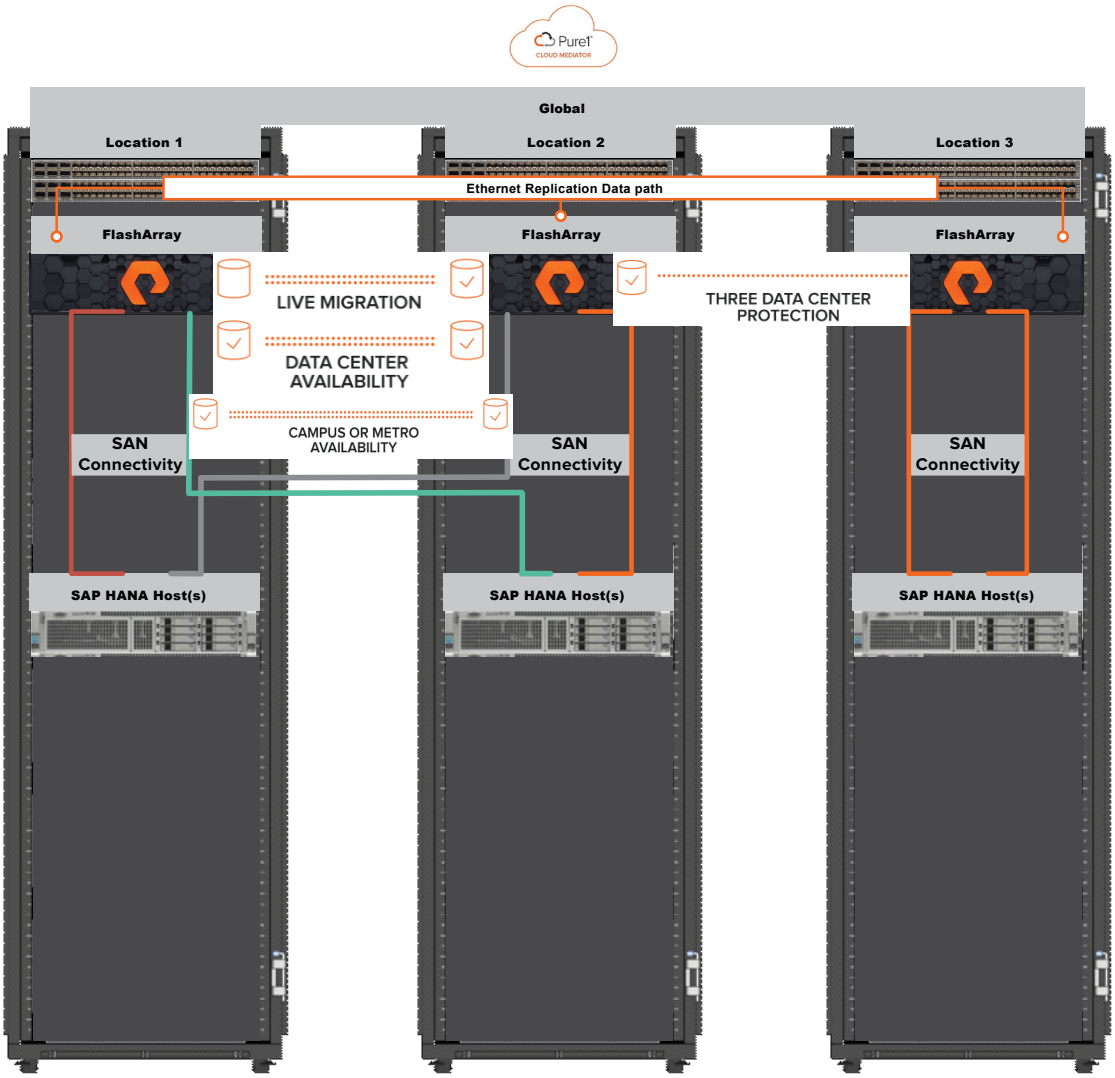

Fig. 6 ActiveCluster deployed in a localised area and remote asynchronous replication for a third site.

ActiveCluster can be further extended from localised high availability scenarios (storage availability and storage replication with a secondary SAP HANA system by adding a remote FlashArray to the ActiveCluster group and having a standby SAP HANA host attached to it. Deploying this scenario requires that the localized high availability group operate synchronous replication over ethernet between each FlashArray and asynchronous replication is then used for replication to the remote system. Pure1 is used to mediate the availability and status of the storage replication relationships and can provide information on which location needs to be used by business applications and operations. An onsite virtual machine can be used as a mediator if the use of Pure1 is not possible.

Business Continuity Solutions (Multi-Site Disaster Recovery)

One to Many snapshot replication

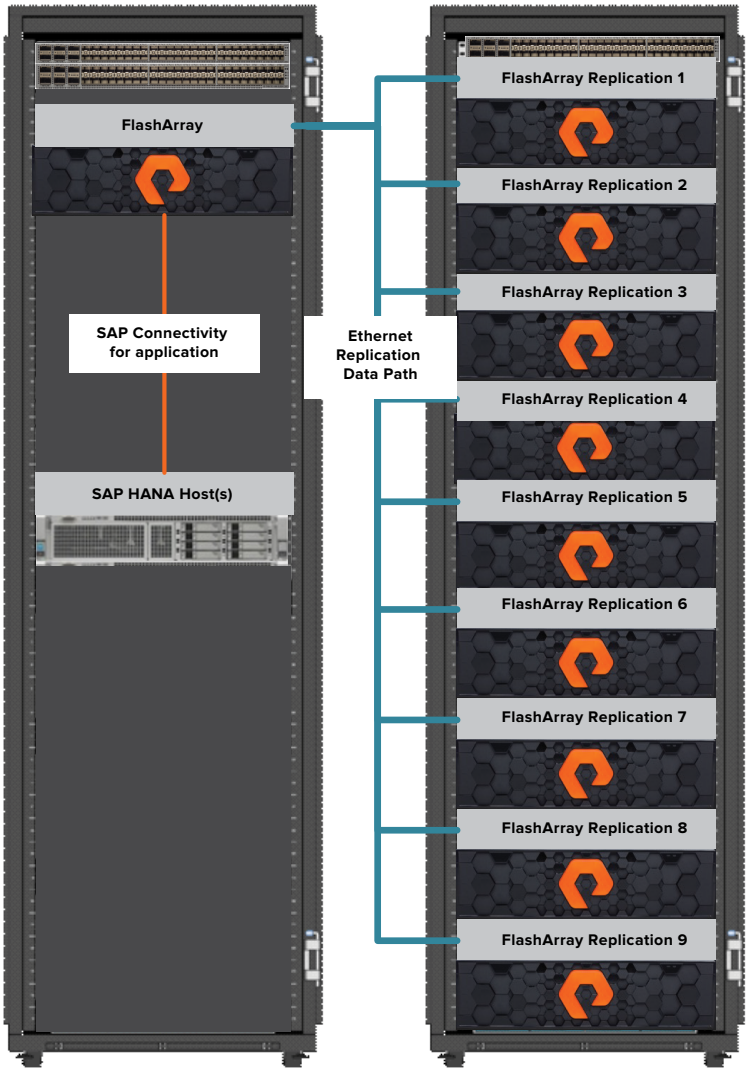

Fig. 7 A single FlashArray system replicating snapshots and volumes to as many as 9 other FlashArrays.

On-Premise FlashArray and hybrid Cloud Block Store deployments can be configured to replicate snapshots to multiple other Pure Storage Purity based deployments to enable disaster recovery over multiple sites or availability zones.

Many to One snapshot and volume replication

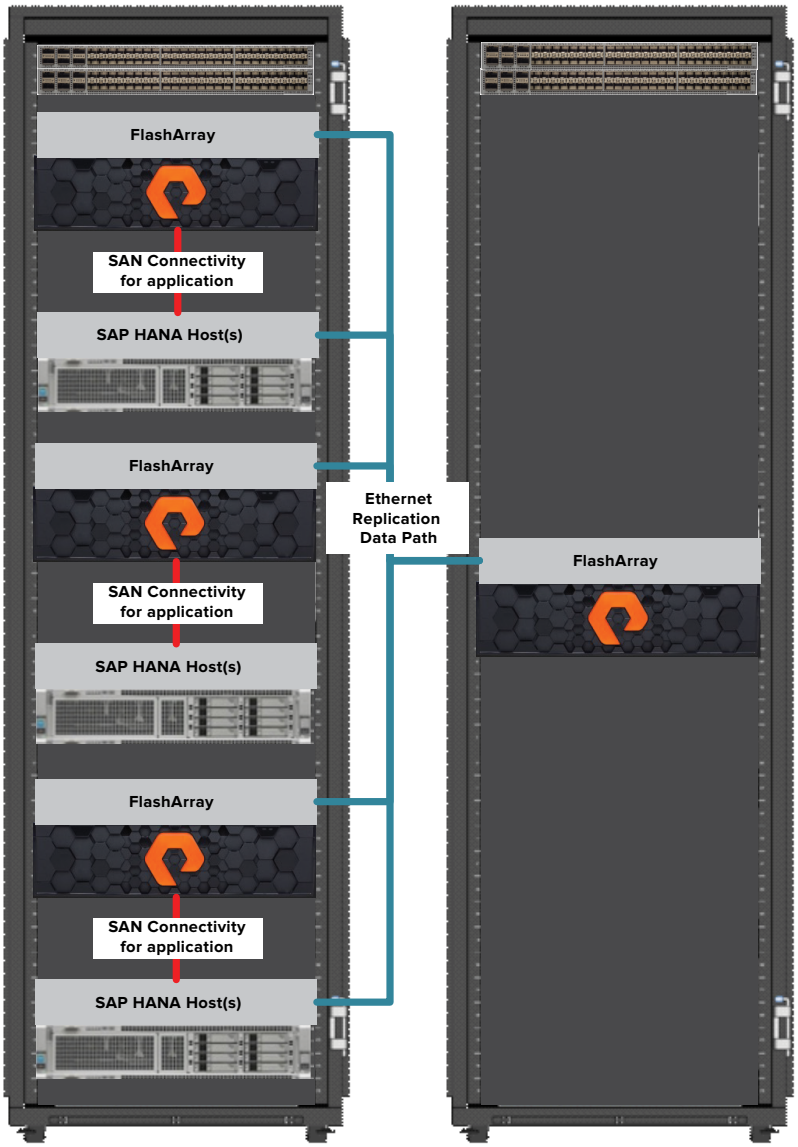

Fig. 8 Many FlashArray systems replicating snapshots and volumes to a single FlashArray.

A single on premise FlashArray or hybrid Cloud Block Store deployment can be used as a single replication target for many other Pure Storage Purity based deployments to centralise business continuity business operations.

Data Protection Solutions for SAP HANA

SAP HANA offers a range of backup types to enable backup and recovery operations. The following backup types are supported by SAP HANA:

Full Backups

Full backups can be performed in one of two ways,

— Backup of the data area streamed off to a local disk, NFS mount or via a supported third-party backup tool using SAP HANA backint.

— Data snapshots where the data persisted in the data area at a particular point in time is created with all of the information required to recover SAP HANA to a consistent state.

Delta backups (incremental and differential backups) can be created after a full backup has been taken (excluding data snapshots) to shorten the amount of time required for a backup window.

Data Snapshots are created in the storage system with minimal impact on database performance and do not consume any additional resources such as memory or compute. Recovery operations from data snapshots are significantly faster than recovery from a streamed backup as it only needs to be made available in the data area of the storage system.

Redo log backups

The architecture of SAP HANA In-memory databases is one which requires the persistence of transaction logs to ensure durability. With each insert, update or delete command executed the operations are serialised and written to the log area. As the log area fills to capacity, in order to ensure both further durability and capacity management, backups of each transaction log need to be made and moved to another storage medium. Once a backup of a transaction log has been created, the space it occupies in the log area can be freed up for further transactions to take its place. If the log area fills to capacity then the SAP HANA database will be unable to process any further transactions and the system will halt until the space is available to create new transaction logs. Transaction logs are rolled up and changes are merged into the persisted data volume every 5 minutes (default) allowing for the data area to be in a consistent state when row or columnar data must be read from storage.

During the recovery of a database both the log area (if it still exists) and any redo log backups can be used to replay transactions which have occurred since the last backup ensuring the database and business data available is in its most recent state.

Backups created using third-part backup tools

Backups created using 3rd party tools use Backint for SAP HANA which is an API focused on integrating backup and recovery, retention and vendor specific parameters directly from SAP HANA's tools. Vendors which develop backint for SAP HANA functionality and support are required to undergo a certification and approval processes by SAP in order to be used in production deployments.

In the instance where the SAP HANA deployment is a virtual machine hosted within a VMware vSphere environment the data protection solution can be done on multiple levels namely data streaming or storage snapshots. However, the data protection solution an organization will choose depends on factors such as SAP HANA deployed on virtual volumes (VVols) or VMFS based storage.

Application Consistent Storage Snapshots

Overview

In order to create application consistent storage snapshots for SAP HANA databases it is important to note that this operation requires the database instance, operating system and storage system to all be synchronised during execution.

Application consistent snapshots can be manually performed using SAP HANA Studio, a secure shell terminal (SSH) signed in with a user who has read/write and execute permissions to the SAP HANA data volume and the Pure Storage FlashArray web user interface. The process can also be automated using programmatic logic interfacing with SAP HANA database SQL, secure shell bash commands and Pure Storage FlashArrays ReST API or the associated software development kit (SDK).

The process for creating application consistent snapshots differs for systems deployed with multiple database containers (a system with a single tenant is still deployed as MDC), single containers, single host (scale up) and distributed host (scale out) systems. Storage snapshots are only supported by SAP from the SAP HANA 2.0 SPSSPS 04 release and onwards.

Within this section the areas explored will be how storage snapshots for SAP HANA systems can be created manually by a user or automatically for a landscape of different systems and system types. The concepts discussed within each deployment type for automation are high level and can be applied to any scripting or high-level programming language with the relevant libraries and support. It is assumed that any SAP HANA Scale up deployment has a single volume for data and log mount points separate from one another and not configured using logical volume manager. Recovery for each deployment type is also explored.

Required libraries/modules for automation

— A library/module which enables the user to SSH into the SAP HANA system with a username and password
  - Posh-SSH for PowerShell (https://www.powershellgallery.com/packages/Posh-SSH/2.0.2)
  - SSH.Net for .Net Languages (https://github.com/sshnet/SSH.NET)
  - JSch for Java® (http://www.jcraft.com/jsch/)
  - Paramiko for Python® (https://github.com/paramiko/paramiko)

— If possible, an SDK to work with the Pure Storage FlashArray ReST API
  - PowerShell SDK for FlashArray – (https://www.powershellgallery.com/packages/PureStoragePowerShellSDK/1.13.1.12)
  - Python Pure Storage REST Client – (https://pypi.org/project/purestorage/)
  - View the ReST API documentation for FlashArray on the FlashArray Web GUI (Help->REST API Guide)

— A library/module to interact with SAP HANA using the hdbsql query language.
  - SAP HANA Client on Windows® (https://help.sap.com/viewer/e7e79e15f5284474b965872bf0fa3d63/2.0.01/en-US/68f5b289fab2427e9580a4524071ba96.html)
  - SAP HANA Client on Linux/UNIX® (https://help.sap.com/viewer/e7e79e15f5284474b965872bf0fa3d63/2.0.01/en-US/006cc8dc05b2404cb6148493f854b7cb.html)

Configuration and operation

SAP HANA Scale Up / Single host system deployment

MANUAL OPERATION

Step 1. Verify the system type

Using **SAP HANA Studio** connected to the system deployed with the SAP HANA Instance, navigate towards the **configuration** page

SAP HANA Studio, Configuration Page.

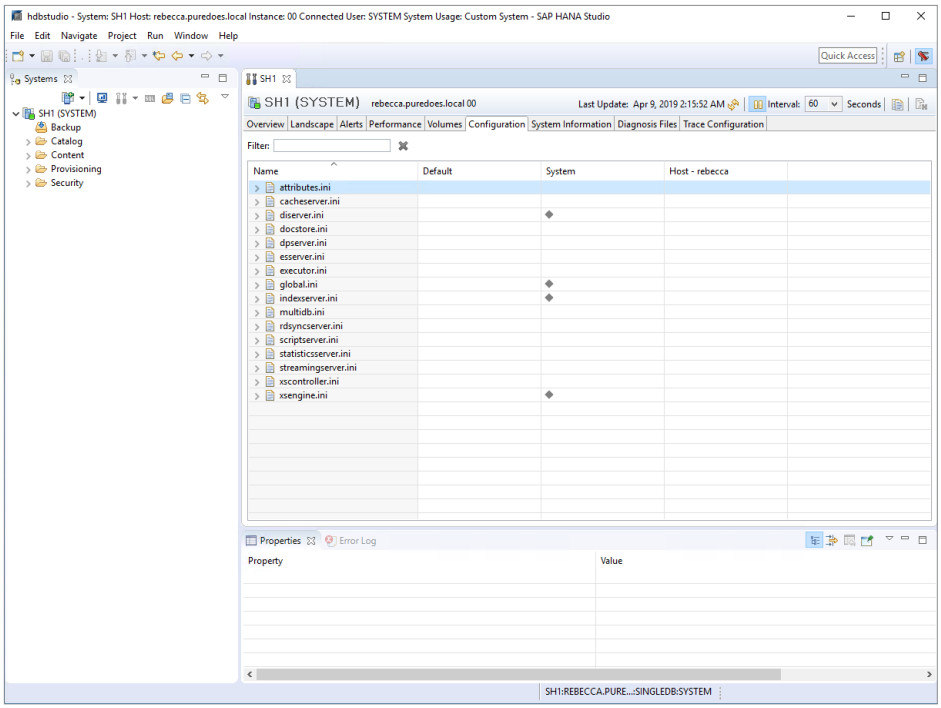

Expand **global.ini** and then further expand the **multidb** section. Look for the key **mode** and observe its value. If this value is “singledb” then the system is a single container and if the value is “multidb” the system is then set to be a multiple container system.

In SAP HANA Studio the global.ini file is expanded.

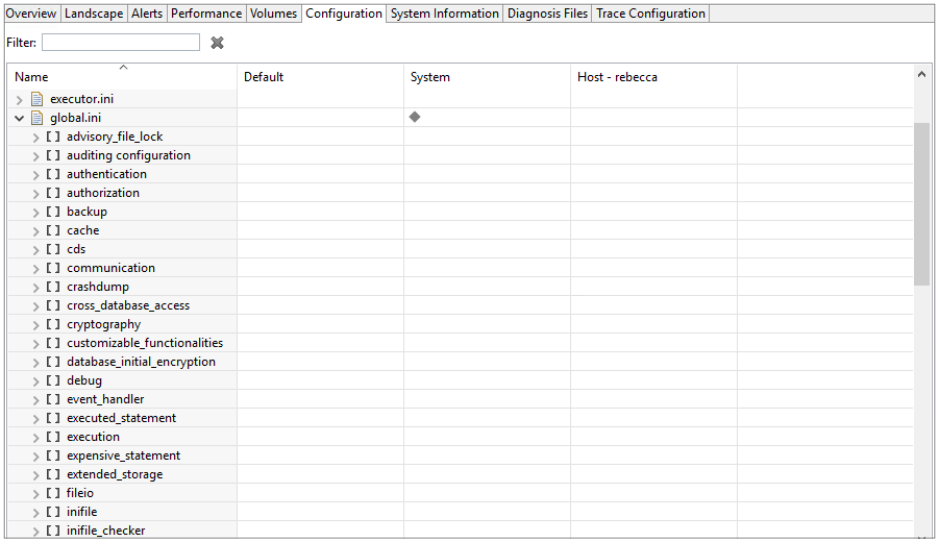

In SAP HANA Studio the multidb section under global.ini shows the system properties.

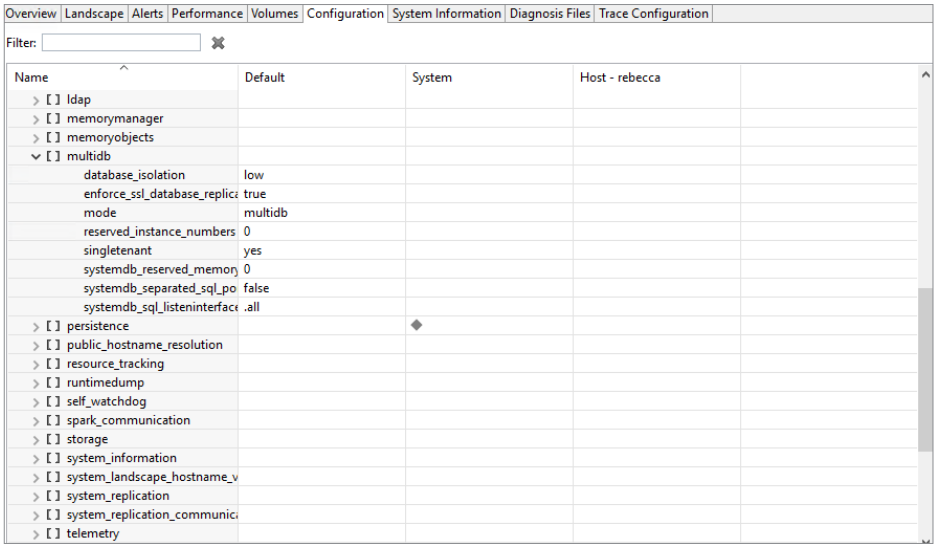

Step 2. Connect to the correct database instance for the creation and management of storage snapshots

If the system is a single database container then ensure that under mode the **Single container** is selected, for a multiple container system ensure that **Multiple containers** is selected, and when available ensure that **System database** is the selected target to connect to. It is important to note that selecting Single container and the mode to connect to will still connect to the single tenant container.

Using SAP HANA Studio connect to a single container system.

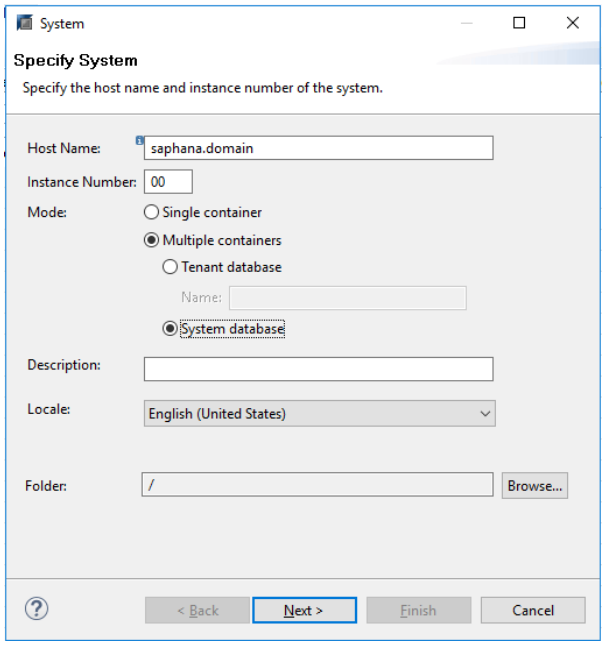

Using SAP HANA Studio connect to the System dataset on a system with multiple containers.

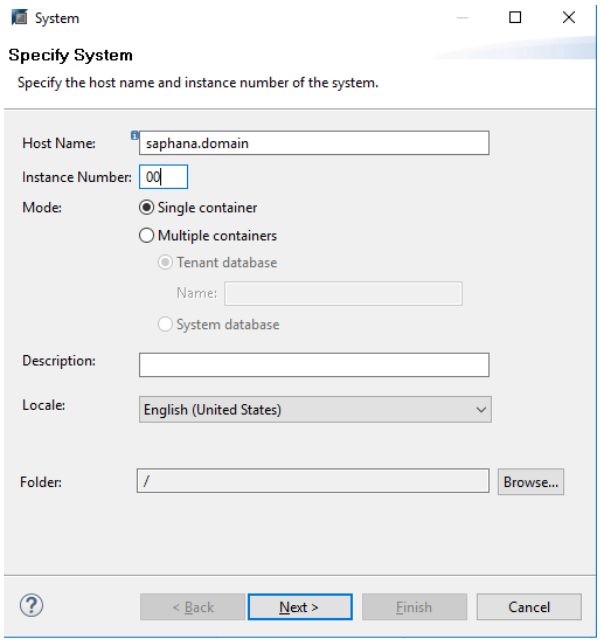

Connected to a single container system.

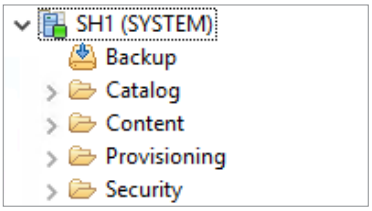

Connected to a multiple container system.

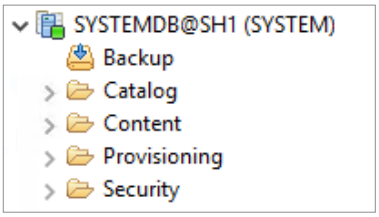

Step 3. Prepare the system for a storage snapshot

This can be done from the backup console or using the side menu by right-clicking on the **Backup** system folder and selecting **Manage Storage Snapshot....**

In SAP HANA Studio right click on the "Backup" system folder and select "Manage Storage Snapshot...".

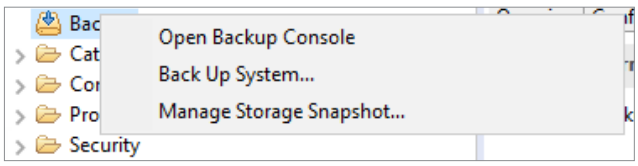

In SAP HANA Studio and the backup console, select the camera next to "Status of Currently Active Data Backup".

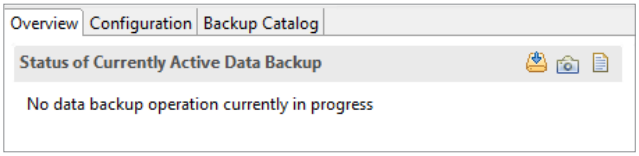

Select Prepare and add a comment if needed, then press the OK button.

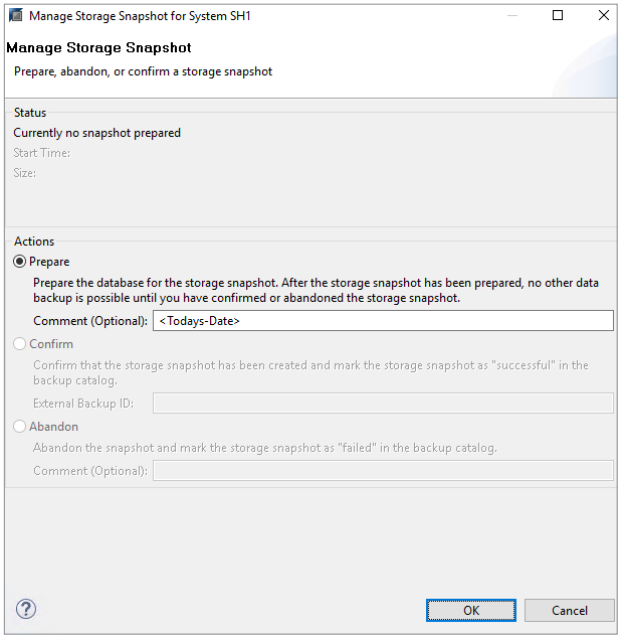

Once the database snapshot is ready then the below will show in the backup console.

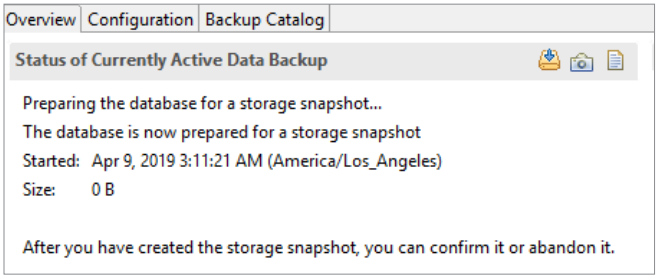

Step 4. Retrieve the SAP HANA prepared storage backup ID using the SQL console in SAP HANA Studio

Use the SQL console for the database instance the prepared snapshot has been created for and run the query:

```
"SELECT BACKUP_ID, COMMENT FROM M_BACKUP_CATALOG WHERE ENTRY_TYPE_NAME = 'data snapshot'
AND STATE_NAME = 'prepared'"
```

This will return the backup ID for the prepared snapshot, which is used in the FlashArray block device snapshot creation as a suffix to link the SAP HANA storage snapshot ID to the Volume snapshot.

In SAP HANA Studio right click on the instance and select "Open SQL Console".

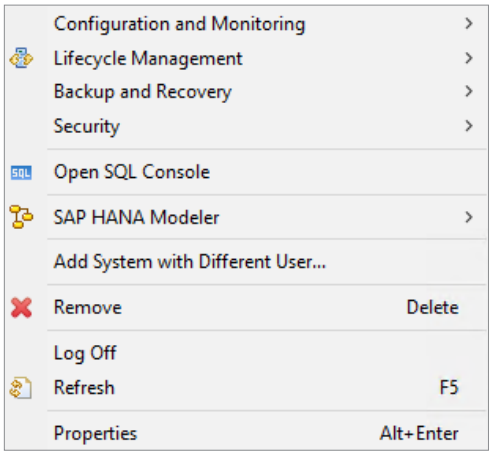

In SAP HANA Studio execute the query to return the Backup_ID of the prepared database snapshot.

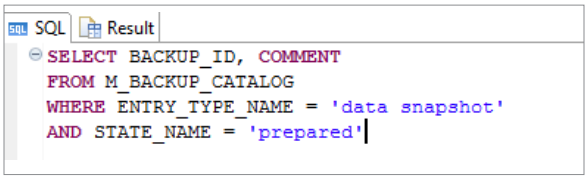

The Backup ID is returned with any comments added to the entry, take note of the Backup.

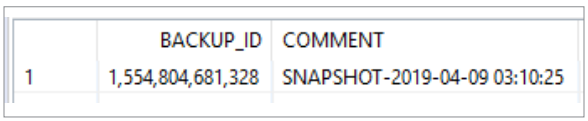

	BACKUP_ID	COMMENT
1	1,554,804,681,328	SNAPSHOT-2019-04-09 03:10:25

Step 5. Freeze the filesystem for the SAP HANA data persistence mount point

Open a terminal (SSH or local to the system) and ensure the prompt is logged in as a user who has read, write and execute permissions on the SAP HANA data persistence mount point. The data persistence mount point can be identified by inspecting the global.ini persistence section for “basepath_datavolumes”. Take the value offered by the global.ini file and remove the database name from the path to only get the base path mount point (this assumes that the SAP HANA systems has been setup with a mount point and single volume each for log and data). We will then use the “fsfreeze” Linux utility to halt any IO to the volume and ensure consistency.

In SAP HANA Studio the multidb section under global.ini shows the system properties.

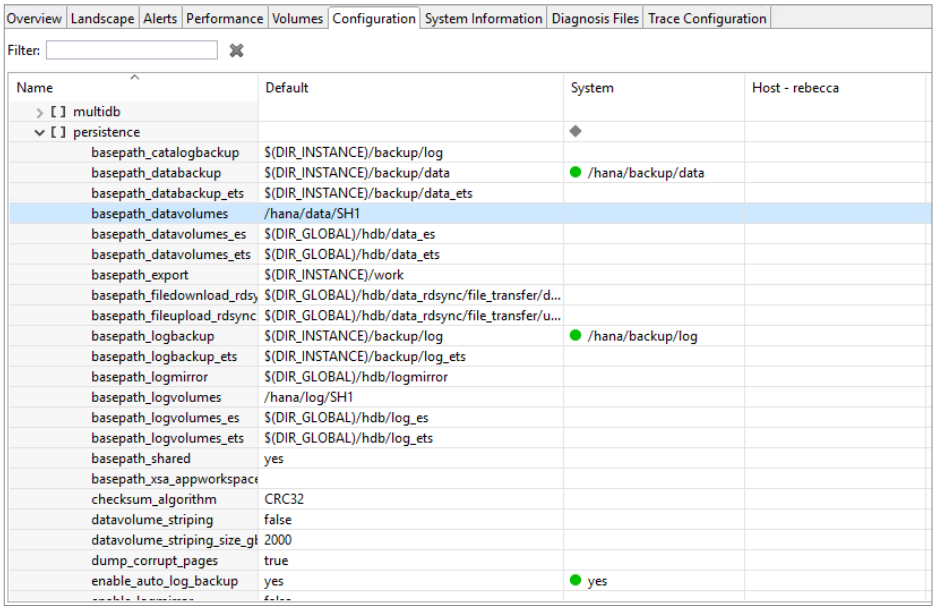

Freeze the filesystem of the data persistence mount point using the fsfreeze utility

```
:~ # /sbin/fsfreeze --freeze /hana/data
```

Step 6. Take a snapshot of the block device in the FlashArray

This step is shown using the web based graphical user interface to operate the FlashArray storage device. It is assumed the user can identify the block volume which matches the SAP HANA persistence data volume. In the user interface navigate to **Storage**, select the **Volumes tab**. Select the volume which corresponds to the SAP HANA persistence data volume. Under **Volume Snapshots**

Pure Storage FlashArray Web Graphical User interface – Main page.

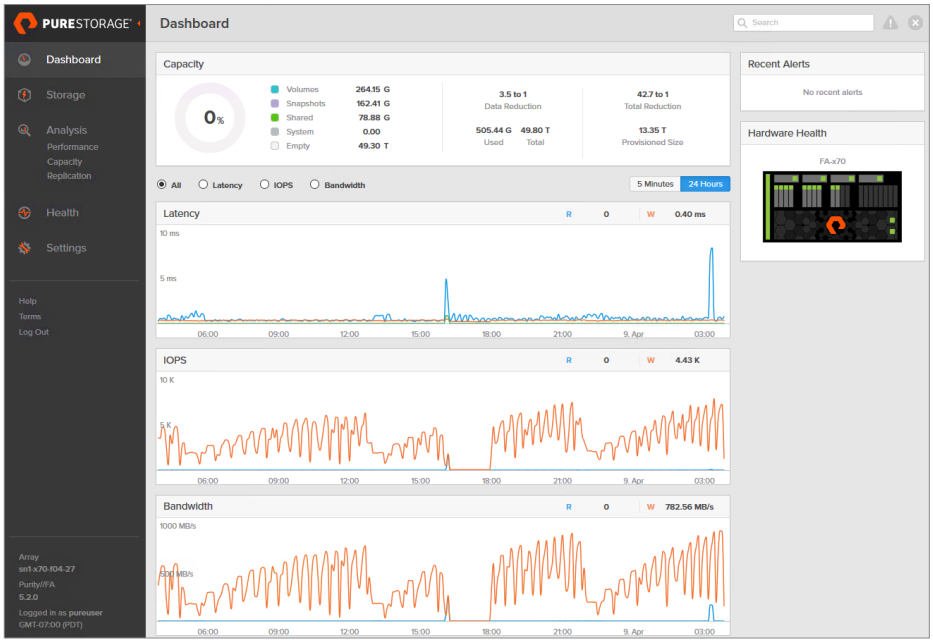

Select Storage and navigate to the "Volumes" section to view all volumes and snapshots.

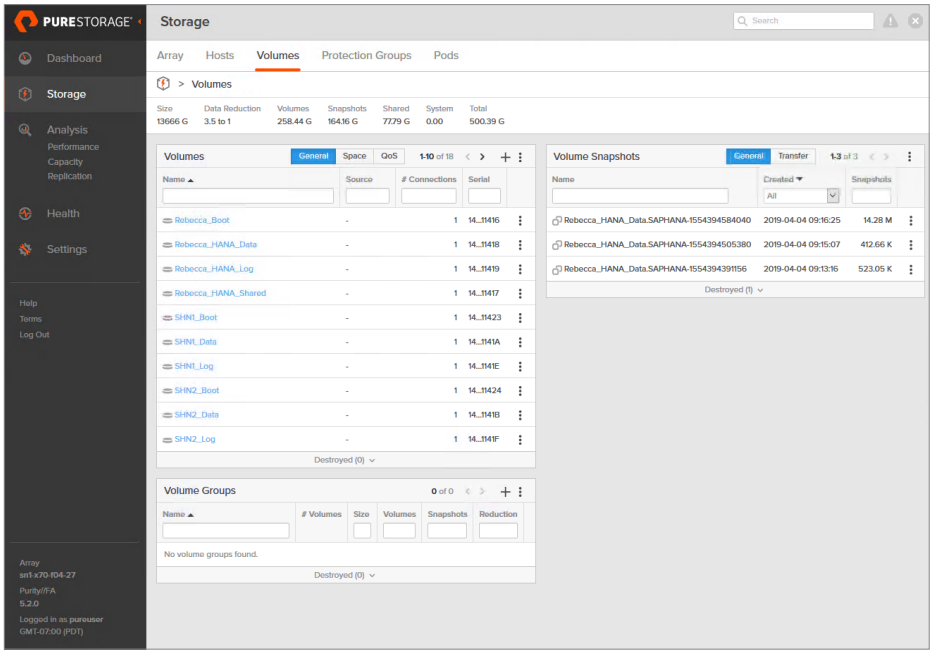

Select the volume which matches the SAP HANA Data persistence volume and select the "+" next to Volume Snapshots.

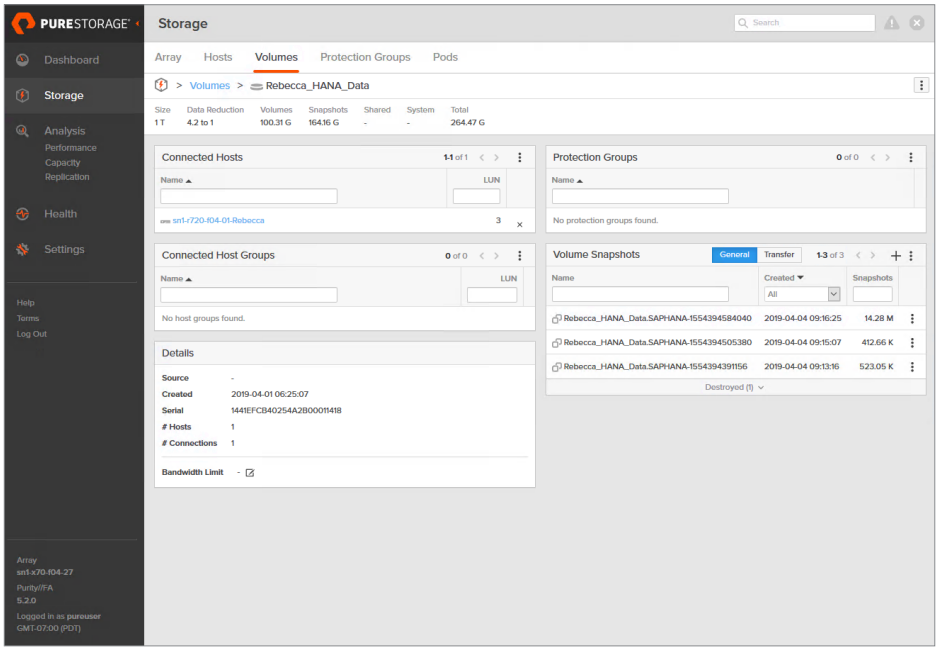

Create the snapshot with the suffix "SAPHANA-<Backup_ID from SAP HANA prepared snapshot>

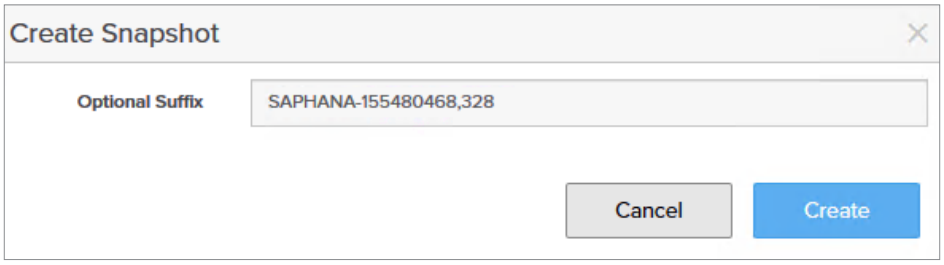

The snapshot is created and listed under the Volume snapshots.

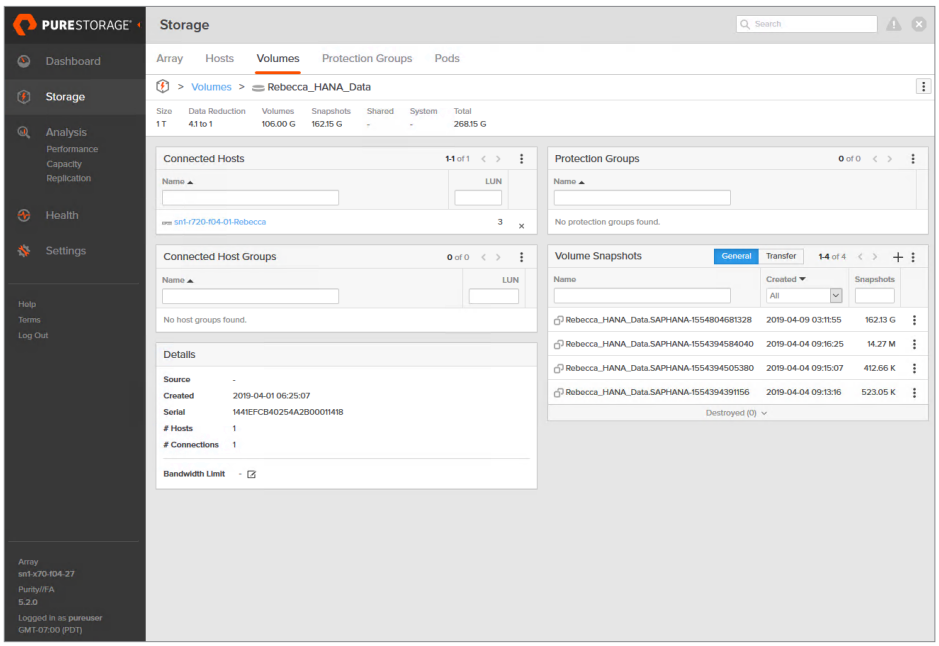

Step 7. Unfreeze the filesystem for the SAP HANA data persistence mount point

Open a terminal (SSH or local to the system) and ensure the prompt is logged in as a user who has read, write and execute permissions on the SAP HANA data persistence mount point. The data persistence mount point can be identified by inspecting the global.ini persistence section for “basepath_datavolumes”. Take the value offered by the global.ini file and remove the database name from the path to only get the base path mount point (this assumes that the SAP HANA systems has been setup with a mount point and single volume each for log and data). We will then use the “fsfreeze” Linux utility to resume IO to the volume and allow the database to continue operation.

Unfreeze the filesystem of the data persistence mount point using the fsfreeze utility

```
:~ # /sbin/fsfreeze --unfreeze /hana/data
```

Step 8. Confirm or abandon the snapshot

Confirm or abandon the snapshot in SAP HANA studio, allowing the backup to be marked as valid or invalid. If something has not operated as expected then the snapshot should be abandoned. To confirm the snapshot an External Backup ID must be supplied, the Backup_ID originally offered by the prepared snapshot and used as a suffix for the Block volume storage snapshotr is used for the value.

Confirm the Snapshot and supply the External Backup ID, and press OK.

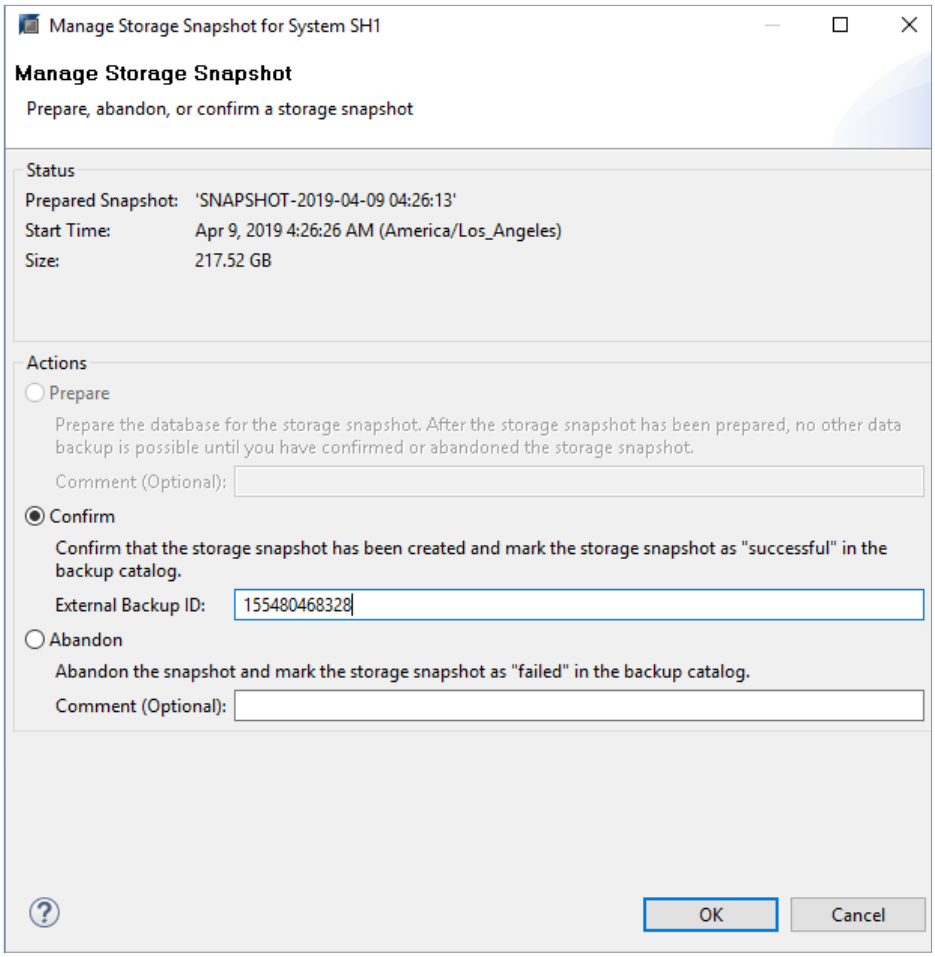

The backup now shows in the backup catalog as complete.

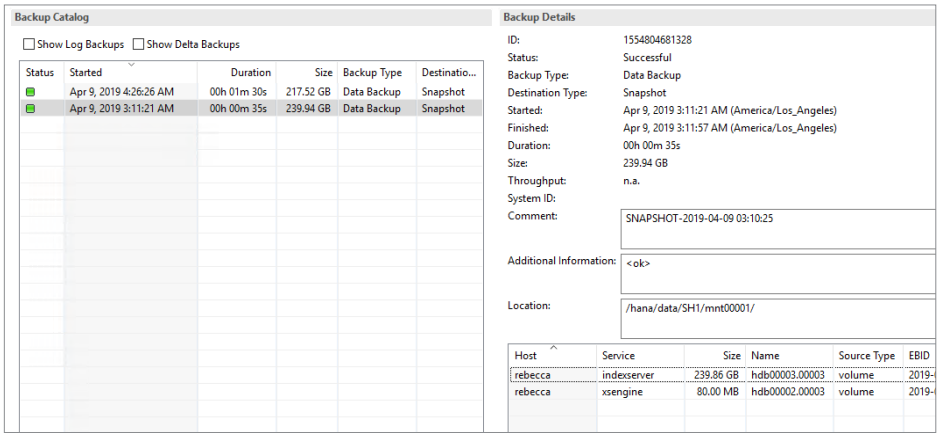

AUTOMATED OPERATION

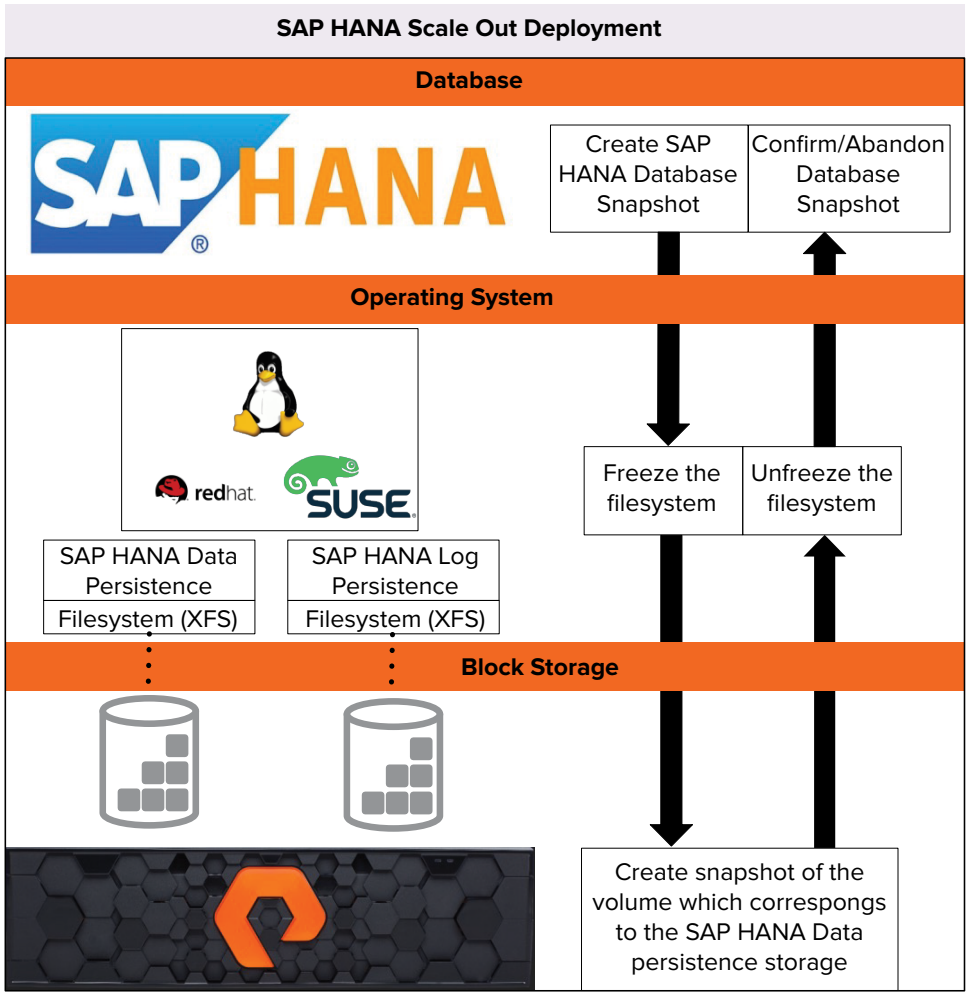

Fig. 9 Workflow to create a storage snapshot for a single host SAP HANA database.

Step 1. Check the SAP HANA System mode

Using the connection string

```
Driver={HDBODBC}; ServerNode=<HostAddress>:3 <InstanceNumber> 15; UID=
<Database User>; PWD=<DatabasePassword>;
```

Connect to the SAP HANA database and run the following query to determine system mode:

```
SELECT VALUE FROM M_INIFILE_CONTENTS WHERE FILE_NAME = 'global.ini' AND
SECTION = 'multidb' AND KEY = 'mode'
```

```
SELECT VALUE FROM M_INIFILE_CONTENTS WHERE FILE_NAME = 'global.ini' AND SECTION = 'multidb' AND KEY = 'mode'
```

	VALUE
1	multidb

The result will be either "singledb" or "multidb", if the result is "multidb" then change the connection string to:

```
Driver={HDBODBC}; ServerNode=<HostAddress>:3 <InstanceNumber> 13; UID=
<Database User>; PWD=<DatabasePassword>;
```

Changing the port number allows the application to connect to the System database. In this case the port to connect on for instance 00 would change from 30015 to 30013. In the event of the system running in "singledb" mode, continue to use the original connection string.

Step 2. Determine the block storage volume on FlashArray that corresponds to the SAP HANA persistent data volume mount point in Linux

Using the established connection string run the following query to determine the SAP HANA persistence data volume mount point:

```
SELECT VALUE FROM M_INIFILE_CONTENTS WHERE FILE_NAME = 'global.ini' AND SECTION = 'persistence' AND KEY =
'basepath_datavolumes' AND VALUE NOT LIKE '$%'
```

```
SELECT VALUE FROM M_INIFILE_CONTENTS WHERE FILE_NAME = 'global.ini' AND SECTION = 'persistence' AND KEY =
'basepath_datavolumes' AND VALUE NOT LIKE '$%'
```

VALUE
/hana/data/SH1

The value returned will include the database name at the end (e.g. SH1 will correspond to /hana/data/SH1). The mount point needed to interact with is the directory above the database name (e.g. /hana/data/SH1 becomes /hana/data/).

An SSH connection needs to be created to query the operating system for the device serial number for the SAP HANA data persistence volume, once the command line is available for reading and writing we run the “df -h” command to view all mounted volumes and mount points as well as retrieving the device mapper or storage device (sd) the mount point is mapped to. The output of df -h then needs to be piped and “grep” used to isolate the specific entry for the required volume. It is possible to query the /etc/fstab for the same information but the contents of fstab may not always be what the system is working on at that point in time.

```
df -h | grep <SAP HANA Data Persistence mount point>
```

```
Rebecca:~ # df -h | grep /hana/data
/dev/mapper/3624a93701441efcb40254a2b00011418-part1  1.0T  470G  554G  46% /hana/data
```

Once the device the mountpoint corresponds to has been isolated, udevadm is then queried for the serial number known as “DM_SERIAL” in its output using the command:

```
udevadm info --query=all --name<device name> | grep DM_SERIAL
```

```
Rebecca:~ # udevadm info --query=all --name=/dev/mapper/3624a93701441efcb40254a2b00011418-part1 | grep DM_SERIAL
E: DM_SERIAL=3624a93701441efcb40254a2b00011418
```

Using the serial number returned for the device, it is possible to match it up to the block storage volume a FlashArray. Note that the block volume serial number will be all of the characters after “3624a9370”

Step 3. Prepare the database snapshot and retrieve the backup ID for it

Using the established connection string to execute the query needed to prepare a database snapshot.

```
BACKUP DATA FOR FULL SYSTEM CREATE SNAPSHOT COMMENT ‘SNAPSHOT-<Snapshot Time>
```

To retrieve the backup ID, execute the following query:

```
SELECT BACKUP_ID, COMMENT FROM M_BACKUP_CATALOG WHERE
ENTRY_TYPE_NAME = ‘data snapshot’ AND STATE_NAME = ‘prepared’
```

	BACKUP_ID	COMMENT
1	1,554,804,681,328	SNAPSHOT-2019-04-09 03:10:25

Step 4. Freeze the filesystem

An SSH connection needs to be created with the operating system on which the SAP HANA instance is installed to execute command line arguments. To freeze the filesystem the fsfreeze utility will be used as it supports EXT3/4, ReiserFS, JFS and XFS. The mount point retrieved in step 2 will be used during the freeze operation.

```
/sbin/fsfreeze –freeze <path to mount point> /sbin/fsfreeze –freeze <path to mount point>
```

```
:~ # /sbin/fsfreeze --freeze /hana/data
```

Step 5. Query the relevant FlashArray for a list of its volumes and search them for the serial number contained with the serial number returned by step 2, then create a snapshot once the volume has been located

```
$Array = New-PfaArray -EndPoint $FlashArrayAddress -username $User -Password
$Password -IgnoreCertificateError
$Volumes = Get-PfaVolumes -Array $Array
```

Example of PowerShell using the Pure Storage PowerShell SDK

```
$VolumeSnapshot = New-PfaVolumeSnapshots -Array $Array -Sources $volume.name -Suffix $SnapshotSuffix
```

Once the correct volume has been found create a snapshot with a specified Snapshot suffix

```
return $VolumeSnapshot.serial
```

Return the serial number for the volume to be used as a reference when confirming or abandoning the snapshot

Step 6. Unfreeze the filesystem

An SSH connection needs to be created with the operation system on which the SAP HANA instance is installed to execute command line arguments. To unfreeze the filesystem, use the same mount point used in step 3.

```
/sbin/fsfreeze --unfreeze <path to mount point>
```

```
:~ # /sbin/fsfreeze --unfreeze /hana/data
```

Step 7. Confirm or abandon the database snapshot

Using the connection string established in step 1, confirm or abandon the snapshot using hdbsql commands:

```
BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID <External Backup ID> SUCCESSFUL;
```

Confirm the snapshot if all of the previous steps executed successfully.

```
BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID <Prepared database snapshot backup ID>
UNSUCCESSFUL <additional comments>;
```

Abandon the snapshot if one of the previous steps did not execute successfully.

RECOVERY PROCESS

The recovery process for a storage snapshot will restore the SAP HANA instance and all of the data to the point in time at which the storage snapshot was taken. It is important to note that the recovery of a system using storage snapshots with multiple database containers (MDC) is only supported from SAP HANA 2.0 SPS04.

SAP HANA Studio Step 1. Prepare to Recover System Database (MDC systems only)

Navigate to the Systems inventory and right-click on the SystemDB connection for the relevant database and navigate to the Backup and Recovery sub menu then select “Recover System Database…”. A prompt to shut down the relevant system will be shown as recovery can only be done when it is offline.

Navigate to the systems inventory , right click the SystemDB and navigate to Backup and Recovery > Recover System Database…

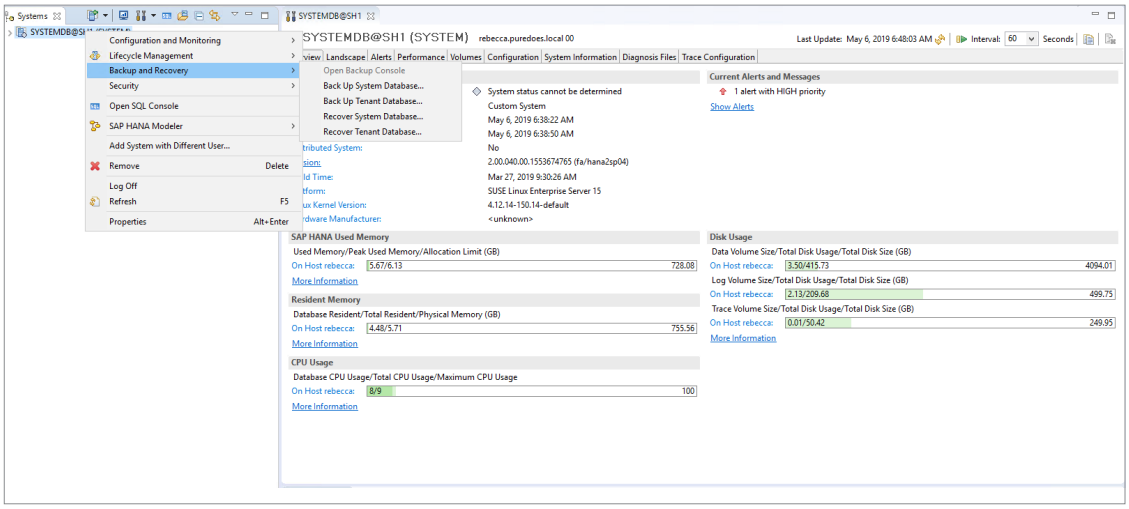

Enter the relevent values for the <sid>adm user created during initial SAP HANA installation.

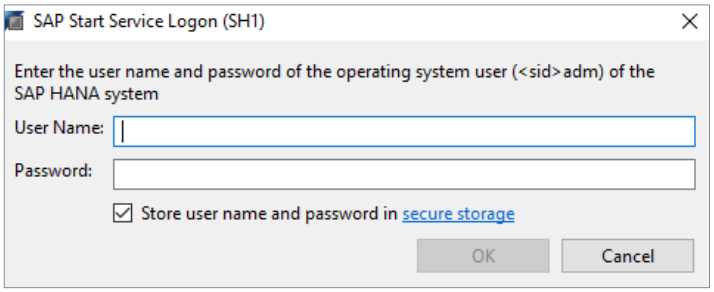

SAP HANA studio will display a prompt to shut down the system.

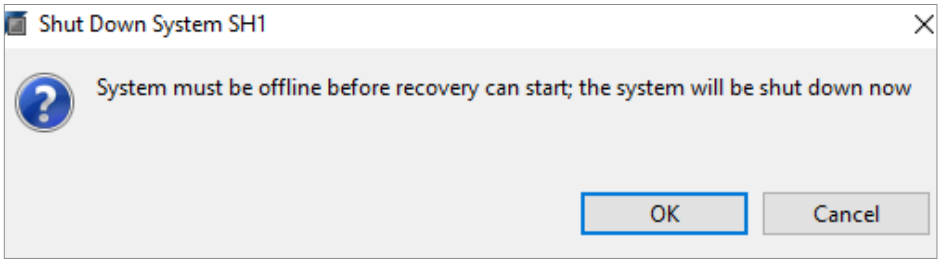

The graphical process gives the user an option of choosing a recovery type.

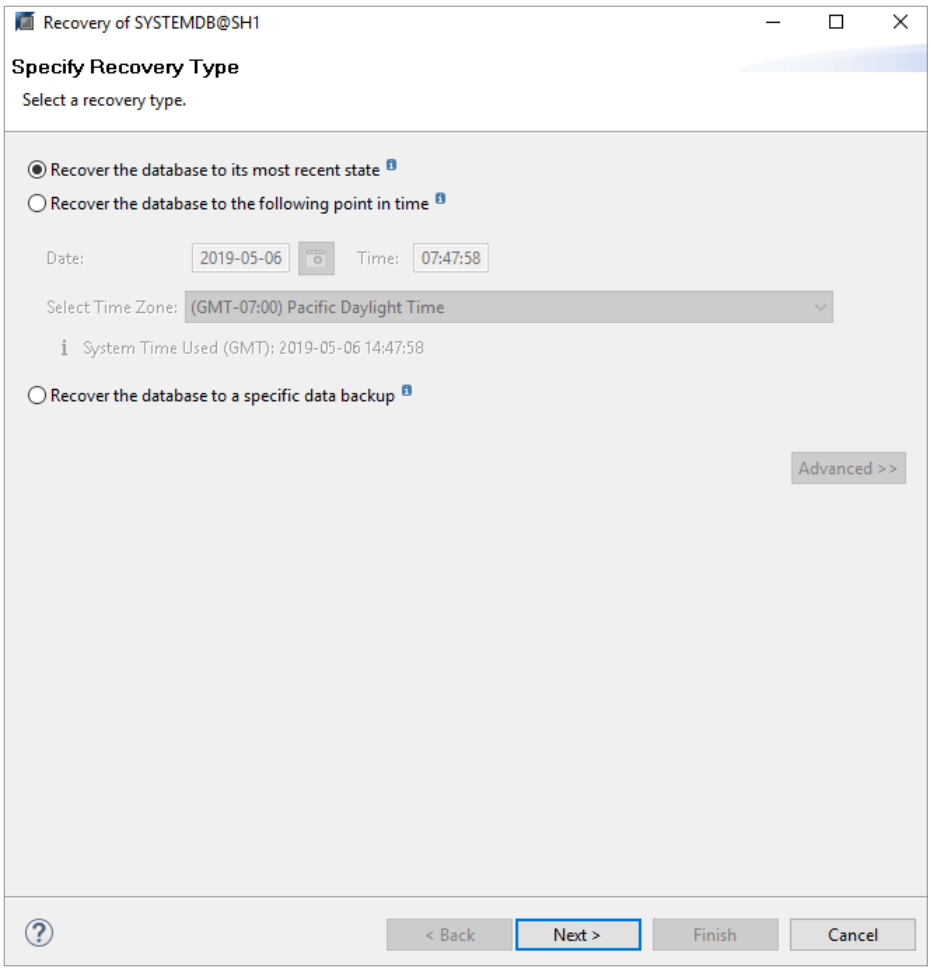

The location of the backup catalog is normally populated, only change this under specialized circumstances where the recovery catalog is in a different location.

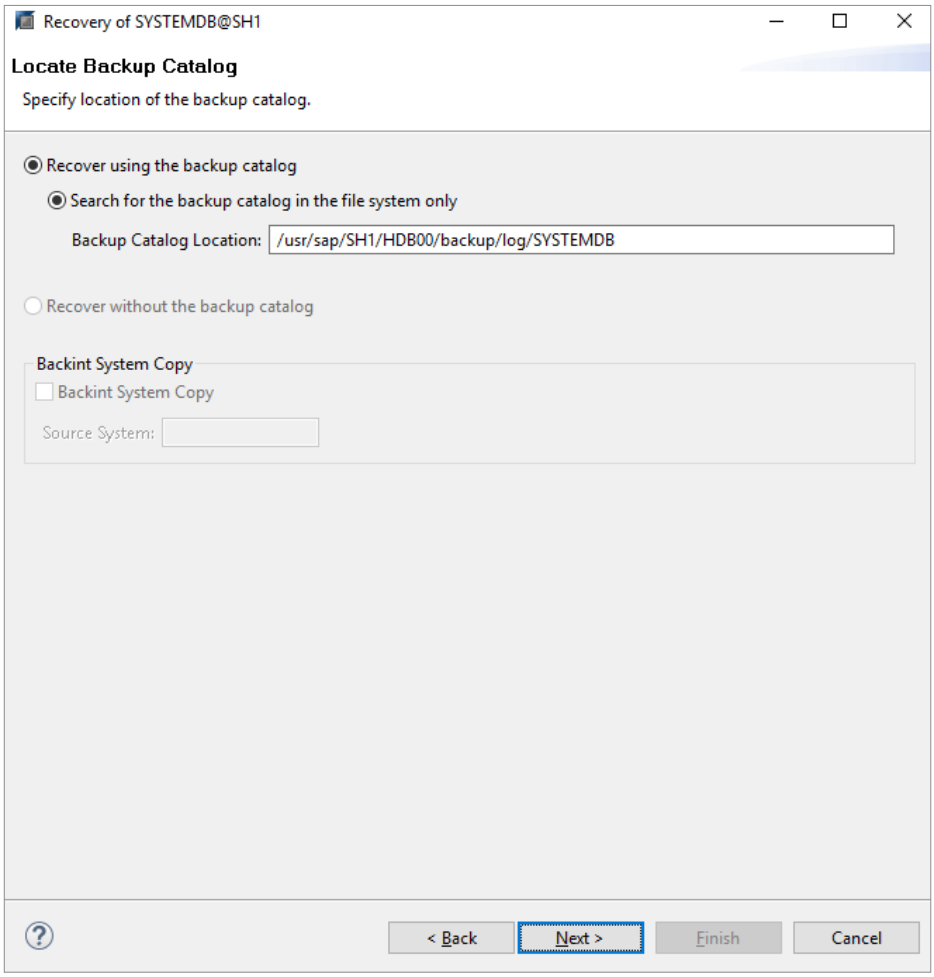

All records in the catalog will be shown with the relevant information, take note of the Backup ID for the required recovery point.

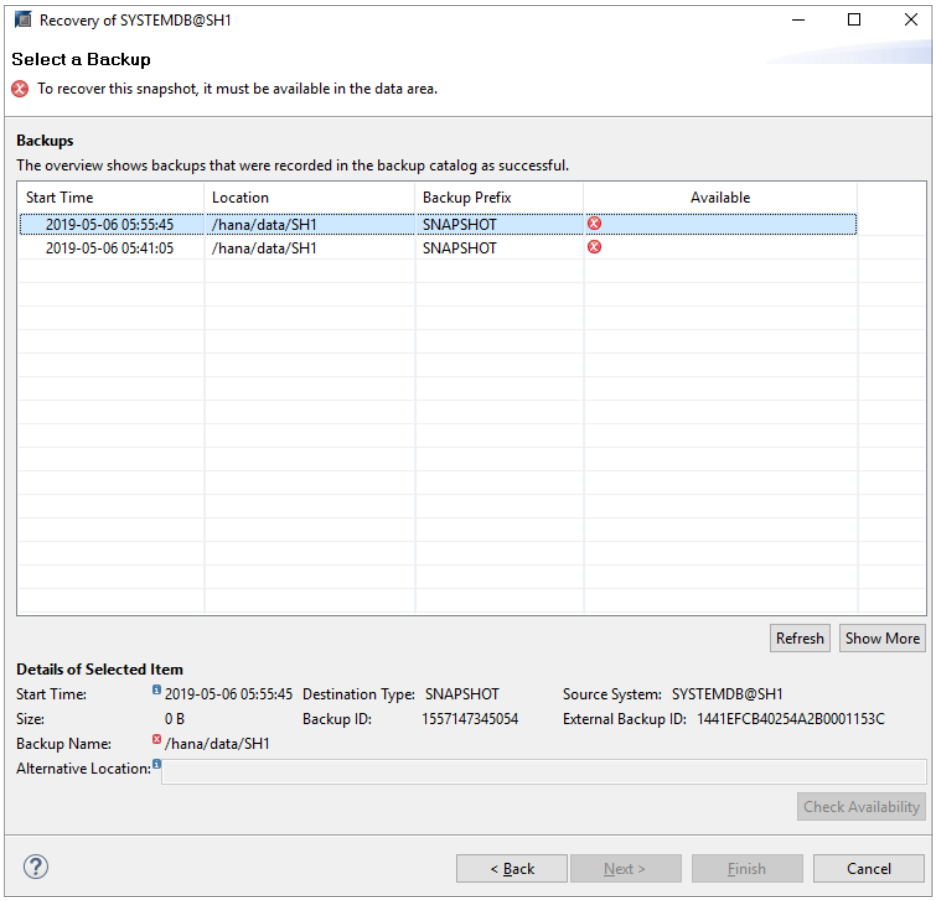

SAP HANA Studio Step 1a. Recover block storage volume

Recover the relevant storage volume using the FlashArray web user interface. This operation can also be completed using the command line interface or ReST API.

Important: The persistence data volume must be unmounted from the operating system before a snapshot is restored. This can be done using the "umount" command in a terminal or SSH connection.

SSH command to unmount the volume at the specified location.

```
:~ # umount /hana/data
```

Navigate to the FlashArray web user interface and select "Storage" from the sidebar and then navigate to the "Volumes" section.

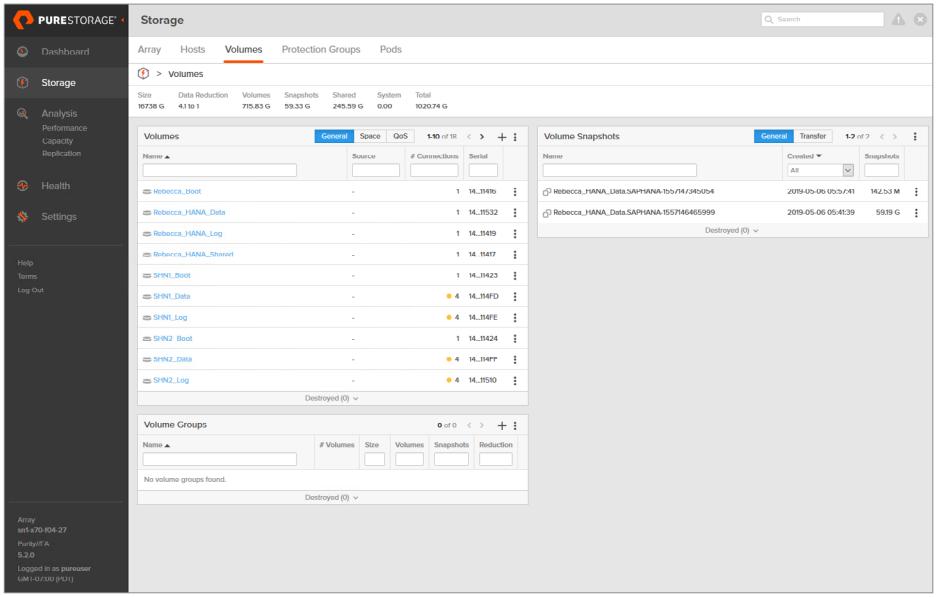

Navigate to the block volume on which the relevant hosts SAP HANA data volume is located.

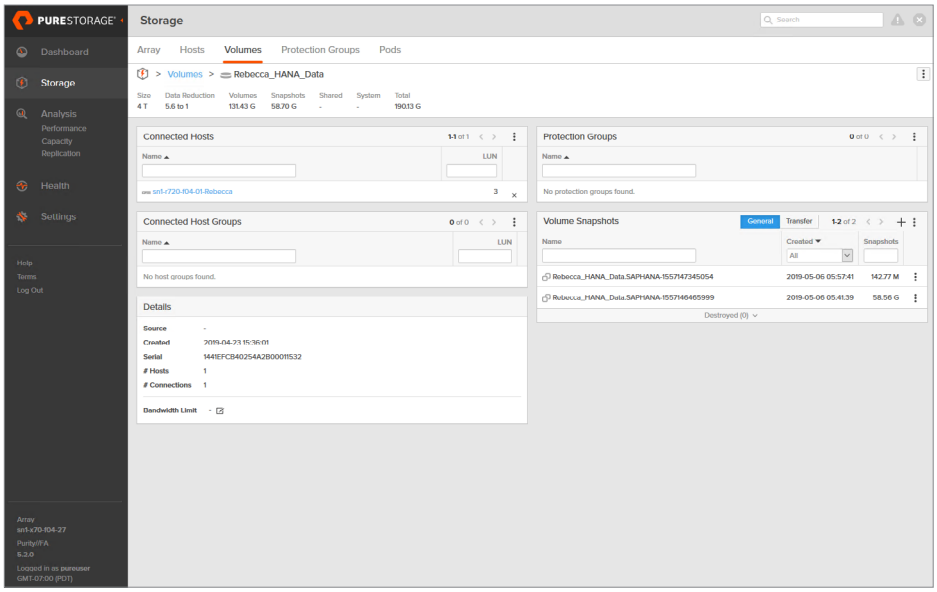

Important: The snapshot which matches the "Backup ID" in the catalog is appended with the exact same value, this is the volume which should be restored. Note that the addition of this value is done during storage snapshot creation by the user.

Select the 3 vertical dots to bring up the context menu and select Restore...

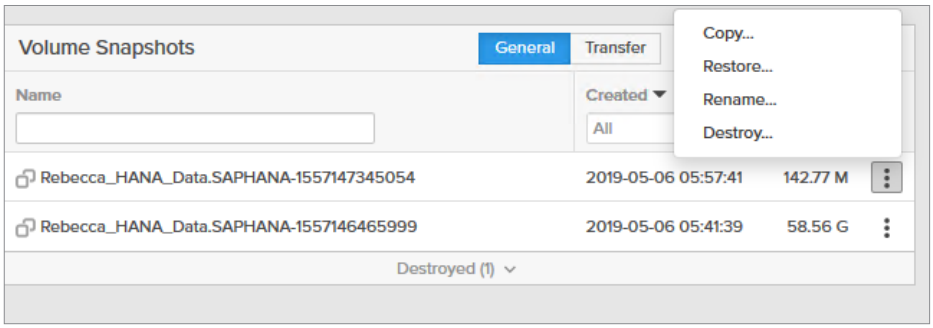

Check the snapshot to be restored matches the requirements for this recovery scenario and then confirm the operation by selecting "Restore".

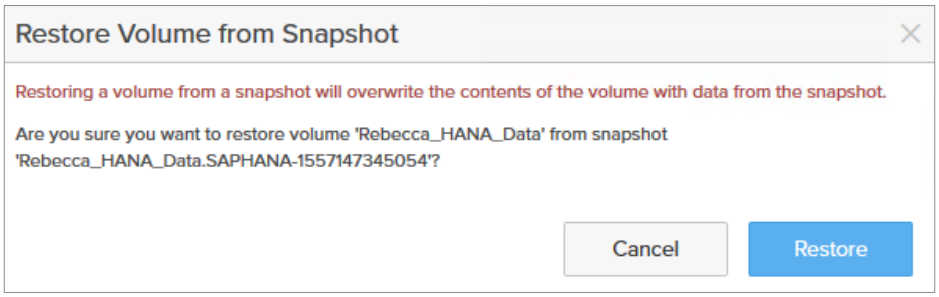

Return to the SSH connection and remount the volume to the SAP HANA data location. This assumes that the correct values in /etc/fstab are present.

```
:~ # mount /hana/data/
```

SAP HANA Studio Step 1b. Recover the System Database

Select “Refresh” from the Recovery window in SAP HANA studio and the relevant recovery point should be shown as available, then select “Next“ to proceed.

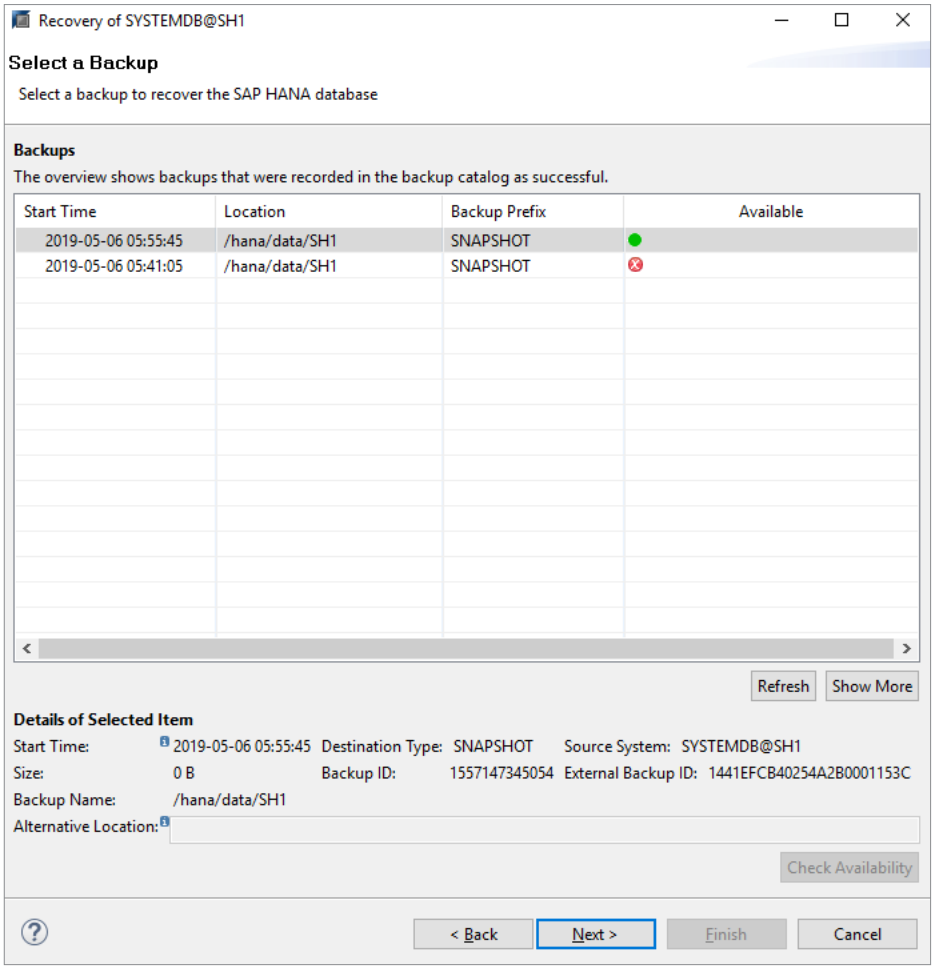

Review the relevant settings for "Initialize Log Area", "Use Delta Backups" and "Install New License Key". In some scenarios these selections will be greyed out depending on the recovery type.

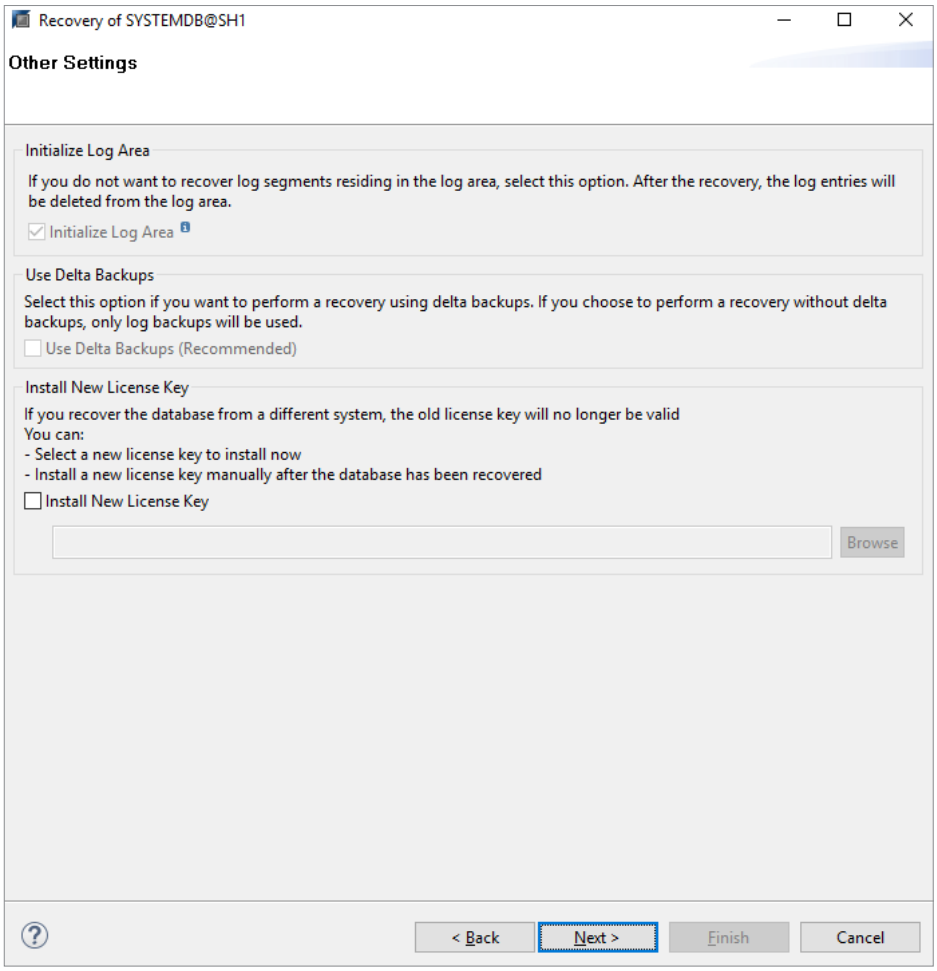

Review all of the recovery settings and proceed by selecting "finish".

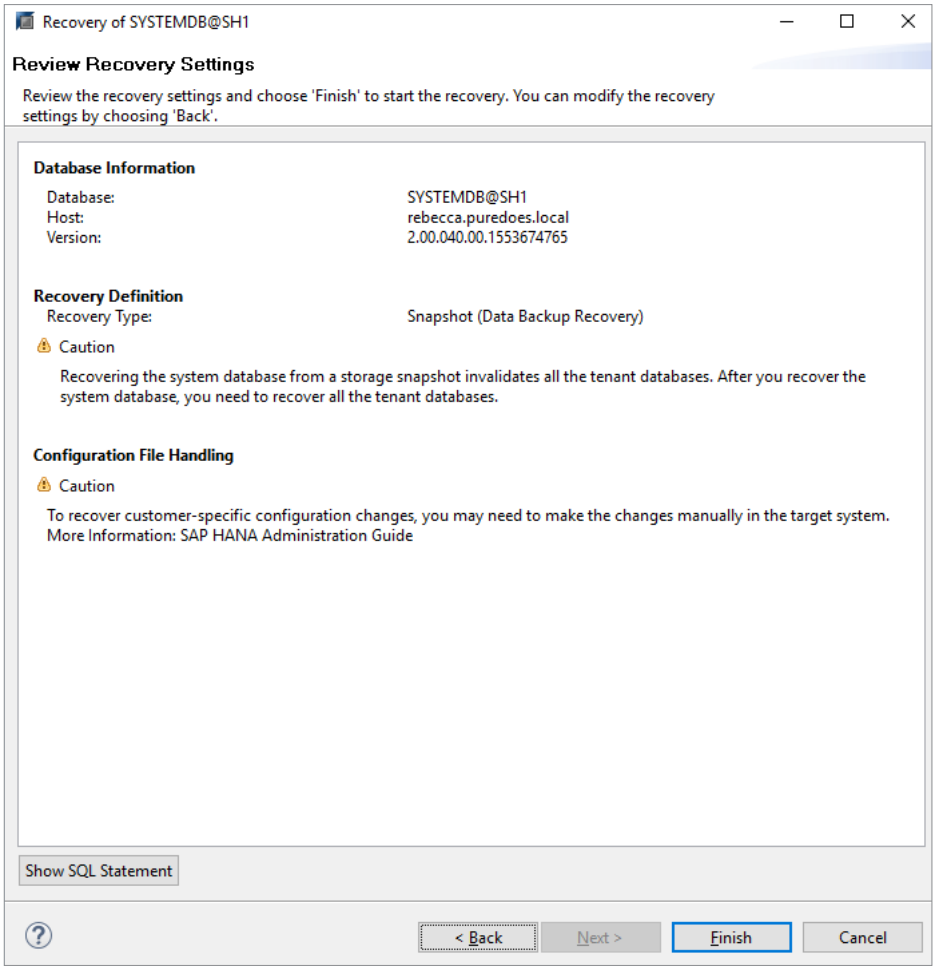

SAP HANA Studio Step 2. Recover Tenant Database

During the recovery of the System database all tenant databases are invalidated and need to be individually recovered by repeating the below steps.

Important: Follow from here if using single container SAP HANA 1.0 system. Notable differences are the absence of “System” and “Tenant” terminology.

Important: In SAP HANA 2.0 SPS04 and onwards repeat these steps for each additional tenant database.

Important: Ensure that Step 1a is completed, while the database is offline, for SAP HANA 1.0 Single tenant systems before proceeding with the below.

Navigate to the main database view in SAP HANA Studio, right click on the relevant database connection, navigate to “Backup and Recovery” and then Select “Recover Tenant Database…”. When using a single container SAP HANA 1.0 system this will simple state “Recover database…”

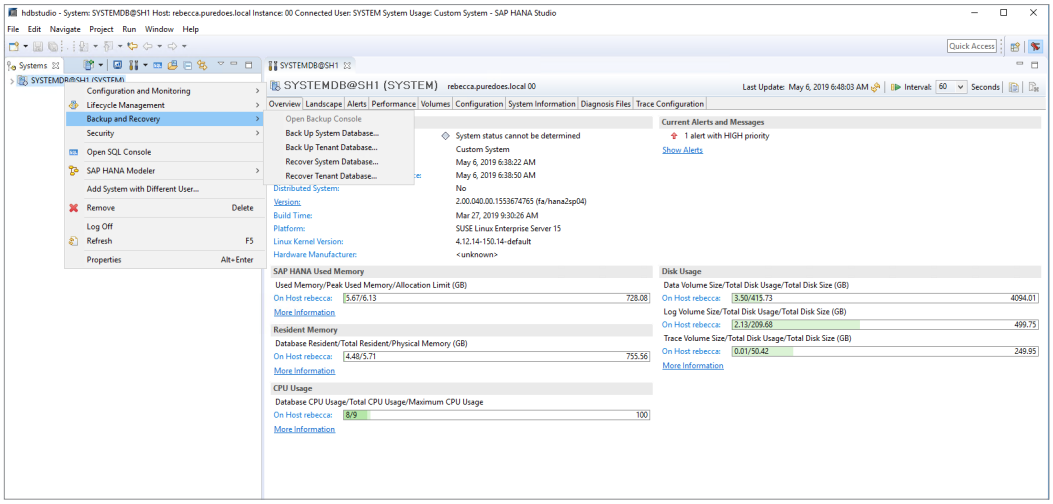

(MDC Only) Select the tenant database to restore. In SAP HANA 2.0 SPS03 and previous releases only a single tenant system can be restored.

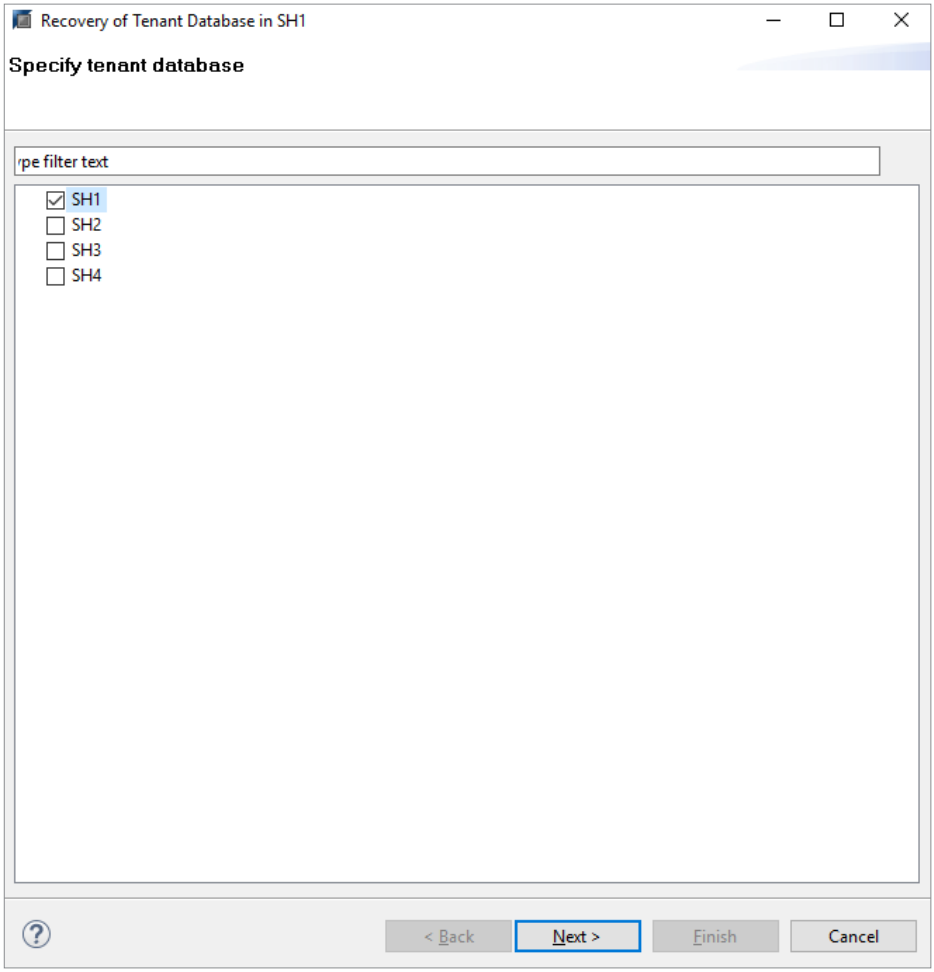

The location of the backup catalog is normally populated, only change this under specialized circumstances where the recovery catalog is in a different location.

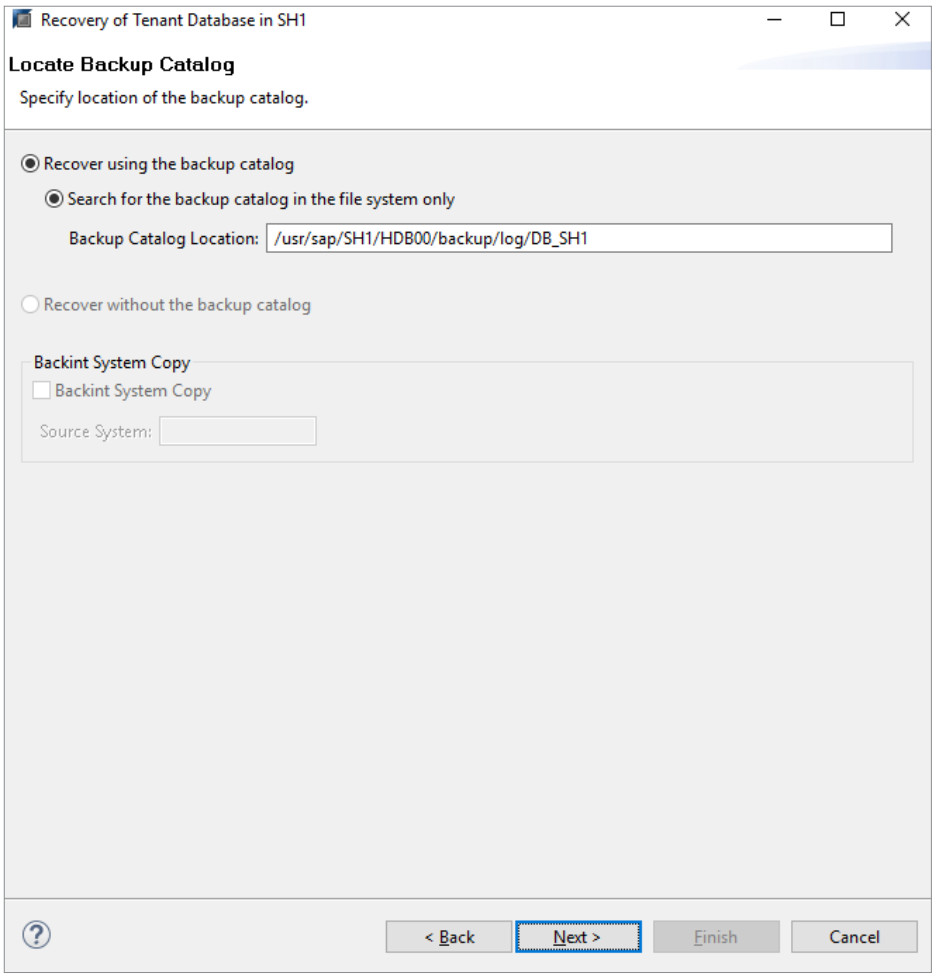

Accept the prompt to shut down the database.

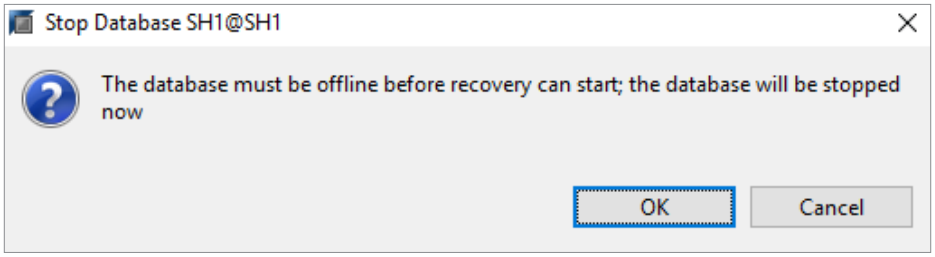

The relevant recovery point should be shown as available, then select “Next “to proceed.

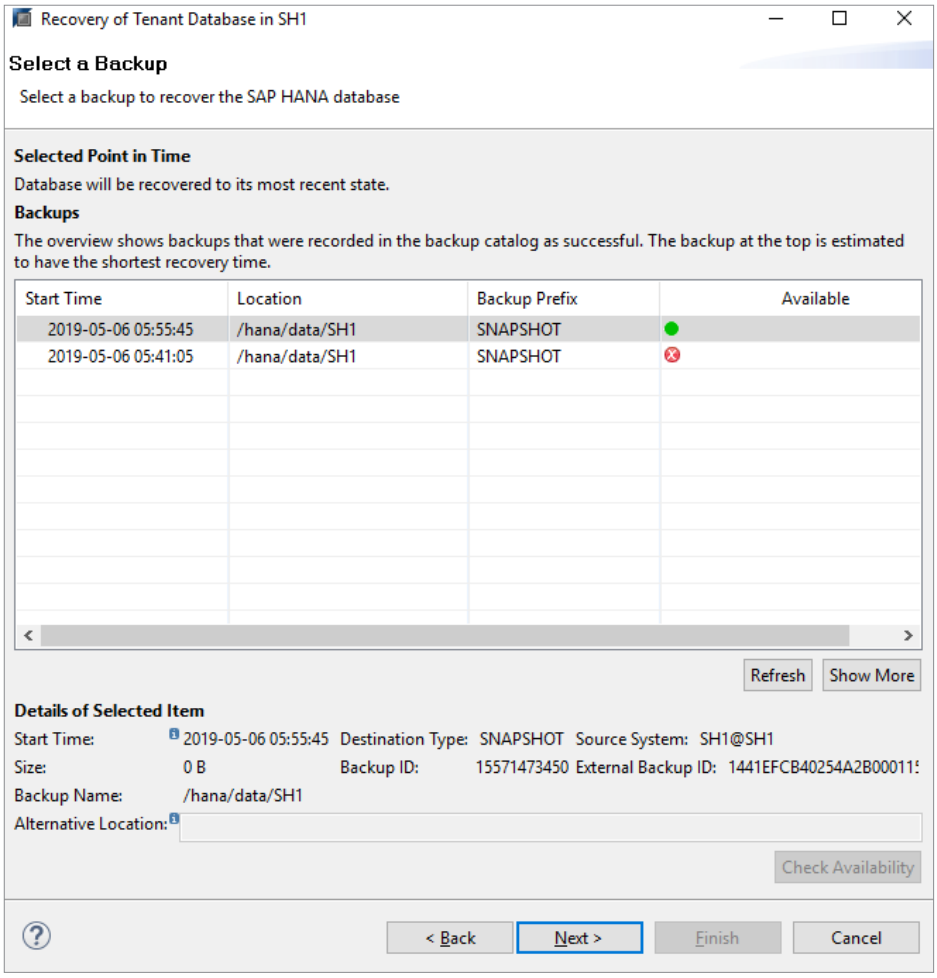

The log backup location is typically populated with the system default values, only change these values in specialised circumstances.

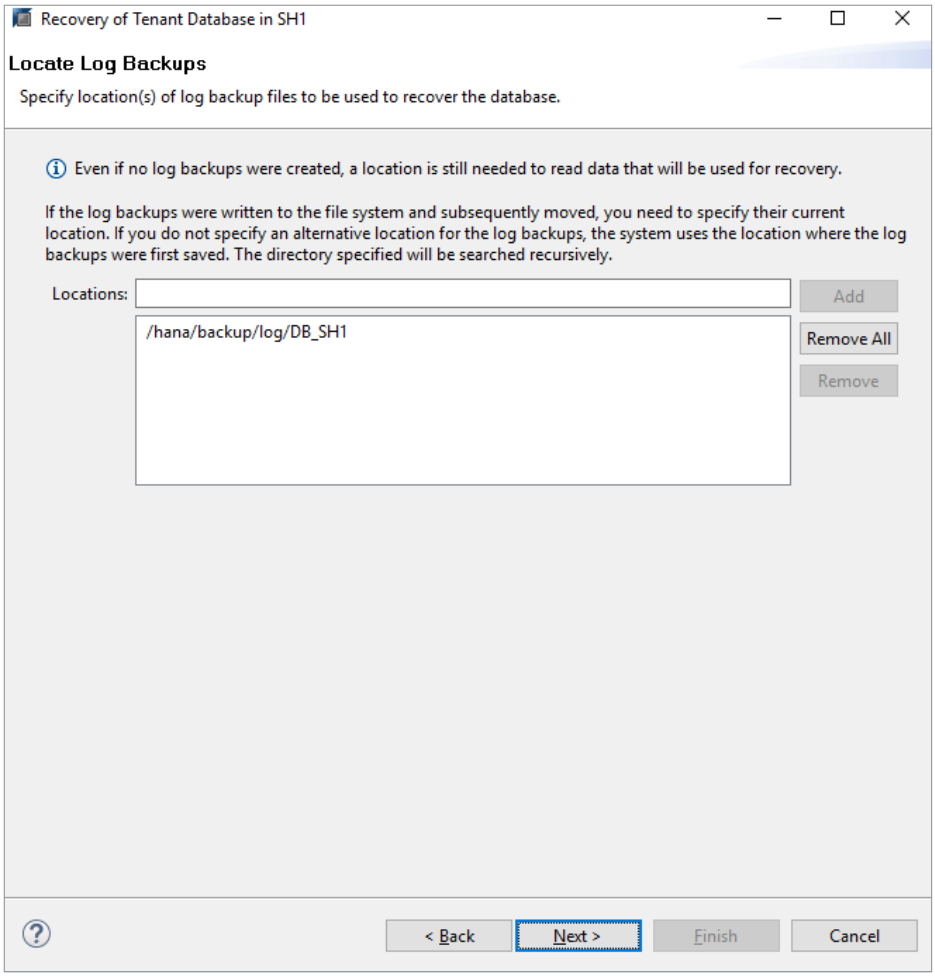

Review the relevant settings for "Check the availibility of delta and log backups", "Initialize Log Area" , "Use Delta Backups" and "Install New License Key". In some scenarios these selections will be greyed out depending on the recovery type.

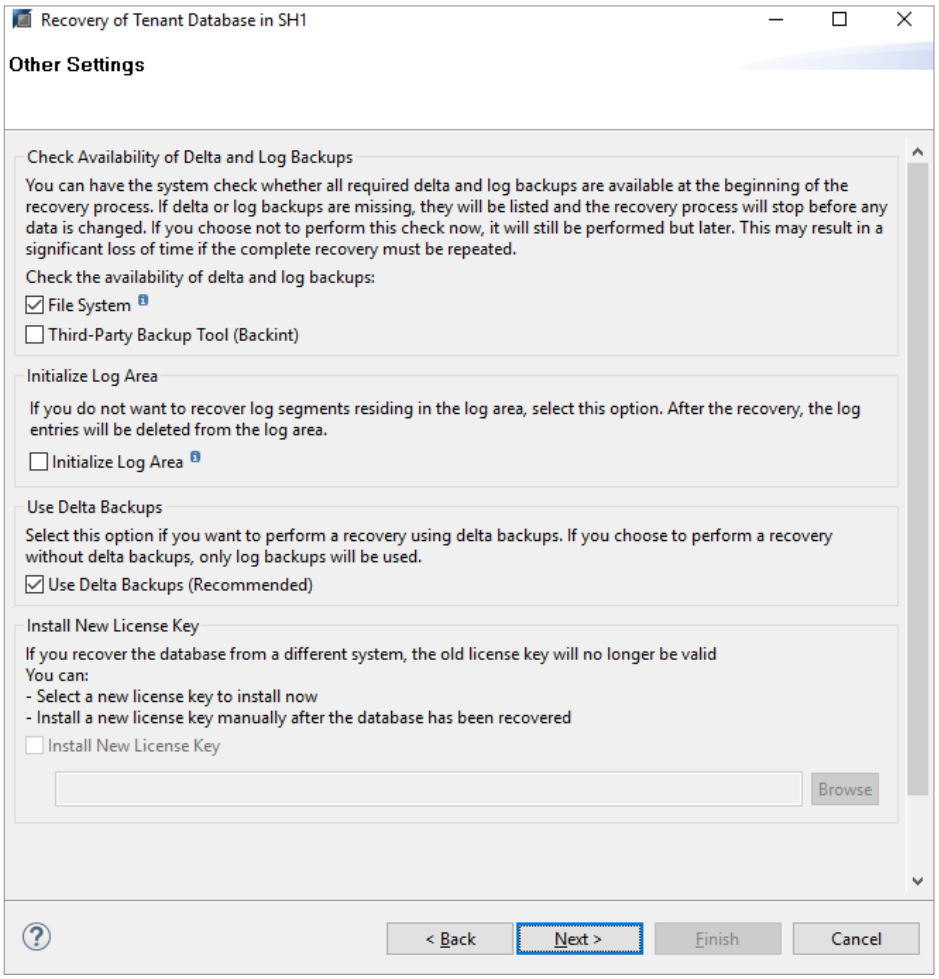

Review all of the recovery settings and proceed by selecting "finish".

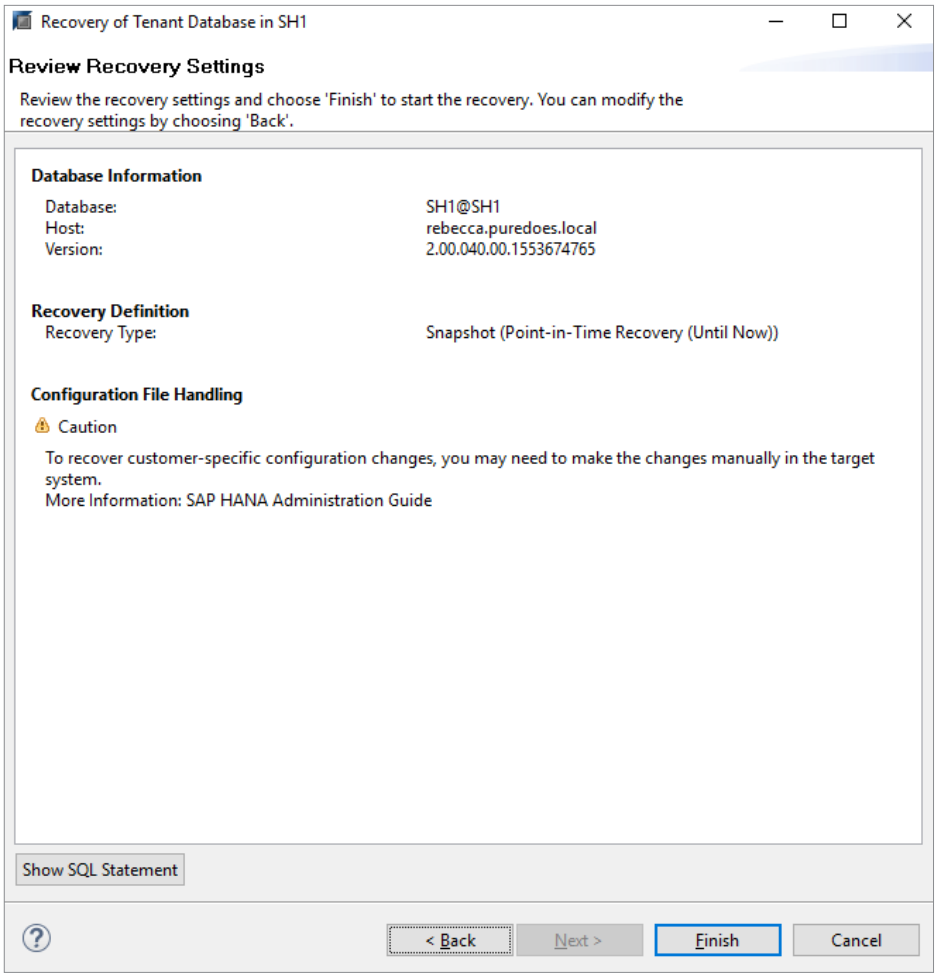

When recovery is completed the status will be shown.

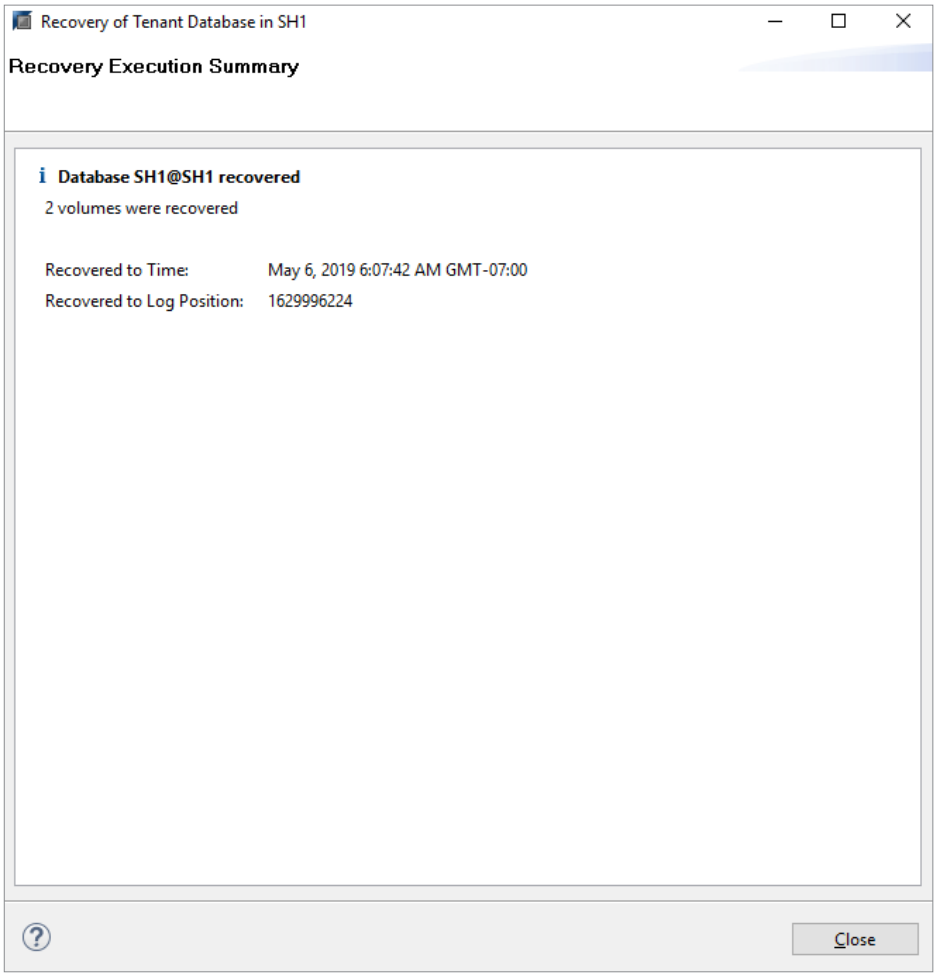

SAP HANA Cockpit Step 1. Prepare to Recover System Database (MDC systems only)

Navigate to the SAP HANA Cockpit web interface and ensure that all of the relevant information is available for it to interface with the relevant deployment to be restored. SAP HANA Systems can be started and stopped from the cockpit interface, where the database will need to be stopped before a recovery can proceed.

Ensure SAP Control credentials are entered and that the resource for HANA cockpit is connected to the SystemDB. Then select the relevant resource for restore operations.

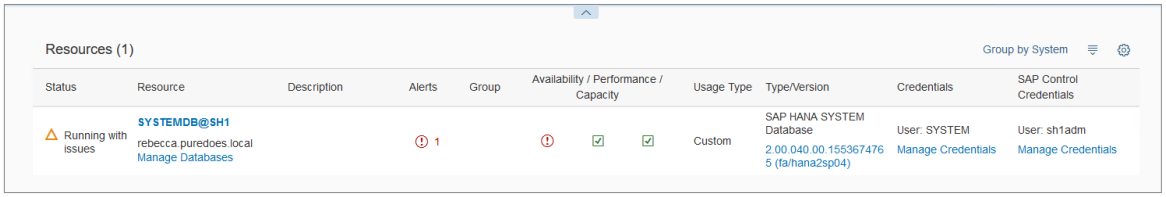

From SAP HANA cockpit choose the overall database status tile and select "stop system".

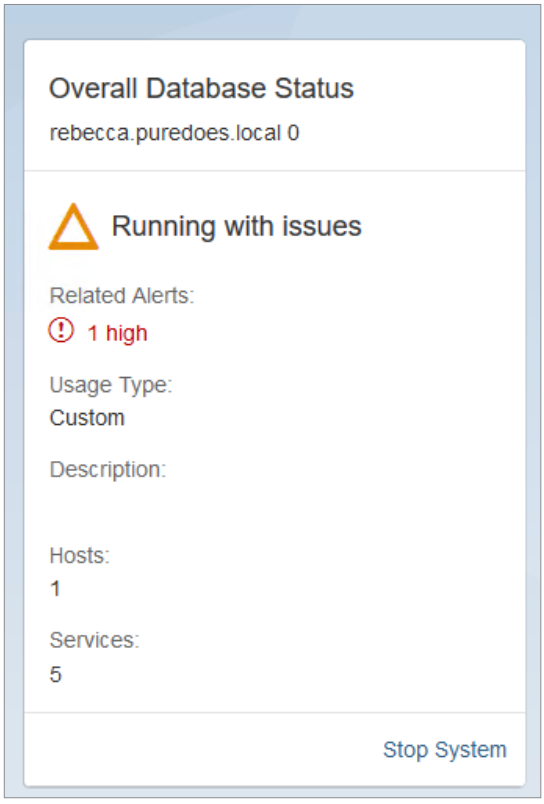

Stop the system using either “softly” or “immediately”.

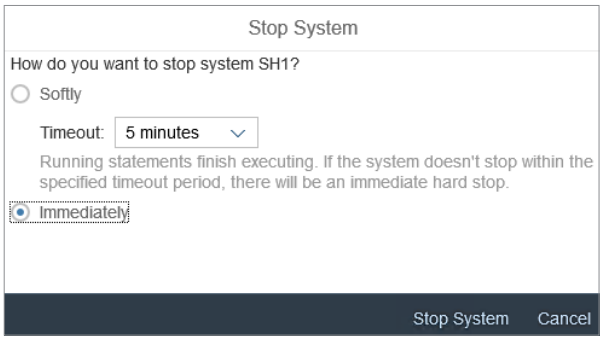

Monitor the system processes until all have been shut down.

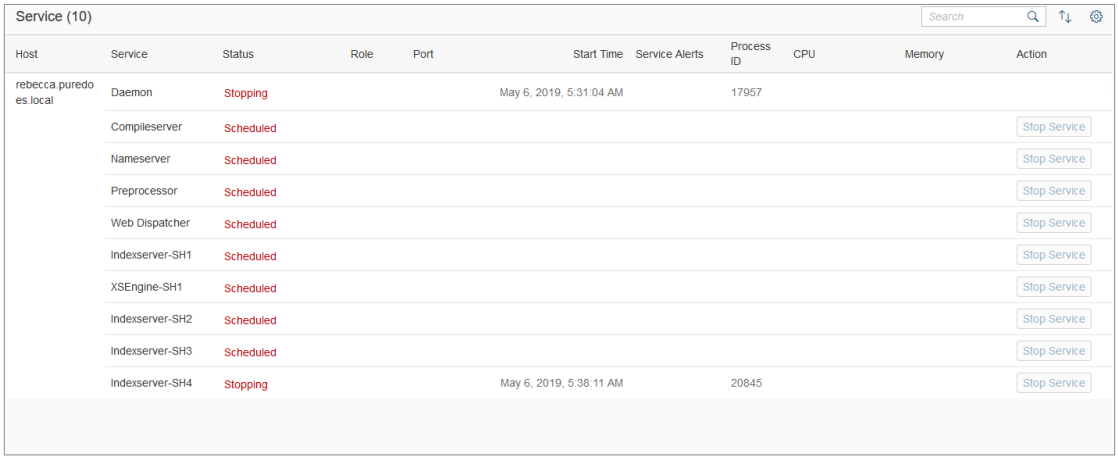

Once all of the system has shut down, select the relevant resource.

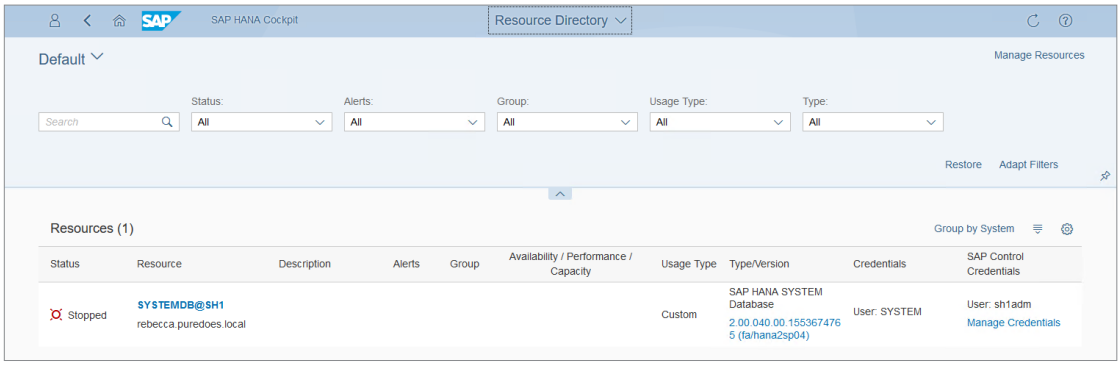

In the "Database Administration" tile, select "Recover database".

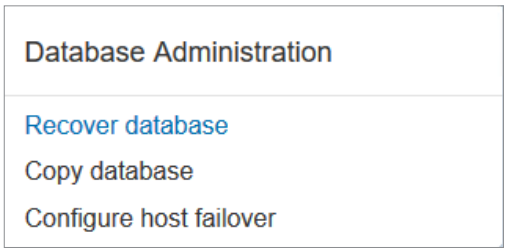

Select the recovery target type.

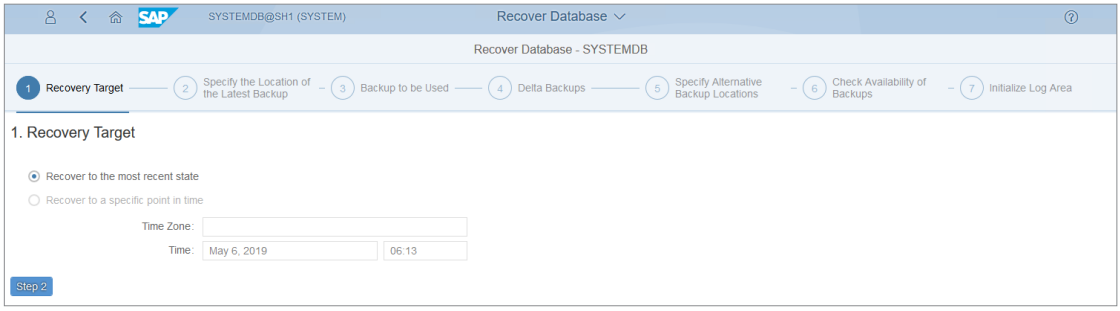

Specify the location of the latest backup catalog, only change these values under specialised circumstances.

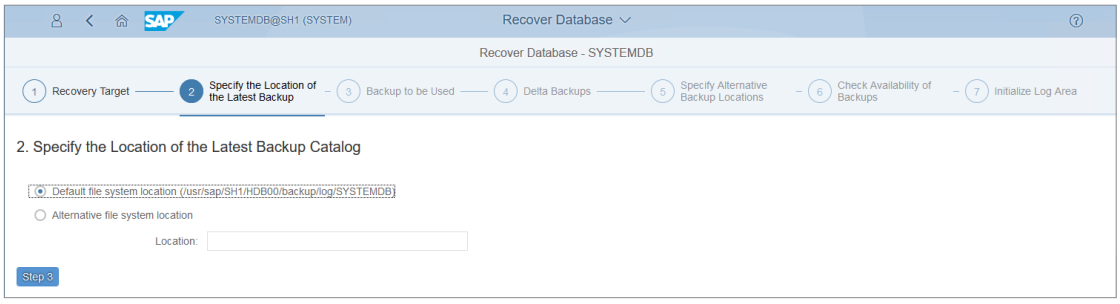

Note the "Backup ID" of the backup to be used and its availability.

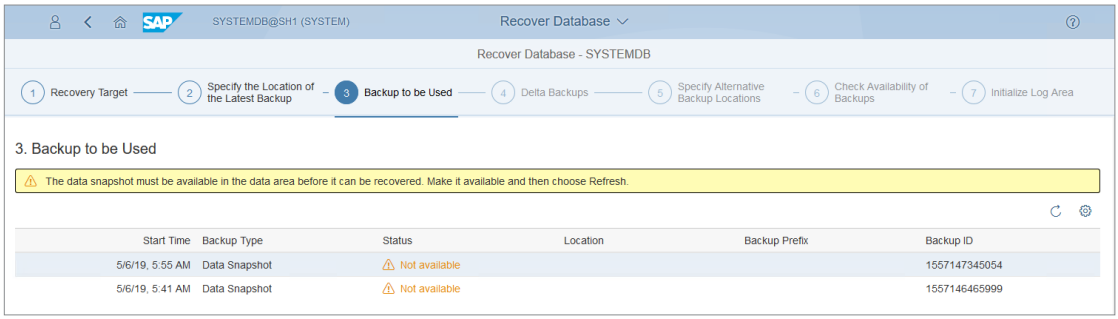

SAP HANA Cockpit Step 1a. Recover block storage volume

Recover the relevant storage volume using the FlashArray web user interface. This operation can also be completed using the command line interface or ReST API.

Important: The persistence data volume must be unmounted from the operating system before a snapshot is restored. This can be done using the "umount" command in a terminal or SSH connection.

SSH command to unmount the volume at the specified location.

```
:~ # umount /hana/data
```

Navigate to the FlashArray web user interface and select "Storage" from the sidebar and then navigate to the "Volumes" section.

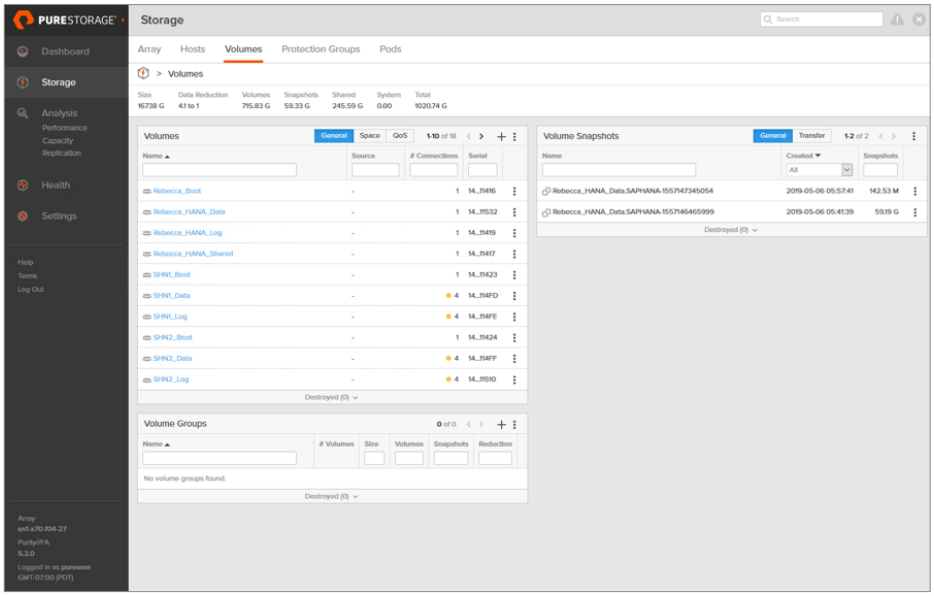

Navigate to the block volume on which the relevant hosts HANA data volume is located.

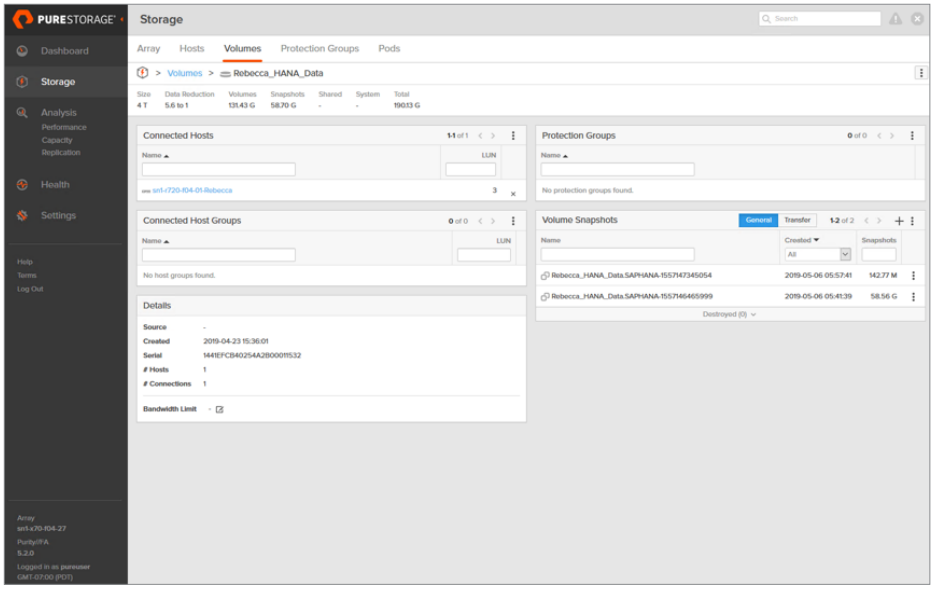

Important: The snapshot which matches the “Backup ID” in the catalog is appended with the exact same value, this is the volume which should be restored. Note that the addition of this value is done during storage snapshot creation by the user.

Select the 3 vertical dots to bring up the context menu and select Restore…

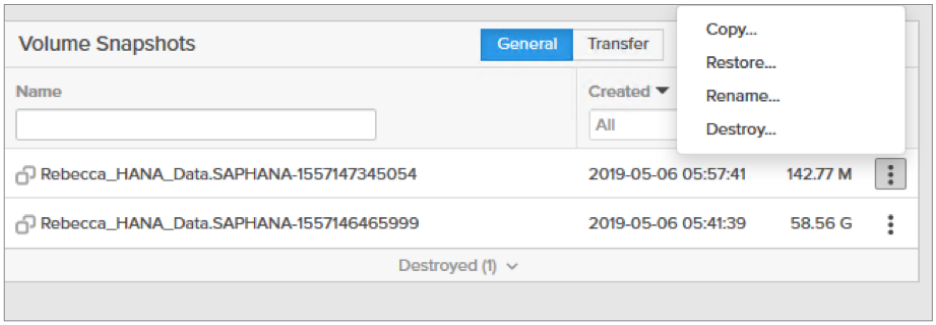

Check the snapshot to be restored matches the requirements for this recovery scenario and then confirm the operation by selecting "Restore".

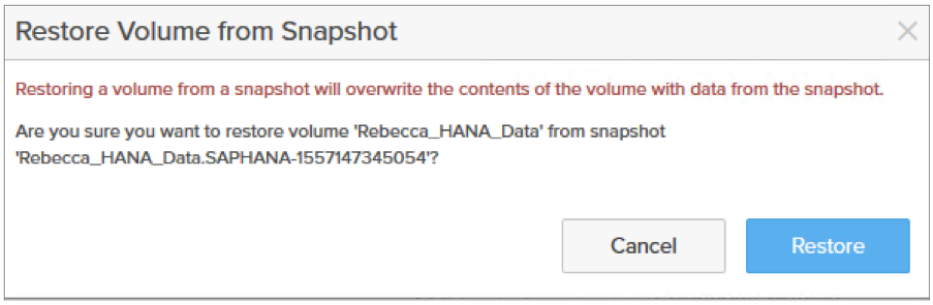

Return to the SSH connection and remount the volume to the HANA data location. This assumes that the correct values in /etc/fstab are present.

```
:~ # mount /hana/data/
```

SAP HANA Cockpit Step 1b. Recover the System Database

Select the refresh icon in the top right-hand corner under Backups to be Used to rescan the backup catalog.

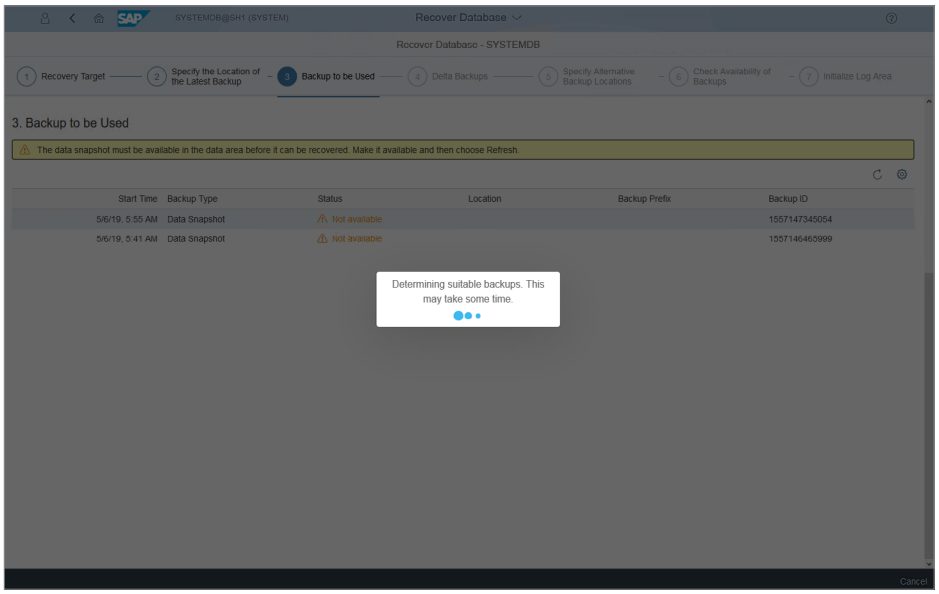

Once the relevant recovery point has been located it should show as "Available", select "Step 4" to proceed.

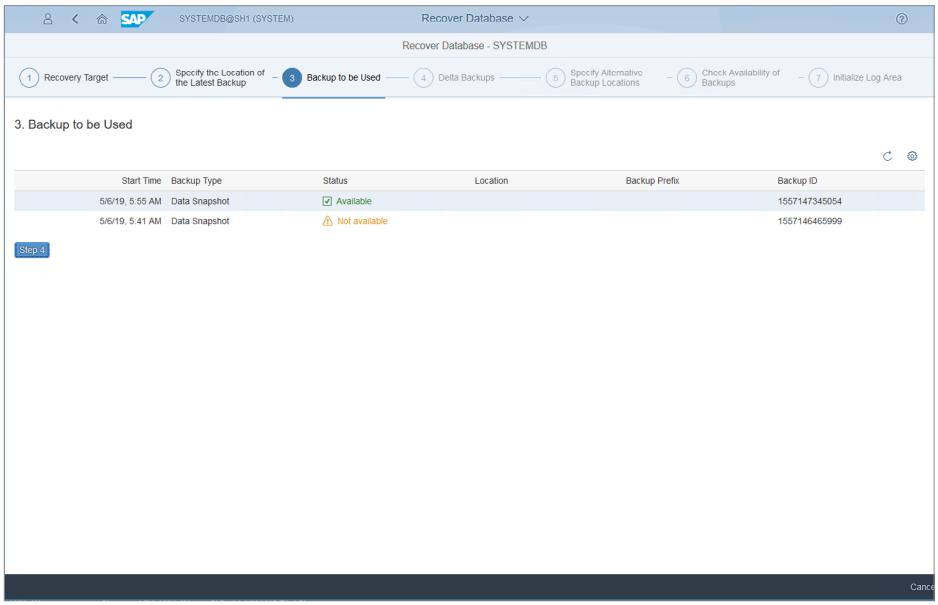

Select the preferred value for Delta Backups.

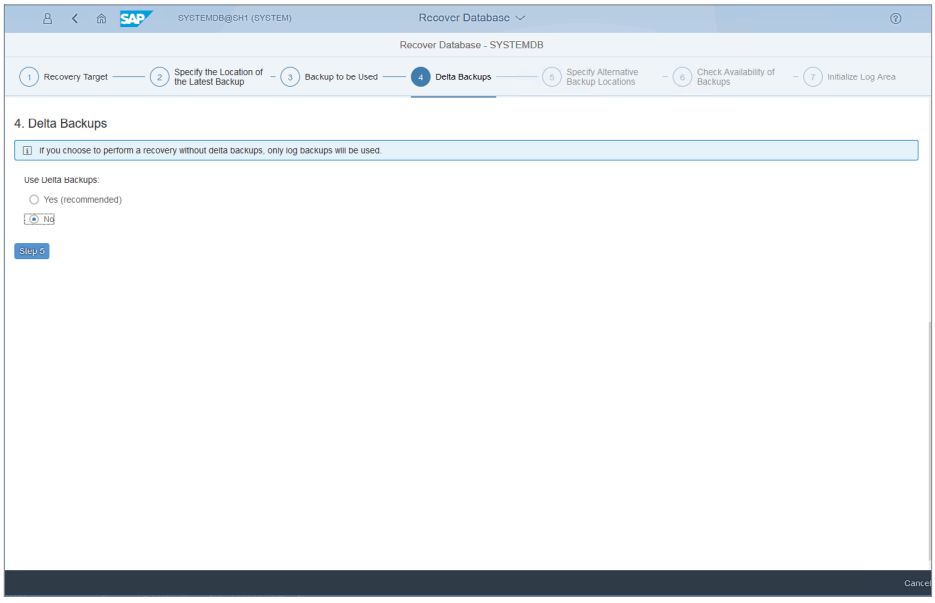

Specify any alternative backup locations.

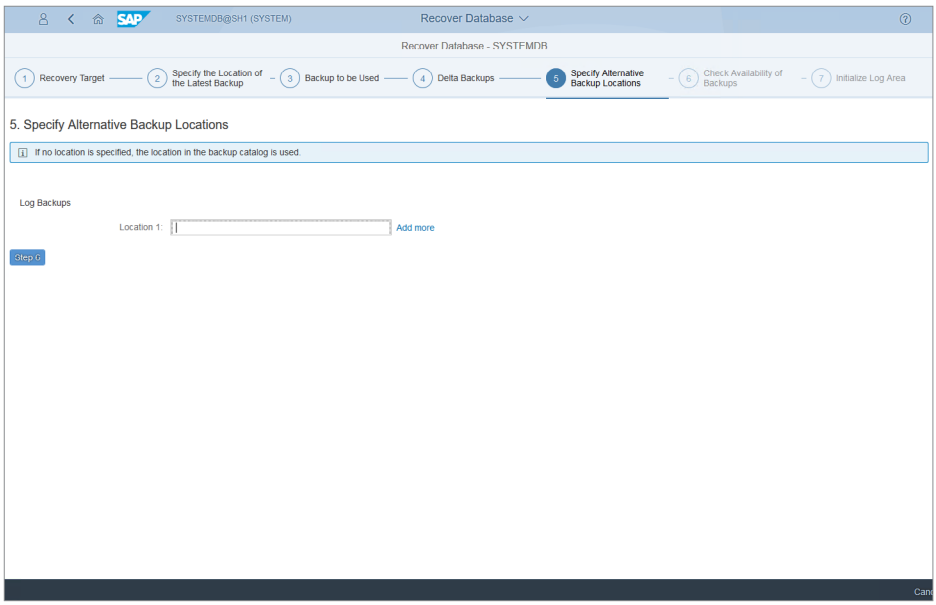

Specify if the availability of backups should be checked.

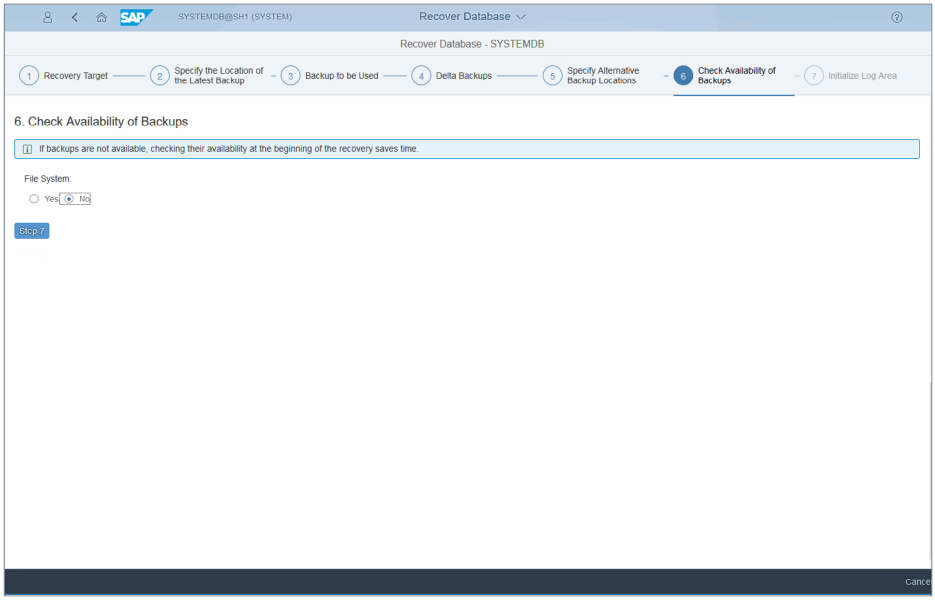

Specify if the log area should be initialized, doing so invalidates any logs or log backups made after the recovery point.

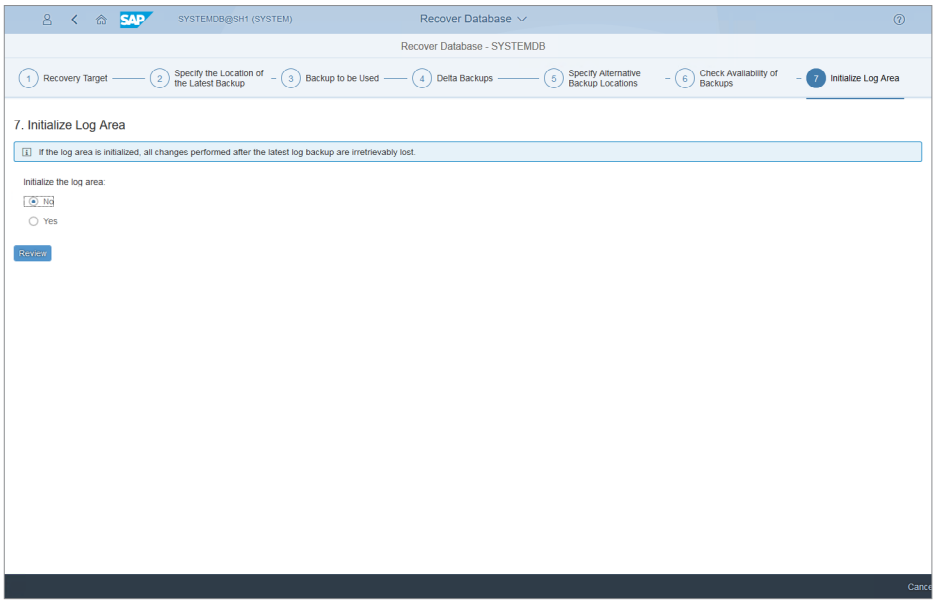

Review the information for the recovery operation and then select "Start Recovery".

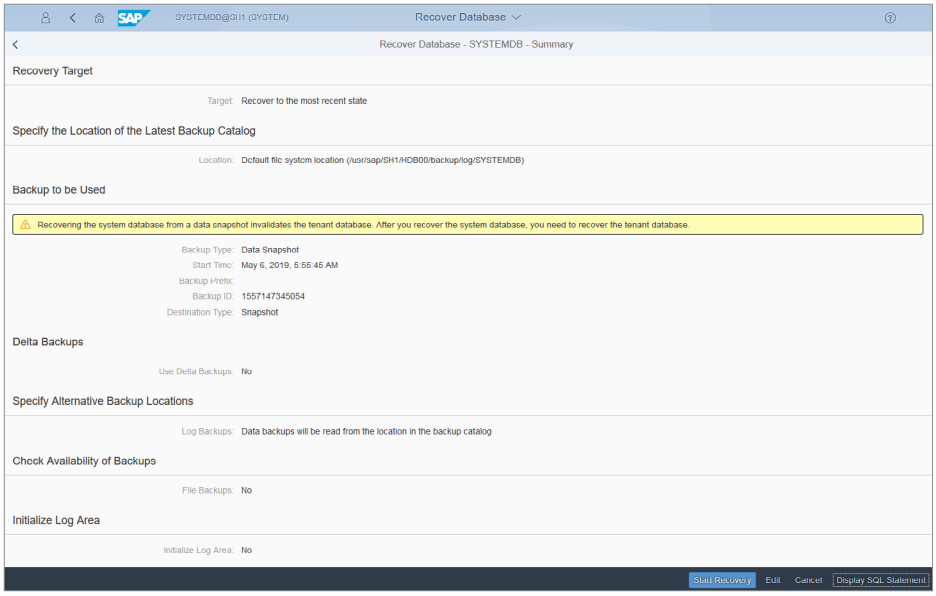

While restoring the System Database the status can be observed from the same view.

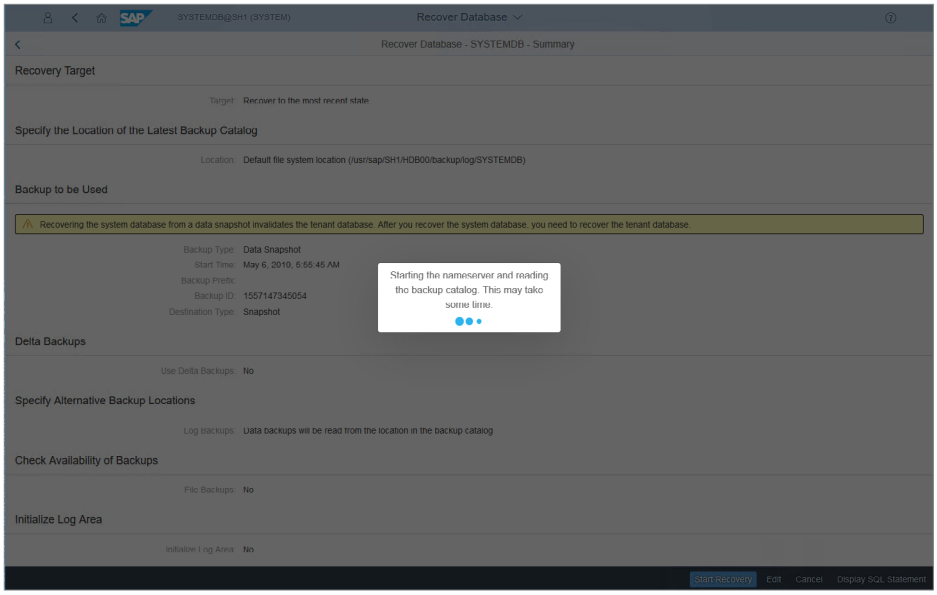

SAP HANA Cockpit Step 2. Recover Tenant Database

During the recovery of the System database all tenant databases are invalidated and need to be individually recovered by repeating the below steps.

Important: Follow from here if using single container SAP HANA 1.0 system. Notable differences are the absence of “System” and “Tenant” terminology.

Important: In SAP HANA 2.0 SPS04 and onwards repeat these steps for each additional tenant database.

Important: Ensure that Step 1a is completed, while the database is offline, for SAP HANA 1.0 Single tenant systems before proceeding with the below.

Once the System Database has been recovered select "Manage Databases" from the SystemDB database view.

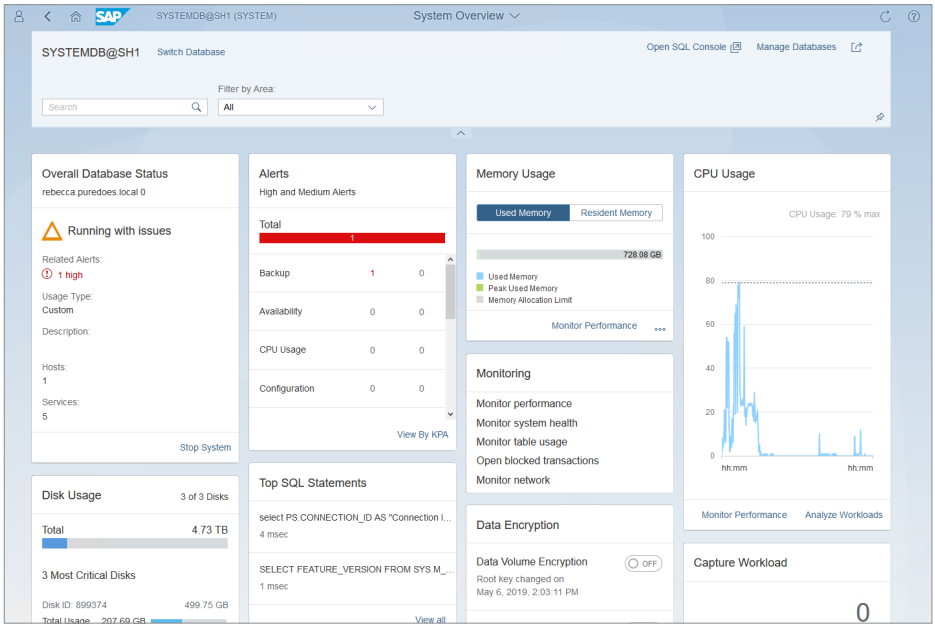

All of the tenants will be shown, but in an offline state. In SAP HANA SPS03 and previous releases only a single tenant system can be recovered.

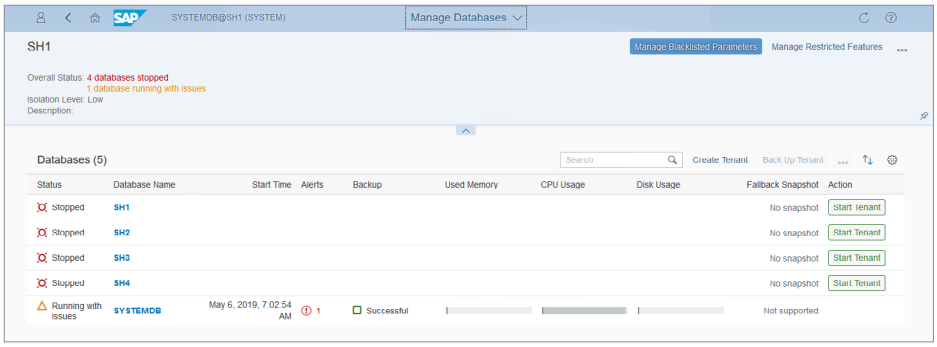

Highlight the relevant tenant database to be recovered and then select the three dots in the right-hand corner to bring up the context menu, then select "Recover Tenant".

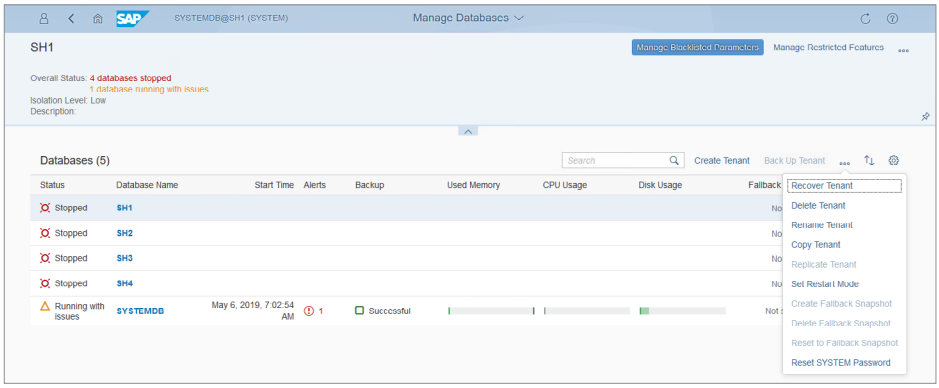

Select the recovery target type.

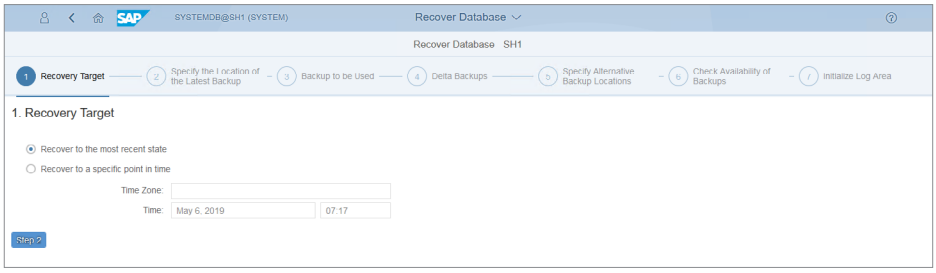

Specify the location of the latest backup catalog, only change these values under specialised circumstances.

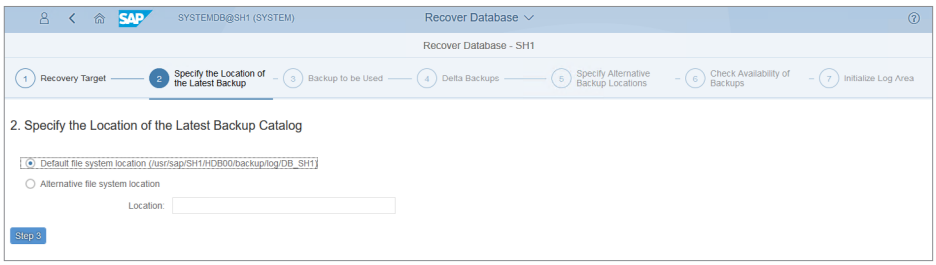

The relevant backup to be used will have the Status of "Available", select this and then select "Step 4".

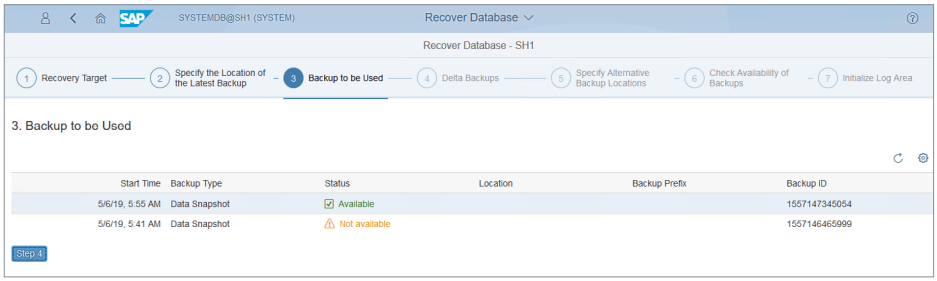

Specify if Delta Backups are to be used.

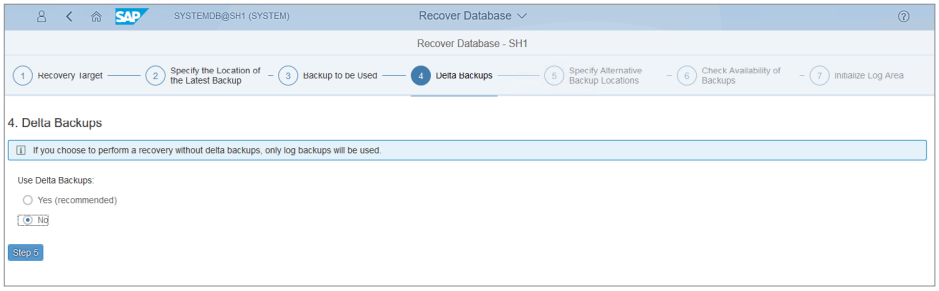

Specify an Alternative Backup location if any exist.

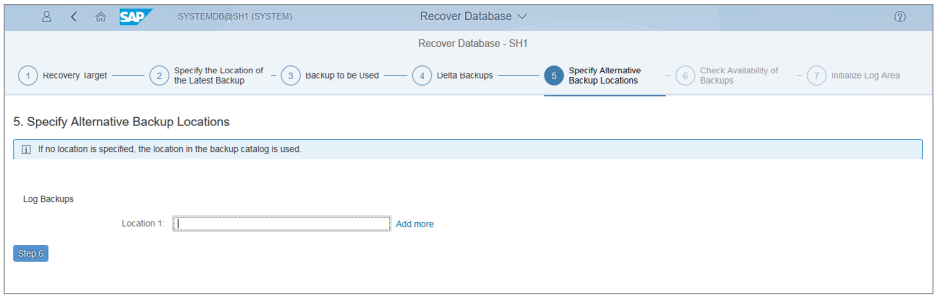

Specify if the log area must be initialized.

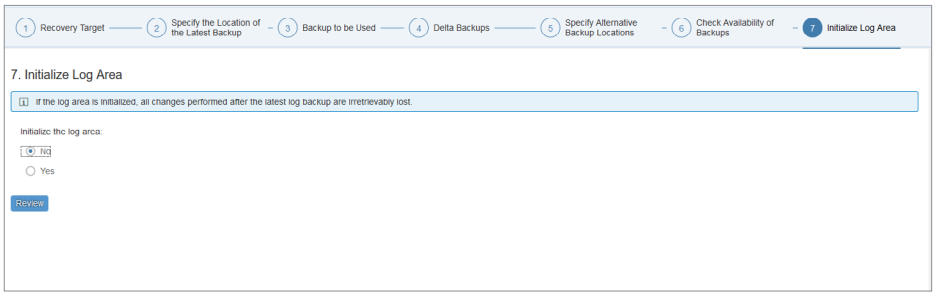

Review the information for the recovery operation and then select "Start Recovery".

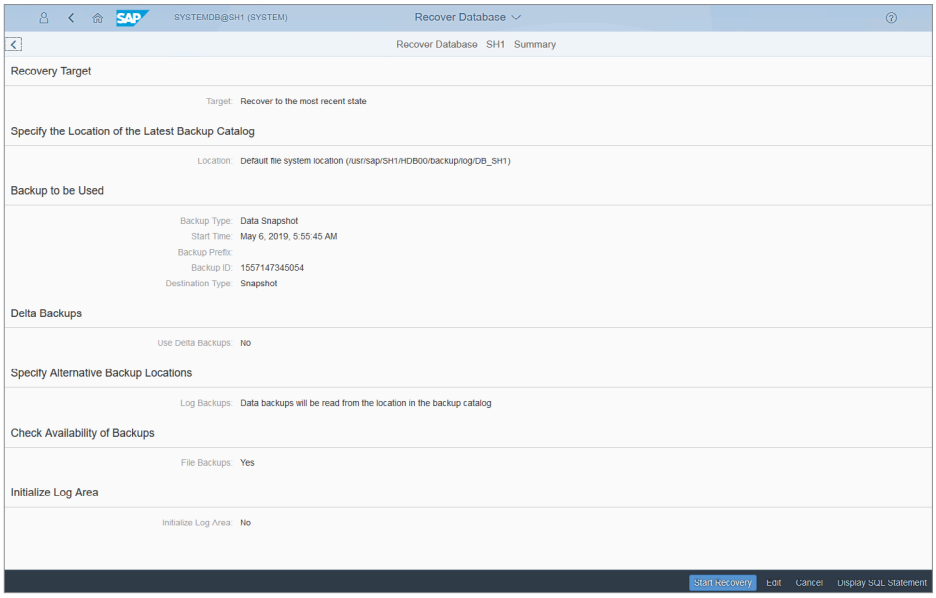

Once the tenant has been successfully recovered the status will be shown.

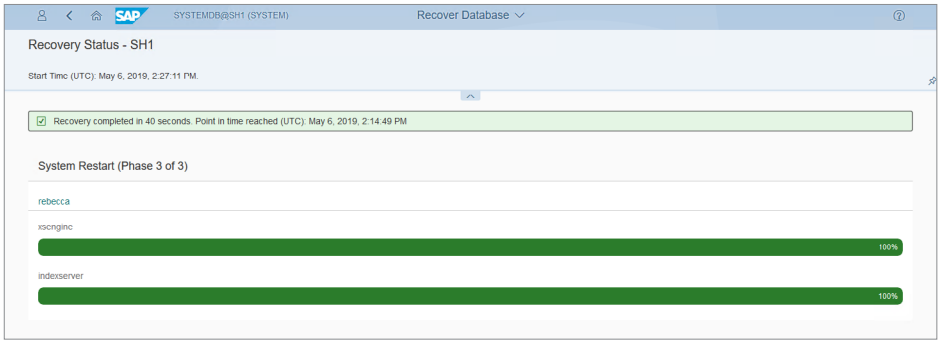

SAP HANA Scale out / Distributed Host deployment

MANUAL OPERATION

Step 1. Verify the system type

Using SAP HANA Studio connected to any one of the systems deployed with the SAP HANA instance, navigate towards the configuration page

SAP HANA Studio, Configuration Page.

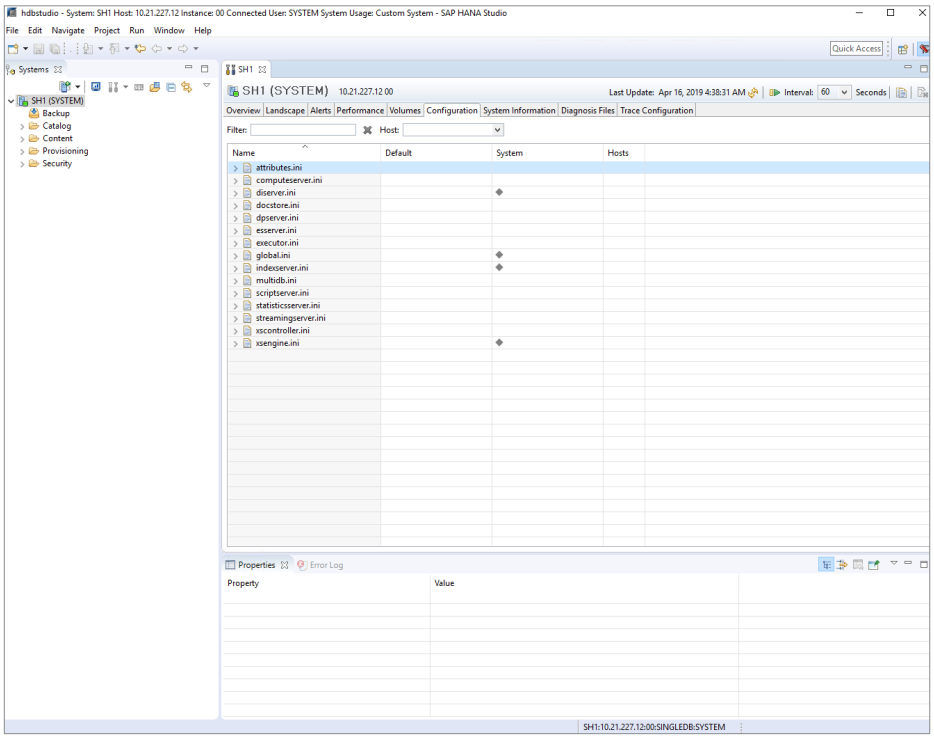

Expand the **global.ini** filename and then further expand the **multidb** section. Look for the key **mode** and observe its value. If this value is “singledb” then the system is a single container and if the value is “multidb” the system is then set to be a multiple container system.

In SAP HANA Studio the global.ini file is expanded.

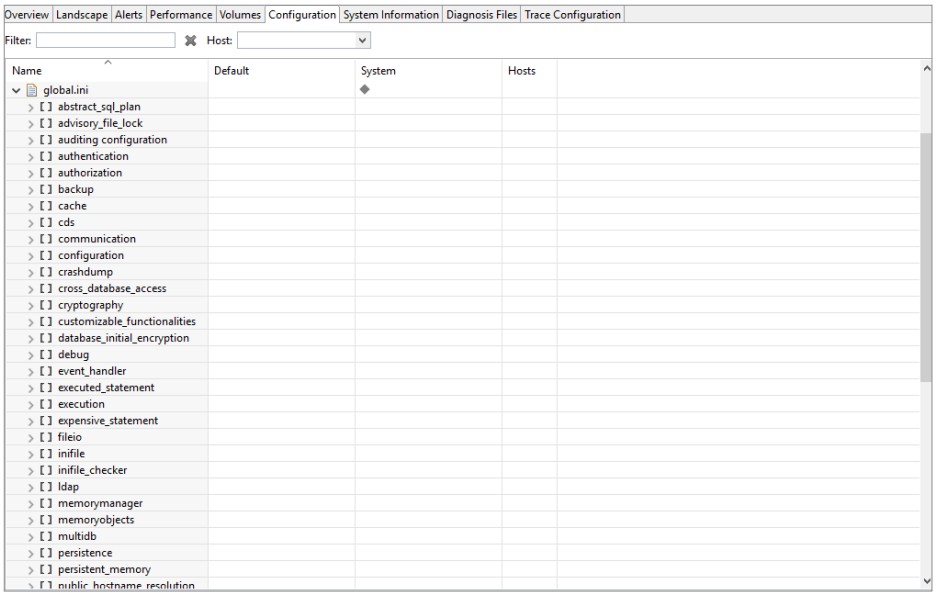

In SAP HANA Studio the multidb section under global.ini shows the system properties.

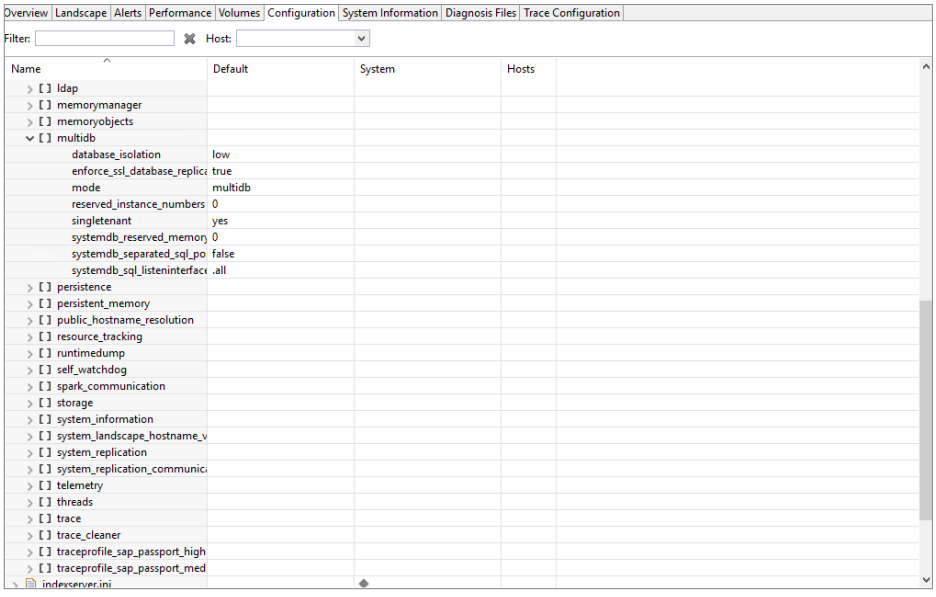

Step 1a. In the event of the database mode being “multidb” find the host on which the SYSTEMDB is running

To identify the host which the SYSTEMDB is running on the nameserver and other non-persistent services needs to be identified. As per the SAP Knowledge Base article (201244) the most notable differences between the list of services running on the SYSTEM DB and tenant database are the nameserver and other non-persistent services are available on the SYSTEMDB. The SYSTEMDB will only ever exist on a single host in a scale out/distributed system environment and cannot be distributed across multiple hosts.

To investigate the services on each system in a scale out/distributed environment the landscape view and services tab is inspected to provide the required information.

The host on which the “nameserver” service and “master” detail are shown, is the system on which the SYSTEMDB is running.

In SAP HANA Studio, under Landscape for the deployment, the SYSTEMDB is running on SHN3.

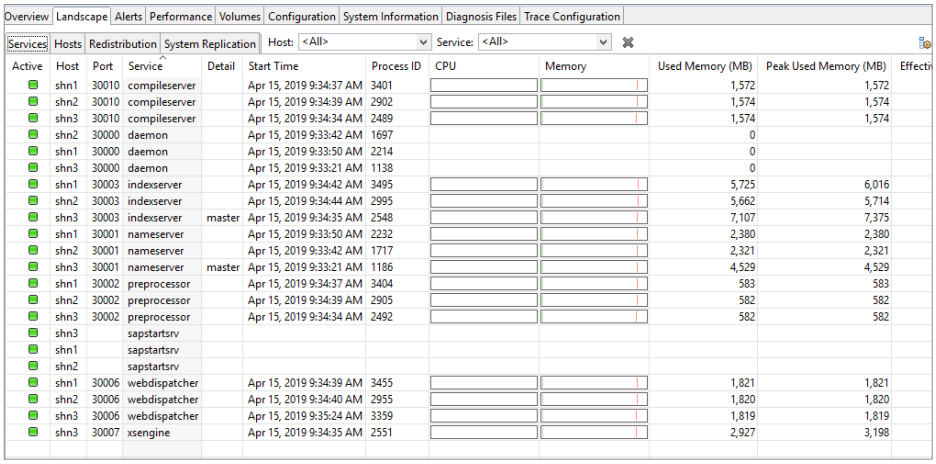
Overview | Landscape | Alerts | Performance | Volumes | Configuration | System Information | Diagnosis Files | Trace Configuration

Services | Hosts | Redistribution | System Replication | Host: <All> | Service: <All>

Active	Host	Port	Service	Detail	Start Time	Process ID	CPU	Memory	Used Memory (MB)	Peak Used Memory (MB)	Effecti
	shn1	30010	compileserver		Apr 15, 2019 9:34:37 AM	3401			1,572	1,572	
	shn2	30010	compileserver		Apr 15, 2019 9:34:39 AM	2902			1,574	1,574	
	shn3	30010	compileserver		Apr 15, 2019 9:34:34 AM	2489			1,574	1,574	
	shn2	30000	daemon		Apr 15, 2019 9:33:42 AM	1697			0		
	shn1	30000	daemon		Apr 15, 2019 9:33:50 AM	2214			0		
	shn3	30000	daemon		Apr 15, 2019 9:33:21 AM	1138			0		
	shn1	30003	indexserver		Apr 15, 2019 9:34:42 AM	3495			5,725	6,016	
	shn2	30003	indexserver		Apr 15, 2019 9:34:44 AM	2995			5,662	5,714	
	shn3	30003	indexserver	master	Apr 15, 2019 9:34:35 AM	2548			7,107	7,375	
	shn1	30001	nameserver		Apr 15, 2019 9:33:50 AM	2232			2,380	2,380	
	shn2	30001	nameserver		Apr 15, 2019 9:33:42 AM	1717			2,321	2,321	
	shn3	30001	nameserver	master	Apr 15, 2019 9:33:21 AM	1186			4,529	4,529	
	shn1	30002	preprocessor		Apr 15, 2019 9:34:37 AM	3404			583	583	
	shn2	30002	preprocessor		Apr 15, 2019 9:34:39 AM	2905			582	582	
	shn3	30002	preprocessor		Apr 15, 2019 9:34:34 AM	2492			582	582	
	shn3		sapstartsrv								
	shn1		sapstartsrv								
	shn2		sapstartsrv								
	shn1	30006	webdispatcher		Apr 15, 2019 9:34:39 AM	3455			1,821	1,821	
	shn2	30006	webdispatcher		Apr 15, 2019 9:34:40 AM	2955			1,820	1,820	
	shn3	30006	webdispatcher		Apr 15, 2019 9:35:24 AM	3359			1,819	1,819	
	shn3	30007	xsengine		Apr 15, 2019 9:34:35 AM	2551			2,927	3,198	

Step 1b. Identify which host each storage device is mounted to it at that point in time

The volumes and services used can be identified by selecting the volumes view and viewing the information for the "indexserver" service where "Data Volume Size" has a value. Each data persistence mount point for the scale out/distributed environment will typically be in the form "{BASE_PATH}/Data/<Database Name>/mntXXXXX" but this can vary between each deployment.

In SAP HANA Studio the relevant volumes can be identified under the "Volumes" view. In this case the mount points and hosts they are attached to are

SHN1: /hana/data/SH1/mnt00002
SHN2: /hana/data/SH1/mnt00003
SHN3: /hana/data/SH1/mnt00001

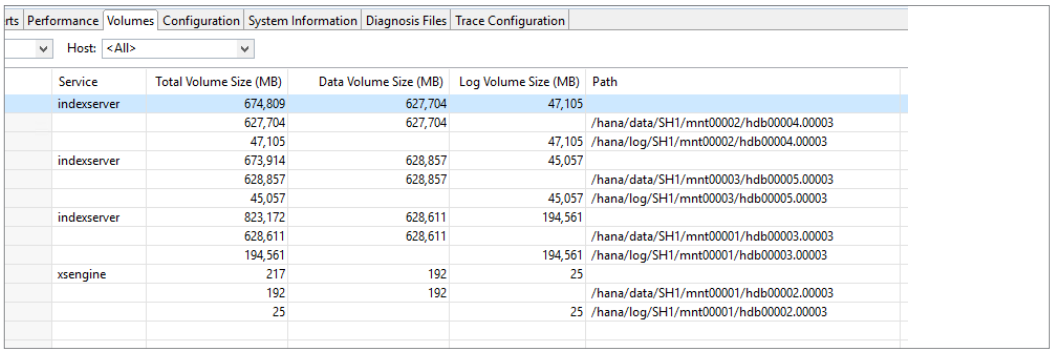

Service	Total Volume Size (MB)	Data Volume Size (MB)	Log Volume Size (MB)	Path
indexserver	674,809	627,704	47,105	
	627,704	627,704		/hana/data/SH1/mnt00002/hdb00004.00003
	47,105		47,105	/hana/log/SH1/mnt00002/hdb00004.00003
indexserver	673,914	628,857	45,057	
	628,857	628,857		/hana/data/SH1/mnt00003/hdb00005.00003
	45,057		45,057	/hana/log/SH1/mnt00003/hdb00005.00003
indexserver	823,172	628,611	194,561	
	628,611	628,611		/hana/data/SH1/mnt00001/hdb00003.00003
	194,561		194,561	/hana/log/SH1/mnt00001/hdb00003.00003
xsengine	217	192	25	
	192	192		/hana/data/SH1/mnt00001/hdb00002.00003
	25		25	/hana/log/SH1/mnt00001/hdb00002.00003

Step 2. Connect to the correct database instance for the creation and management of storage snapshots

If the system is a single database container then ensure that under mode the **Single container** selection is selected, for a multiple container system ensure that **Multiple containers** is selected, and when available ensure that **System database** is the selected target to connect to. It is important to note that selecting **Single container** and the mode to connect to will still connect to the single tenant container.

Using SAP HANA Studio connect to a single container system.

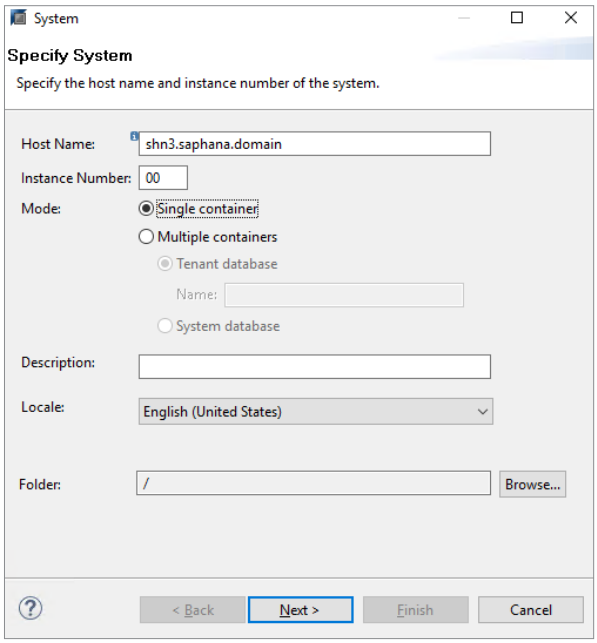

Using SAP HANA Studio connect to the System dataset on a system with multiple containers.

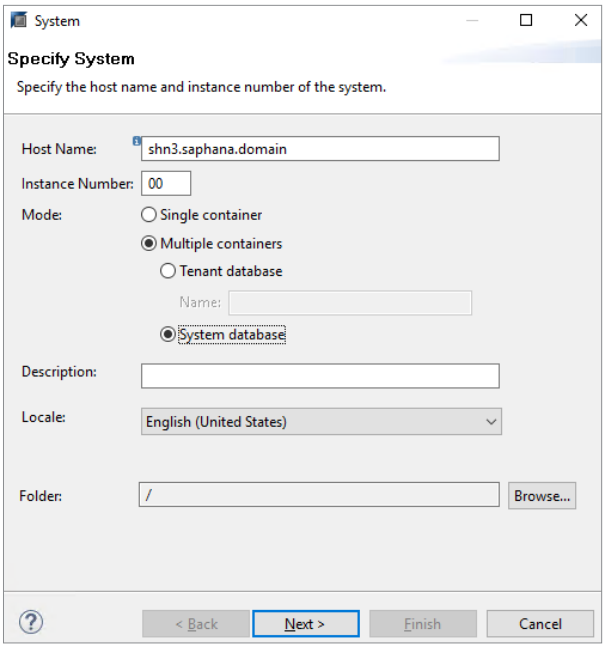

Connected to a single container system.

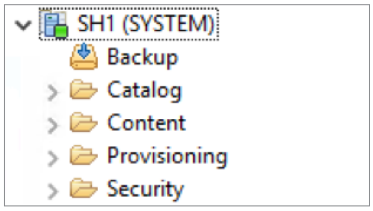

Connected to a multiple container system.

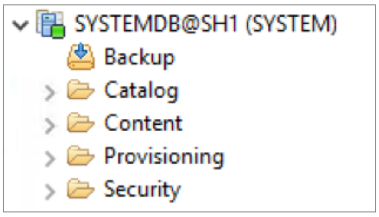

Step 3. Prepare the system for a storage snapshot

This can be done from the backup console or using the side menu by right-clicking on the **Backup** system folder and selecting **Manage Storage Snapshot....**

In SAP HANA Studio right-click on the "Backup" system folder and select "Manage Storage Snapshot...".

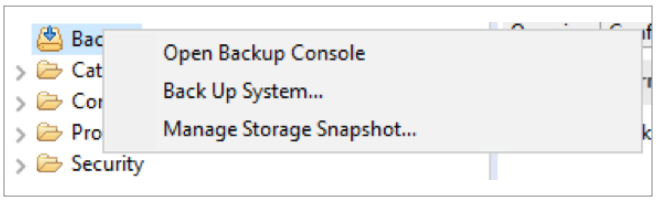

In SAP HANA Studio and the backup console, select the camera next to "Status of Currently Active Data Backup".

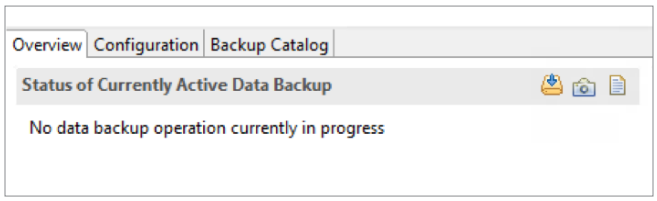

Select Prepare and add a comment if needed, then press the OK button.

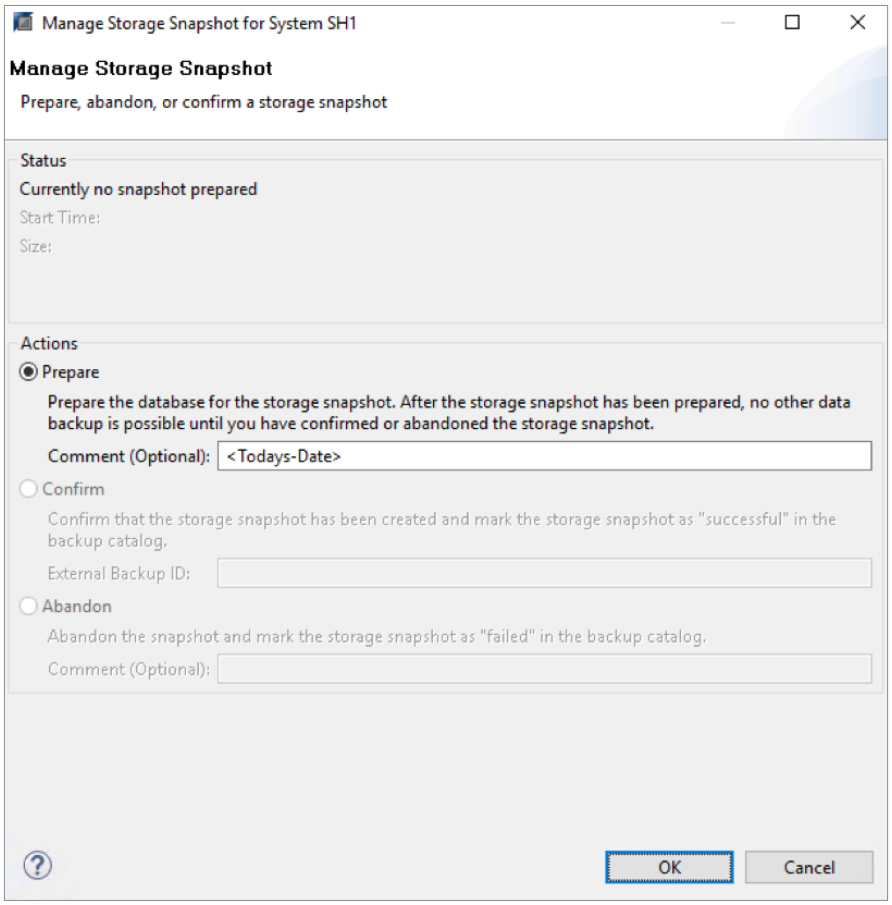

Once the database snapshot is ready then the below will show in the backup console.

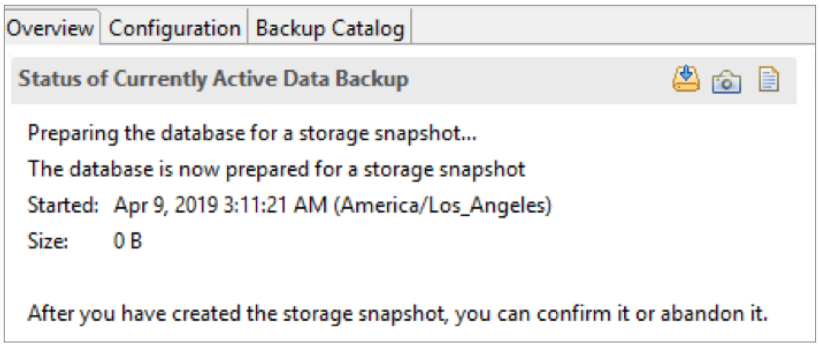

Step 4. Retrieve the SAP HANA prepared storage backup ID using the SQL console in SAP HANA Studio

Use the SQL console for the database instance the prepared snapshot has been created for and run the query:

```
"SELECT BACKUP_ID, COMMENT FROM M_BACKUP_CATALOG WHERE ENTRY_TYPE_NAME = 'data snapshot'
AND STATE_NAME = 'prepared'"
```

This will return the backup ID for the prepared snapshot, which is used in the FlashArray block device snapshot creation as a suffix to link the SAP HANA storage snapshot ID to the Volume snapshot.

In SAP HANA Studio right-click on the instance and select "Open SQL Console".

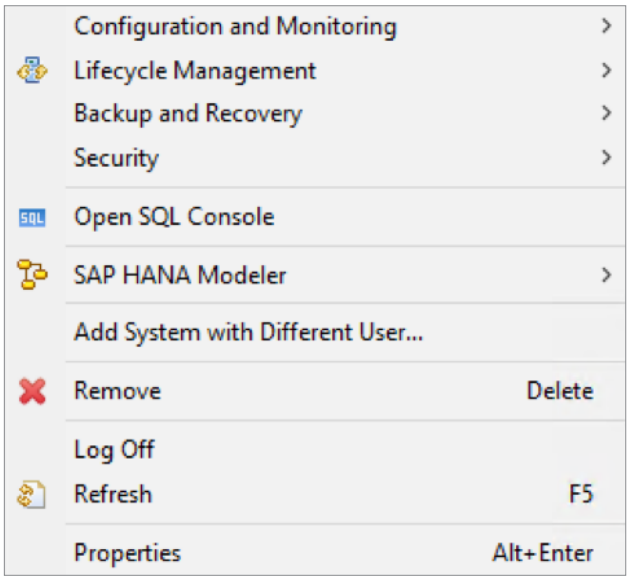

In SAP HANA Studio execute the query to return the Backup_ID of the prepared database snapshot.

SQL | Result

```
SELECT BACKUP_ID, COMMENT
FROM M_BACKUP_CATALOG
WHERE ENTRY_TYPE_NAME = 'data snapshot'
AND STATE_NAME = 'prepared'
```

The Backup ID is returned with any comments added to the entry, take note of the Backup.

	BACKUP_ID	COMMENT
1	1,554,804,681,328	SNAPSHOT-2019-04-09 03:10:25

Step 5. Freeze the filesystem for the SAP HANA data persistence mount points on each host

Open a terminal (SSH or local to the system) and ensure the prompt is logged in as a user who has read, write and execute permissions on the SAP HANA data persistence mount point. The data persistence mount points and each host they are attached to were identified in step 1b. The “fsfreeze” Linux utility will be used to halt any IO to the volume and ensure consistency.

Freeze the filesystem of each data persistence mount point using the fsfreeze utility.

```
SHN1:~ # /sbin/fsfreeze --freeze /hana/data/SH1/mnt00002
SHN2:~ # /sbin/fsfreeze --freeze /hana/data/SH1/mnt00003
SHN3:~ # /sbin/fsfreeze --freeze /hana/data/SH1/mnt00001
```

Step 6. Take a snapshot of the block device in the FlashArray

This step is shown using the web based graphical user interface to operate the FlashArray storage device. It is assumed the user can identify the block volume which matches each SAP HANA persistence data volume. In the user interface navigate to the “Storage” section, select “Volumes”. Select the volume which corresponds to each SAP HANA persistence data volume. Under “Volume Snapshots”

Pure Storage FlashArray Web Graphical User interface – Main page.

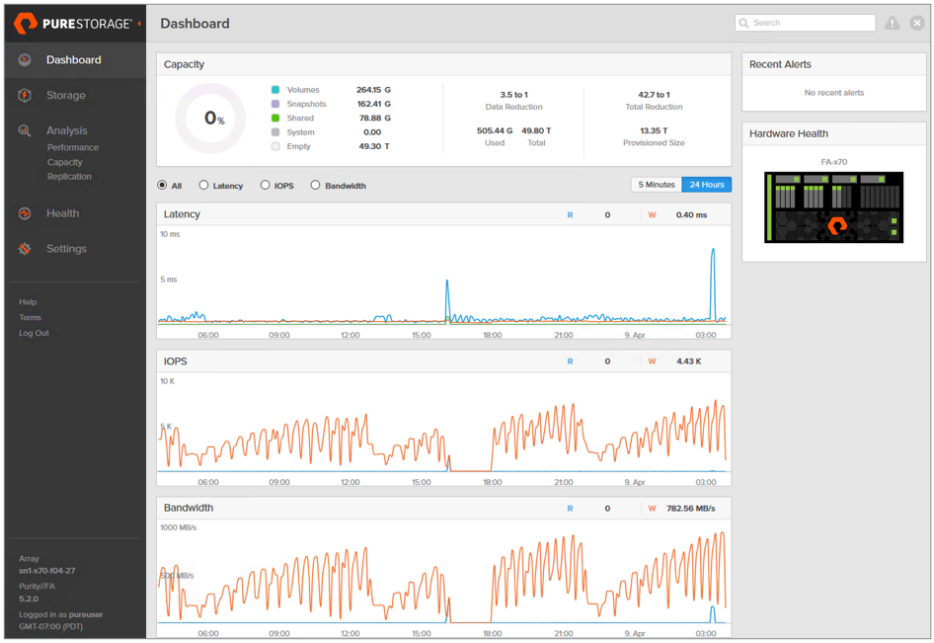

Select Storage and navigate to the "Volumes" section to view all volumes and snapshots.

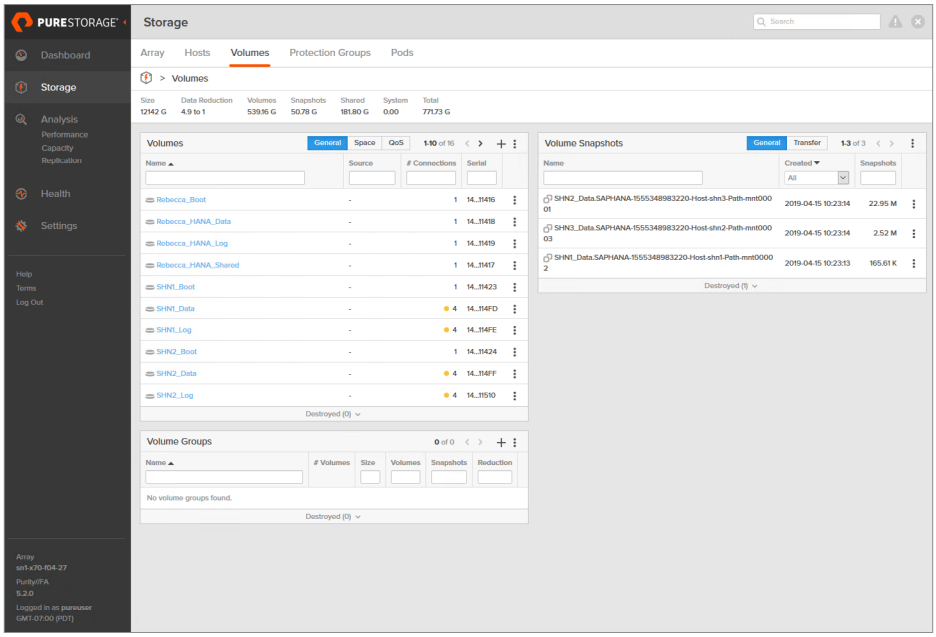

Select the volume which matches the SAP HANA Data persistence volume and select the "+" next to Volume Snapshots.

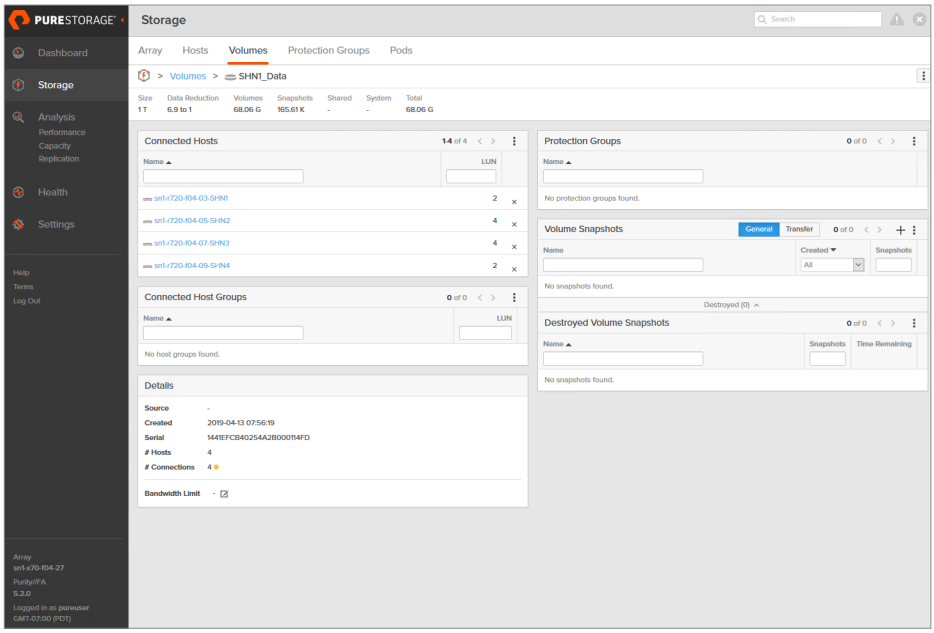

Create the snapshot with the suffix "SAPHANA-<Backup_ID from SAP HANA prepared snapshot>-Host-<Host>-Path-<mounted Path At time of Snapshot>

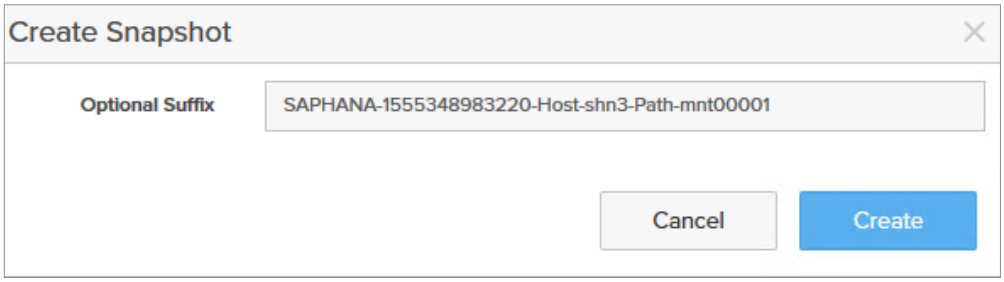

The snapshot is created and listed under the Volume snapshots.

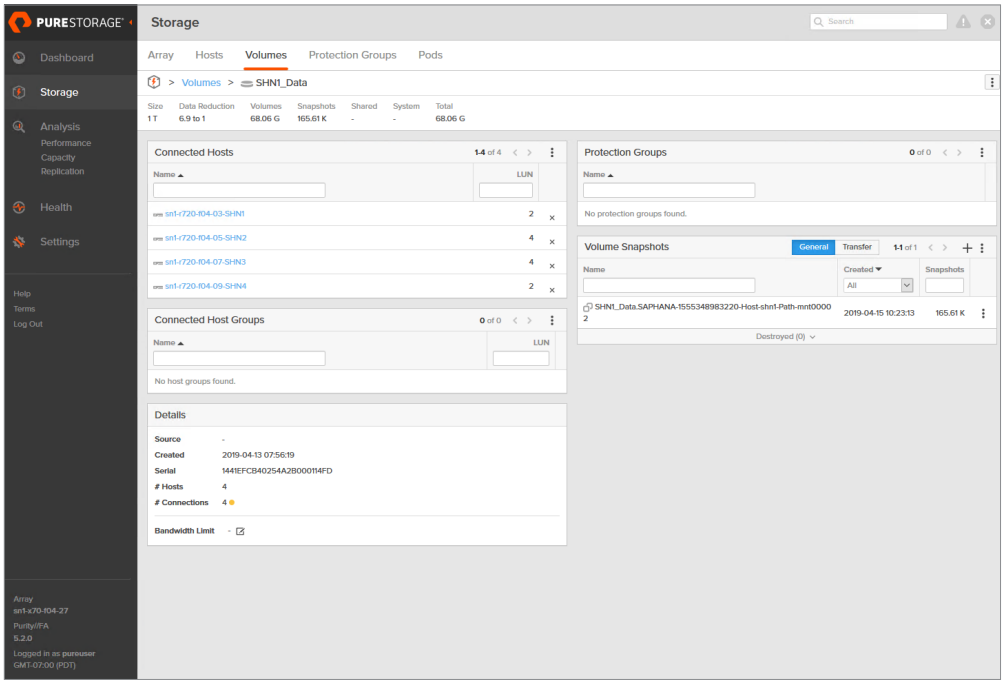

Step 7. Unfreeze the filesystem for each SAP HANA data persistence mount point

Open a terminal (SSH or local to the system) and ensure the prompt is logged in as a user who has read, write and execute permissions on each SAP HANA data persistence mount point. We will then use the “fsfreeze” Linux utility to resume IO to the volume and allow the database to continue operation.

Unfreeze the filesystem of each data persistence mount point using the fsfreeze utility

```
SHN1:~ # /sbin/fsfreeze --unfreeze /hana/data/SH1/mnt00002
SHN2:~ # /sbin/fsfreeze --unfreeze /hana/data/SH1/mnt00003
SHN3:~ # /sbin/fsfreeze --unfreeze /hana/data/SH1/mnt00001
```

Step 8. Confirm or abandon the snapshot

Confirm or abandon the snapshot in SAP HANA studio , allowing the backup to be marked as valid or invalid. If something has not operated as expected then the snapshot should be abandoned. To confirm the snapshot an External Backup ID must be supplied. The Backup_ID originally offered by the prepared snapshot and used as a suffix for the Block volume storage snapshot is used for the value.

Confirm the Snapshot and supply the External Backup ID, and press OK.

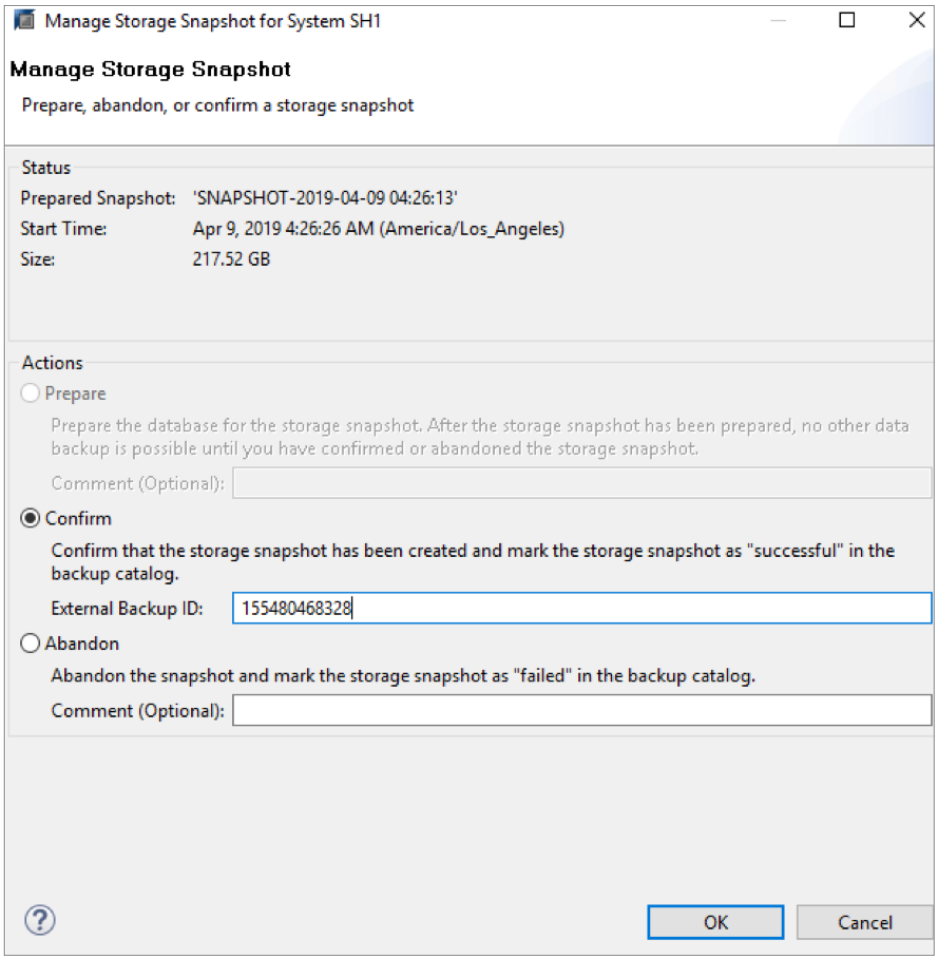

The backup now shows in the backup catalog as complete.

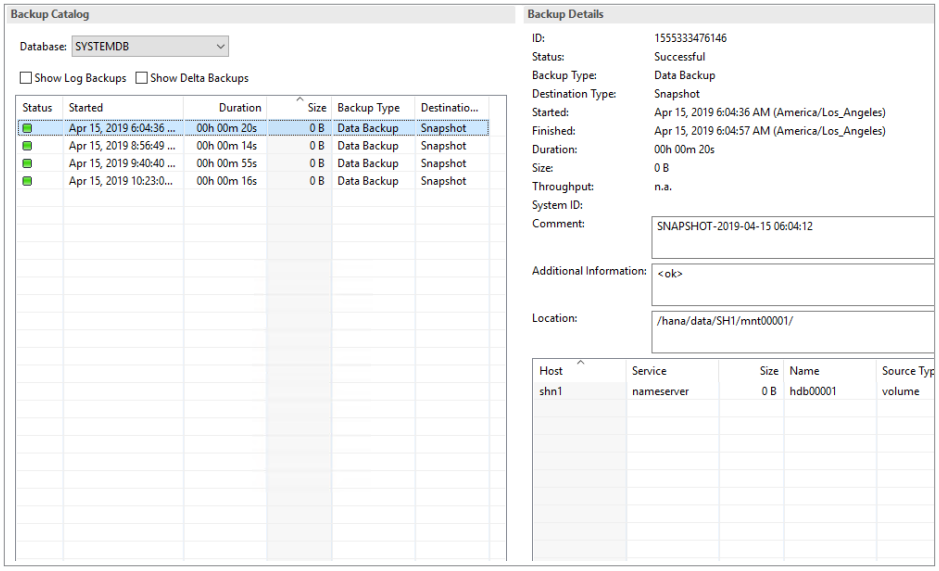

AUTOMATED OPERATION

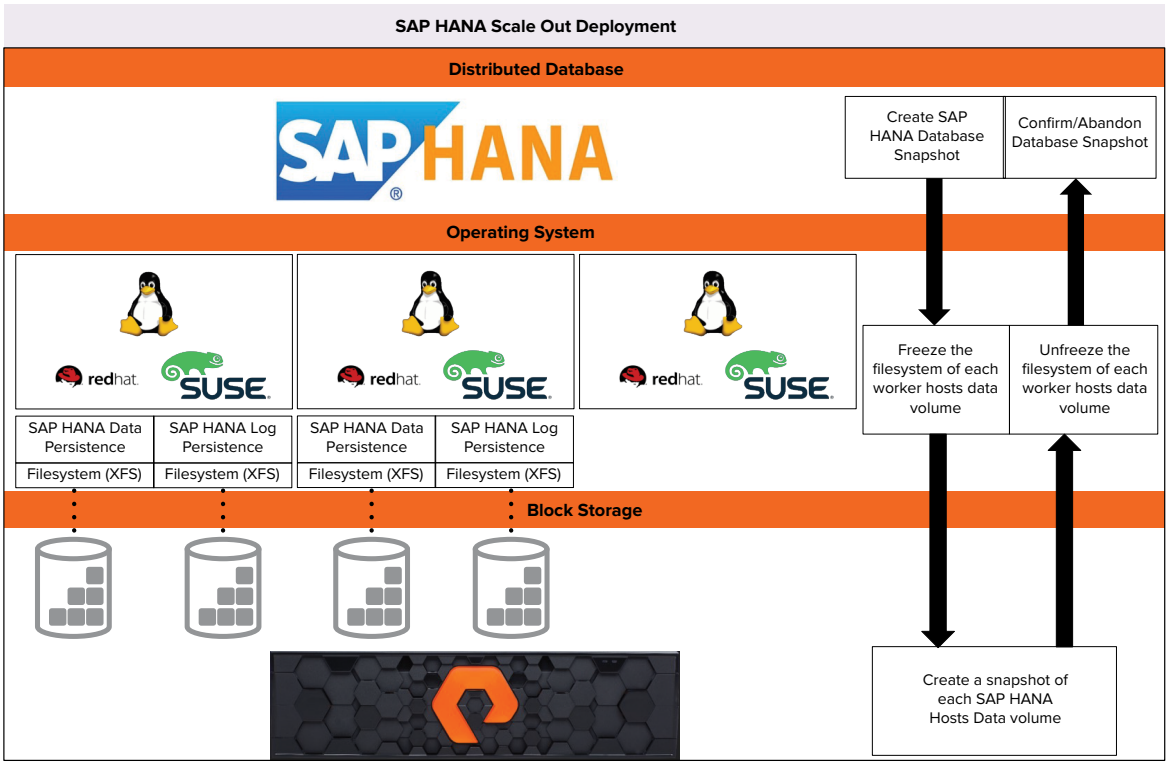

Fig. 10 Workflow to create a storage snapshot for a distributed system/multi-host SAP HANA database.

Step 1. Check the SAP HANA System mode

Using the resilient connection string:

```
Driver={HDBODBC}; ServerNode=<Host-01-Address>:3 <InstanceNumber> 15, < Host-02-Address >:3 <InstanceNumber> 15, < Host-03-Address >:3 <InstanceNumber> 15; UID= <Database User>; PWD=<DatabasePassword>;
```

Connect to the SAP HANA database and run the following query to determine system mode:

```
SELECT VALUE FROM M_INIFILE_CONTENTS WHERE FILE_NAME = 'global.ini' AND SECTION = 'multidb' AND KEY = 'mode'
```

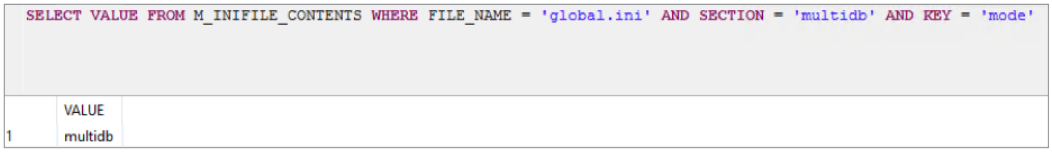

	VALUE
1	multidb

The result will be either "singledb" or "multidb", if the result is "multidb" then a further query needs to be run to identify which host the SYSTEMDB is running on:

```
"SELECT HOST FROM SYS.M_SERVICES WHERE DETAIL = 'master' AND SERVICE_NAME = 'nameserver'"
```

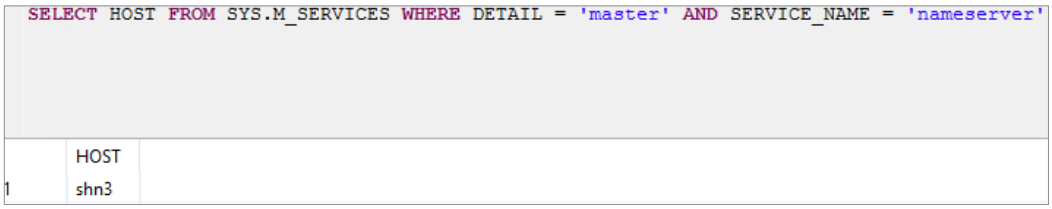

	HOST
1	shn3

Once the host on which the host is running has been isolated change the connection string to the following:

```
Driver={HDBODBC}; ServerNode=<SystemDBHost>:3 <InstanceNumber> 13; UID= <Database User>; PWD=<DatabasePassword>;
```

Changing the port number allows the application to connect to the system database on the specified host. In this case the port to connect on for instance 00 would change from 30015 to 30013. In the event of the system running in “singledb” mode, continue to use the original connection string.

Step 2. Identify which host each volume is attached to and the associated mount point for each of the SAP HANA persistent data volumes

Using the established connection string run the following query to determine all of the SAP HANA persistence data volume mount points:

```
SELECT HOST, STORAGE _ID, PATH, KEY, VALUE
FROM SYS.M_ATTACHED_STORAGES WHERE KEY = ‘WWID’
AND PATH LIKE (SELECT CONCAT(VALUE,’%’) FROM M_INIFILE_CONTENTS
```

```
WHERE FILE_NAME = ‘global.ini’
AND SECTION = ‘persistence’
AND KEY = ‘basepath_datavolumes’
AND VALUE NOT LIKE ‘$%’)
```

```
SELECT HOST,STORAGE_ID,PATH,KEY,VALUE
FROM SYS.M_ATTACHED_STORAGES WHERE KEY = 'WWID'
AND PATH LIKE (SELECT CONCAT(VALUE,'%') FROM M_INIFILE_CONTENTS
 WHERE FILE_NAME = 'global.ini'
 AND SECTION = 'persistence'
```

	HOST	STORAGE_ID	PATH	KEY	VALUE
1	shn1	3	/hana/data/SH1/mnt00002	WWID	3624a93701441efcb40254a2b000114fd
2	shn2	7	/hana/data/SH1/mnt00003	WWID	3624a93701441efcb40254a2b00011511
3	shn3	11	/hana/data/SH1/mnt00001	WWID	3624a93701441efcb40254a2b000114ff

Step 3. Prepare the database snapshot and retrieve the backup ID for it

Using the established connection string to execute the query needed to prepare a database snapshot.

```
BACKUP DATA FOR FULL SYSTEM CREATE SNAPSHOT COMMENT 'SNAPSHOT-<Snapshot Time>
```

To retrieve the backup ID, execute the following query:

```
SELECT BACKUP_ID, COMMENT FROM M_BACKUP_CATALOG WHERE ENTRY_TYPE_NAME = 'data snapshot'
AND STATE_NAME = 'prepared'
```

	BACKUP_ID	COMMENT
1	1,554,804,681,328	SNAPSHOT-2019-04-09 03:10:25

Step 4. Freeze the filesystem for each mountpoint attached to each host

An SSH connection needs to be created with the operating system on which the SAP HANA instance is installed to execute command line arguments. To freeze the filesystem the fsfreeze utility will be used as it supports EXT3/4, ReiserFS, JFS and XFS. The mount points and hosts retrieved in step 2 will be used during the freeze operation.

```
/sbin/fsfreeze –freeze <path to mount point>
```

```
SHN1:~ # /sbin/fsfreeze --freeze /hana/data/SH1/mnt00002
SHN2:~ # /sbin/fsfreeze --freeze /hana/data/SH1/mnt00003
SHN3:~ # /sbin/fsfreeze --freeze /hana/data/SH1/mnt00001
```

Step 5: Query the relevant FlashArray for a list of its volumes and search them for the serial number contained with the values returned by step 2, then create a snapshot once each volume has been located

Example of PowerShell using the Pure Storage PowerShell SDK

```
$Array = New-PfaArray -EndPoint $FlashArrayAddress -username $User -Password $Password -IgnoreCertificateError
$Volumes = Get-PfaVolumes -Array $Array
```

Once the correct volume has been found create a snapshot with a specified Snapshot suffix

```
$VolumeSnapshot = New-PfaVolumeSnapshots -Array $Array -Sources $volume.name -Suffix $SnapshotSuffix
```

Step 5. Unfreeze the filesystem for each mountpoint attached to each host

An SSH connection needs to be created with the operation system on which the SAP HANA instance is installed to execute command line arguments. To unfreeze the filesystem, use the same mount points and hosts in step 3.

```
/sbin/fsfreeze --unfreeze <path to mount point>
```

```
SHN1:~ # /sbin/fsfreeze --unfreeze /hana/data/SH1/mnt00002
SHN2:~ # /sbin/fsfreeze --unfreeze /hana/data/SH1/mnt00003
SHN3:~ # /sbin/fsfreeze --unfreeze /hana/data/SH1/mnt00001
```

Step 6. Confirm or abandon the database snapshot

Using the connection string established in step 1, confirm or abandon the snapshot using hdbsql commands:

Confirm the snapshot if all of the previous steps executed successfully.

```
BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID <External Backup ID> SUCCESSFUL; BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID <External Backup ID> SUCCESSFUL;
```

Abandon the snapshot if one of the previous steps did not execute successfully.

```
BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID <Prepared database snapshot backup ID> UNSUCCESSFUL <additional comments>;
```

RECOVERY PROCESS

The recovery process for a storage snapshot will restore the SAP HANA instance and all of the data to the point in time at which the storage snapshot was taken. It is important to note that the recovery of a system using storage snapshots with multiple database containers (MDC) is only supported from SAP HANA 2.0 SPS04.

SAP HANA Studio Step 1. Prepare to Recover System Database (MDC systems only)

Navigate to the Systems inventory and right-click on the SystemDB connection for the relevant database and navigate to the Backup and Recovery sub menu then select “Recover System Database…”. A prompt to shut down the relevant system will be shown as recovery can only be done when it is offline.

Navigate to the systems inventory , right-click the SystemDB and navigate to Backup and Recovery > Recover System Database…

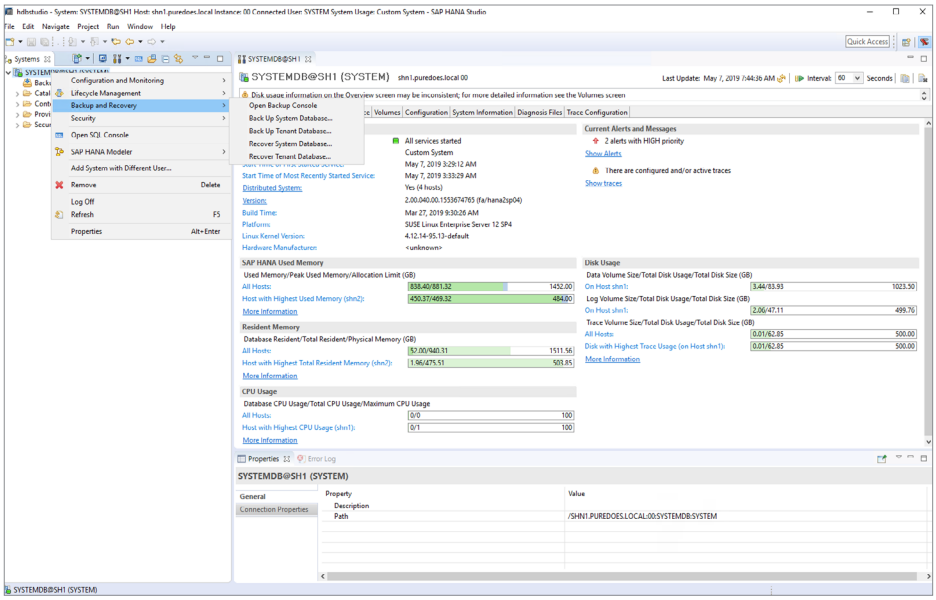

Enter the relevent values for the <sid>adm user created during initial SAP HANA installation.

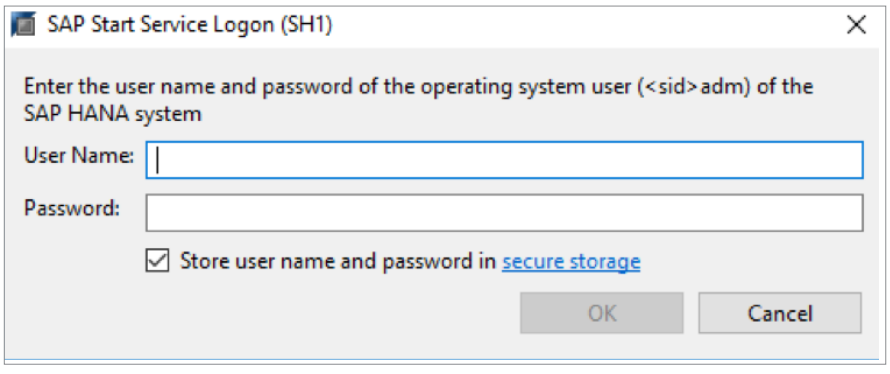

SAP HANA studio will display a prompt to shut down the system.

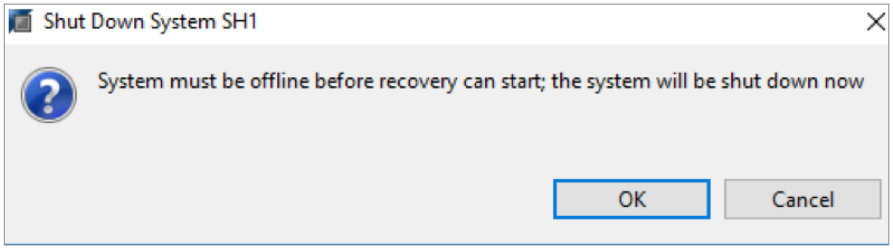

The graphical process gives the user an option of choosing a recovery type.

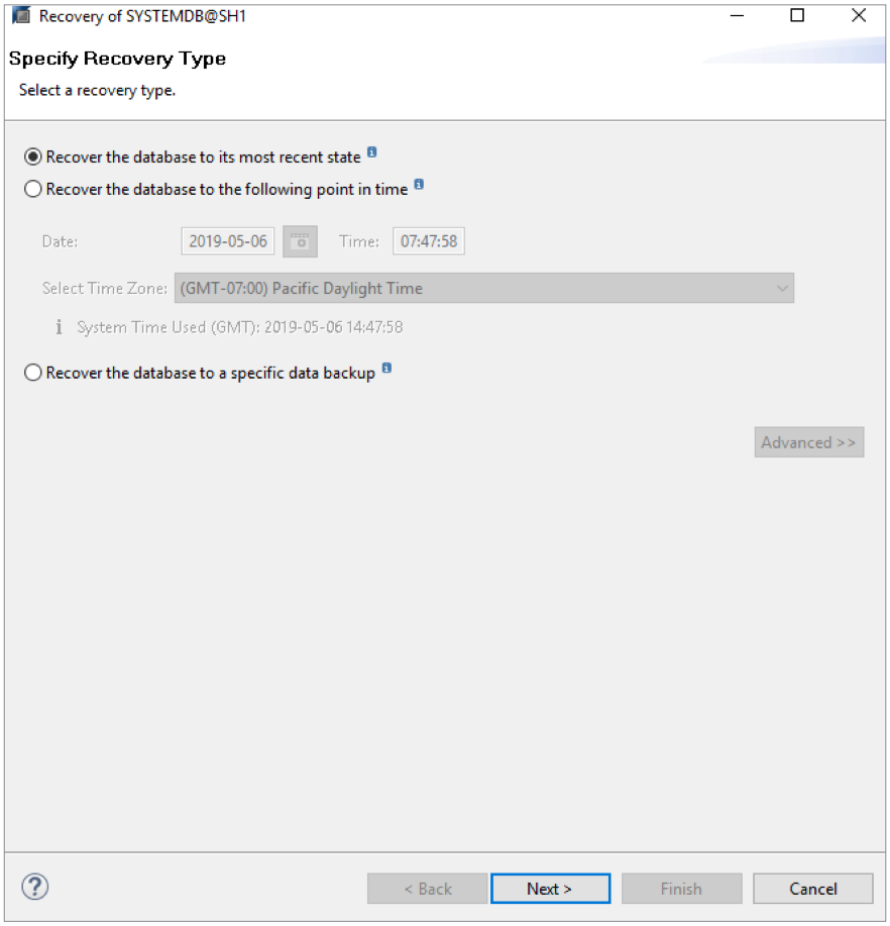

The location of the backup catalog is normally populated, only change this under specialized circumstances where the recovery catalog is in a different location.

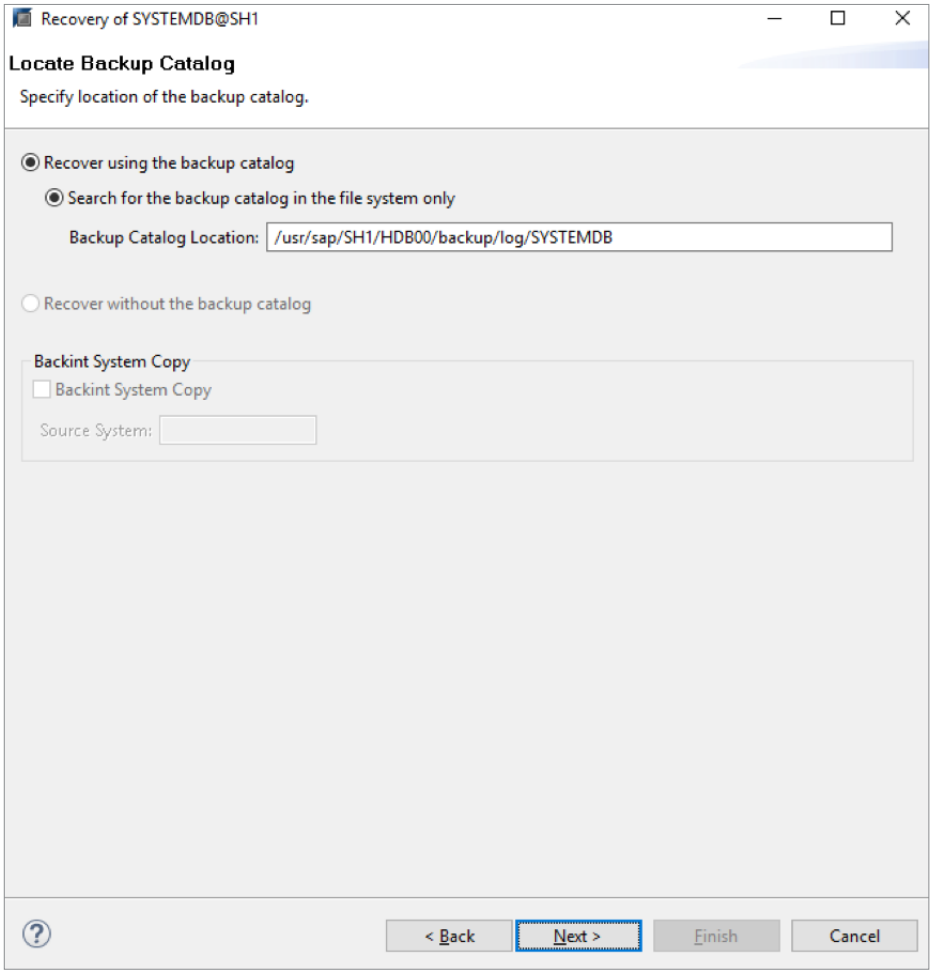

All records in the catalog will be shown with the relevant information, take note of the Backup ID for the required recovery point.

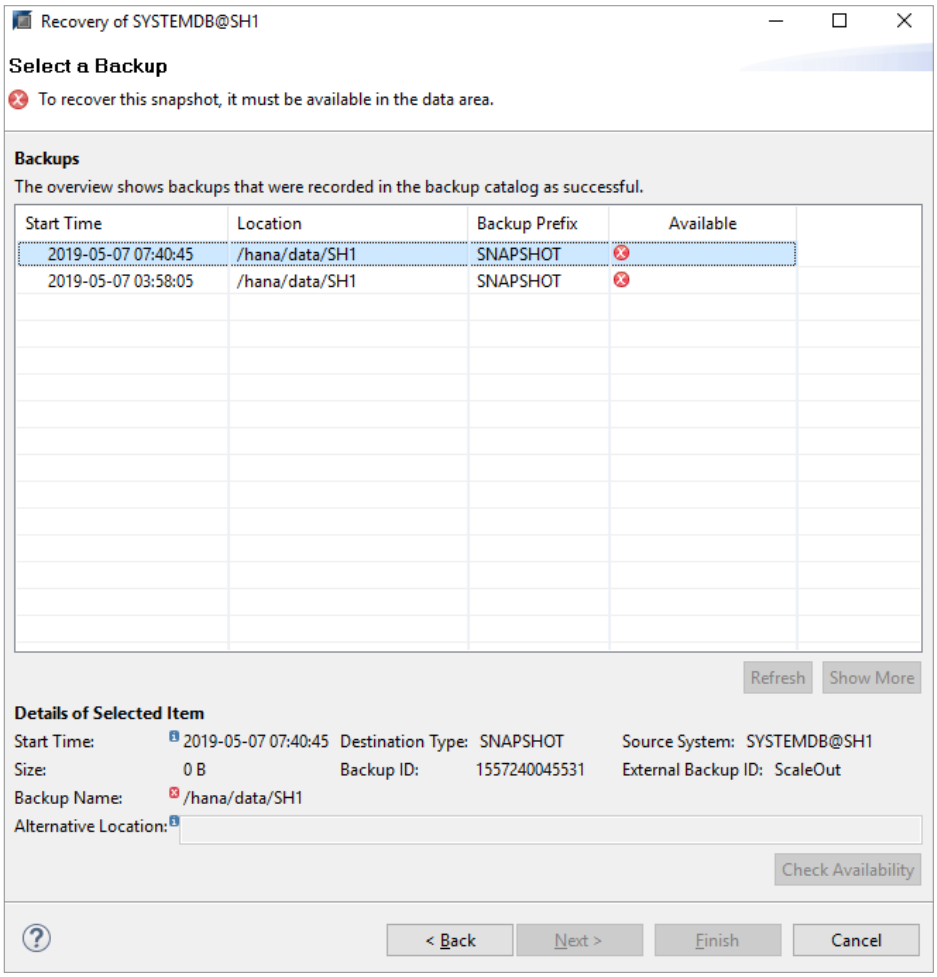

SAP HANA Studio Step 1a. Recover block storage volume

Recover the relevant storage volume using the FlashArray web user interface. This operation can also be completed using the command line interface or ReST API.

Important: The persistence data volumes must be unmounted from the operating system before a snapshot is restored. This can be done using the “umount” command in a terminal or SSH connection.

In the FlashArray web user interface navigate to "Storage" and select the "Volumes" section.

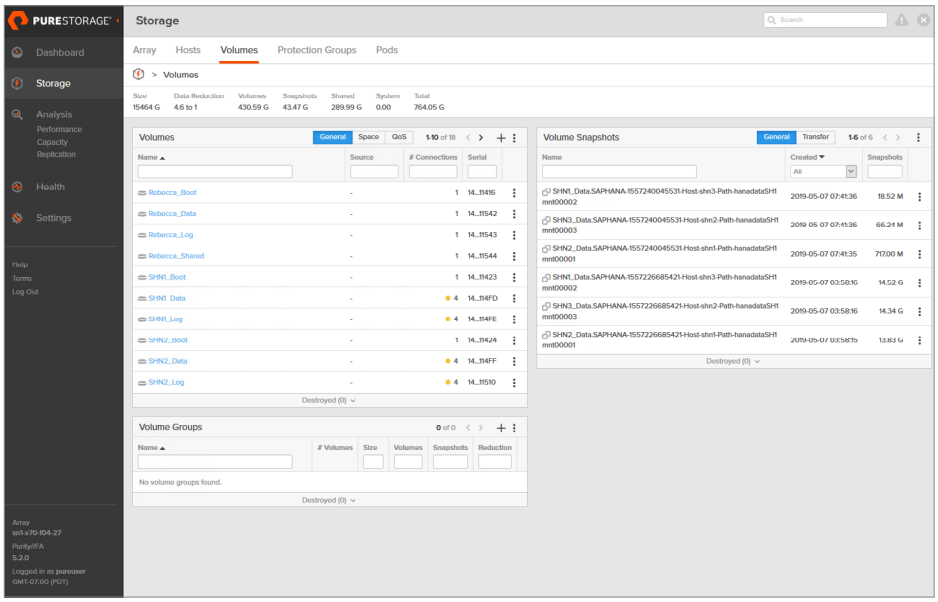

Important: The snapshot which matches the "Backup ID" in the catalog is appended with the exact same value, this is the volume which should be restored. Note that the addition of this value is done during storage snapshot creation by the user.

Identify the snapshots relevant to the Backup_ID of the recovery point.

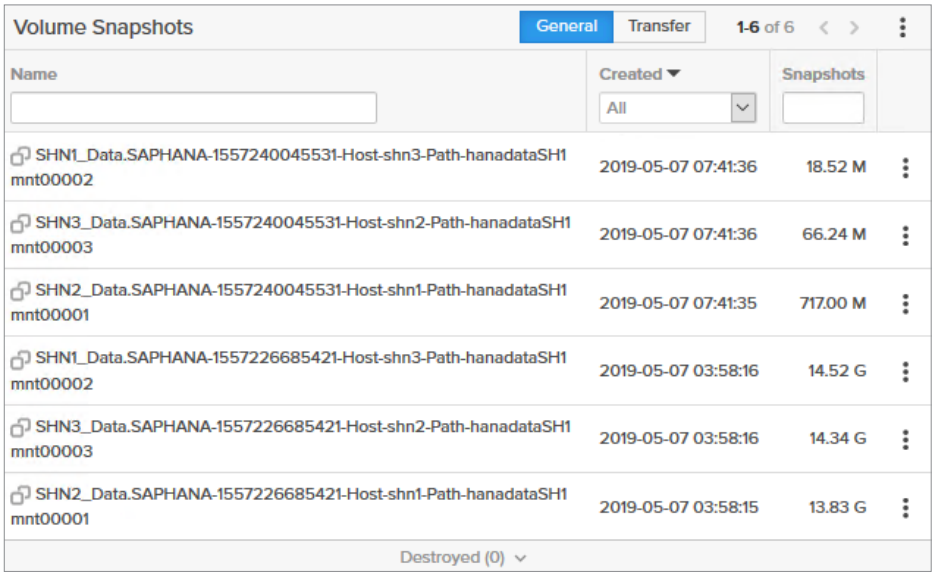

Unmount each persistence data volume from the mount point on each host

```
SHN1:~ # umount /hana/data/SH1/mnt00001
SHN2:~ # umount /hana/data/SH1/mnt00003
SHN3:~ # umount /hana/data/SH1/mnt00002
```

Select the 3 dots on the right-hand side of the relevant snapshot" and select "Restore" for each required block device. A prompt will appear for each snapshot, check the information and then confirm the restore operation.

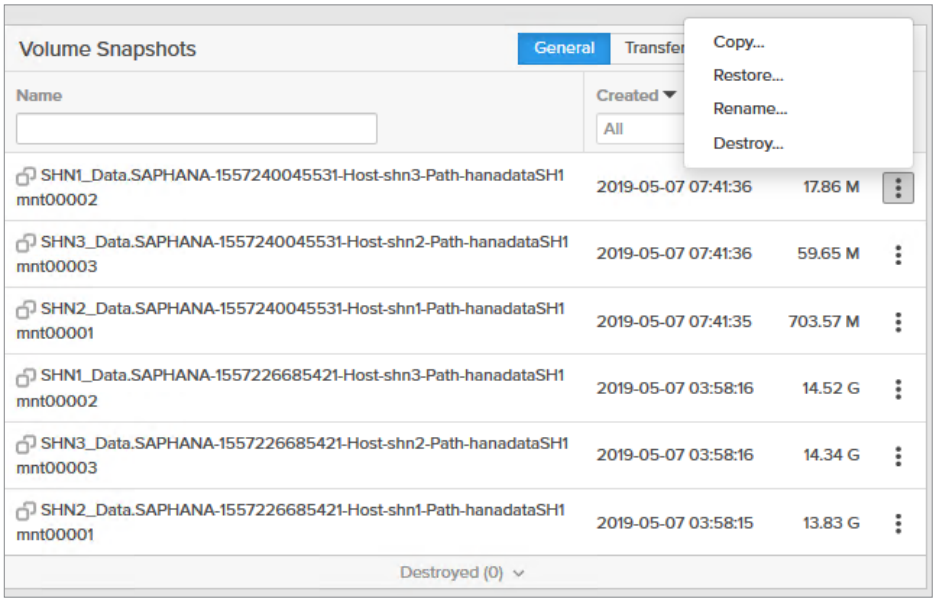

Verify that the volumes have restored correctly.

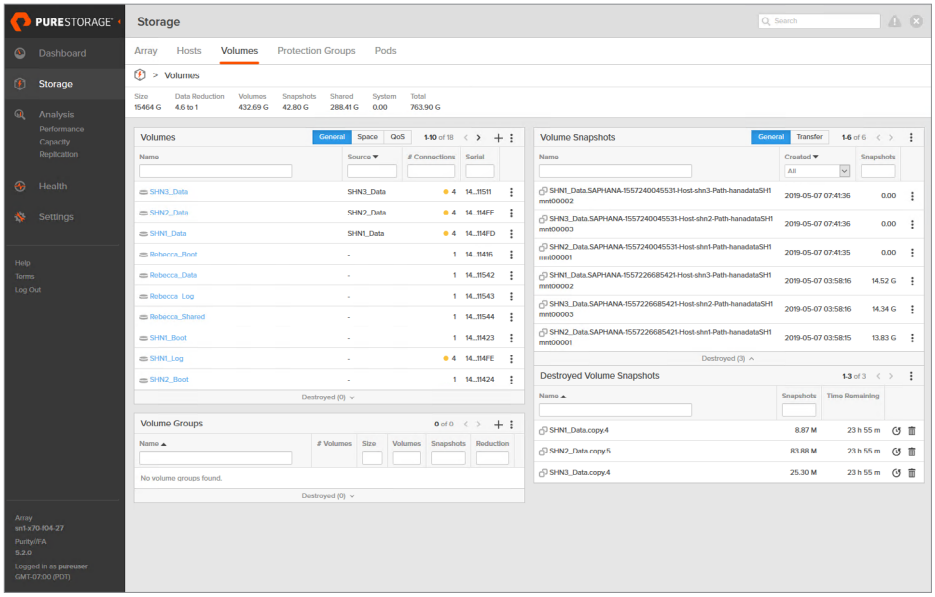

Navigate to each volume and check the details match what is required of the SAP HANA scale out system.

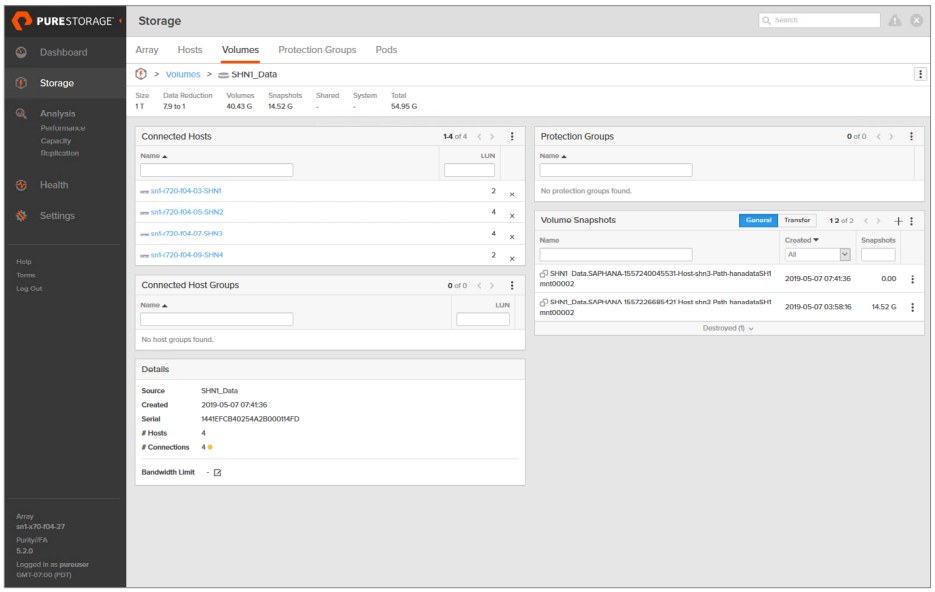

Check the details for each volume required to attach to any host in the distributed host setup.

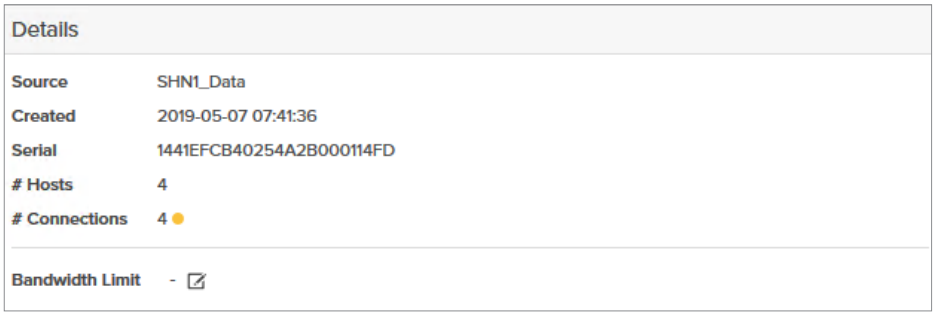

Check the details for each volume required to attach to any host in the distributed host setup.

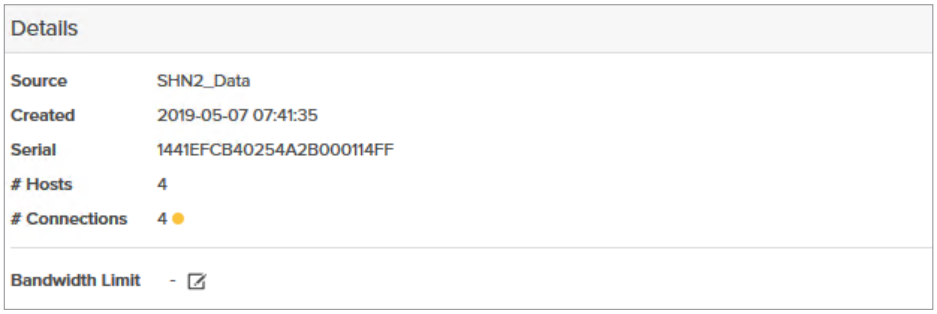

Check the details for each volume required to attach to any host in the distributed host setup.

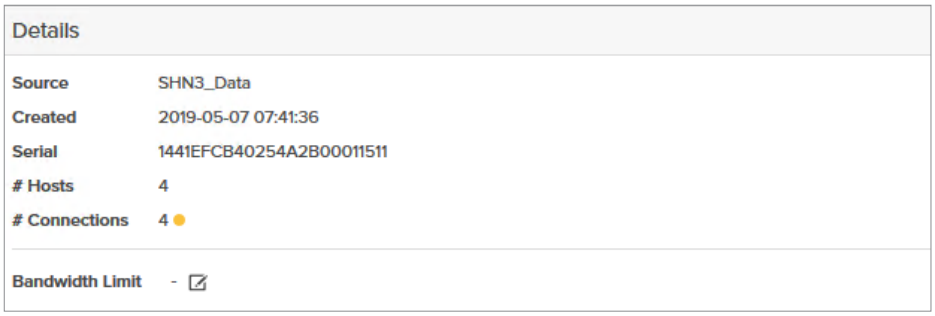

SAP HANA Studio Step 1b. Recover the System Database

Select "Refresh" once all of the steps in Step 1a are completed and wait for the user interface to refresh. Once the required recovery point is shown to be available select "Next".

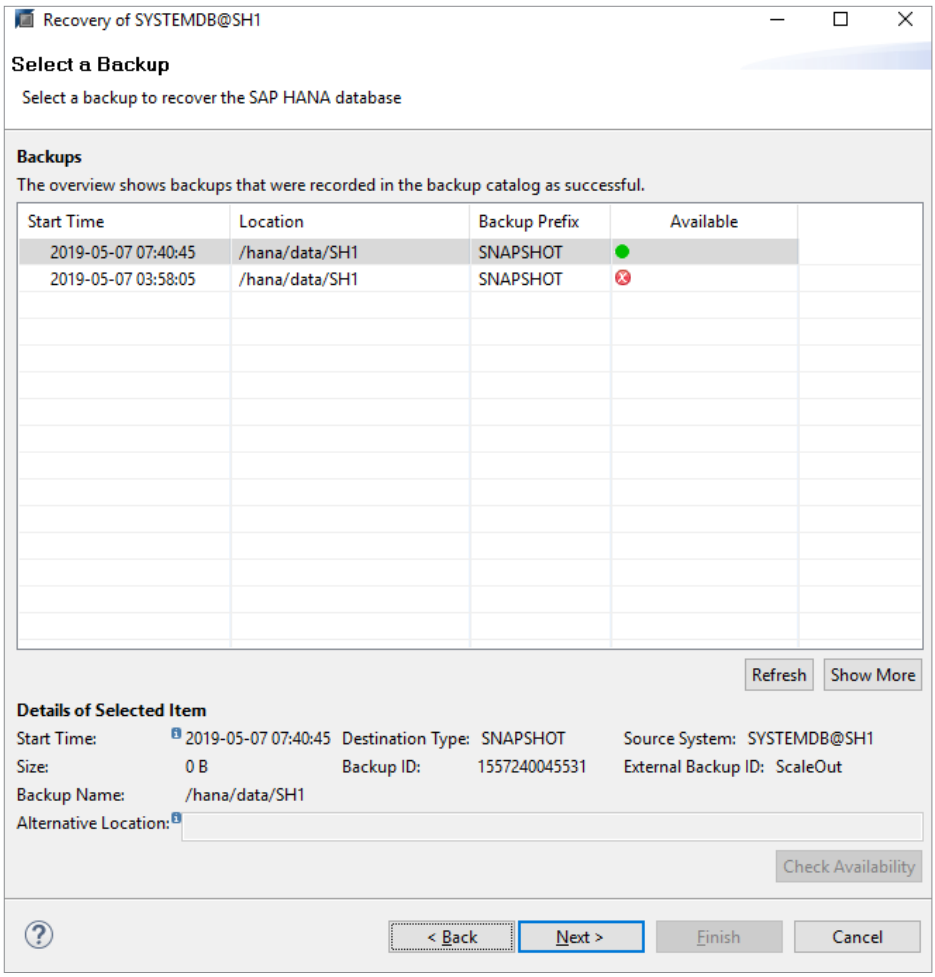

Review the relevant settings for "Initialize Log Area" , "Use Delta Backups" and "Install New License Key". In some scenarios these selections will be greyed out depending on the recovery type.

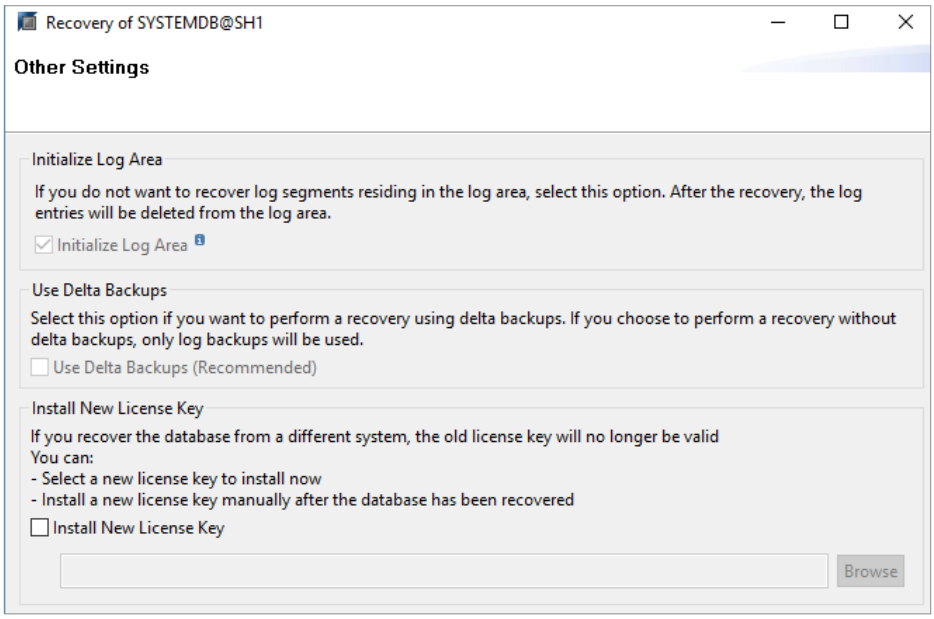

Review all of the recovery settings and proceed by selecting "finish".

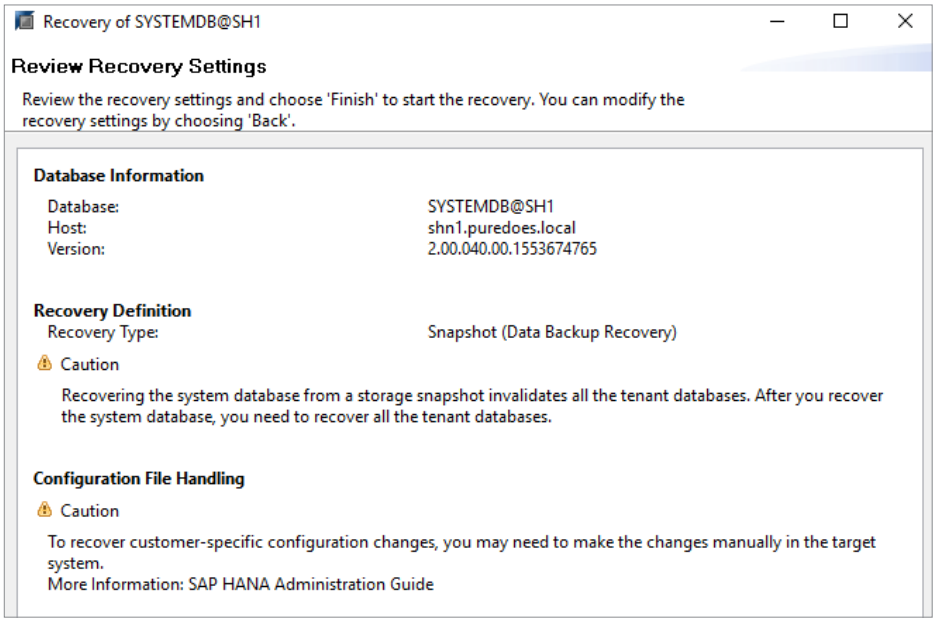

Monitor the recovery process.

Verify the system was restored successfully.

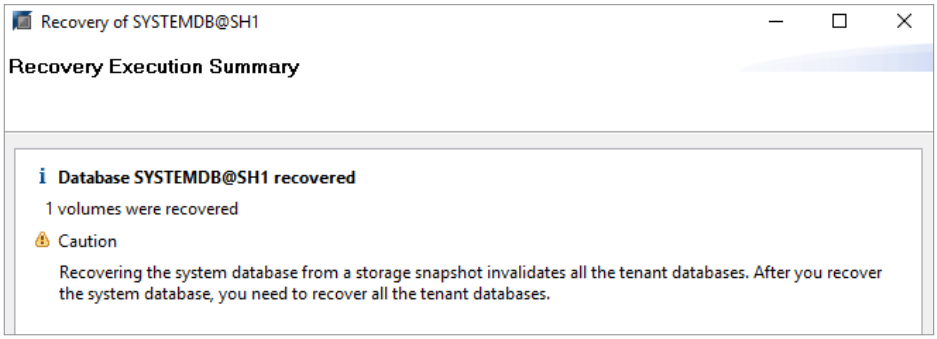

SAP HANA Studio Step 2. Recover Tenant Database

During the recovery of the System database all tenant databases are invalidated and need to be individually recovered by repeating the below steps.

Important: Follow from here if using single container SAP HANA 1.0 system. Notable differences are the absence of “System” and “Tenant” terminology.

Important: In SAP HANA 2.0 SPS04 and onwards repeat these steps for each additional tenant database.

Important: Ensure that Step 1a is completed, while the database is offline, for SAP HANA 1.0 Single tenant systems before proceeding with the below.

Navigate to the main database view in SAP HANA Studio, right-click on the relevant database connection, navigate to "Backup and Recovery" and then Select "Recover Tenant Database…". When using a single container SAP HANA 1.0 system this will simply state "Recover database…"….".

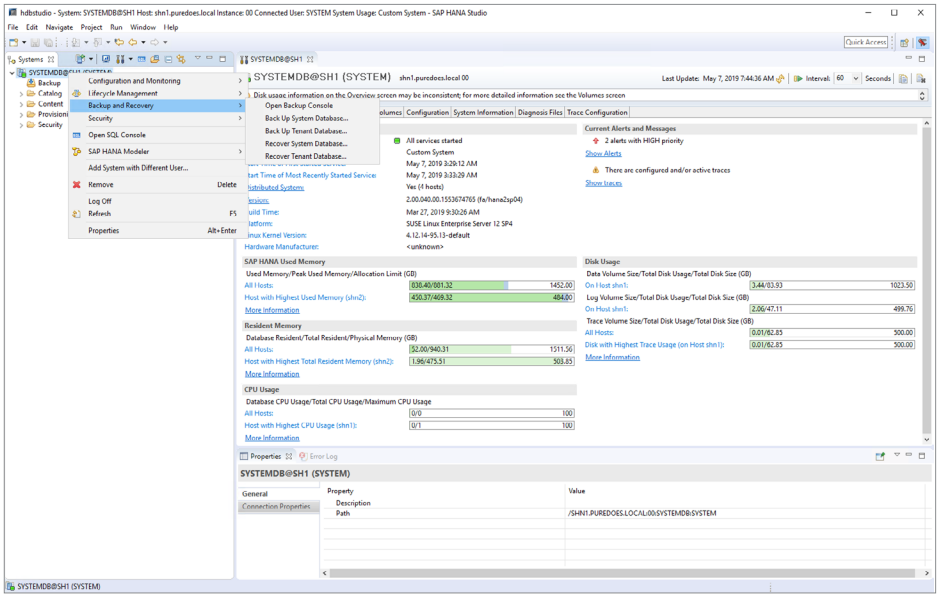

(MDC Only) Select the tenant database to restore. In SAP HANA 2.0 SPS03 and previous releases only a single tenant system can be restored.

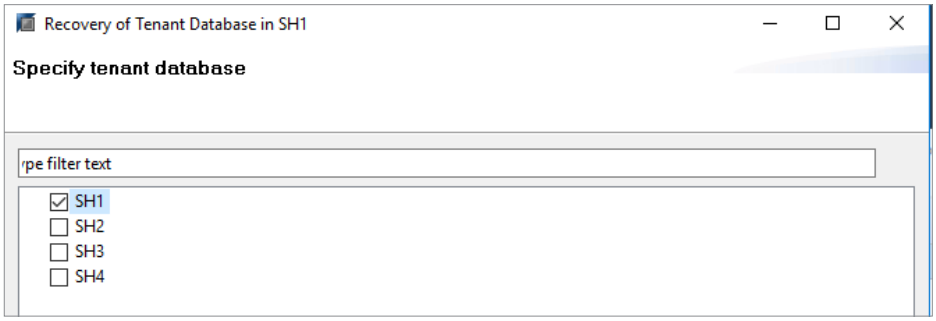

Specify the Recovery Type.

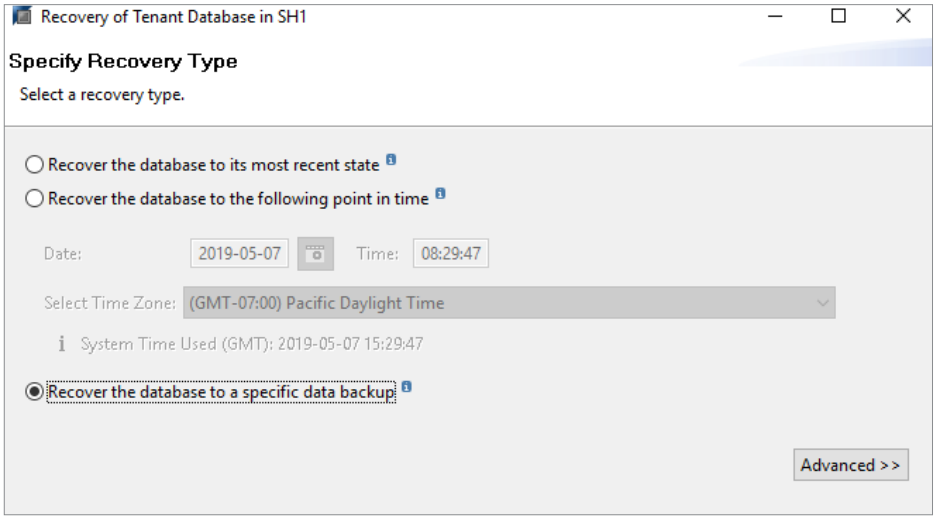

The location of the backup catalog is normally populated, only change this under specialized circumstances where the recovery catalog is in a different location.

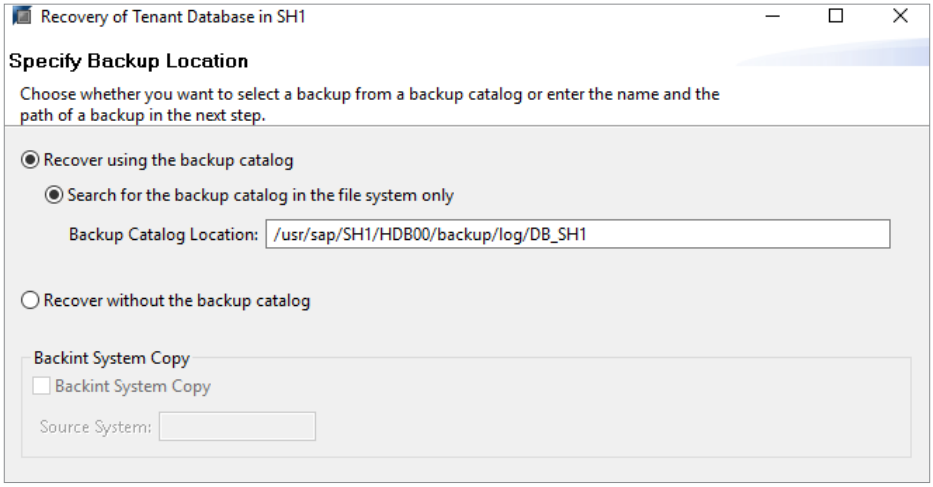

Accept the prompt to shut down the database.

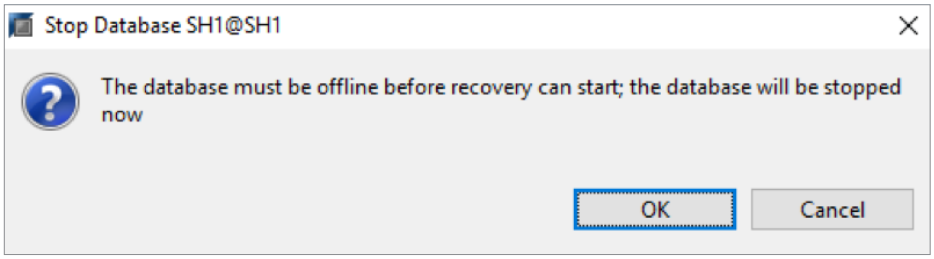

The relevant recovery point should be shown as available, then select “Next “to proceed.

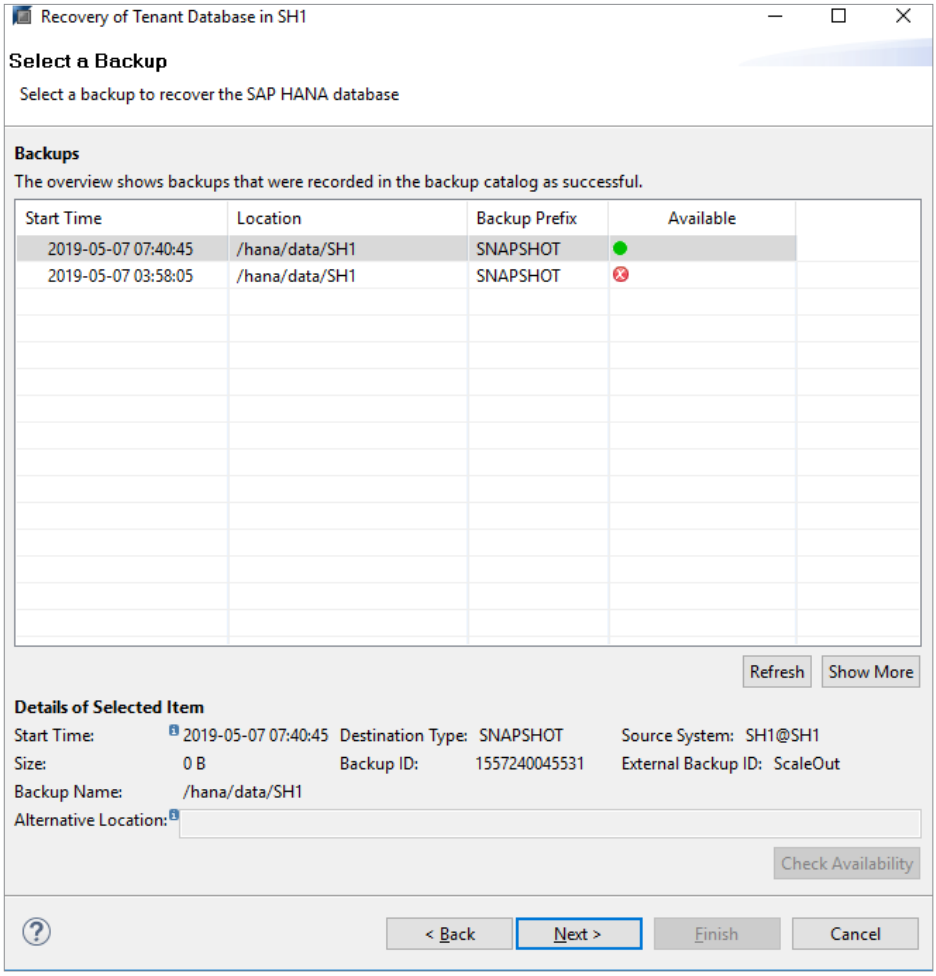

Review the relevant settings for "Initialize Log Area", "Use Delta Backups" and "Install New License Key". In some scenarios these selections will be greyed out depending on the recovery type.

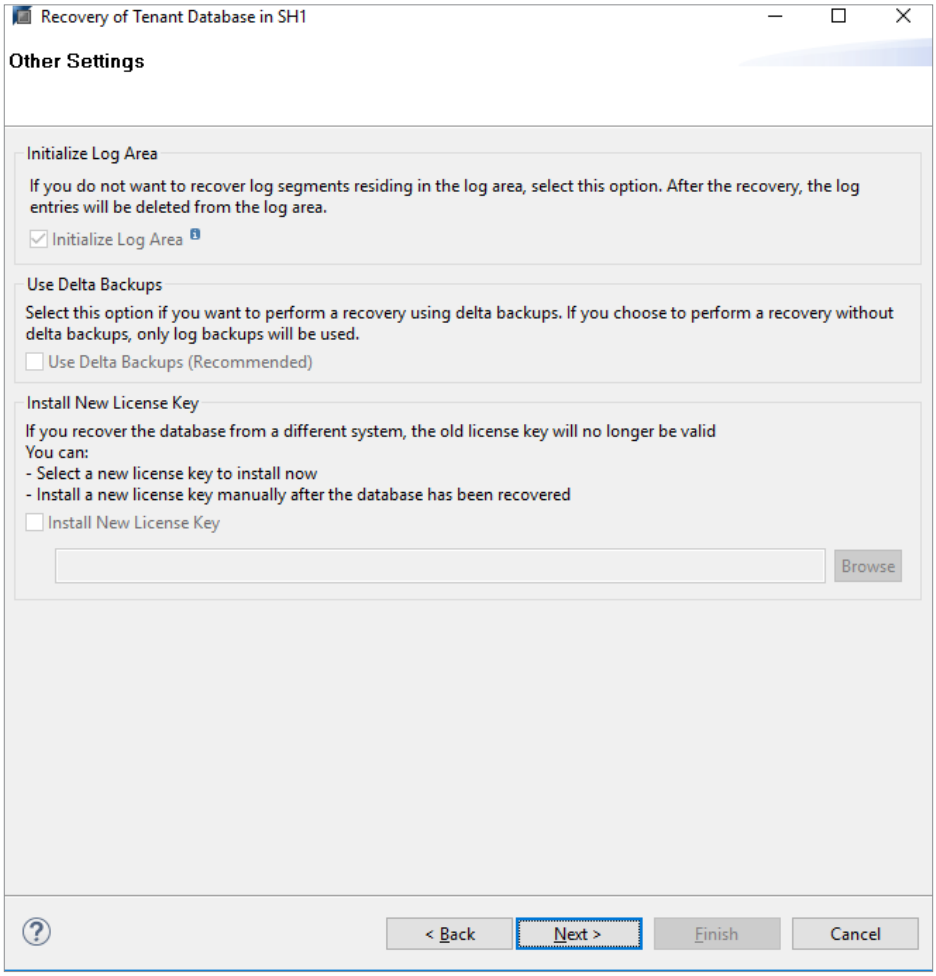

Monitor the data recovery operation.

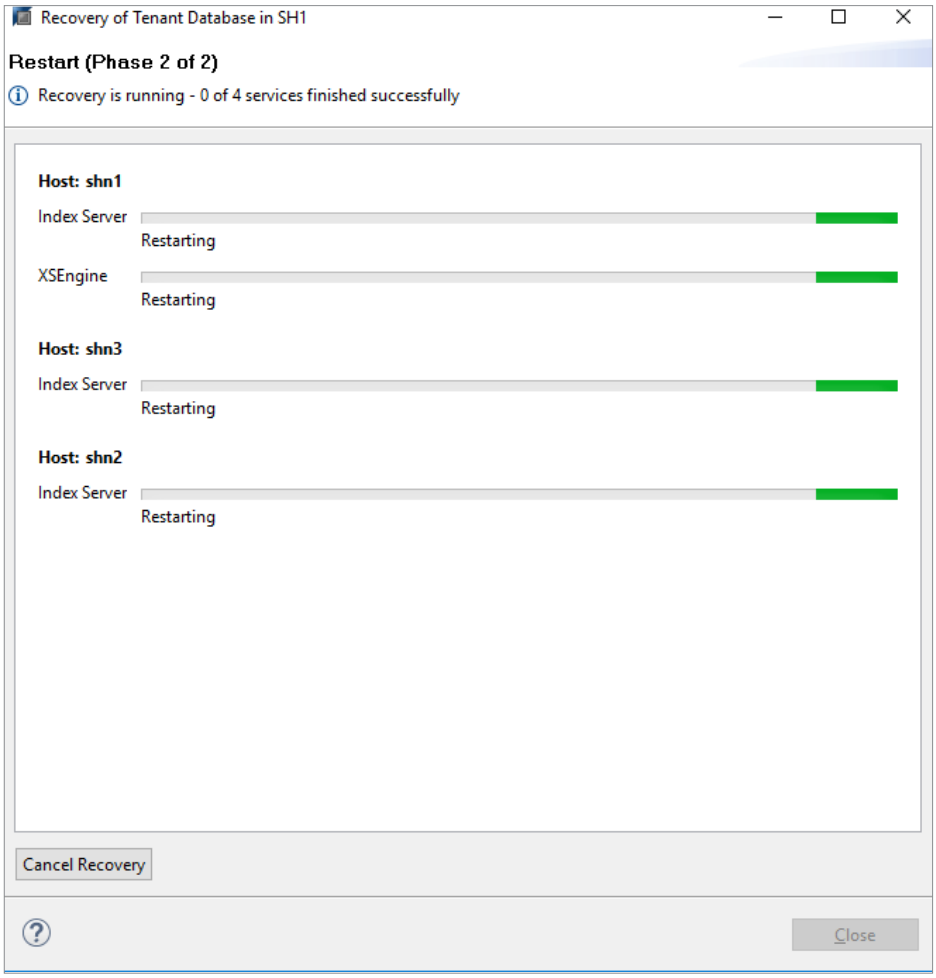

SAP HANA Cockpit Step 1. Prepare to Recover System Database (MDC systems only)

Navigate to the SAP HANA Cockpit web interface and ensure that all of the relevant information is available for it to interface with the relevant deployment to be restored. SAP HANA Systems can be started and stopped from the cockpit interface, where the database will need to be stopped before a recovery can proceed.

Ensure SAP Control credentials are entered and that the resource for HANA cockpit is connected to the SystemDB. Then select the relevant resource for restore operations.

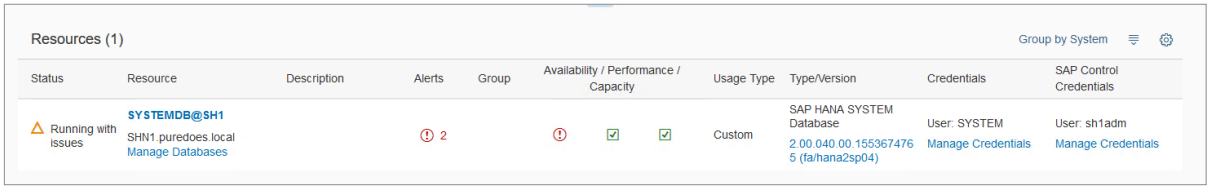

From SAP HANA cockpit choose the overall database status tile and select "stop system".

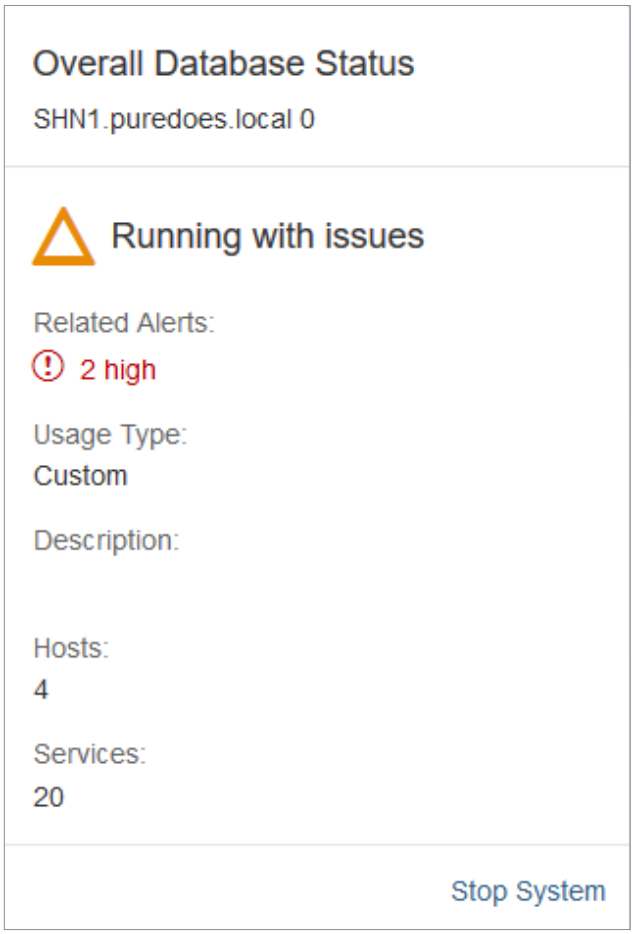

Stop the system using either “softly” or “immediately”.

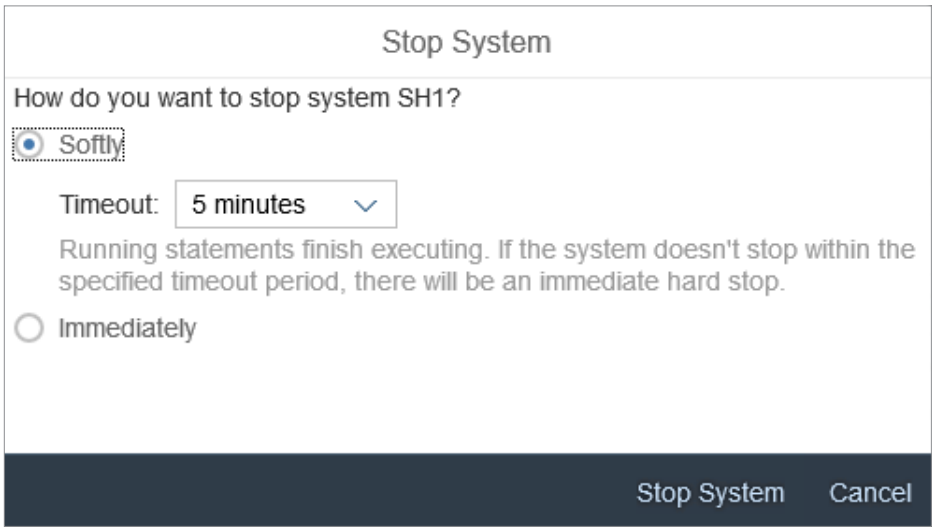

Monitor the system processes until all have been shut down.

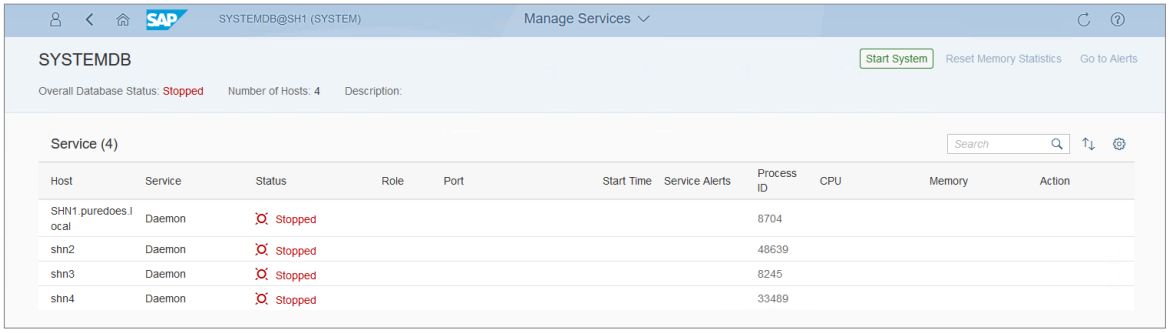

Navigate towards the SystemDB resource page.

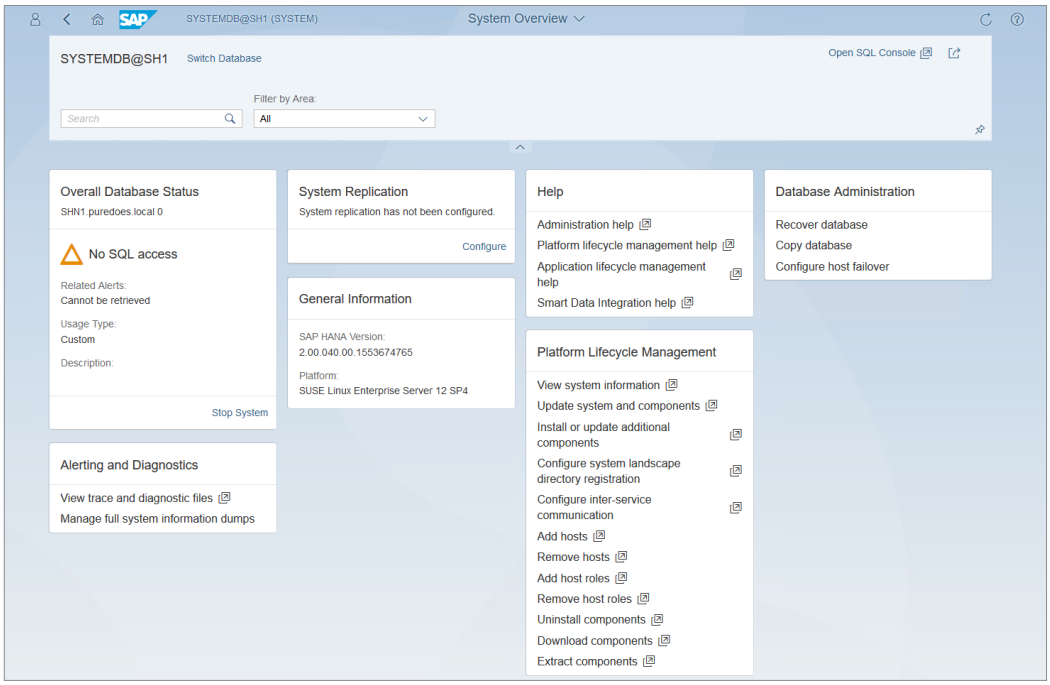

In the "Database Administration" tile, select "Recover database".

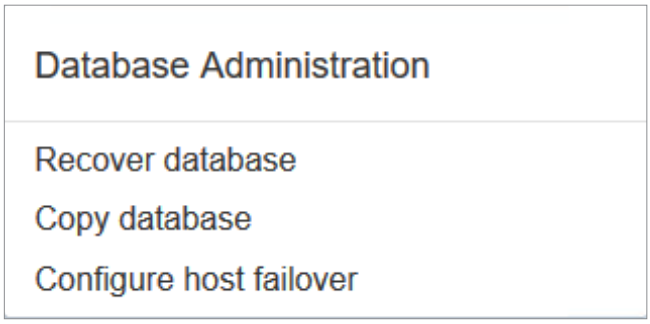

Select the recovery target type.

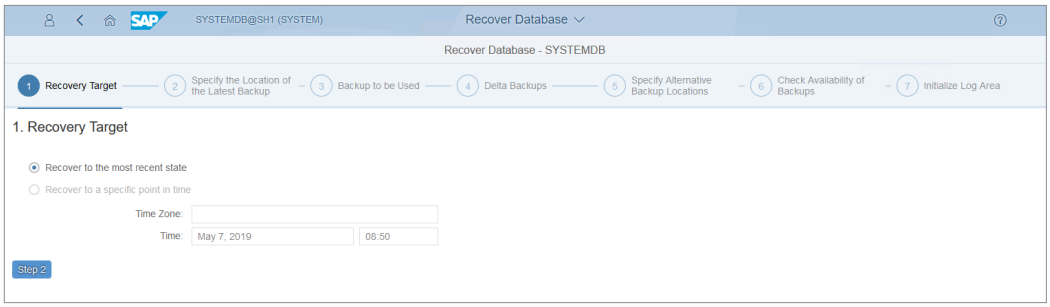

Specify the location of the latest backup catalog, only change these values under specialised circumstances.

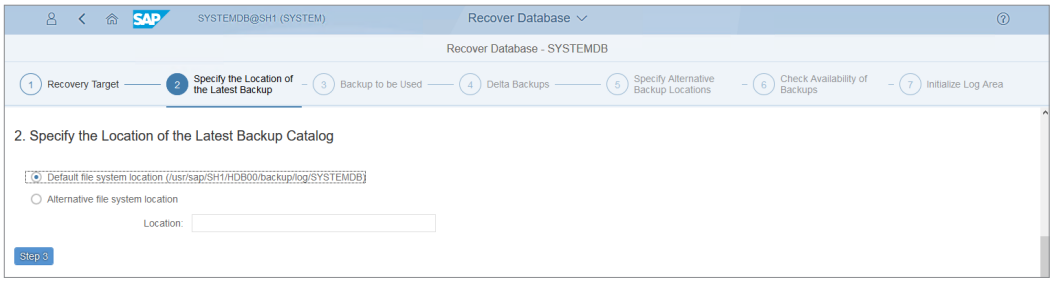

Note the "Backup ID" of the backup to be used and its availability.

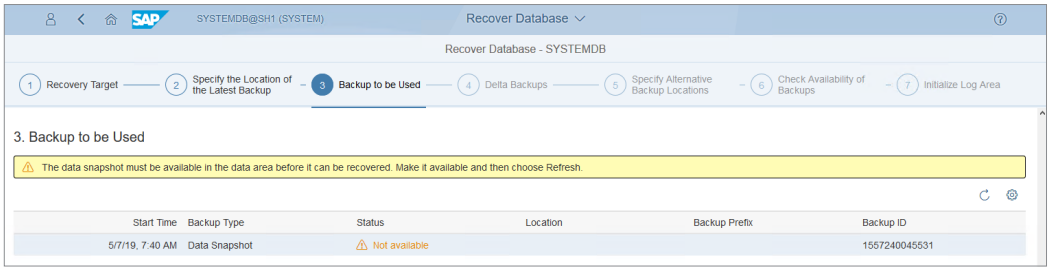

SAP HANA Studio Step 1a. Recover block storage volume

Recover the relevant storage volume using the FlashArray web user interface. This operation can also be completed using the command line interface or ReST API.

Important: The persistence data volumes must be unmounted from the operating system before a snapshot is restored. This can be done using the "umount" command in a terminal or SSH connection.

In the FlashArray web user interface navigate to "Storage" and select the "Volumes" section.

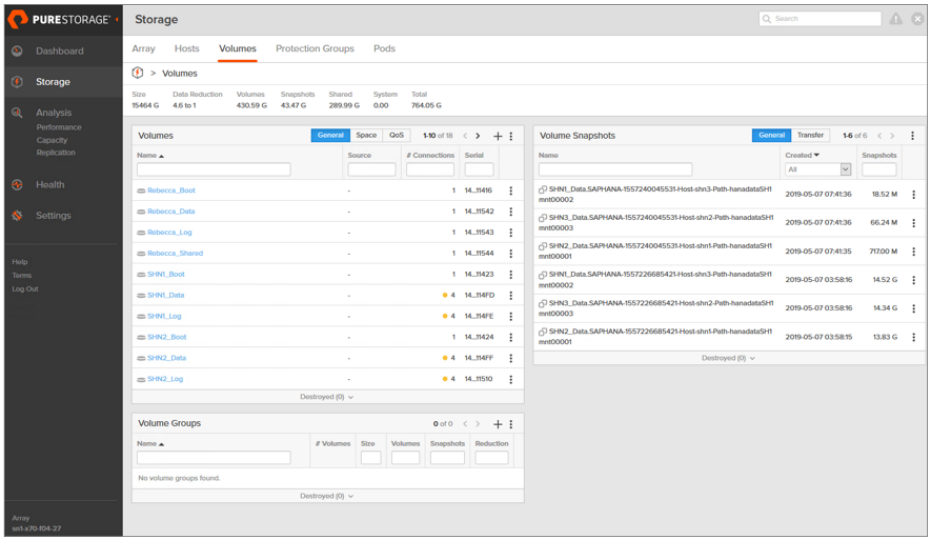

Important: The snapshot which matches the "Backup ID" in the catalog is appended with the exact same value, this is the volume which should be restored. Note that the addition of this value is done during storage snapshot creation by the user.

Identify the snapshots relevant to the Backup_ID of the recovery point.

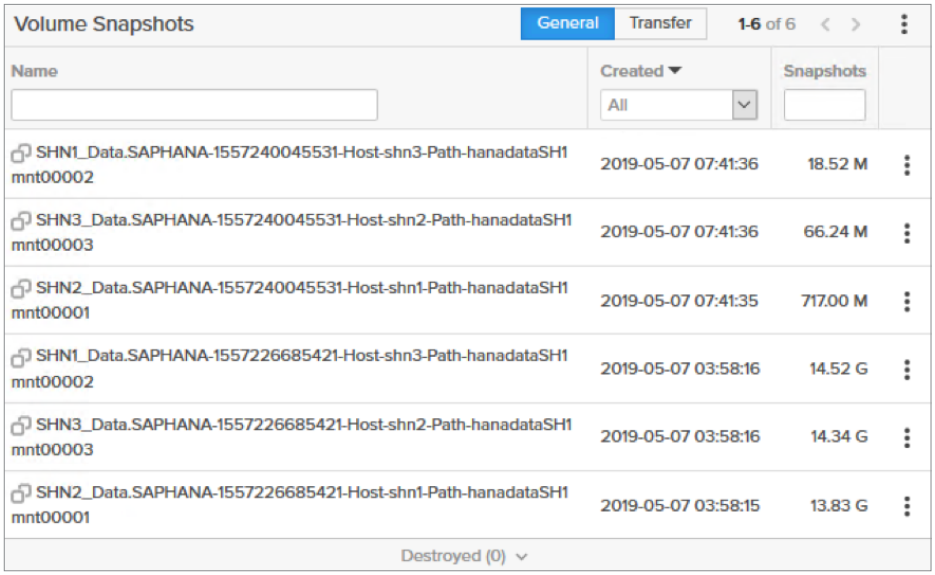

Unmount each persistence data volume from the mount point on each host.

```
SHN1:~ # umount /hana/data/SH1/mnt00001
SHN2:~ # umount /hana/data/SH1/mnt00003
SHN3:~ # umount /hana/data/SH1/mnt00002
```

Select the 3 dots on the right-hand side of the relevant snapshot" and select "Restore" for each required block device. A prompt will appear for each snapshot, check the information and then confirm the restore operation.

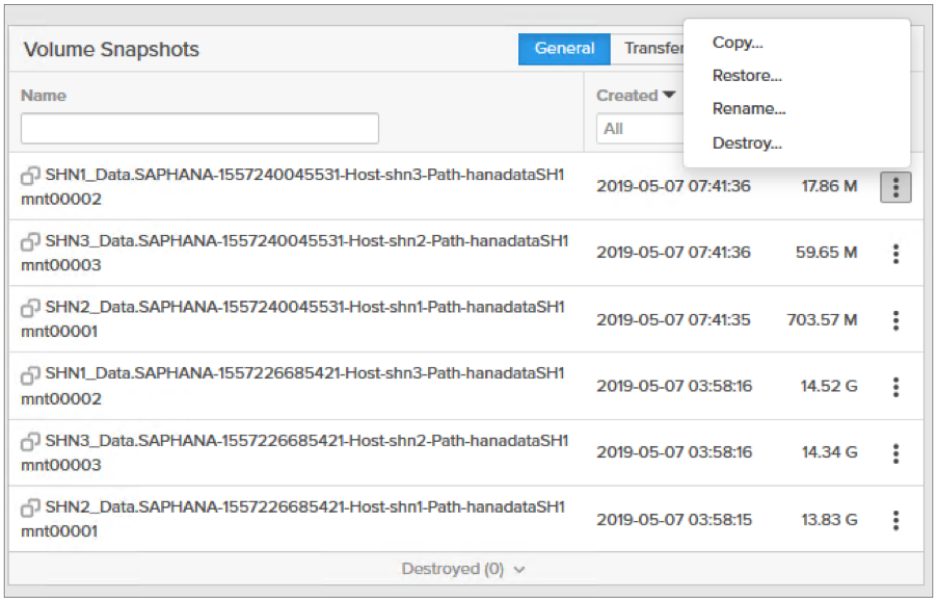

Verify that the volumes have restored correctly.

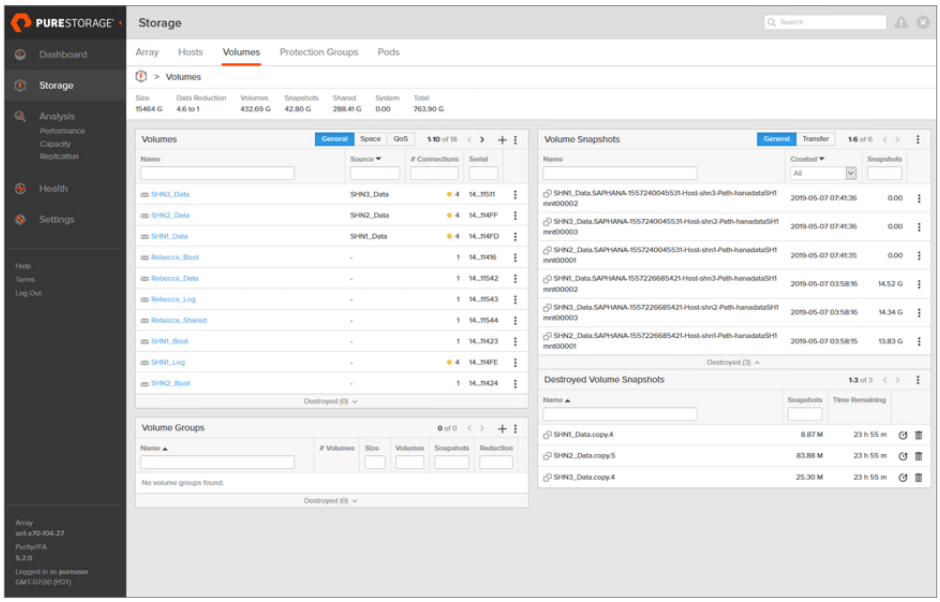

Navigate to each volume and check the details match what is required of the SAP HANA scale out system.

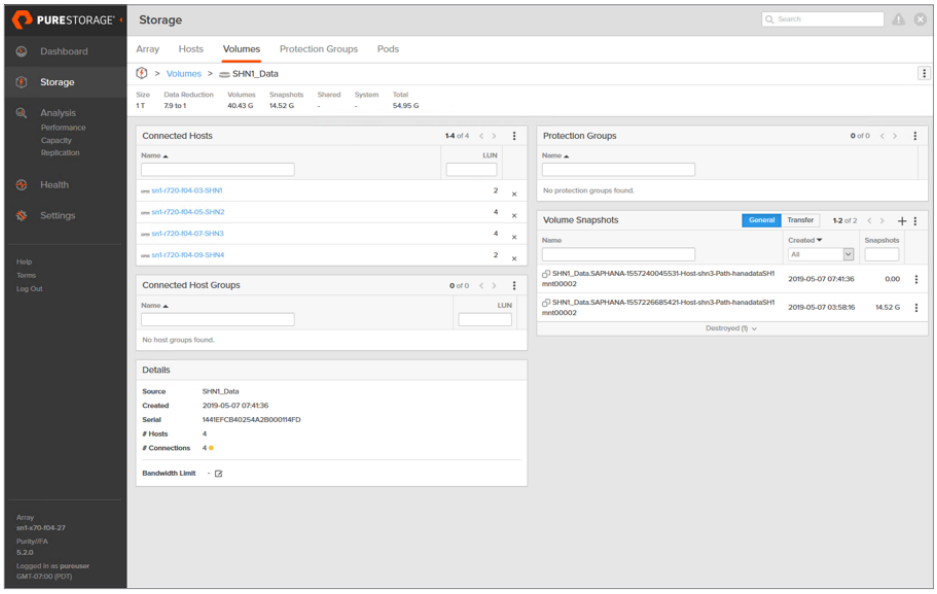

Check the details for each volume required to attach to any host in the distributed host setup.

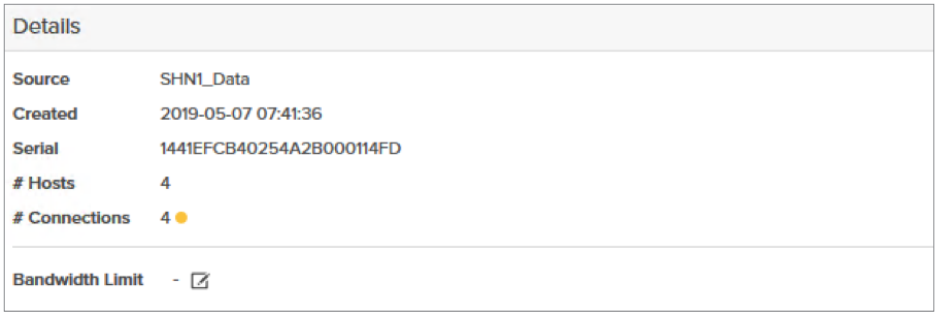

Check the details for each volume required to attach to any host in the distributed host setup.

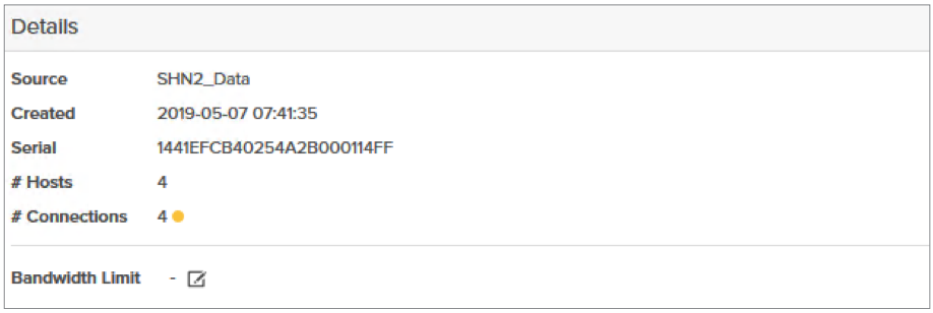

Check the details for each volume required to attach to any host in the distributed host setup.

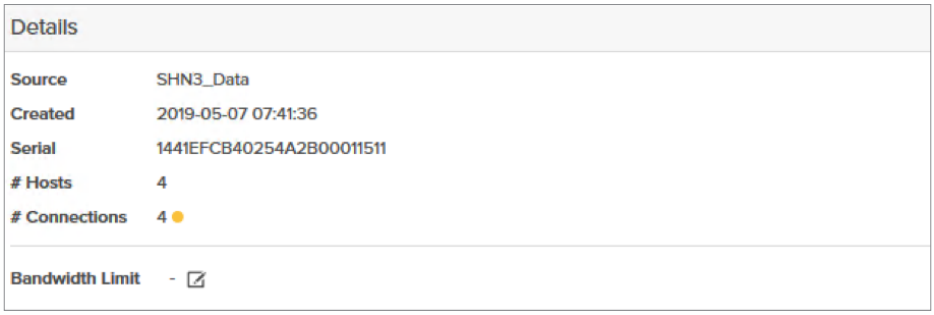

SAP HANA Cockpit Step 1b. Recover the System Database

Select the refresh icon in the top right-hand corner under Backups to be Used to rescan the backup catalog.

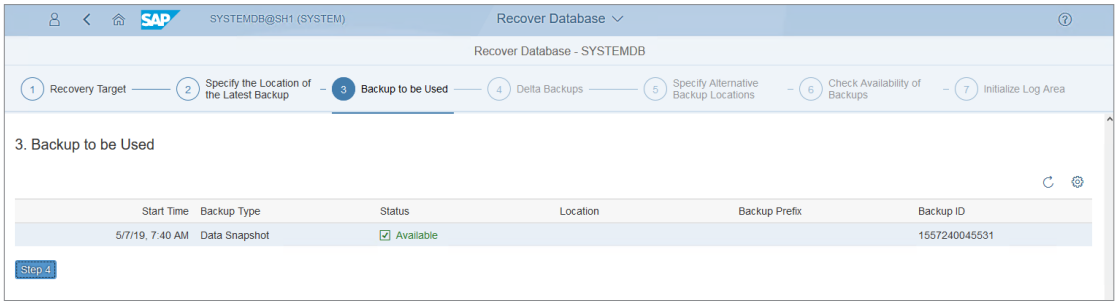

Select the preferred value for Delta Backups.

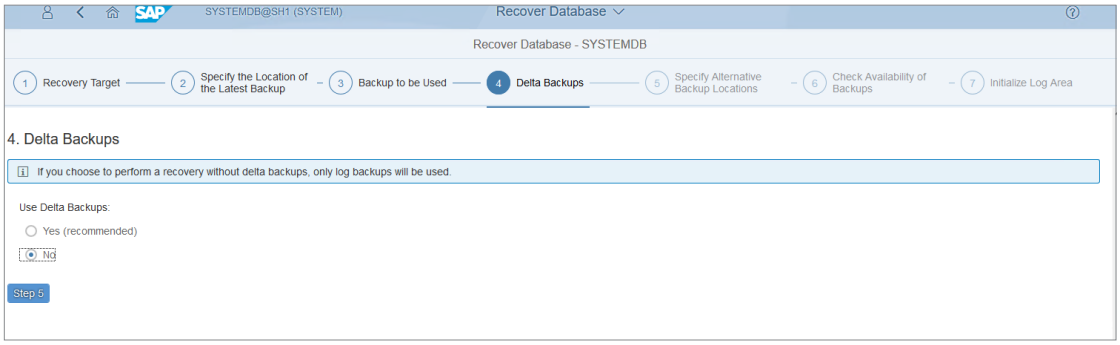

Specify any alternative backup locations.

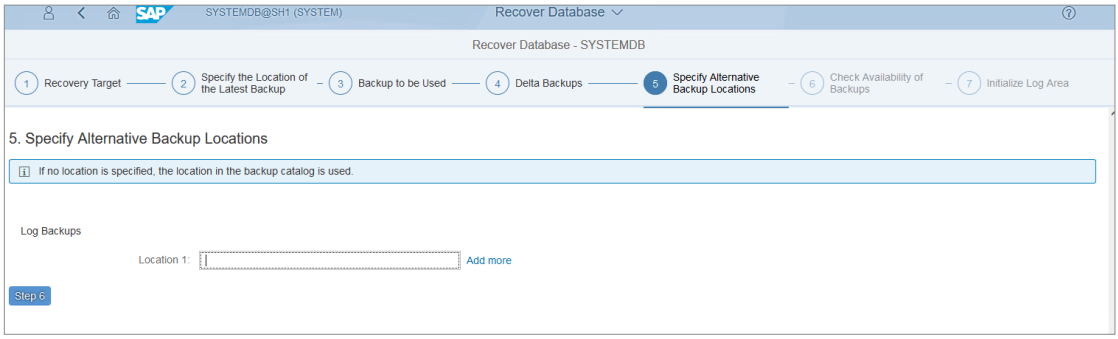

Specify if the availability of backups should be checked.

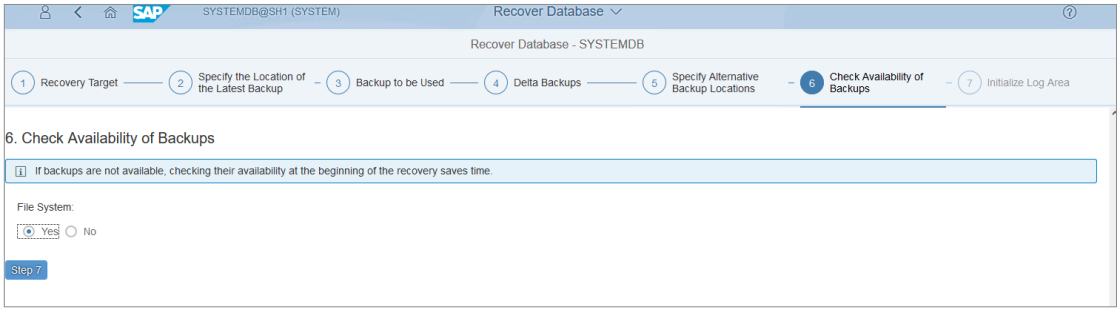

Specify if the log area should be initialized, doing so invalidates any logs or log backups made after the recovery point.

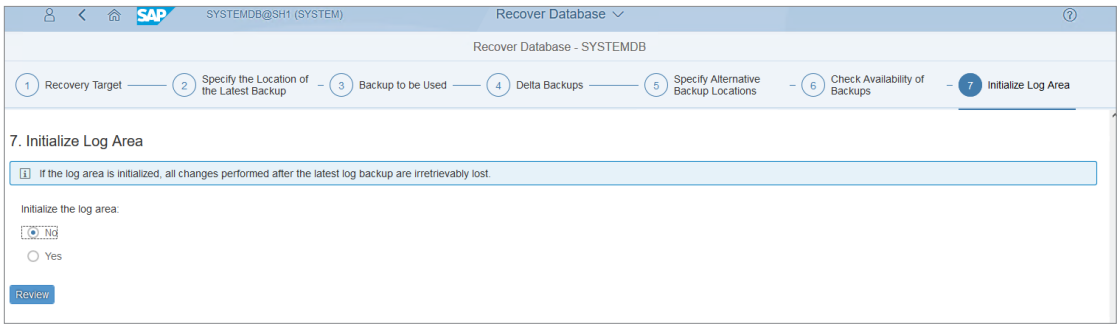

Review the information for the recovery operation and then select “Start Recovery”.

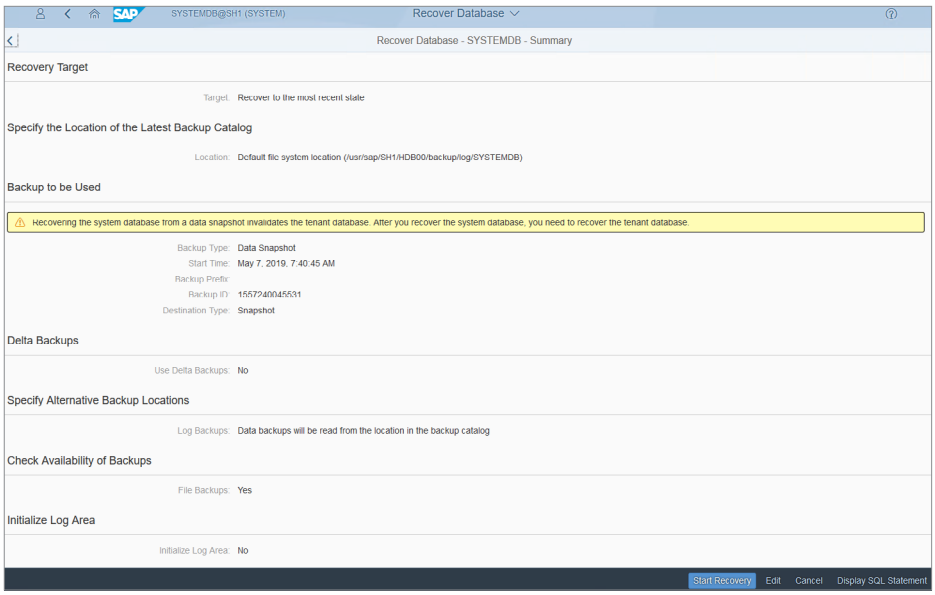

While restoring the System Database the status can be observed from the same view.

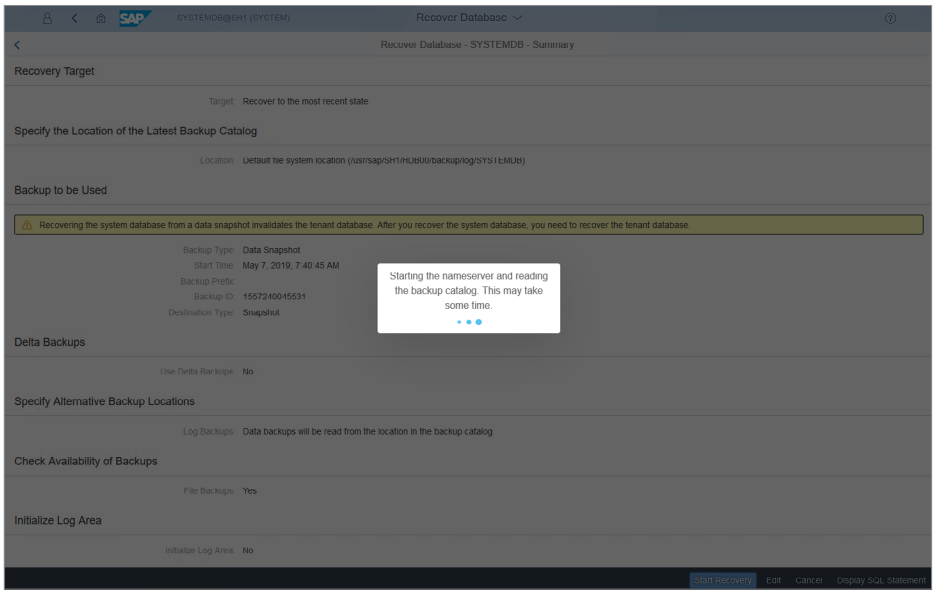

SAP HANA Cockpit Step 2. Recover Tenant Database

During the recovery of the System database all tenant databases are invalidated and need to be individually recovered by repeating the below steps.

Important: Follow from here if using single container SAP HANA 1.0 system. Notable differences are the absence of “System” and “Tenant” terminology.

Important: In SAP HANA 2.0 SPS04 and onwards repeat these steps for each additional tenant database.

Important: Ensure that Step 1a is completed, while the database is offline, for SAP HANA 1.0 Single tenant systems before proceeding with the below.

Once the System Database has been recovered select “Manage Databases” from the SystemDB database view. All of the tenants will be shown, but in an offline state. In SAP HANA SPS03 and previous releases only a single tenant system can be recovered.

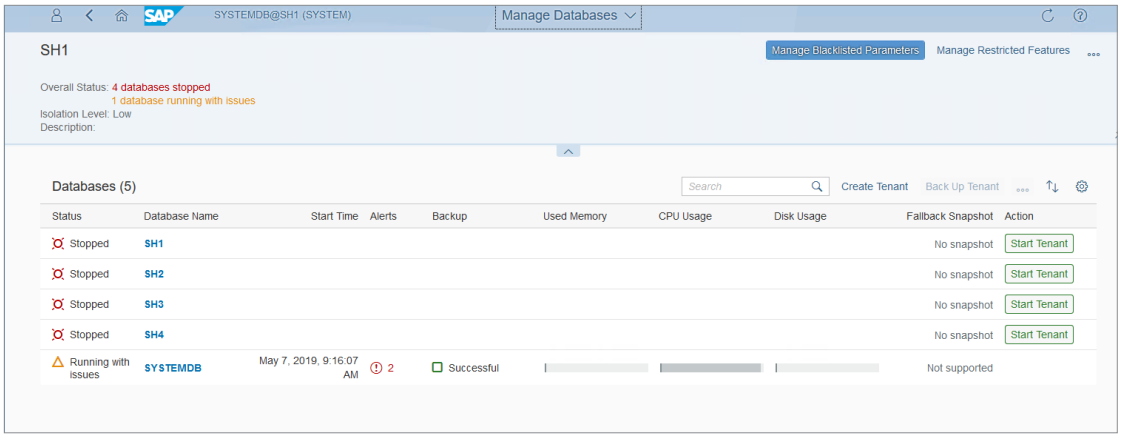

Highlight the relevant tenant database to be recovered and then select the three dots in the right-hand corner to bring up the context menu, then select “Recover Tenant”.

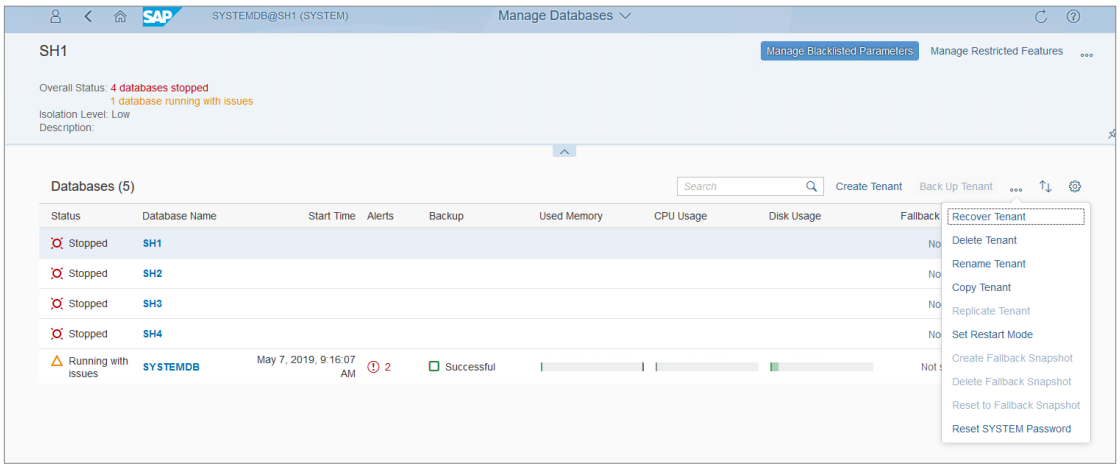

Select the recovery target type.

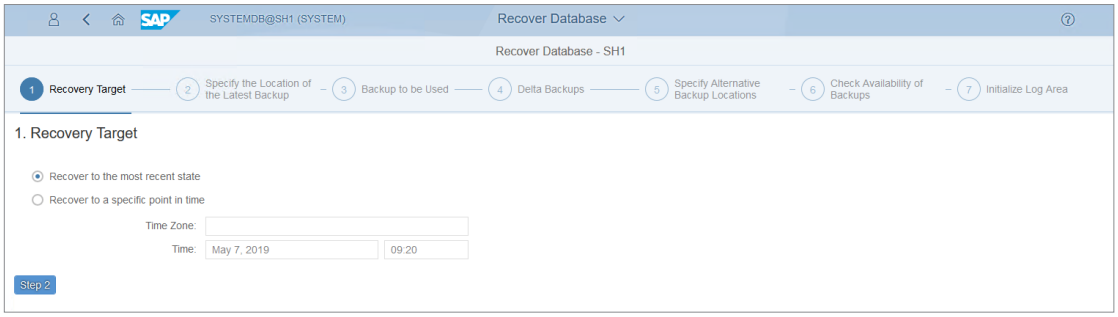

Specify the location of the latest backup catalog, only change these values under specialised circumstances.

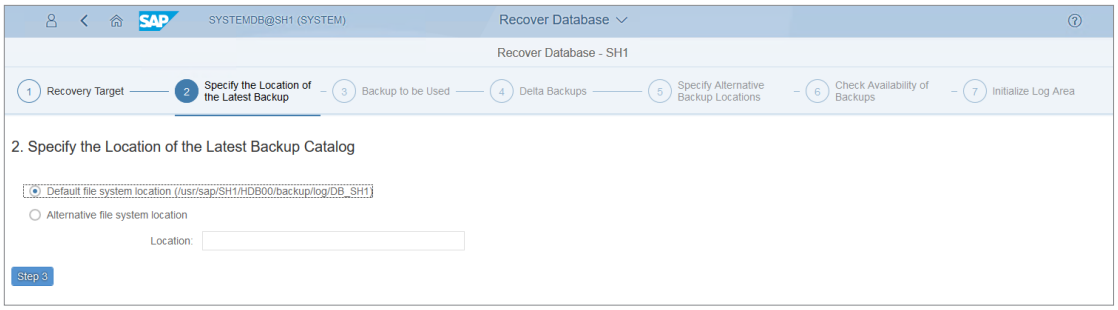

The relevant backup to be used will have the status of "Available", select this and then select "Step 4".

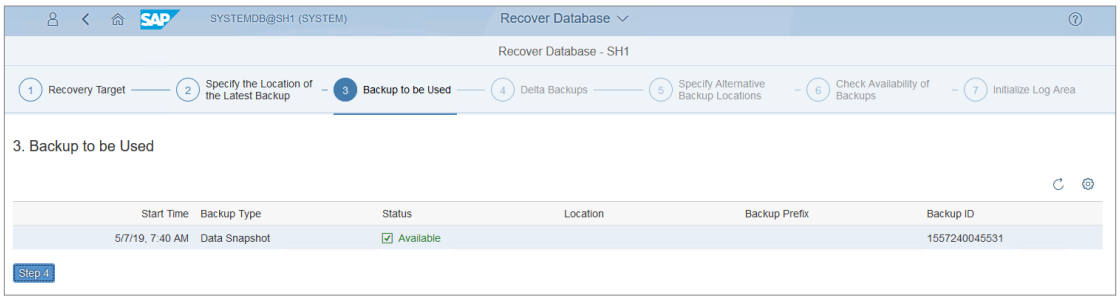

Specify if Delta Backups are to be used.

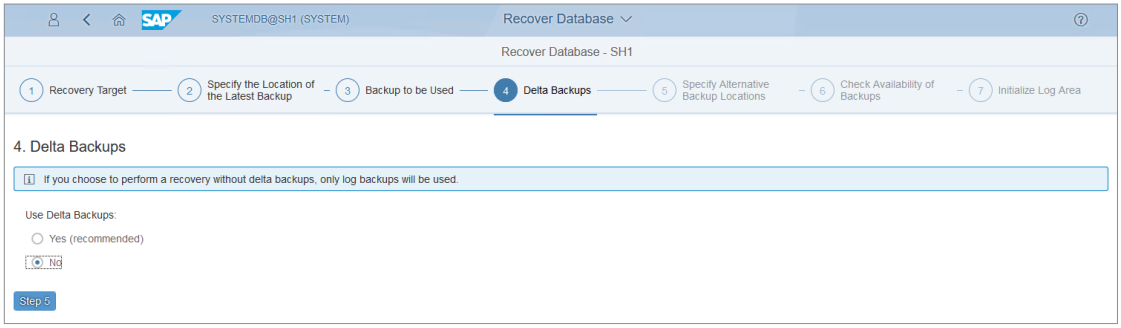

Specify an Alternative Backup location if any exist.

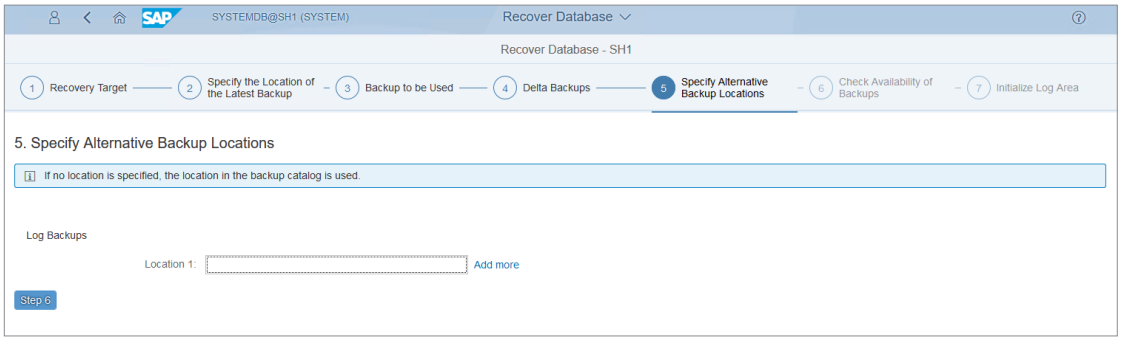

Specify if the Availability of Backups must be checked.

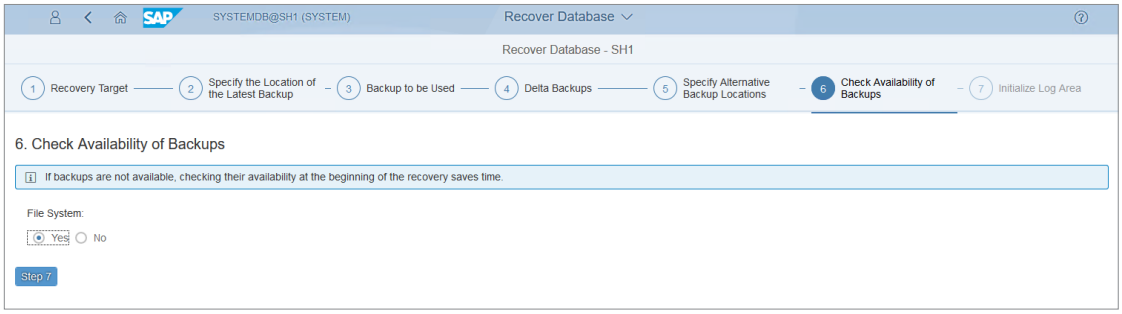

Specify if the log area must be initialized.

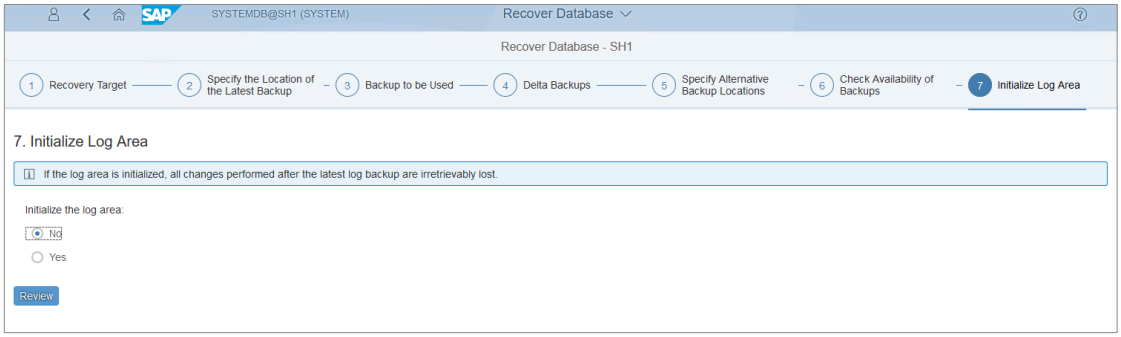

Review the information for the recovery operation and then select “Start Recovery”.

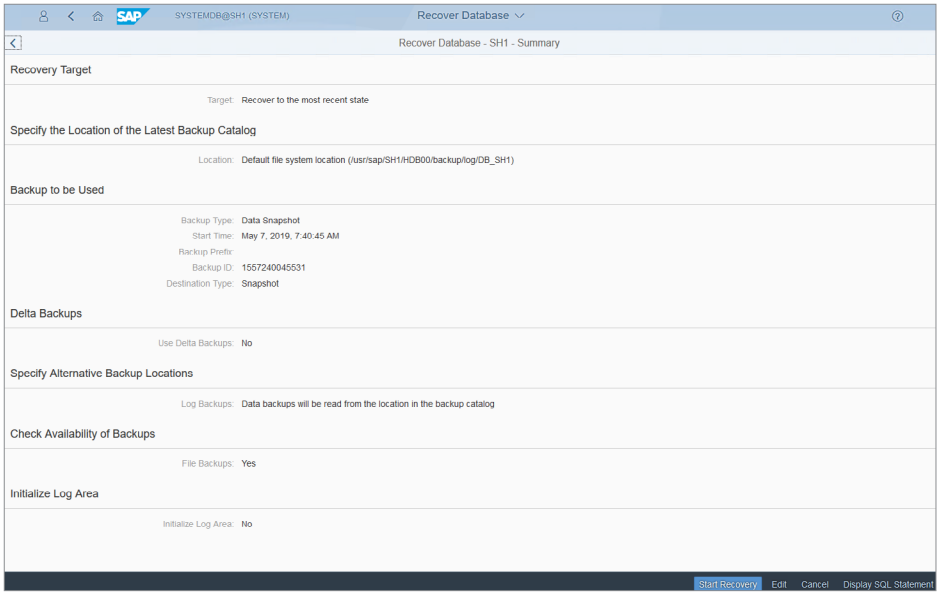

Once the tenant has been successfully recovered the status will be shown.

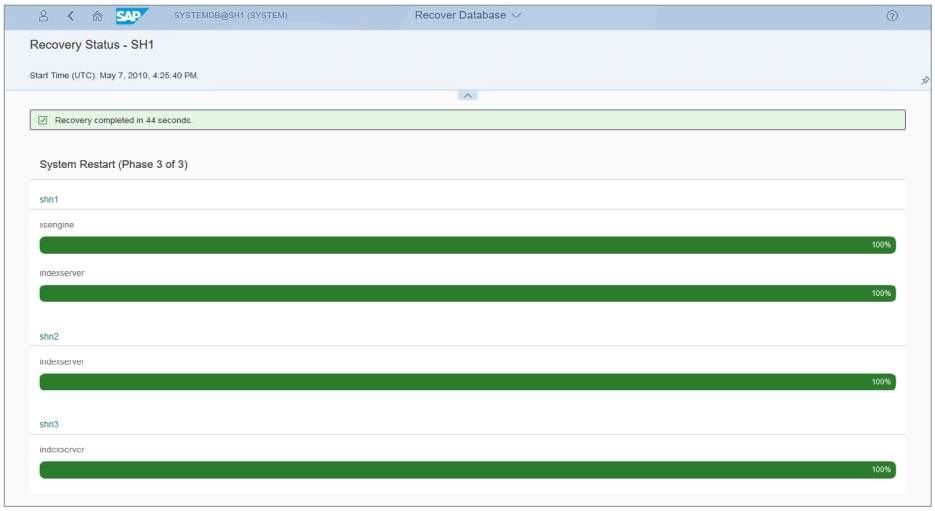

SAP HANA Deployment on VMware ESXi Hypervisors

In many scenarios it is appropriate to use snapshots of the virtual machine and to recover by simply restoring this snapshot. In other specialised scenarios it is possible to combine the recovery catalog present in SAP HANA with virtual machine snapshots and storage snapshots. Best practices for using snapshots in a VMware vSphere environment indicate that virtual machine snapshots should not be used as backups, instead ensure that snapshots are used in conjunction with an independent software vendor for backup to move the virtual machine snapshot to a different location. Virtual machine snapshots should not be retained for more than 72 hours due to storage space management and performance degradation.

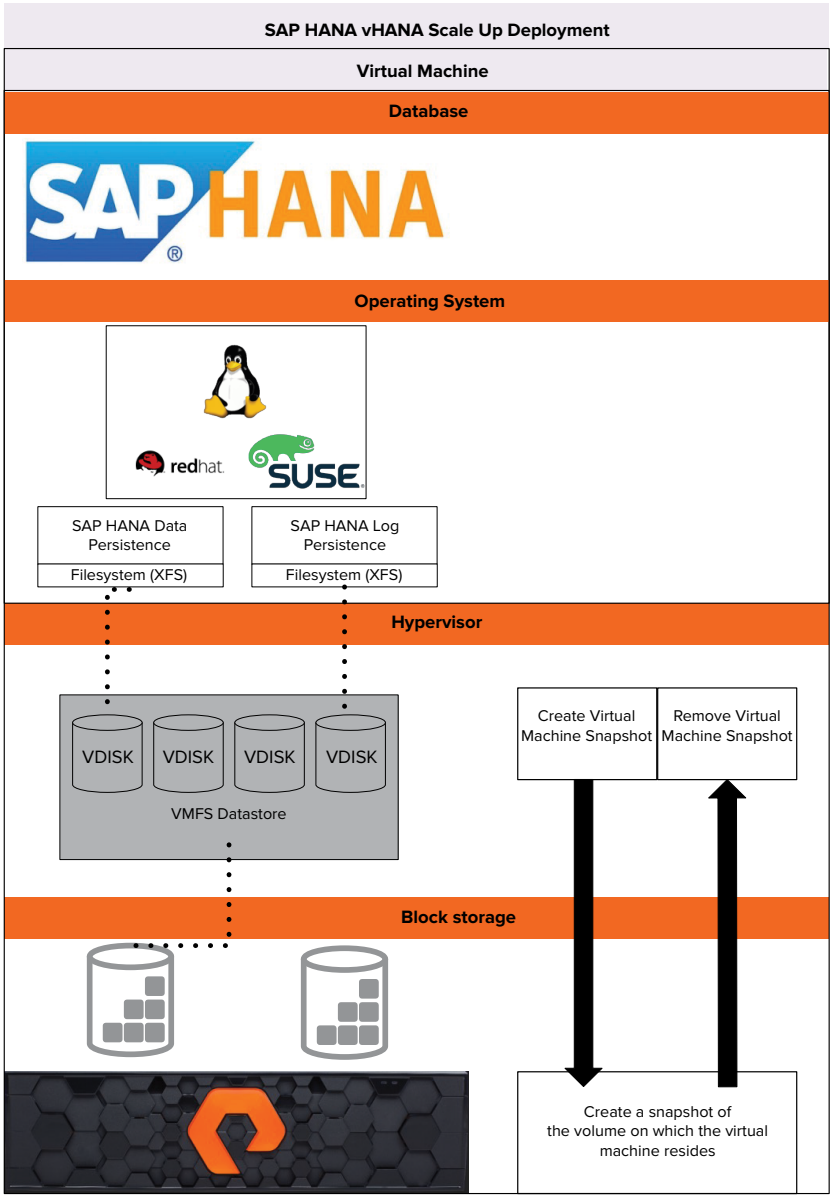

Fig. 11 Workflow to create a storage snapshot for SAP HANA database deployed on VMware vSphere.

With FlashArray and the Purity operating environment, storage snapshots containing a previously snapshotted virtual machine can be considered a point-in-time recovery point, or if the storage snapshot is transported using Snap-to-NFS or Snap-to-S3 then it can be considering a backup.

VIRTUALISED SAP HANA PROTECTION USING VMWARE VSPHERE SNAPSHOTS

In order to create a virtual machine snapshot a user with the necessary permissions should be logged into the vSphere web client. Navigate to the VMs and Templates view, select the relevant SAP HANA system and right-click to bring up the context menu and the select “Take Snapshot…”

VMware vSphere web client with the context menu for a SAP HANA virtual machine.

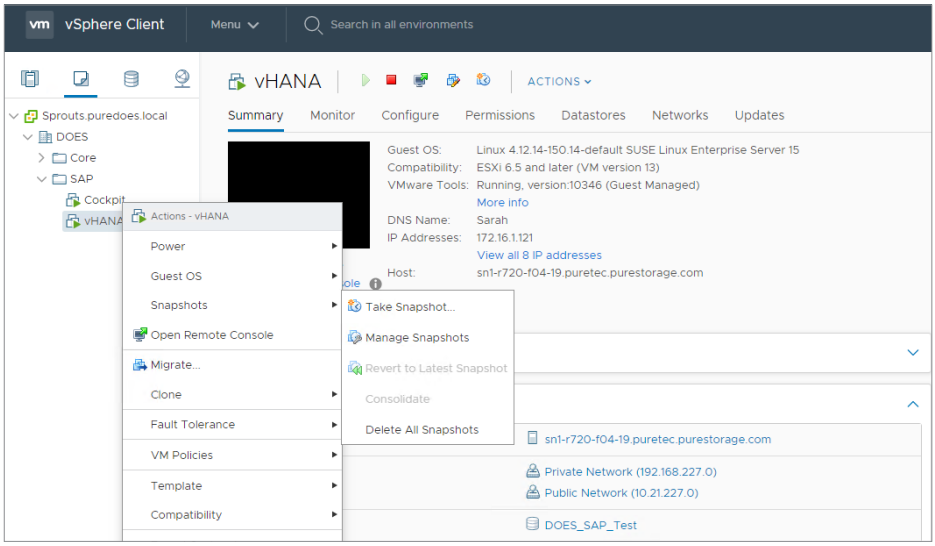

A prompt will show waiting for the relevant information to be entered for a snapshot to be created. If a stateful snapshot of the virtual machine needs to be created ensure that “Snapshot the virtual machines memory” is selected, otherwise unselect it. Creating a snapshot of the virtual machine’s memory can take longer and impact performance.

Prompt waiting for the user to enter specifics about the SAP HANA virtual machine.

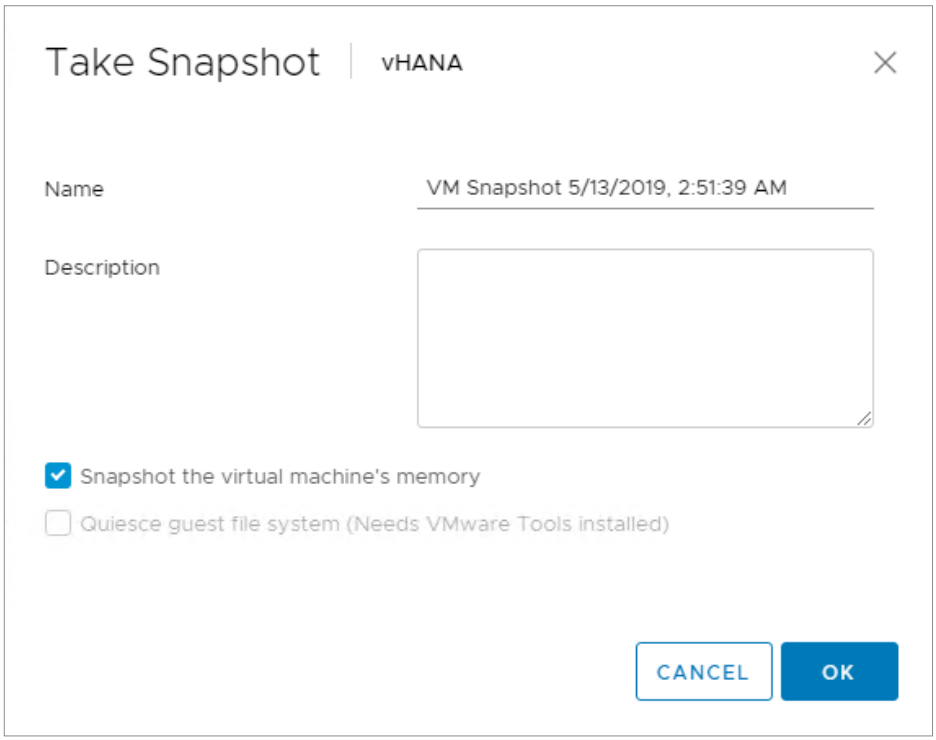

Once the snapshot is created navigate to the FlashArray web user interface and create a storage snapshot of the relevant block storage volume.

In the FlashArray web user interface navigate to the Storage view and select Volumes, then select the volume(s) on which the virtual machine resides.

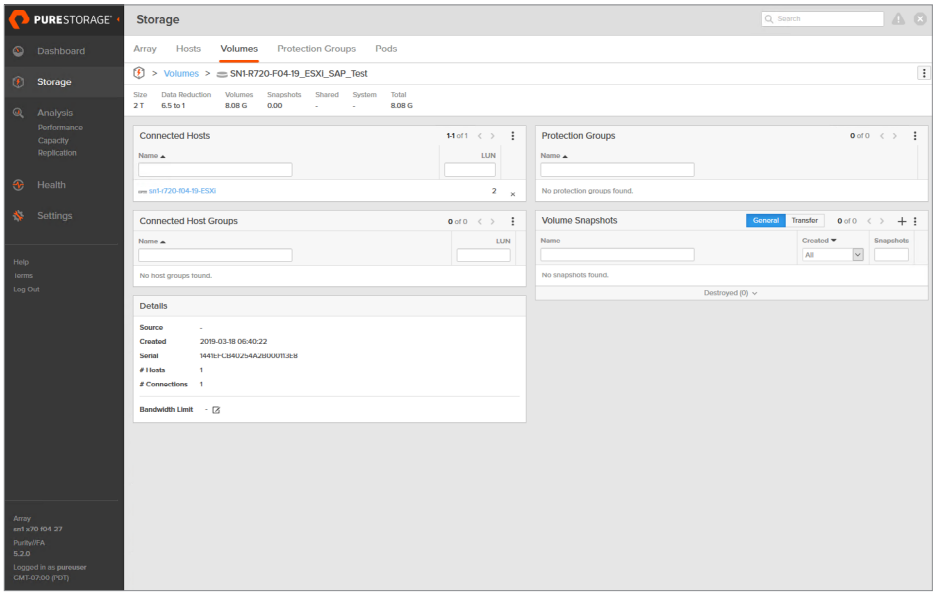

Select the three dots in the upper right-hand corner of Volume snapshots and select "Create…"….".

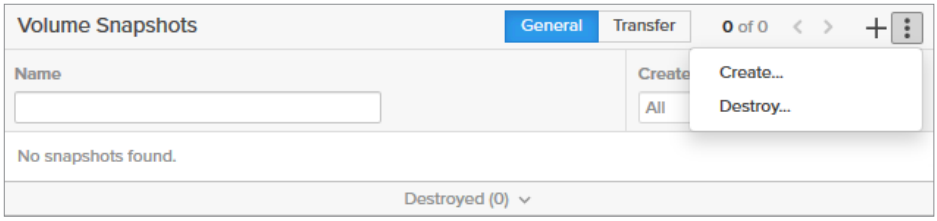

Give the snapshot an optional suffix and then select "Create".

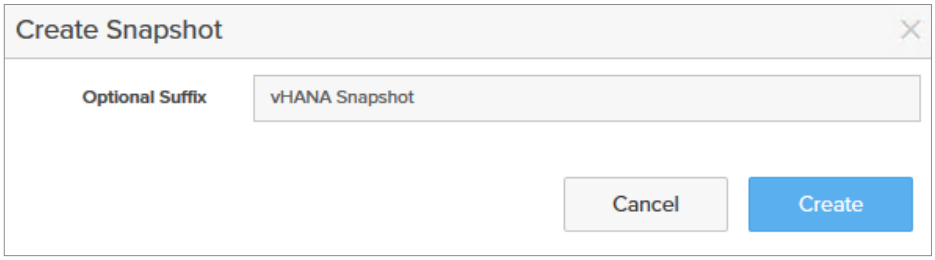

Once the block volume snapshot has been created the virtual machine snapshot must be discarded. This can be done using the vSphere web client, selecting the VMs and templates view and then bringing up the context menu for the vHANA virtual machine and selecting "Snapshots" and then further selecting "Delete All Snapshots".

In the vSphere web client bring up the context menu for the SAP HANA virtual machine and delete all of the snapshots.

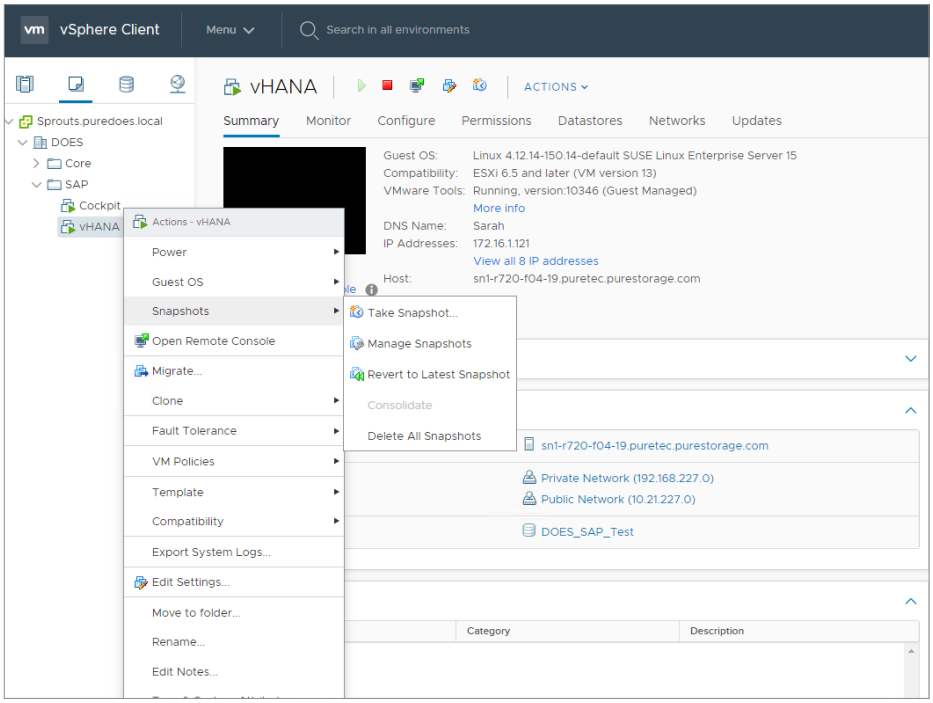

Select OK from the prompt to delete all of the snapshots for the SAP HANA virtual machine.

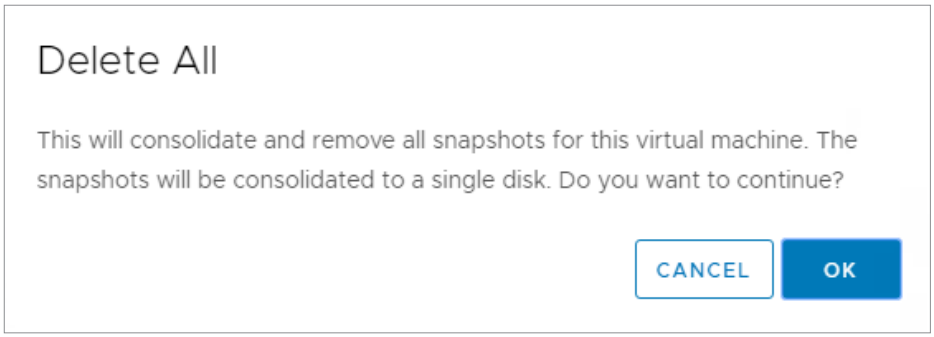

The exact same workflow as set out in the section “Data Protection Solutions for SAP HANA, SAP HANA Scale Up/ Single Host System” can be applied to accomplish application consistent storage snapshots, with the exception that the volume(s) on which the SAP HANA instance is deployed must be identified beforehand. This applies to any implementation of virtual volumes (VVOLS) as well but differs that each virtual disk attached to the virtual machine will have a corresponding block volume in the FlashArray.

Crash Consistent Storage Snapshots

Overview

The creation of a storage snapshot without any application or system coordination to create a record of the operation and ensure that no data is written to the storage device while the operation executes is typically named a crash consistent storage snapshot. It is possible to use this snapshot method to create a point in time recovery point on the storage array and then use it to create system copies or rollback points during development operations. Due to the complex nature of SAP HANA scale out/ distributed host systems only scale up/single host crash consistent snapshots are explored.

The recommended way of creating storage snapshots for production systems would be to use application consistent storage snapshots, crash consistent snapshots do not create a catalog entry for the recovery point nor does it guarantee consistency for a point in time recovery.

Configuration and operation

Creating crash consistent snapshots

The recovery point, when created, must at minimum create a snapshot for both the log and data volumes. Including the /hana/shared (install path) volume is also possible but does not affect recovery. It is optional to have the instance and any tenant database(s) online while creating the storage snapshot. If a degree of data consistency is required then it is advised that “ALTER SYSTEM SAVEPOINT” be run from the database and any tenant databases present on the instance.

In the FlashArray web user interface navigate to the Storage view and select Volumes.

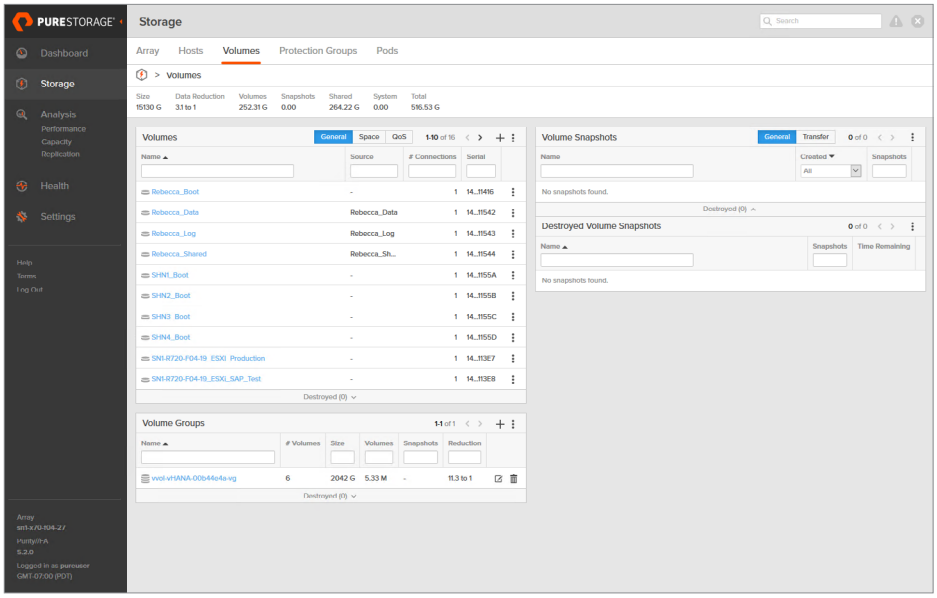

In the volumes view select the vertical ellipsis in the top right-hand corner and select "Create Snapshots..."

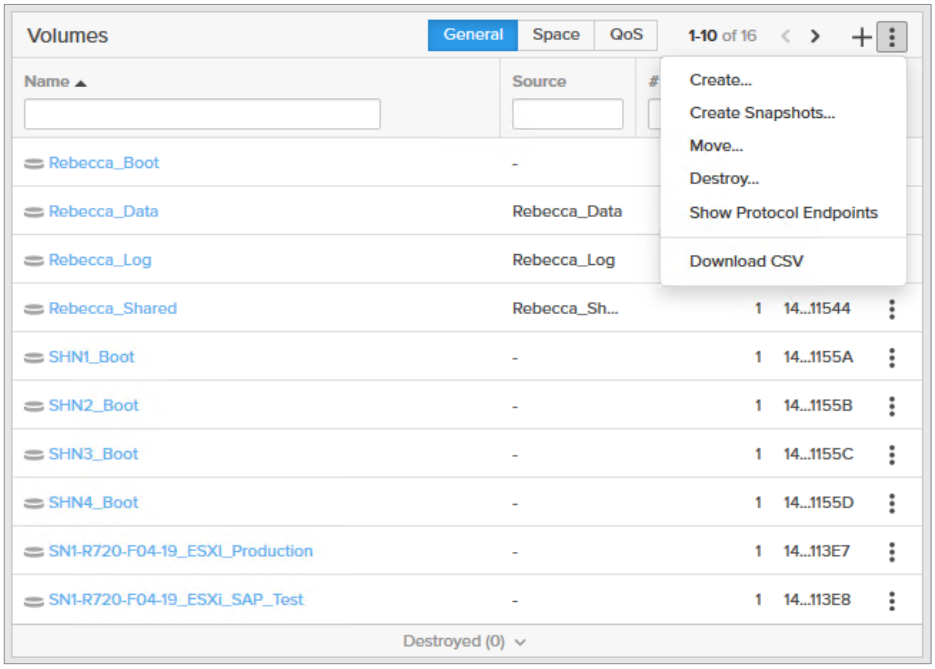

In the Create Volume Snapshots view select the log and data volumes for the SAP HANA instance, then select "Create".

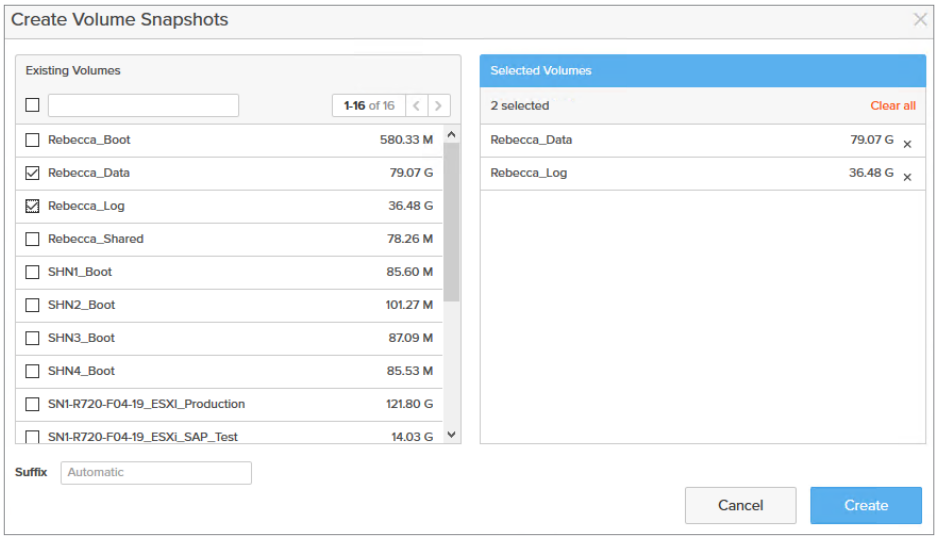

Once the block volume snapshots have been created, they will be shown in the Volume Snapshots view.

RECOVERING FROM CRASH CONSISTENT SNAPSHOTS

In order to restore the SAP HANA instance from a crash consistent storage snapshot, any existing instances must first be shut down.

Step 1: Shut down existing SAP HANA instance

SAP HANA Studio:

In SAP HANA Studio, using the systems view right-click on the relevant instance and select the "Configuration and Monitoring" submenu and then select "Stop System…"….".

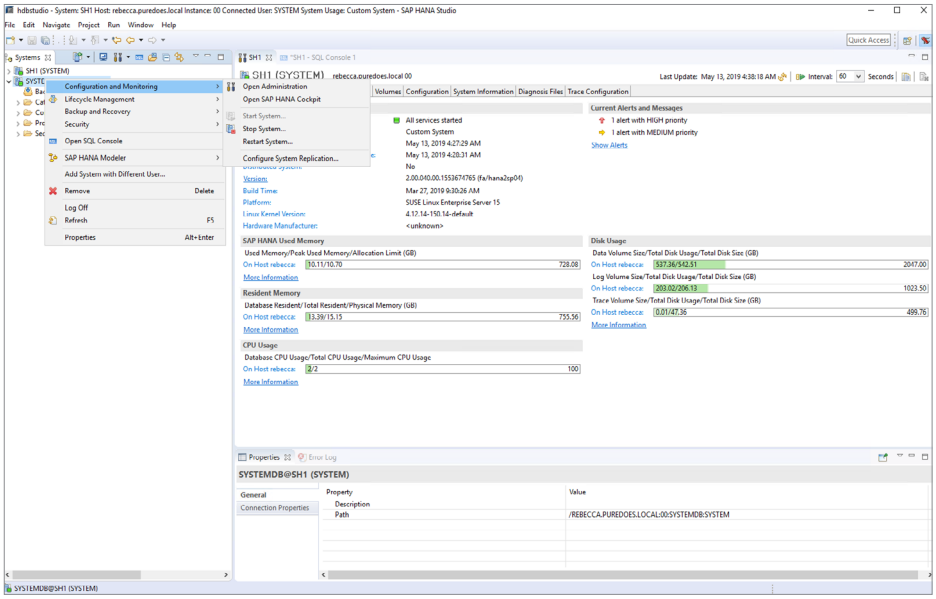

SAP HANA Cockpit:

Under the resources view in SAP HANA cockpit identify the relevant resource. Select the resource to proceed to the system overview.

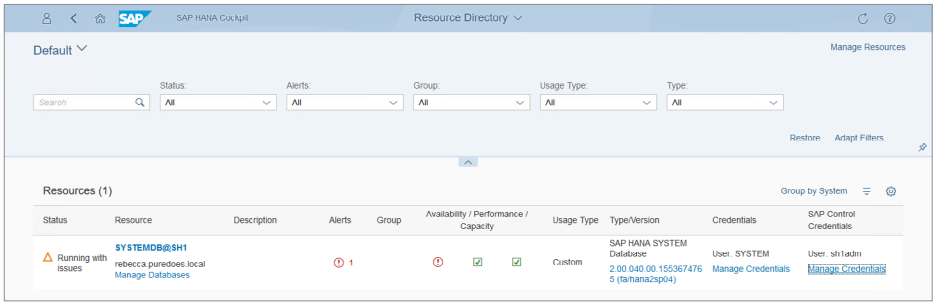

Ensure the relevant resource has the correct SAP Control login credentials.

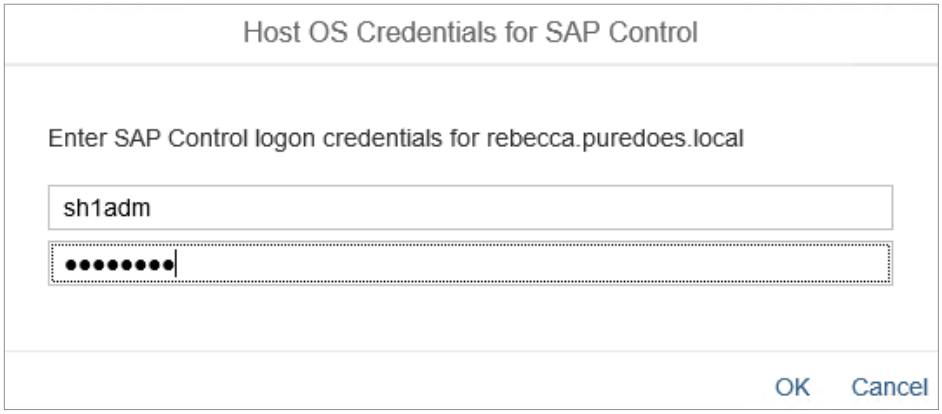

In the system overview identify the “Overall Database Status” tile.

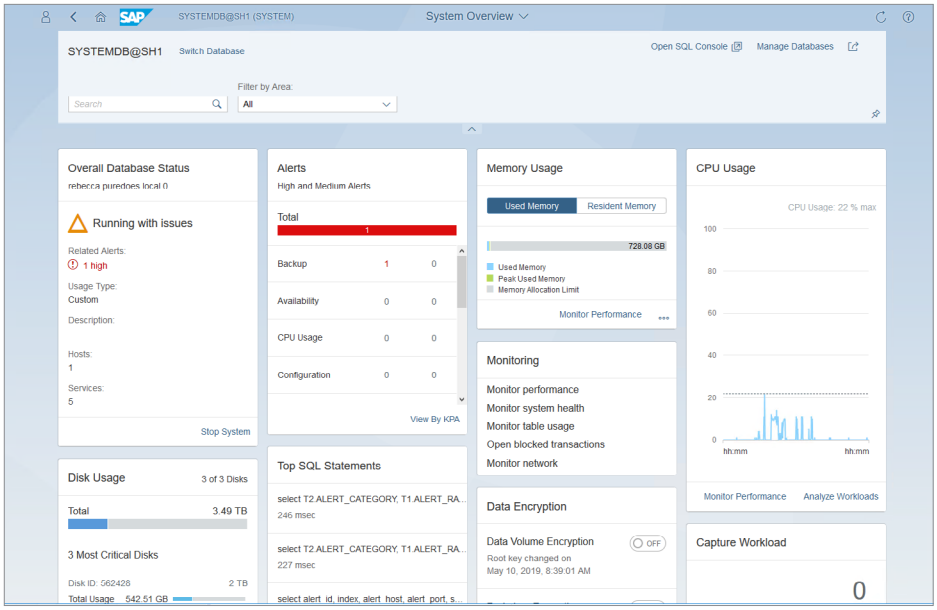

On the “Overall Database Status” tile, select “Stop System”.

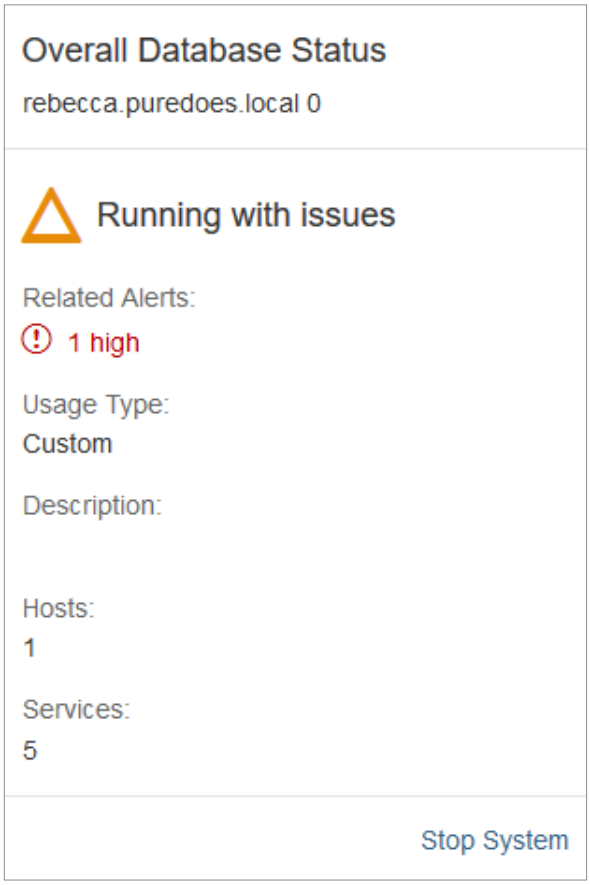

Command Line:

Logged in as <sid>adm :

In the command line logged in as <sid>adm execute "sapcontrol -nr <instance number> -function Stop"

```
/usr/sap/SH1/home> sapcontrol -nr 00 -function Stop
```

Logged in as a root or other user execute "/"/usr/sap/hostctrl/exe/sapcontrol -nr <instance number> -function Stop"

```
/usr/sap/hostctrl/exe/sapcontrol -nr 00 -function Stop
```

Step 2: Umount data and log persistence volumes from the operating system

Umount the data volume using the "umount" command.

```
umount /hana/data
```

Unmount the log volume using the "umount" command.

```
umount /hana/log
```

Step 3: Restore block volume snapshots in the FlashArray user interface

In the Volumes view, identify the relevant volumes to restore and select the top vertical ellipsis to bring up the context menu and select "Restore".

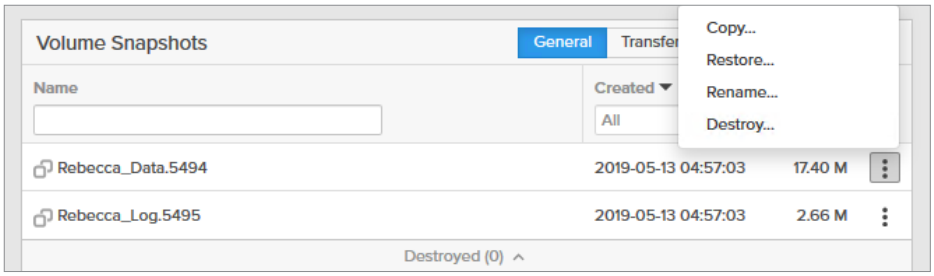

Check that the volume snapshot name presented by the context menu matches what is required.

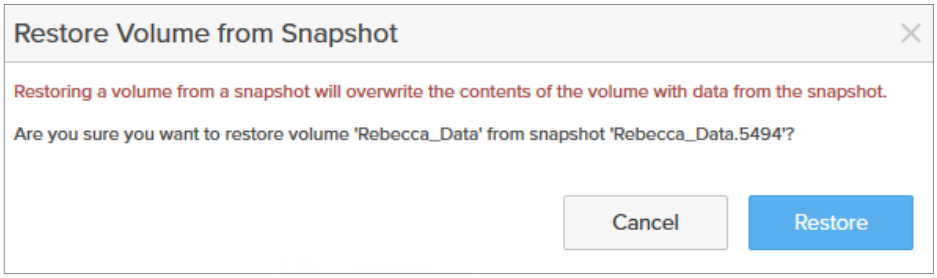

Repeat the operation for any other required volumes. Ensure both the log and data volume are done at minimum by the end of this step.

Step 4: Remount the data and persistent log volumes

Mount the data volume using the "mount" command.

```
mount /hana/data/
```

Mount the log volume using the "mount" command.

```
mount /hana/log/
```

Step 5: Start the SAP HANA instance

SAP HANA Studio:

In SAP HANA Studio, using the systems view right-click on the relevant instance and select the "Configuration and Monitoring" submenu and then select "Start System…"….".

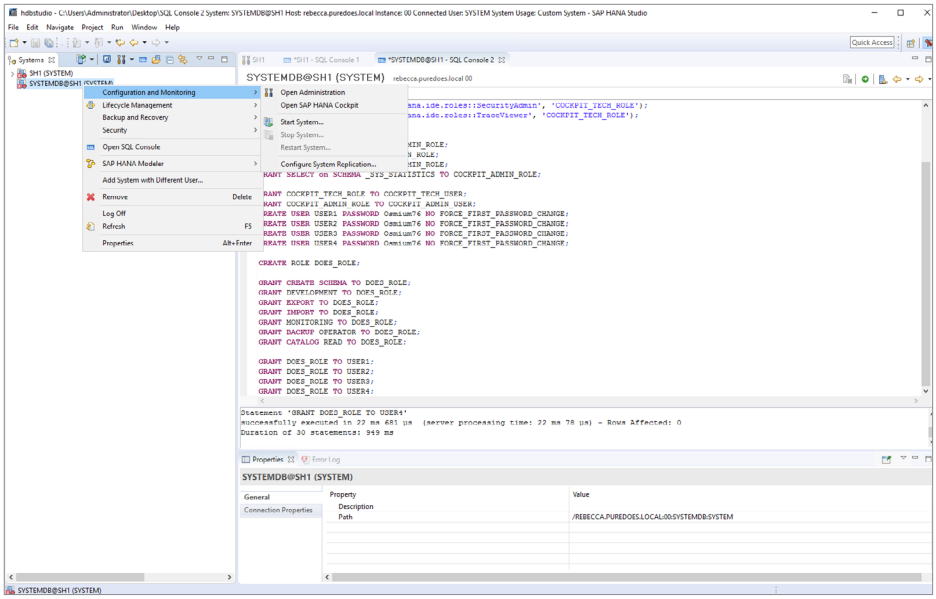

SAP HANA Cockpit:

In the System overview for the relevant resource find the "Overall Database Status" tile and select "Start System".

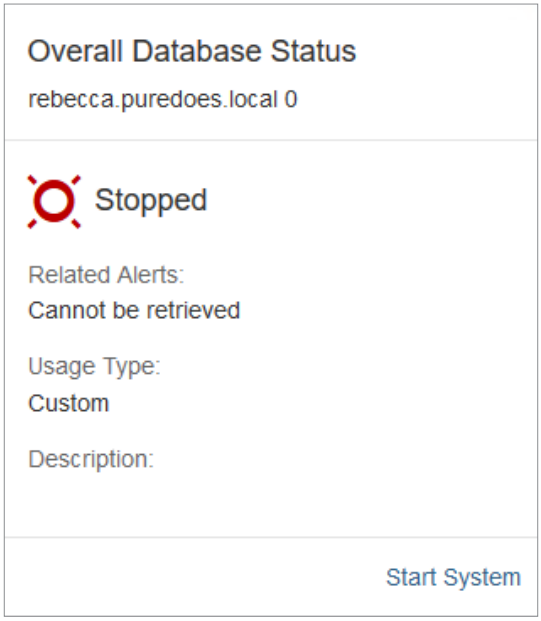

Command Line:

Logged in as <sid>adm:

In the command line logged in as <sid>adm execute "sapcontrol -nr <instance number> -function Start"

```
sapcontrol -nr 00 -function Start
```

Logged in as root

Logged in as a user with the correct permissions execute "/usr/sap/hostctrl/exe/sapcontrol -nr <instance number> -function Start"

```
/usr/sap/hostctrl/exe/sapcontrol -nr 00 -function Start
```

Portable Snapshot Technology

Once a snapshot has been created on any FlashArray//X, FlashArray//M or cloud block store instance it is considered to only be a recovery point and cannot be used for backup and recovery purposes. For snapshots to be appropriate for data protection purposes, block volume data must be moved to a separate location for recoverability to be ensured in the event of critical system loss and eventual failure.

The Purity operating environment offers several distinct mechanisms to offload volume snapshots to a separate location:

— Purity Snap- To- NFS

— Purity CloudSnap™ to S3 on AWS

— Purity FlashRecover (discussed in the section “Business Continuity Solutions for SAP HANA – Multi-Site Disaster Recovery)

Portable snapshot technology is the ability for any snapshot created by FlashArray or Cloud Block Store instances to be transported to a third-party storage target. Once the snapshot has been transported it can then be restored to any other FlashArray or Cloud Block Store instance. No additional backup software, cloud gateways or software licenses are required to make use of this functionality.

The technology functions using an incremental forever model of data movement, minimising the amount of data moved over an interface and limiting backup windows. This is achieved by running an initial full backup, followed by incremental forever snapshots (maintaining data compression) after which every subsequent snapshot is compared against the next, only moving changed blocks of data.

Snapshot offload features makes use of an embedded application deployment system on the FlashArray called Purity Run. Releases prior to Purity 5.2 require Pure Storage Support to install and configure this solution alongside the Snap to NFS feature. In Purity 5.2 and onwards Purity Run is user configurable but Pure Storage Support must be involved for the installation of Snap-to-NFS and Purity CloudSnap.

Snapshot management is accomplished through the use of a Protection Group. This is a collection of common volumes, hosts, targets and policies within which snapshot and replication schedules for any volume members can be applied.

Possible schedules and polices for snapshots and replication.

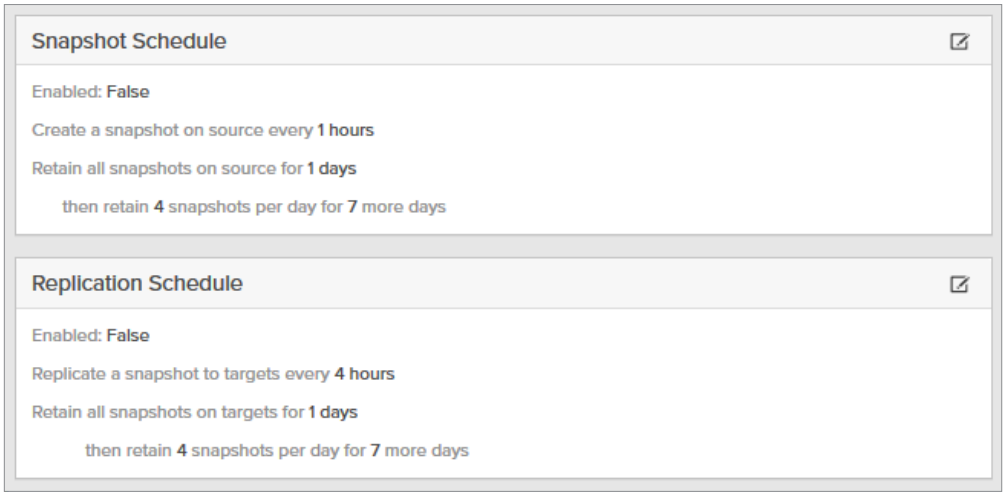

Configuration for offload functionality

Step 1. Install the app from the app catalog

In the FlashArray web user interface navigate to the **Software** tab under **Settings**. This view allows for the app catalog to be viewed and any installed applications are managed from here. In the app catalog install the “offload” application and wait for it to complete. The same operation can be accomplished from the command line by executing “puresw app install offload” and checking the status of the install with “puresw app list”.

Navigate to the FlashArray web user interface and select the Settings view and then navigate to the Software tab.

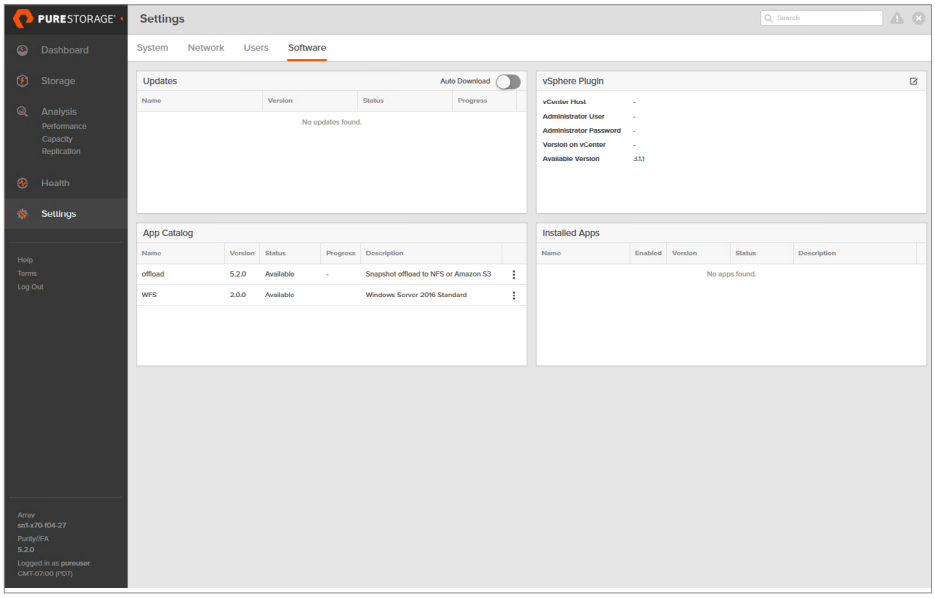

The App catalog tile shows all available applications available for installation, take note of the “offload” app as this is used for Snap-to-NFS functionality.

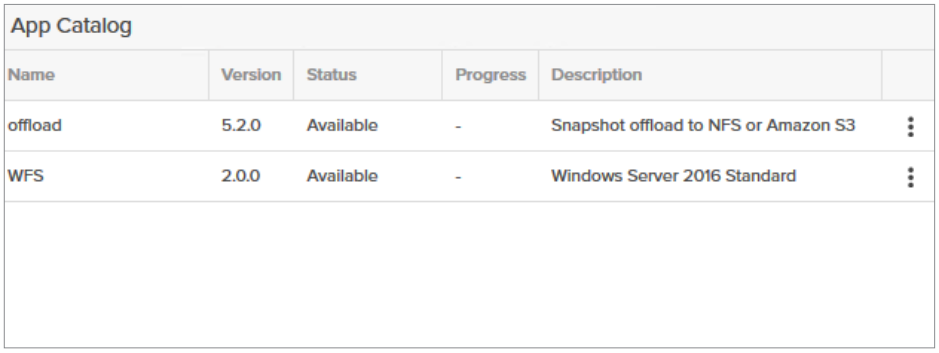

Select the vertical ellipsis for the "offload" application and select "Install".

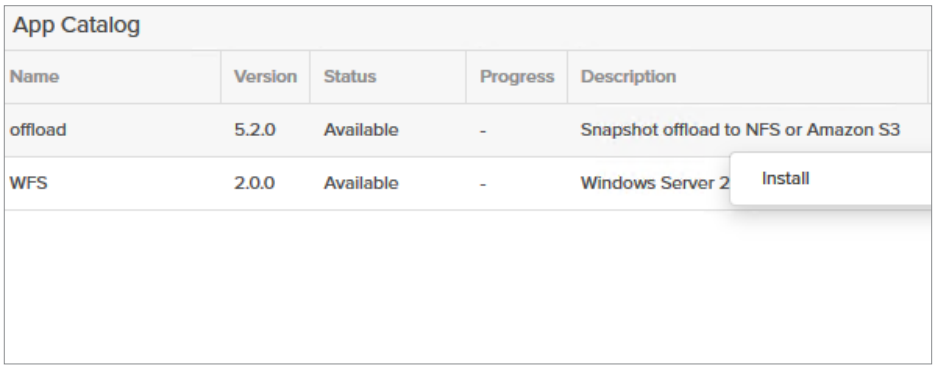

Confirm that the offload app is to be installed by selecting "Install".

Monitor the progress of the install by viewing the App Catalog.

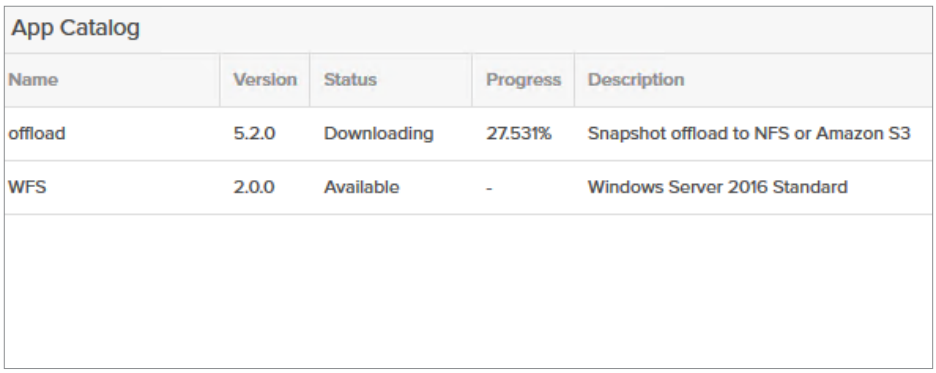

Once the "offload" app has completed installing it will show in the Installed Apps tile. Note that the Enabled and Status columns reading as "false" and "unhealthy" is normal at this stage.

Installed Apps

Name	Enabled	Version	Status	Description
offload	false	5.2.0	unhealthy	Snapshot offload to NFS or Amazon S3

Step 2. Assign network connectivity for the application to use

The command line must be used for this step. A virtual interface must be created for the offload app and then a physical interface on the FlashArray is assigned to it. Any connections assigned to the virtual interface must be enabled in order for communication to work.

Create the virtual interface on the replbond (if present).

```
purenetwork create vif @offload.data0 --subinterfacelist replbond
```

If the replbond is not present create the virtual interface using the physical ports (eth2 or eth3 only).

```
purenetwork create vif @offload.data0 --subinterfacelist ct0.eth2,ct1.eth2
```

Set the network address and associated information for the virtual interface.

```
purenetwork setattr --address 10.21.227.205 --netmask 255.255.255.0 --gateway 10.21.227.1 @offload.data0
```

Enable the physical ports and the virtual interface.

```
purenetwork enable ct0.eth2
purenetwork enable ct1.eth2
purenetwork enable @offload.data0
```

Step 3. Enable the offload app

Enabling the offload app allows the feature to be used and checks that all of the required services are healthy.

Select the vertical ellipsis in the top right-hand corner and select "Enable".

Once the app has started and all features are healthy it will show Enabled as "true" and Status as "healthy".

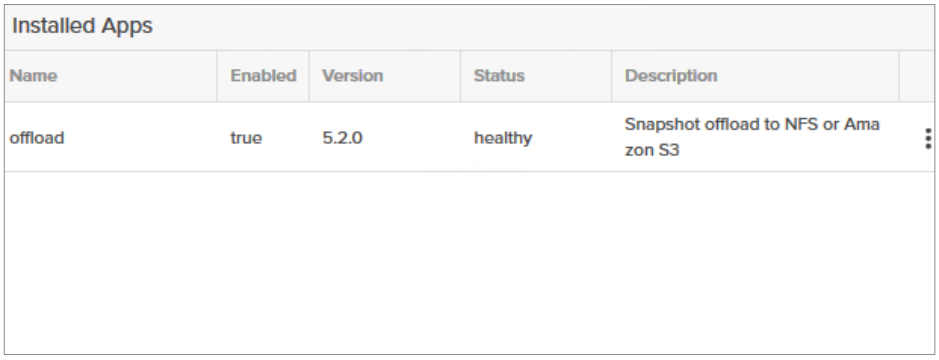

Purity Snap-To-NFS

Overview

Purity Snap-to-NFS allows for snapshots to be transported to a heterogeneous network file system target, natively managed via the FlashArray web interface, command line interface or over the ReST API. This solution allows organizations to back-up to generic inexpensive NFS servers for long-term storage or use existing infrastructure as a data protection target.

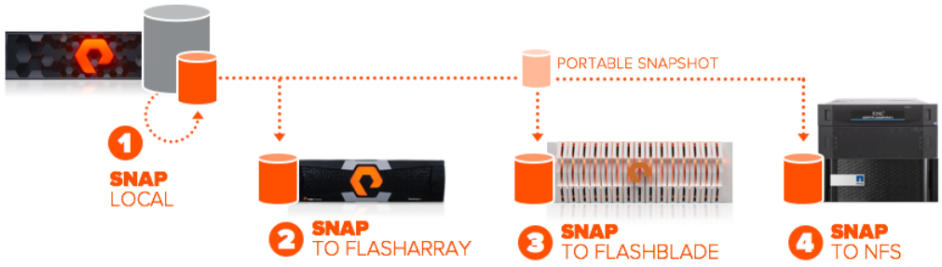

Fig. 12 Portable Snapshots offer the flexibility for various backup targets to be used.

This feature is only available on Purity 5.1 and above for FlashArray models //M20 and above, and //X20 and above.

Configuration and operation

An NFS server supporting NFS v3.X or 4.X is required to serve as the offload target. This NFS server must have network connectivity to the FlashArray replication ports that are assigned to the offload application. The FlashArray must have read/write/execute permissions to the NFS share and authentication must be done locally on the NFS server. .T..

Step 1. Connect to NFS target

With the app running and enabled, the NFS target server needs to be connected to. This can be done by navigating to **Storage** and selecting the **Array** tab and identifying the section for **Offload targets**. When connecting to an offload target the information required for the prompt will be a unique name, the NFS server address, the mount point of the NFS share and any specific mount options.

In the Storage view navigate to the Array tab.

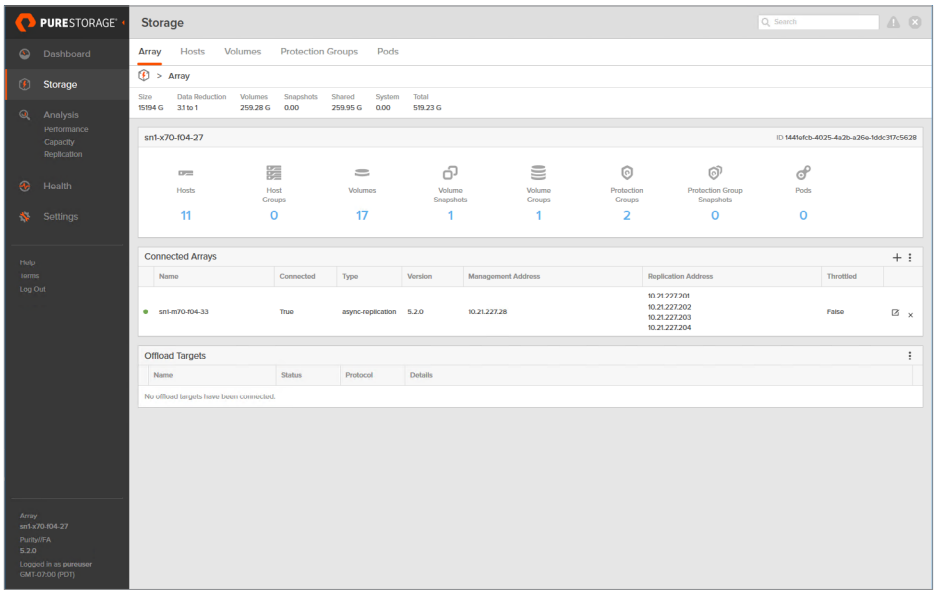

In the Array tab identify the "Offload Targets" section and select the vertical ellipsis to bring up the options and then select "Connect to NFS offload Target".

Enter in the information for a target name, the address of the NFS server, the mount point exported from the NFS server as an NFS share and any specific mount options required.

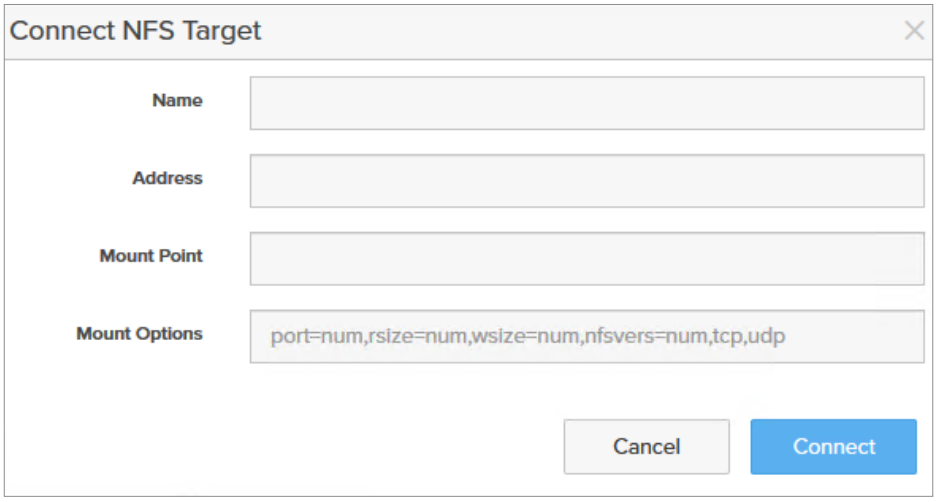

Once all of the relevant information is provided select "Connect".

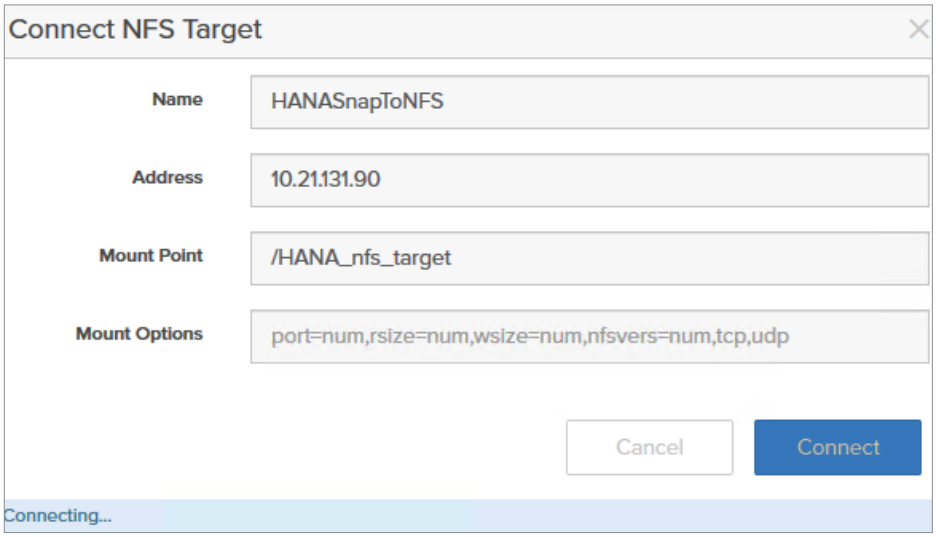

Once the connection has been established it will be shown in the Offload Targets section.

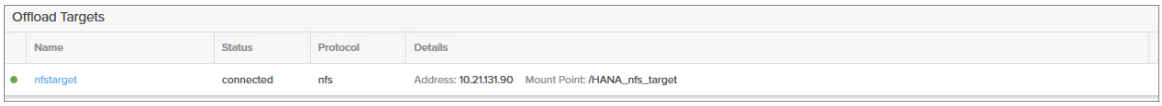

Step 25. Offload snapshots to the NFS target

When offloading snapshots to a target they must be a part of a protection group. A protection group consist of one or many volumes and have different policies applied to it for snapshot schedules and retention or replication schedules. The **Protection Groups** tab can be found in **Storage**.

In the Storage view navigate to the Protection Groups Tab. In this example there is already a protection group created with the name "HANA" with the SAP HANA data persistence volume added as a member.

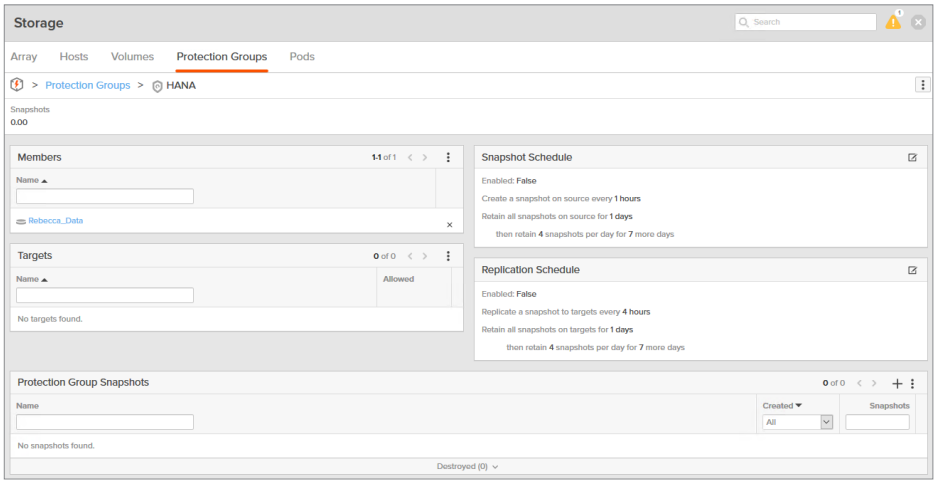

In the section for Members and Targets, select the vertical ellipsis in the upper right-hand corner of Targets and select "Add".

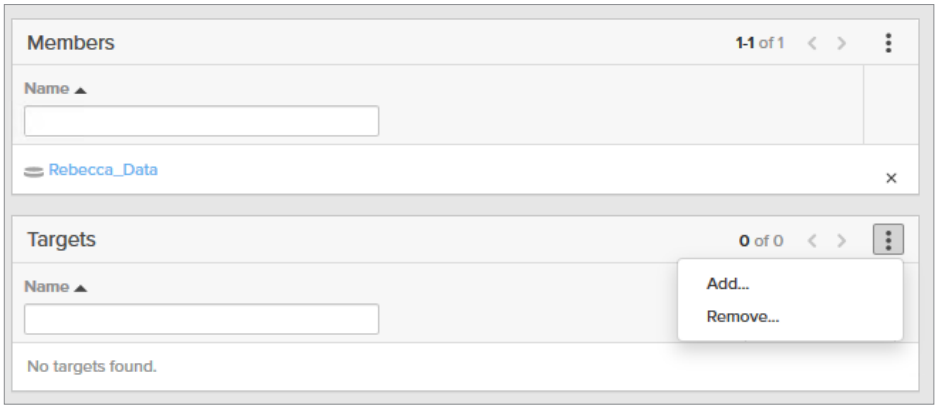

Ensure the required target is selected and then select "Add".

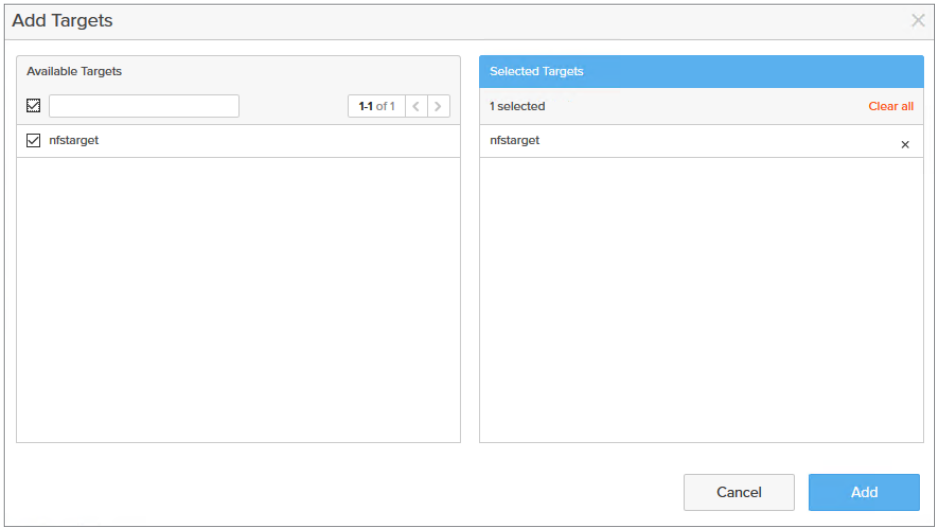

In the Protection Group Snapshots section select the "+" in the top right-hand corner to create a snapshot.

Give the snapshot an optional suffix and then ensure that "Replicate Now" is selected, then select Create.

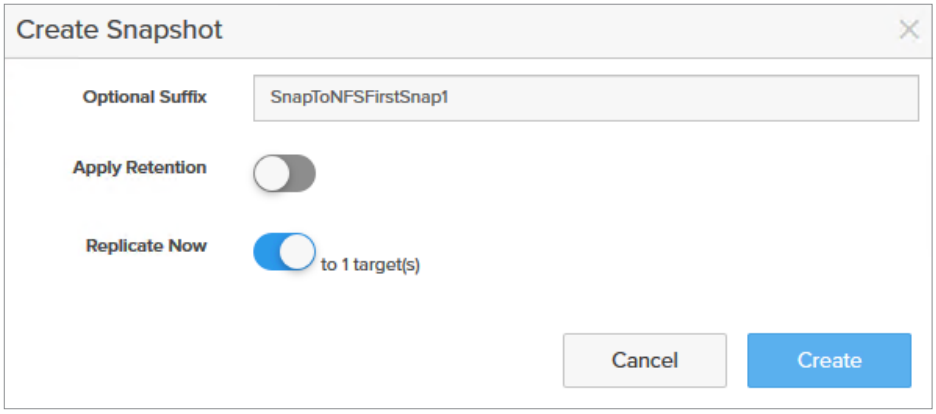

The snapshot is then created and will offload to the NFS share.

Viewing "offload targets" found in the Array tab within the Storage view will show the offload progress.

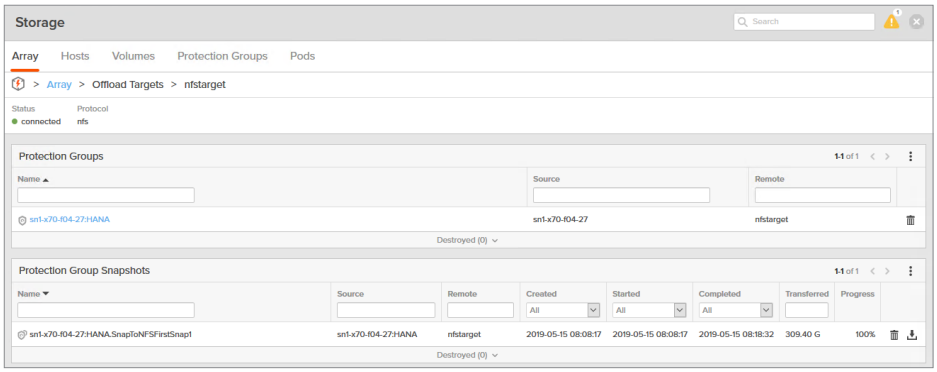

Purity CloudSnap to AWSS3

Overview

Purity CloudSnap allows snapshots to be offloaded to a cloud vendor, extending data protection on FlashArray to cloud storage services. CloudSnap to AWS S3 enables snapshots to be offloaded directly to an S3 bucket on Amazon Web Services. This solution offers organizations cost and efficiency benefits by using less space in the S3 bucket, minimising network utilization and shortening backup windows.

This feature is only available on Purity 5.2 and above for FlashArray models //M20 and above, and //X20 and above.

Configuration and operation

Step 1. Create Amazon Web Services S3 bucket

In the **S3 Management Console** select **Create bucket**. When the prompt appears give the bucket a name and select the appropriate region for it. Blocking public access is recommended for the bucket to ensure that any data stored in it is secure.

In the S3 Management Console, create a bucket and give it a DNS complaint name and set the region it needs to be located in.

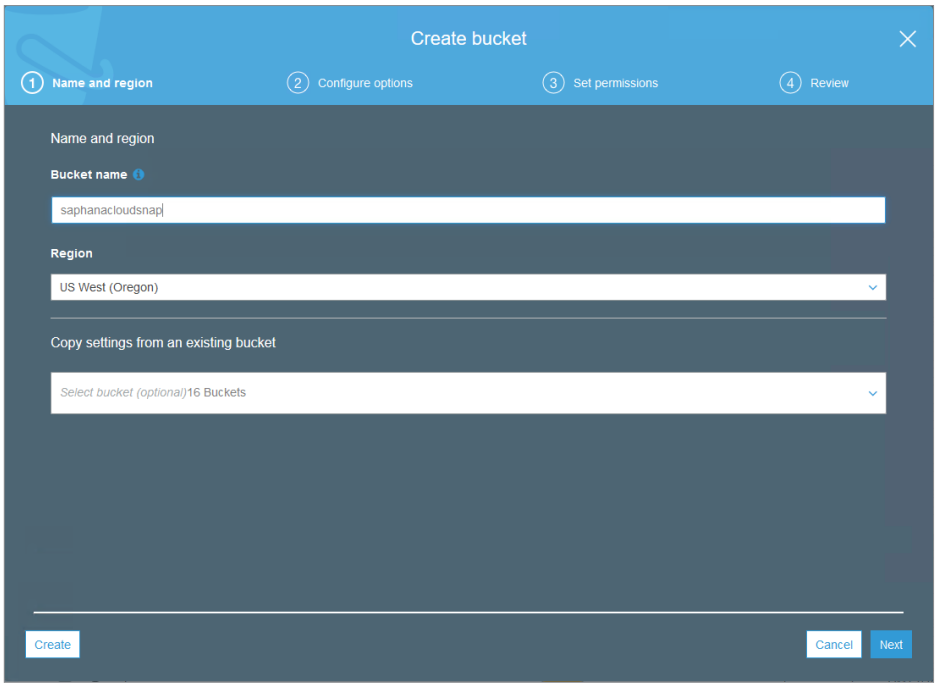

In the configuration options no properties are required for CloudSnap.

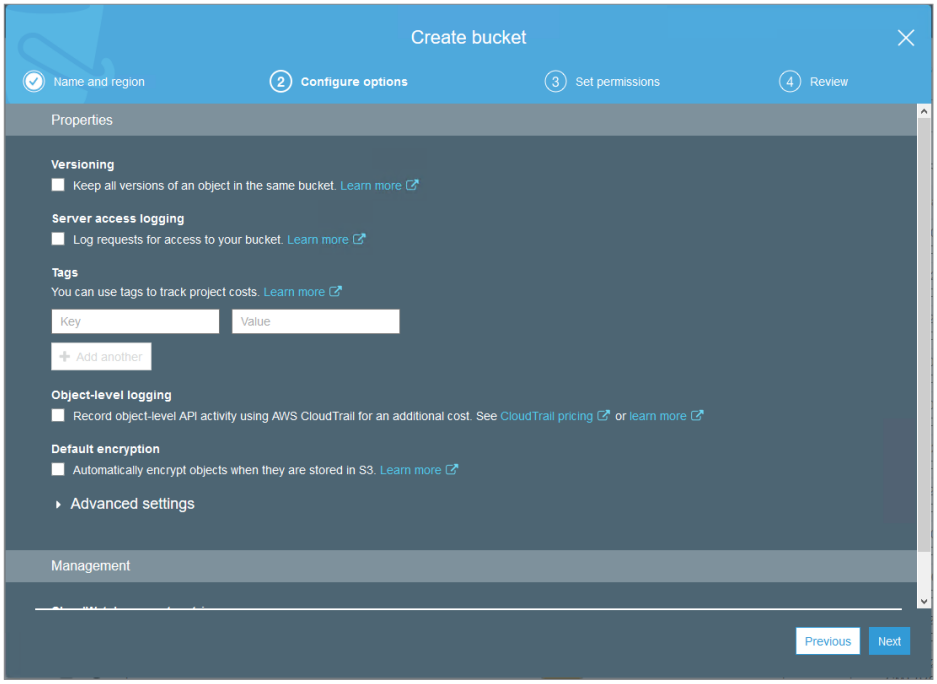

In the permissions ensure all public access is blocked for security reasons.

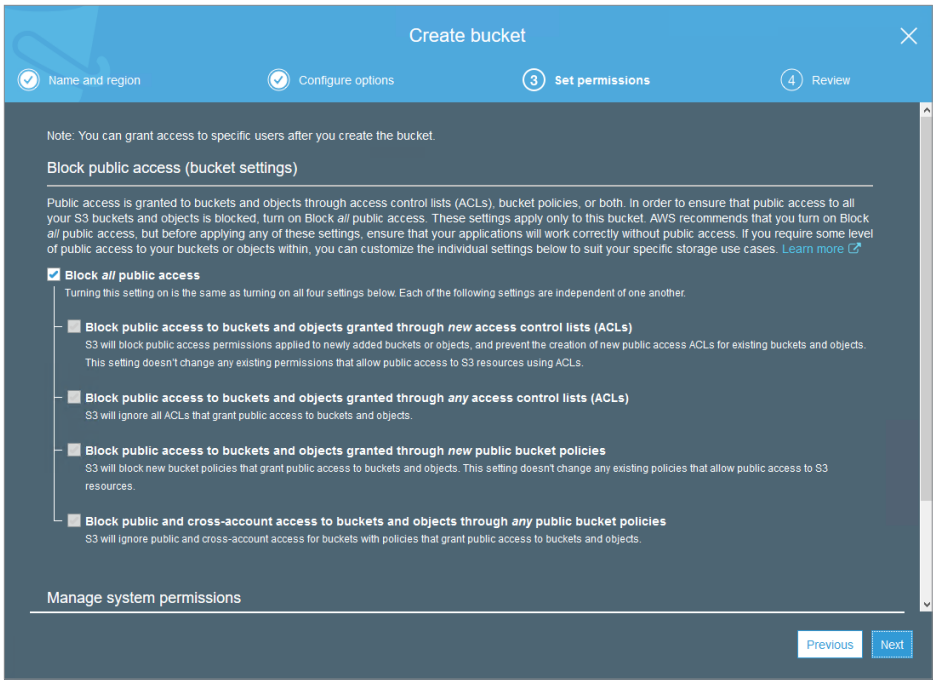

Review configuration for the bucket and select "Create bucket".

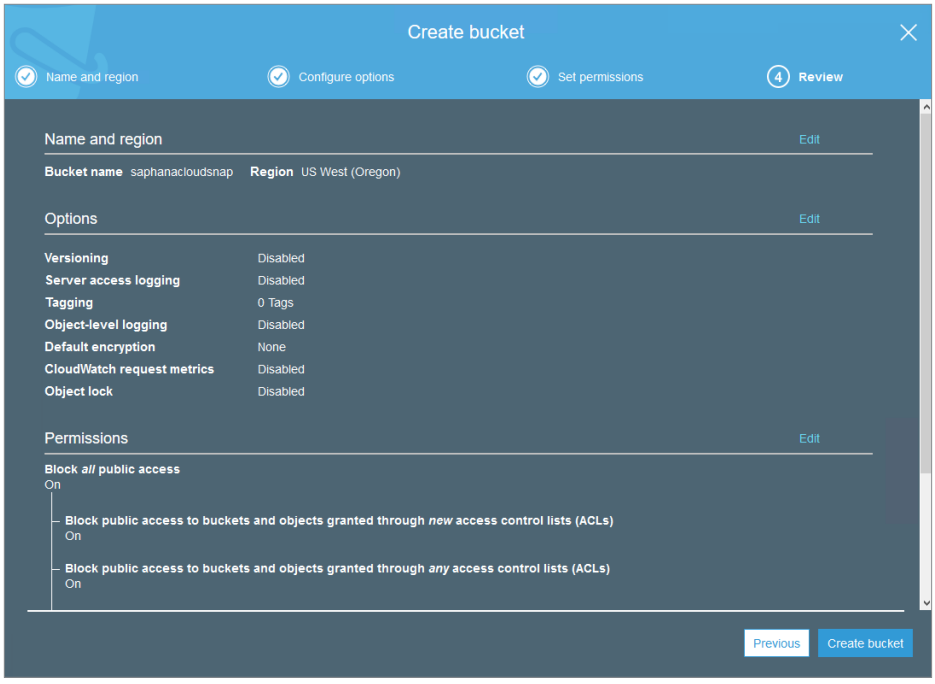

Once the bucket has been created navigate to the management screen, select the "Properties" tab and then select the "Default encryption" tile.

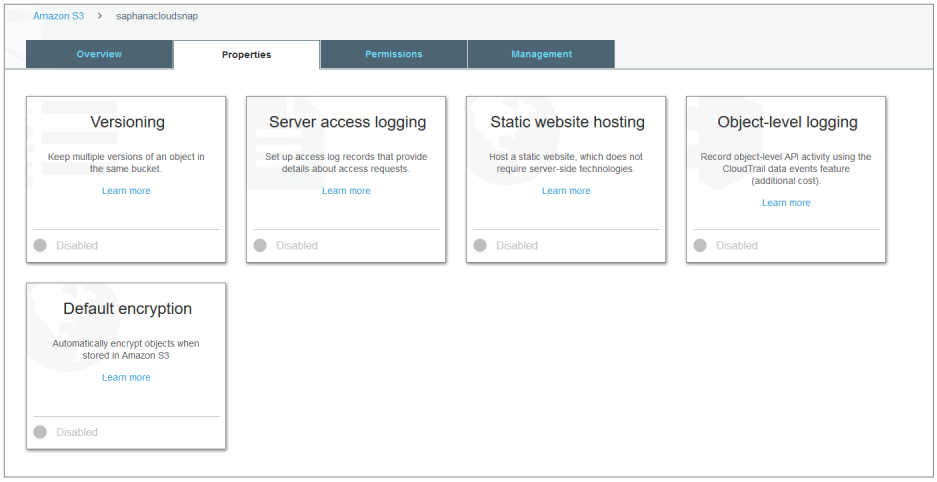

Set the default encryption to "AES-256" and select Save.

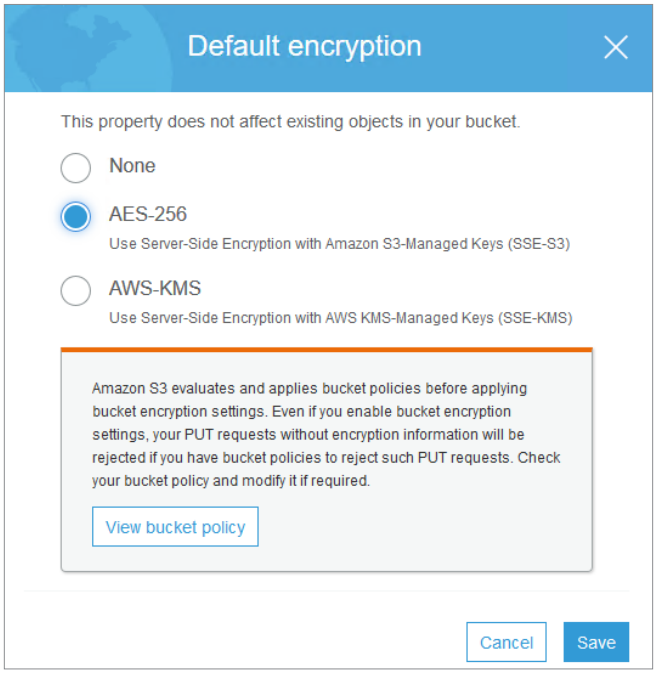

Step 2. Connect to AWS S3 bucket

With the app running and enabled and an S3 bucket created in AWS, the FlashArray needs to be connected to the target. This can be done by navigating to **Storage** and selecting the **Array** tab and identifying the section for **Offload targets**. When connecting to an offload target the information required for the prompt will be a unique name, the Access Key, the bucket name and the Secret Access key.

In the FlashArray web user interface enter the Storage view and navigate to the Array tab.

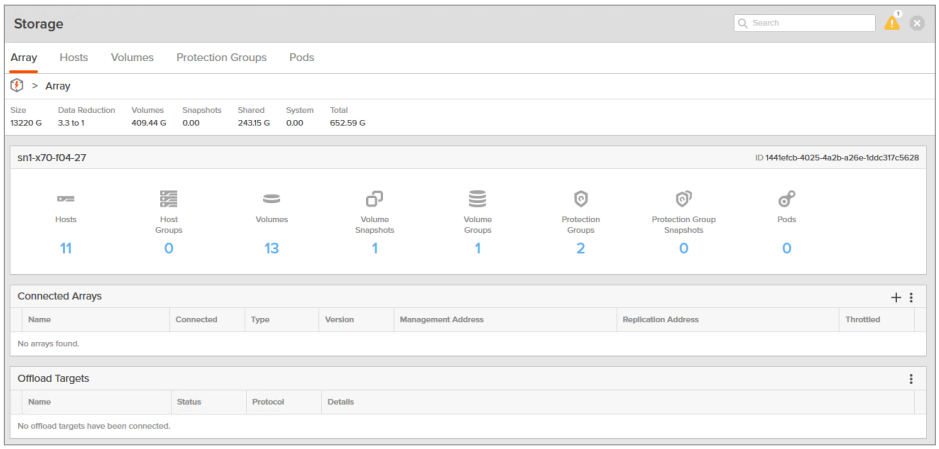

In the Array tab identify the "Offload Targets" section and select the vertical ellipsis to bring up the options and then select "Connect to S3 offload Target".

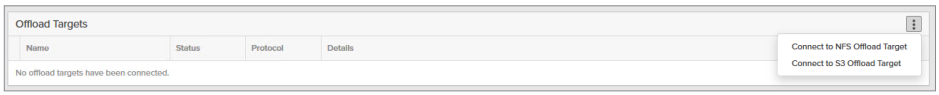

Enter the information for a target name, the Access key ID, Bucket name and Secret Access Key. If the S3 bucket is newly created then Initialize needs to be enabled.

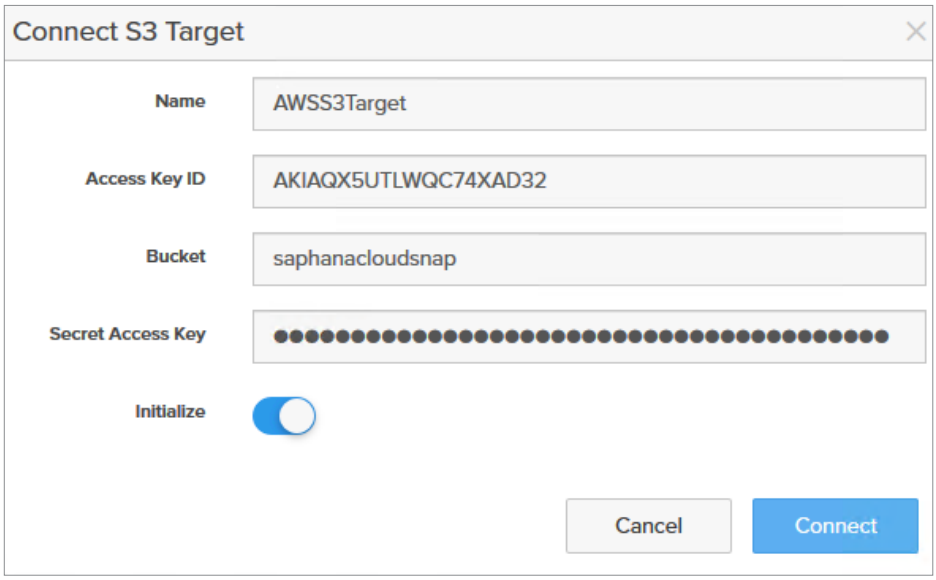

Once the target has been connected, it will be displayed in the Offload Targets section.

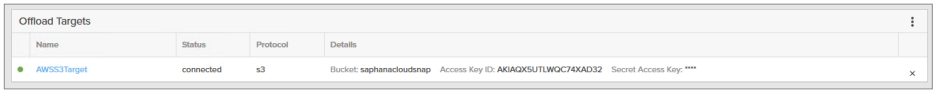

Step 3. Offload snapshots to the AWS S3 target

When offloading snapshots to a target they must be a part of a protection group. A protection group consists of one or many volumes and has different policies applied to it for snapshot schedules and retention or replication schedules. The **Protection Groups** tab can be found in **Storage**.

In the Storage view navigate to the Protection Groups tab. In this example there is already a protection group created with the name "HANA" with the SAP HANA data persistence volume added as a member.

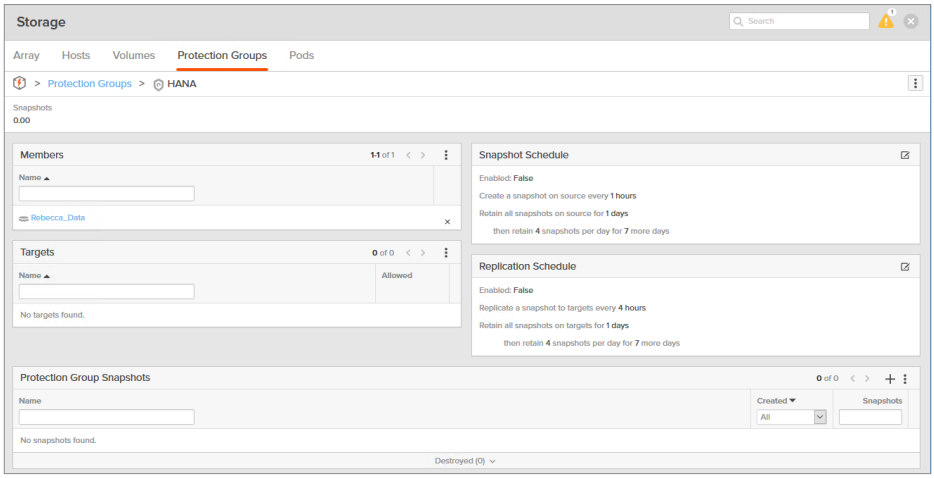

In the section for Members and Targets, select the vertical ellipsis in the upper right-hand corner of Targets and select "Add".

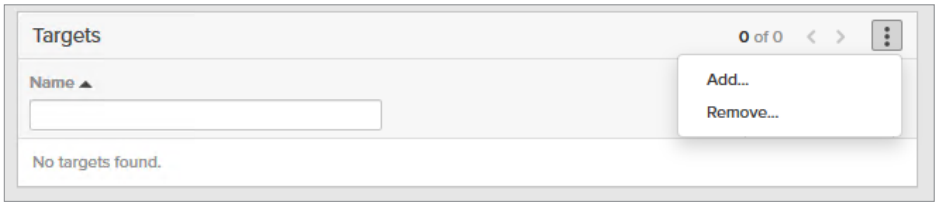

Ensure the required target is selected and then select "Add".

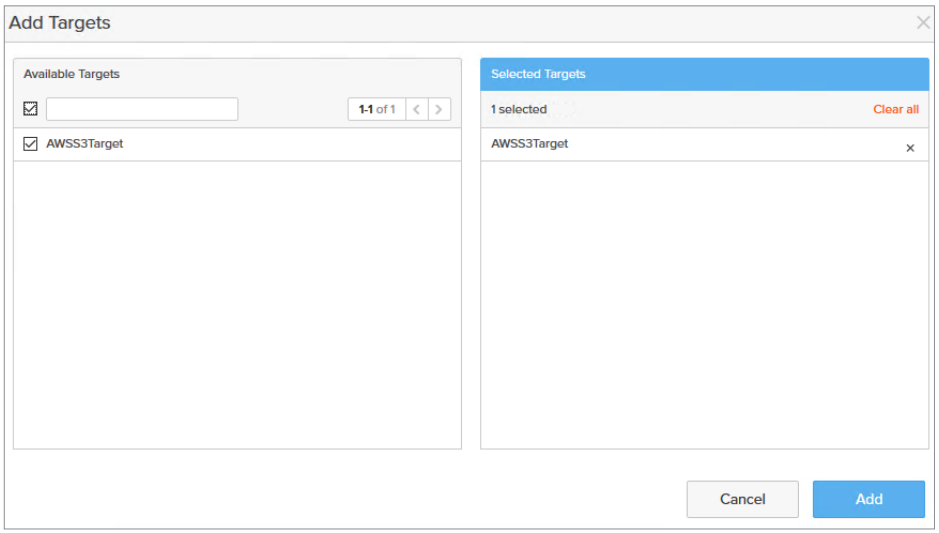

In the Protection Group Snapshots section select the "+" in the top right-hand corner to create a snapshot.

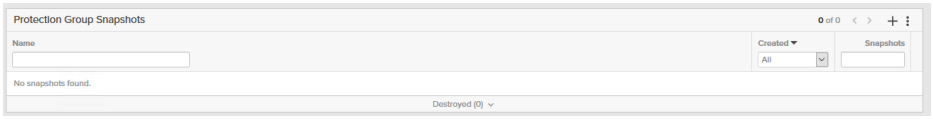

Give the snapshot an optional suffix and then ensure that "Replicate Now" is selected, then select Create.

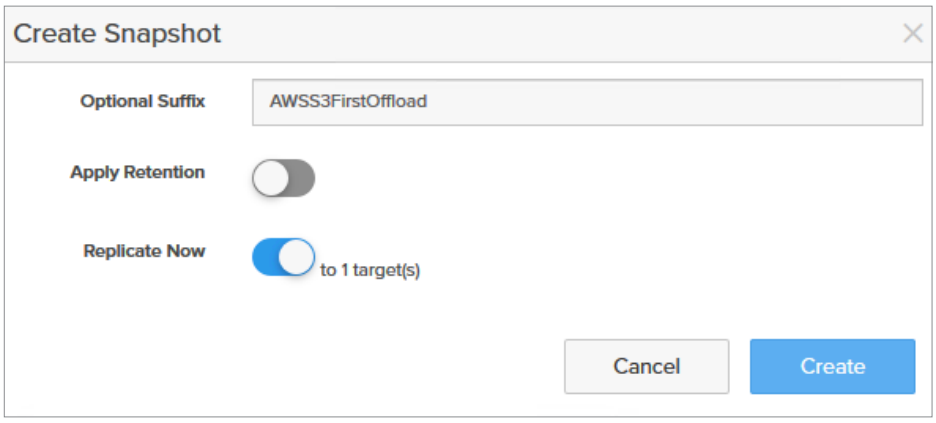

The snapshot is then created and will offload to the AWS S3 bucket.

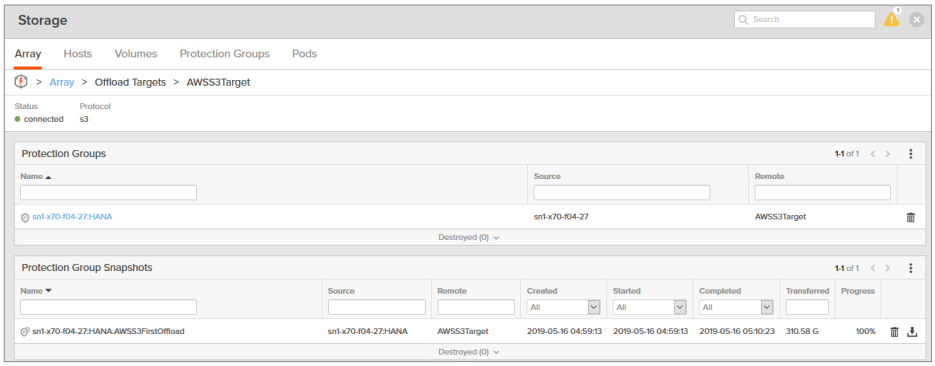

Recovering offloaded snapshots

Recovery of an SAP HANA system using NFS or AWS S3 offloaded snapshots is accomplished by retrieving the block volume snapshot from the target to a relevant FlashArray. Once any volume snapshot has been retrieved from the target it can be used to recover the SAP HANA database with processes detailed in the sections “Application Consistent Storage Snapshots” and “Crash Consistent Storage Snapshots”.

Offloaded snapshots are retrieved in an efficient manner by only retrieving blocks which are not present on the FlashArray, ensuring that a limited amount of bandwidth is used and reducing recovery time.

In the FlashArray web user interface, under the storage view and inside the Array tab, select the NFS or AWS S3 offload target by selecting its name.

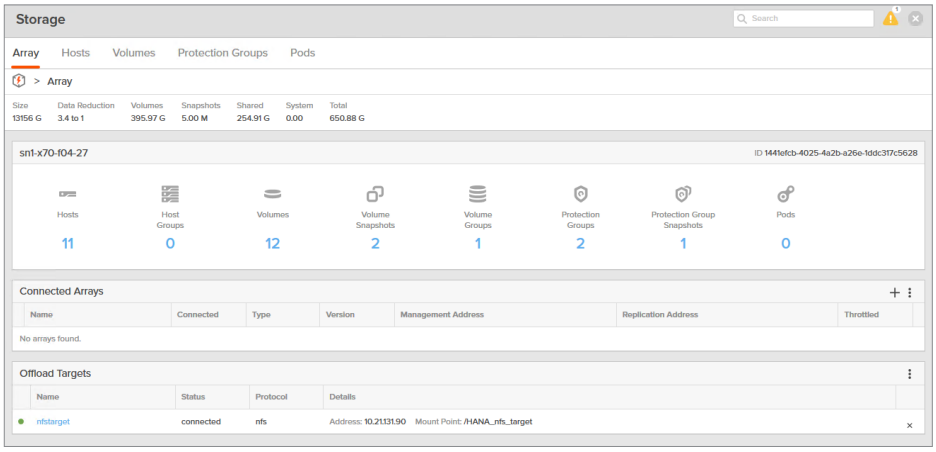

In the properties for the offload target identify the section for "Protection Group Snapshots".

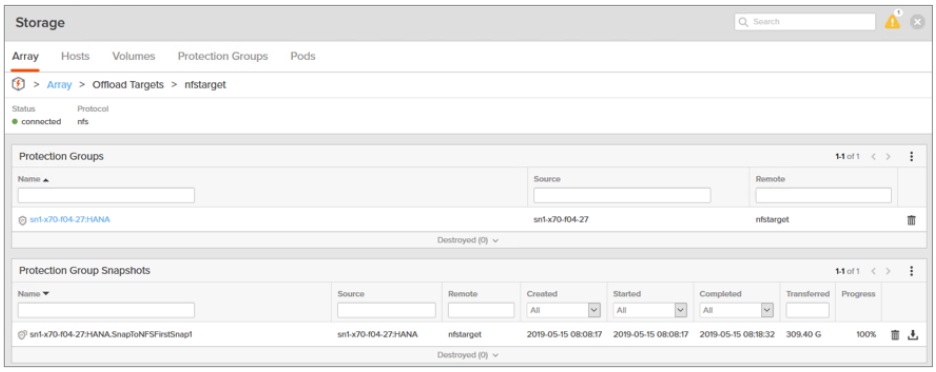

In the section for Protection Group Snapshots, identify the required protection group snapshot to retrieve from the offload target and select the download icon to the right of it.

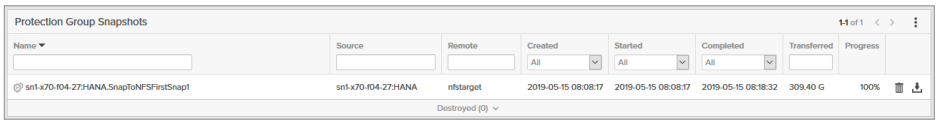

In the prompt for "Get Volume Snapshots" select the volume snapshot that needs to be retrieved and then select "Get".

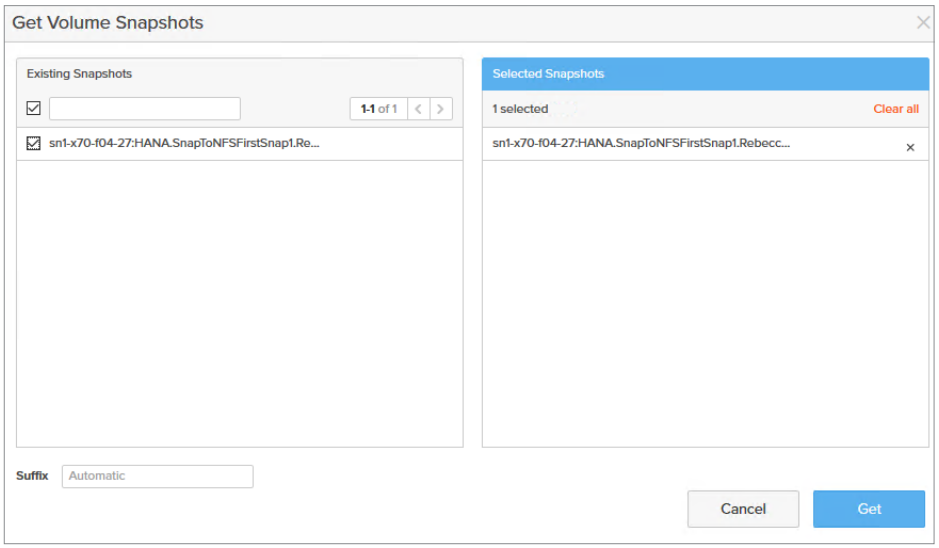

In the summary prompt which appears, select "Go to Volumes page" in the bottom left-hand corner to view the recovered snapshot.

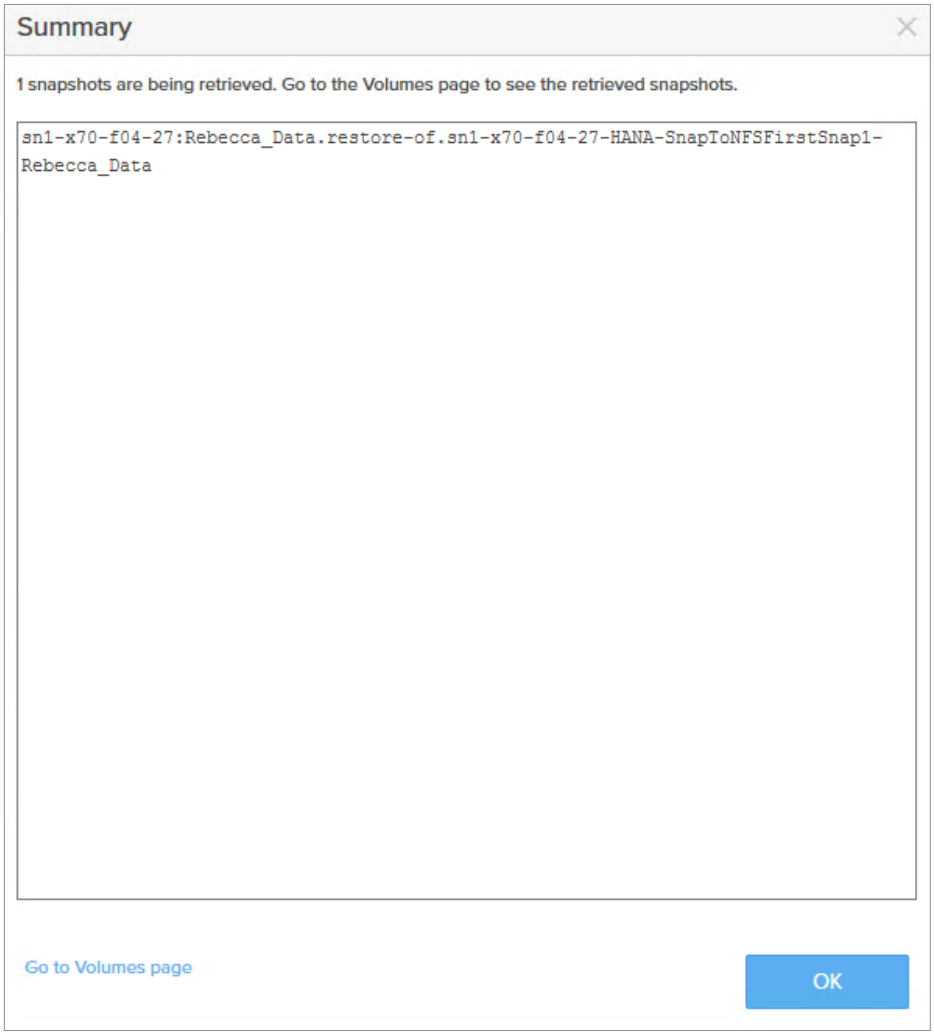

The volume snapshot will show as recovered from the offload target and is available locally on a FlashArray. This can now be copied, restored, renamed or deleted.

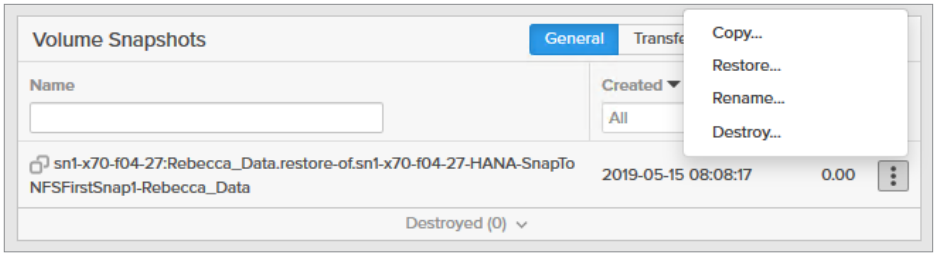

Business Continuity Solutions for SAP HANA

High Availability – ActiveCluster

Overview

Purity ActiveCluster is a flexible Pure Storage technology in the Purity operating environment which can be used as an SAP HANA Storage Replication solution to ensure business continuity in the event of system and site failure scenarios. This solution can be implemented for any SAP HANA system, both single and multiple host, making use of FlashArray in TDI deployments. This is a zero-recovery time objective (RTO) where no time elapses between failure and resolution. The configuration can be extended to provide for a zero-recovery point object (RPO) as well as further increasing operational resiliency (discussed in the section for Multi-Site Disaster Recovery – Adding an asynchronous snapshot replication target to an Active Cluster configuration).

SAP HANA replication can be done either on the system platform (system replication) or underlying storage vendor (storage replication) where each is implemented based on business needs and use case suitability. System replication is a top down (compute to storage) solution replicating the entire platform to another instance, offering organizations a feature set defined by SAP HANA where shadow instances can be preloaded resulting in fast takeovers with little to no impact on performance. Storage replication is a bottom up (storage to compute) solution where data, log and any other block volume data is replicated to remote storage, delivering no preloading capabilities as the whole system is replaced or started fresh on alternative hardware.

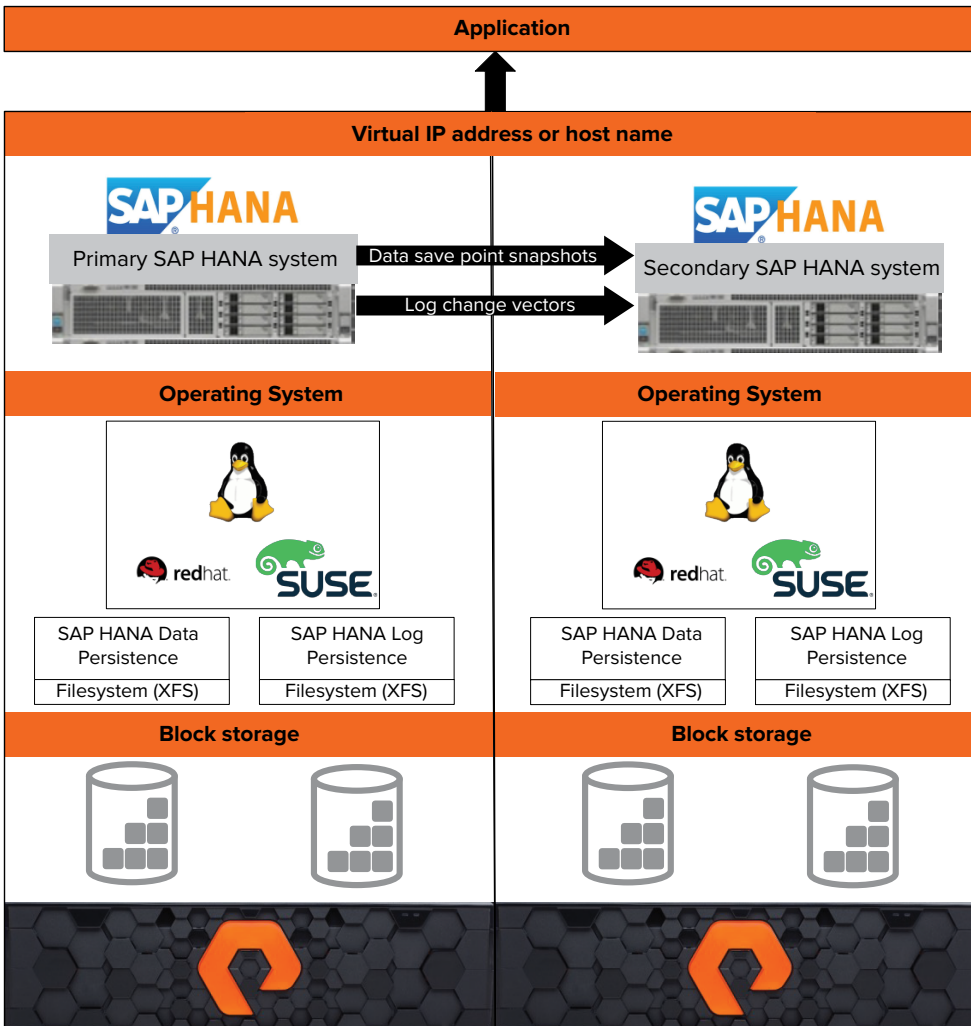

Fig. 13 SAP HANA System replication topology

SAP HANA system replication offers a number of different replication modes, providing high availability and resilience for a number of different scenarios. The terms typically used in system replication are “primary” and “secondary” system, where the primary system is the target for normal operations and data is then sent to the secondary to ensure it continues to be available if the primary is lost. The following replication modes are offered by SAP HANA:

— **Synchronous replication** – The primary system waits until the secondary system has received data and persisted it to disk. In the event of a secondary system failure the primary system will wait until a timeout is reached and then proceeds without replicating data.

— **Synchronous in-memory replication** – The primary system waits until secondary system has received data, but it does not need to be persisted to disk. In the event of a secondary system failure the primary system will wait until a timeout is reached and then proceeds without replicating data.

— **Synchronous full sync replication** – The primary system waits until the secondary system has received data and persisted it to disk. In a failure scenario the primary system will block operations until the secondary system becomes available.

— **Asynchronous replication** – the primary system does not wait for the secondary system to commit data and in the event of the secondary system failing, the primary system will continue to perform operations.

SAP HANA Storage replication with Purity ActiveCluster offers organizations the ability to establish a flexible domain for business continuity through the use of a single user interface without any additional license costs. Storage replication is more suitable for addressing multiple failure domains starting from the storage layer and working towards the platform (SAP HANA database instance) as opposed to system replication which only addresses platform availability and requires each area to be addressed individually.

ActiveCluster is typically configured using two FlashArray storage arrays, defined as “primary”, “secondary” or “tertiary” arrays, but can be extended to three in certain scenarios. The following replication configurations are offered by ActiveCluster:

— **Synchronous replication (Symmetric Active/Active)** – Changes are synchronized and protected in non-volatile memory on both primary and seconds arrays before being acknowledged to the host. Read and write operations can be performed to the same volumes on either primary or secondary array, with optional host-to-array site awareness.

— **Asynchronous replication for snapshots** – Changes are asynchronously written to the secondary (or tertiary) array volumes on a schedule. Any operations performed to the primary (or synchronous active/active) volumes will acknowledge to the host while the asynchronous tertiary array will not be required to complete operations. Asynchronous relationships can be used to convert standby systems to an active state if the primary site(s) are lost. This replication mode is useful for remote site disaster recovery without impacting system performance on the primary site(s).

Pods

ActiveCluster is built on the concept of a **pod**, a stretched storage container that defines a set of objects (hosts, volumes or protection groups) that are replicated together and the arrays they are replicated between. Any FlashArray or cloud block store instance can support multiple pods and any pod that is replicated between them is considered to be “stretched”. A pod is also considered a consistency group where multiple volumes in the same pod are write-order consistent. Pods also provide volume namespaces where different volumes can have the same name if they are in a different pod, allowing the migration of workloads between arrays or consolidating multiple workloads from many arrays onto one without volume name conflicts.

Transparent failover

Transparent failover between arrays when using ActiveCluster is automatic, requiring no intervention from a storage administrator. Failover occurs within standard host I/O timeouts, similar to the way failover occurs between two controllers in one array during non-disruptive hardware or software upgrades. ActiveCluster is designed to provide maximum availability across symmetric active/active storage arrays while preventing a split-brain condition from occurring – split brain being the case where two arrays might serve I/O to the same volume without keeping the data in sync between the two arrays.

Any active/active synchronous replication solution designed to provide continuous availability across two different sites requires a component referred to as a witness, or voter, to mediate failovers while preventing split brain. ActiveCluster includes a simple to use, lightweight, and automatic way for applications to failover transparently, or simply move, between sites in the event of a failure without user intervention: the Pure1 Cloud Mediator.

The Pure1 Cloud Mediator is responsible for ensuring that only one array is allowed to stay active for each pod when there is a loss of communication between the arrays. In the event that the arrays can no longer communicate with each other over the replication interconnect, both arrays will pause I/O and reach out to the mediator to determine which array can stay active for each sync-replicated pod. The first array to reach the mediator is allowed to keep its synchronously replicated pods online. The second array to reach the mediator must stop servicing I/O to its synchronously replicated volumes, in order to prevent split brain. The entire operation occurs within standard host I/O timeouts to ensure that applications experience no more than a pause and resume of I/O

THE PURE1 CLOUD MEDIATOR

A failover mediator must be located in a third site that is in a separate failure domain from either site where the arrays are located. Each array site must have independent network connectivity to the mediator such that a single network outage does not prevent both arrays from accessing the mediator. A mediator should also provide a very lightweight and easy to administer component of the solution. The Pure Storage solution provides this automatically by utilizing an integrated cloud-based mediator. The Pure1 Cloud Mediator provides two main functions:

— Prevent a split-brain condition from occurring, where both arrays are independently allowing access to data without synchronization between arrays.

— Determine which array will continue to service IO to synchronously replicated volumes in the event of an array failure, replication link outage, or site outage.

The Pure1 Cloud Mediator has the following advantages over a typical non-Pure, heavy-handed voter or witness component:

— **SaaS Operational Benefits** – As with any SaaS solution the operational maintenance complexity is removed: nothing to install onsite, no hardware or software to maintain, nothing to configure and support for HA, no security patch updates, etc.

— **Automatically a third site** – The Pure1 Cloud Mediator is inherently in a separate failure domain from either of the two arrays.

— **Automatic Configuration** – Arrays configured for synchronous replication will automatically connect to and use the Pure1 Cloud Mediator.

— **No Misconfiguration** – With automatic and default configuration there is no risk that the mediator might be incorrectly configured.

— **No Human Intervention** – A significant number of issues in non-Pure active/active synchronous replication solutions, particularly those related to accidental split brain, are related to human error. Pure's automated non-human mediator eliminates operator error from the equation.

— **Passive Mediation** – Continuous access to the mediator is not required for normal operations: the arrays will maintain a heartbeat with the mediator. However, if the arrays lose connection to the mediator, they will continue to synchronously replicate and serve data as long as the replication link is active.

ON-PREMISES FAILOVER MEDIATOR

Failover mediation for ActiveCluster can also be provided using an on-premises mediator distributed as an OVF file and deployed as a VM. Failover behaviours are exactly the same as described above. The on-premises mediator simply replaces the role of the Pure1 Cloud Mediator during failover events. The on-premises mediator has the following basic requirements:

— The on-premises mediator can only be deployed as a VM on virtualized hardware. It is not installable as a stand-alone application.

— High Availability for the mediator must be provided by the hosts on which the mediator is deployed. For example, using VMware HA or Microsoft Hyper-V HA Clustering.

— Storage for the mediator must not allow the configuration of the mediator to be rolled back to previous versions. This applies to situations such as storage snapshot restores, or cases where the mediator might be stored on mirrored storage.

— The arrays must be configured to use the on-premises mediator rather than the Pure1 Cloud Mediator.

— The mediator must be deployed in a third site, in a separate failure domain that will not be affected by any failures in either of the sites where the arrays are installed.

— Both array sites must have independent network connections to the mediator such that a failure of one network connection does not prevent both arrays from accessing the mediator

Multipathing

There are different combinations of various multipathing configurations with ActiveCluster each with their own advantages and disadvantages. A Uniform Access Configuration is where all hosts have access to both FlashArrays. When using a uniform access configuration, inter-site network performance and reliability are lower than those of the intra-site storage network, hence it is optimal for hosts at each site to route pod volume I/O to the array local to them and to access the remote array only when the local one is unresponsive.

FlashArray provides an option named preferred array to specify whether a host connection to a volume in a stretched pod is optimized or not.

Most multipathing I/O stacks support Asymmetric Logical Unit Assignment (ALUA), a SCSI feature that allows a host to query an array to determine which of its paths to a logical unit are optimized, meaning a direct path to an array owning the LUN, and which are not. Hosts then route I/O commands solely to optimized paths and use non-optimized paths when the optimized paths are unavailable.

Uniform host access

Uniform host access for any SAP HANA deployment can be configured by allowing two of the arrays to be visible to the compute host, providing an even more resilient setup in the event of any single storage array failure. The underlying volumes required for the SAP HANA system are replicated between the arrays using ethernet connectivity but the SAP HANA platform is not aware of the underlying storage configuration. Synchronous replication is appropriate to be used within a single datacenter or multiple datacenters on the same campus (5ms response latency in a datacenter or 11ms across a metro area) and asynchronous snapshot replication extends this to longer distances with higher latencies.

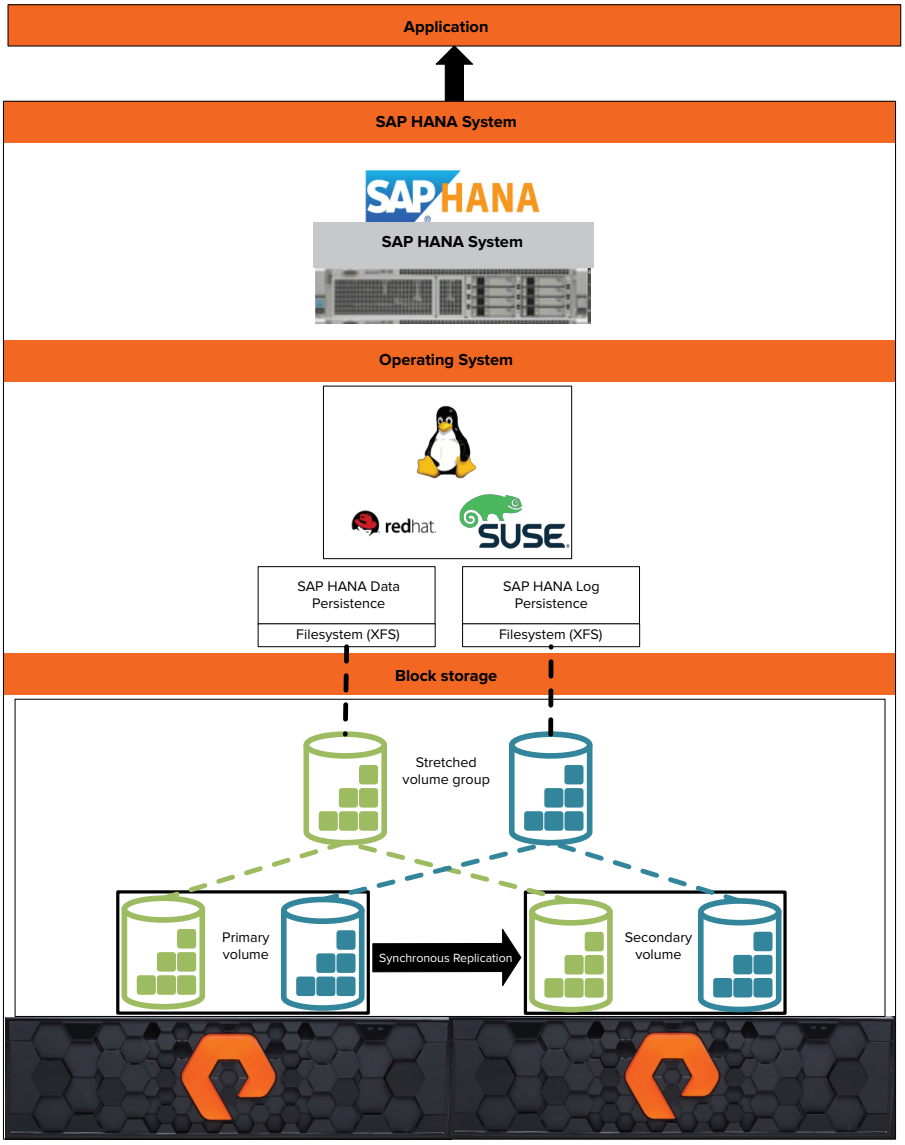

Fig. 14 Uniform storage access for SAP HANA System.

Symmetric Active/Active host access

An Active/Active host access can be configured two separate arrays and compute instances. The arrays will replicate block volume data either synchronously and data can then be read or written by either the primary or secondary system. ActiveCluster makes use of ALUA to expose paths to local hosts as active or optimised to provide the following advantages:

- Volumes in stretched pods are read/write on both arrays. There is no such thing as a passive volume that cannot service both reads and writes.
- The optimized path is defined on a per host-to-volume connection basis using a preferred-array option; this ensures that, regardless of what host a VM or application is running on, it will have a local-optimized path to that volume.

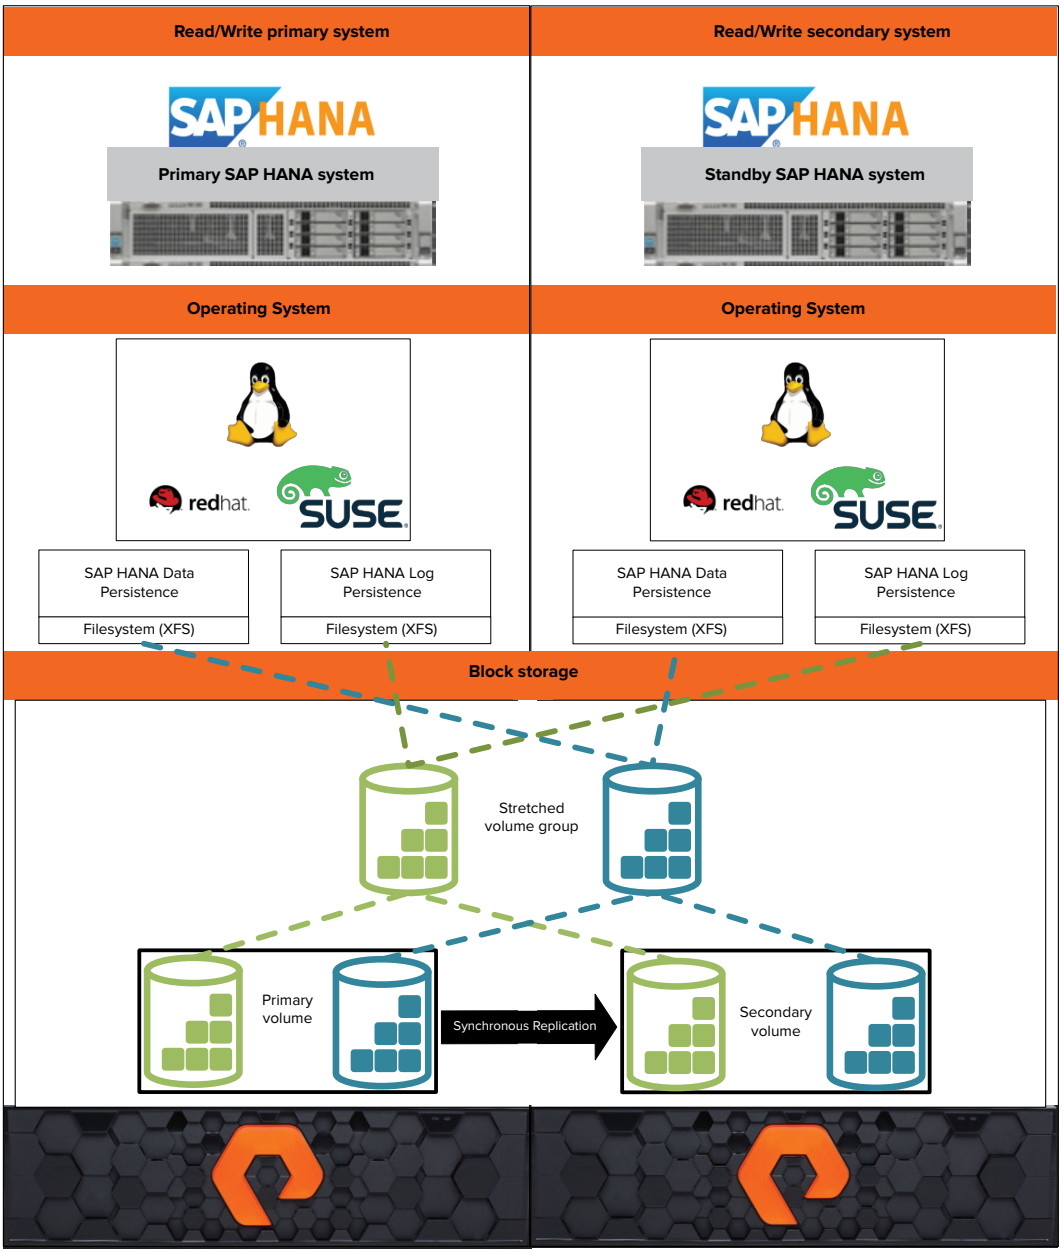

Fig. 15 Symmetric Active/Active configuration.

Configuration and operation

Before processing further, ensure the following environment is configured as set out in the below Requires and Best Practices:

https://support.purestorage.com/FlashArray/PurityFA/Protect/ActiveCluster/ActiveCluster_Requirements_and_Best_Practices

ActiveCluster Glossary of Terms

The following terms have been introduced for ActiveCluster and will be used repeatedly in this document:

- **Pod** – A pod is a namespace and a consistency group. Synchronous replication can be activated on a pod, which makes all volumes in that pod present on both FlashArrays in the pod.
- **Stretching** – Stretching a pod is the act of adding a second FlashArray to a pod. When stretched to another array, the volume data will begin to synchronize, and when complete all volumes in the pod will be available on both FlashArrays.
- **Unstretching** – Unstretching a pod is the act of removing a FlashArray from a pod. This can be done from either FlashArray. When removed, the volumes and the pod itself are no longer available on the FlashArray from which they were removed.
- **Restretching** – When a pod is unstretched, the other array (the array unstretched from) will keep a copy of the pod in the trash can for 24 hours. This allows the pod to be quickly re-stretched without having to resend all data, if re-stretched within 24 hours.

Step 1. Connect to an array and create a synchronous connection

Connecting to another FlashArray requires that a connection key is acquired from the corresponding array attempting to be connected to. This key is used as a secure way of ensuring only authorised individuals can connect to the other array. The other array can be connected to using the virtual IP address of fully qualified domain name of the remote array. ActiveCluster can be configured to use synchronous replication and with a third array use asynchronous snapshot replication.

In the FlashArray web user interface for both the local and remote system, in the Storage section navigate to the Array tab.

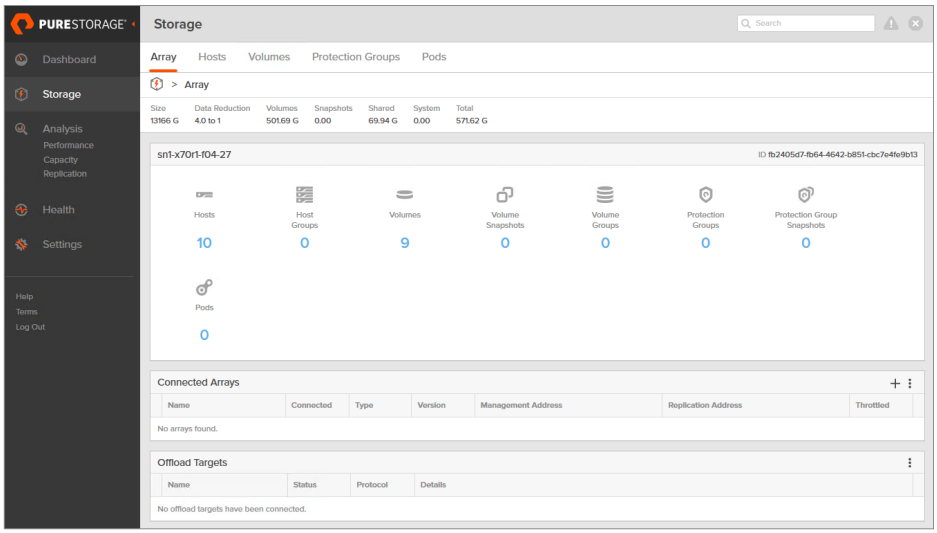

In the interface for the array to be connected to, identify the section for connected arrays and select the 3 vertical ellipses and select "Get Connection Key".

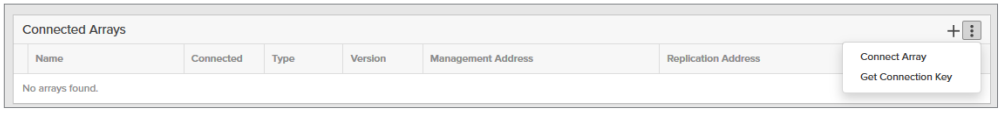

A prompt will show up with the connection key for that array, Select Copy and then navigate to the array to connect from.

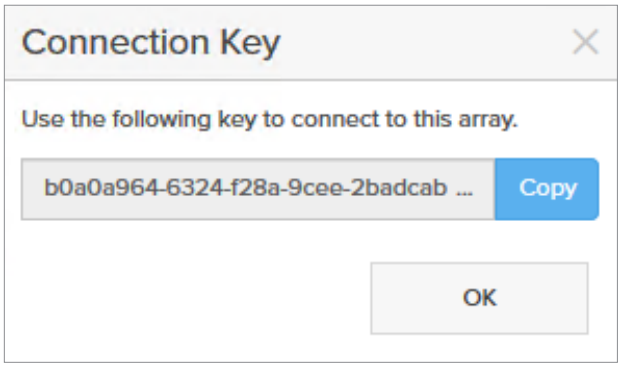

In the user interface for the FlashArray to connect from, return to the Connected Arrays section, select the three vertical ellipses and select "Connect Array".

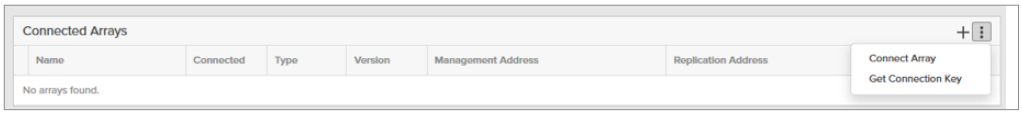

A prompt will show, enter the relevant information for Management Address, Type (Synchronous replication), the connection key and any replication addresses required. Select Connect when all of the relevant information has been entered.

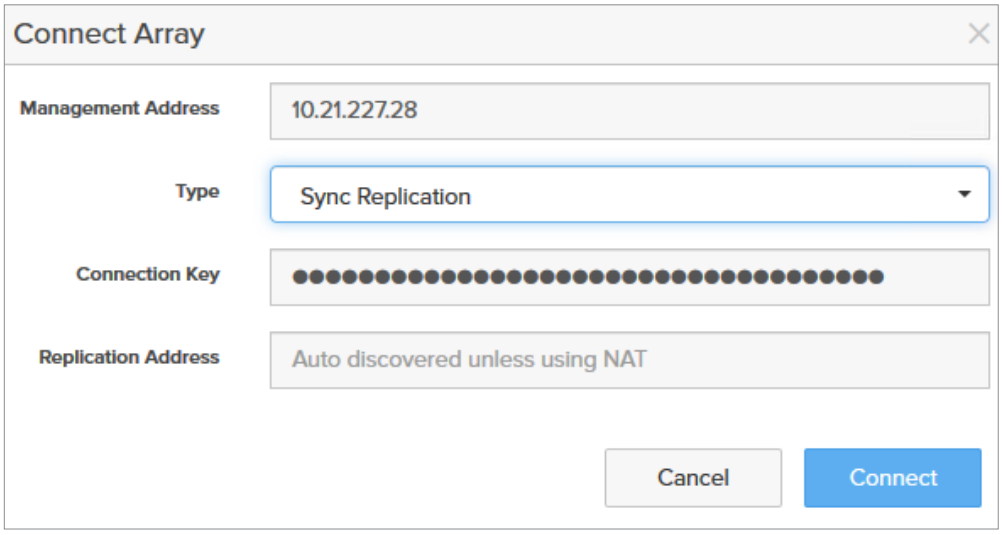

Once connected, the connected array will show under the relevant section.

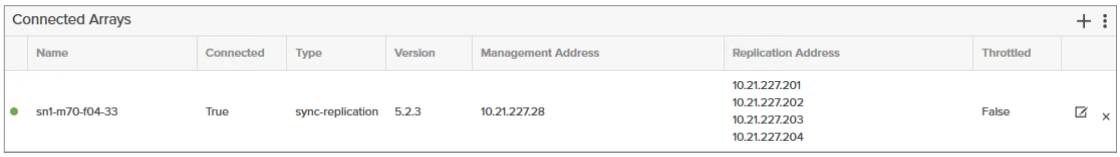

Step 2. Create a POD and add volumes to it

Creating a POD and adding volumes to it is known more commonly as a consistency group. Pods are stretched, unstretched and failover over together. A POD will have its own unique name and any volumes added to that POD will become a part of that namespace.

The default configuration for a new pod used by ActiveCluster is to use the Cloud Mediator – no configuration is required other than to ensure the management network from the FlashArray is redundant and to have IP access to the mediator.

It is important to note once a pod has been "stretched", pre-existing volumes cannot be added to it until it is "un-stretched". Alternatively, once a pod has been stretched, only new volumes can be created in the pod. Therefore, if you would like to add existing volumes to the pod, return to the beginning of this step and then stretch the pod.

Any existing volumes added to the pod should already be exported to a host, new volumes will need to be exported to a host after creation.

In the storage view navigate to the Pods tab.

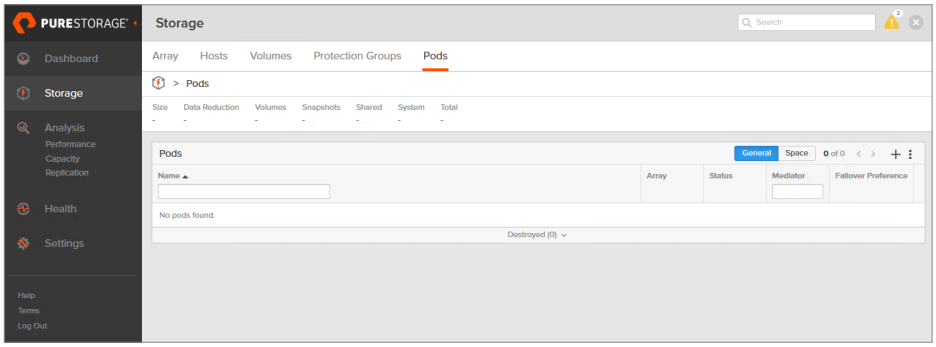

Identify the section for Pods and the three vertical ellipses and then select "Create...".

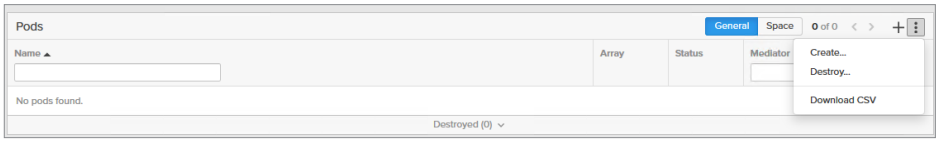

In the prompt which appears give the Pod a name and then select Create.

The POD will show up in the Pods section with the default mediator, pure storage, which is Pure1.

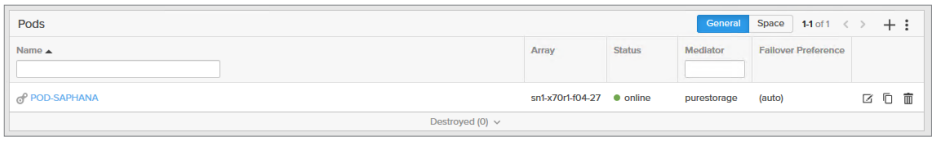

Selecting the Pod will show the management view for it including sections for Arrays, Volumes, Protection Groups, Volume Snapshots, Protection Group Snapshots and any miscellaneous details.

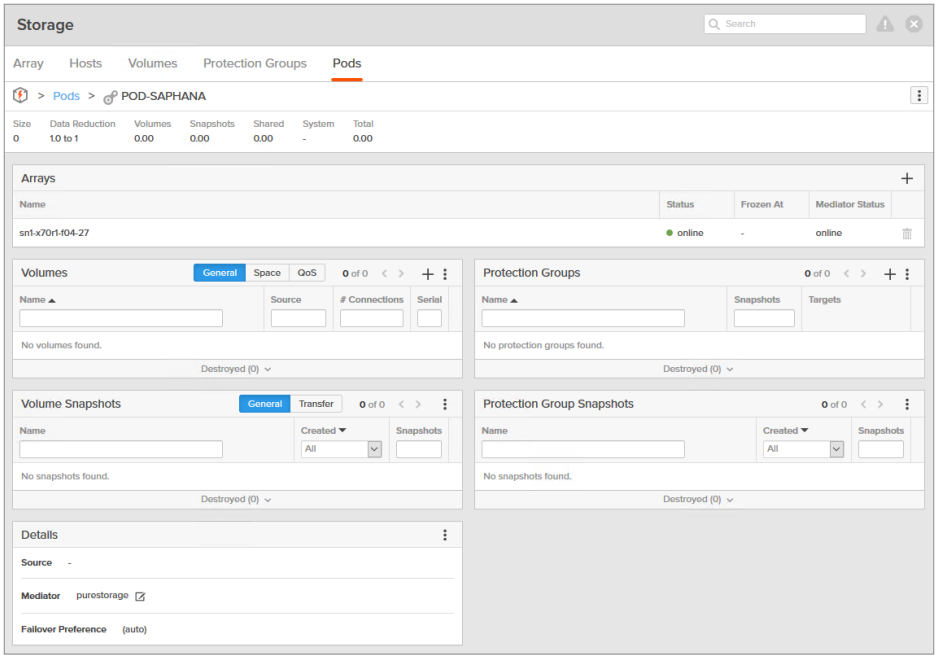

To add volumes to the pod, identify the volumes section and select the three vertical ellipses. To create a new volume in the pod, select "Create…", to move existing volumes into the pod select "Move In…".

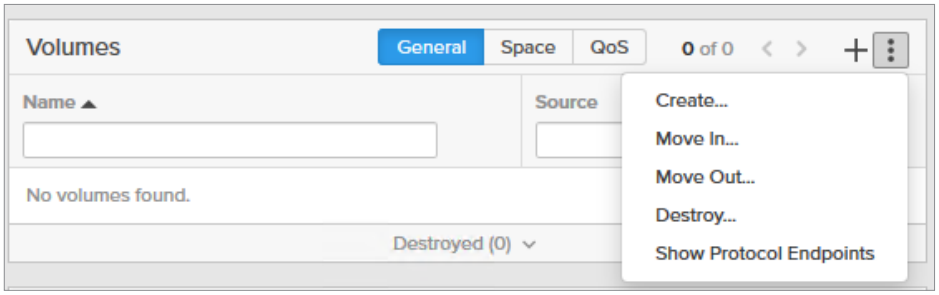

When creating a volume to add to the pod enter an appropriate name, provisioned size and unit, any bandwidth limits and select Create.

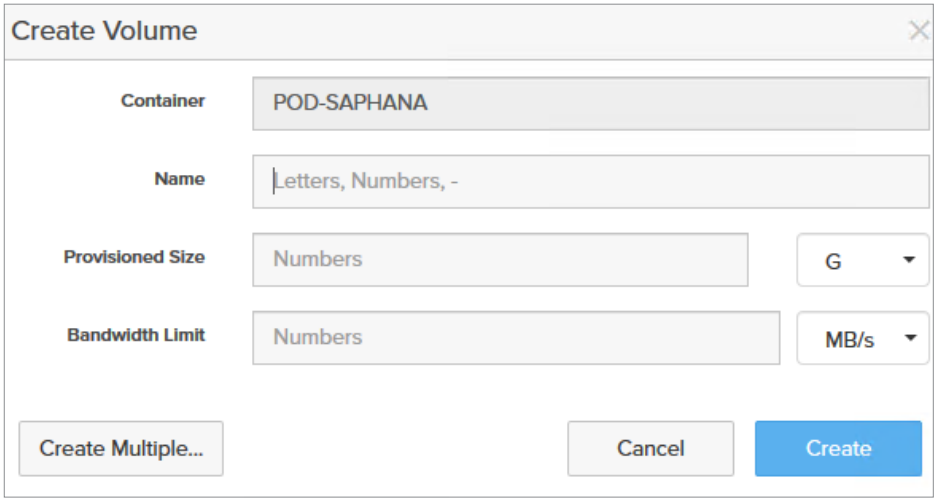

When moving volumes into the pod, select the relevant volumes on the left-hand window and then select "Move" to migrate the volumes to the right-hand window.

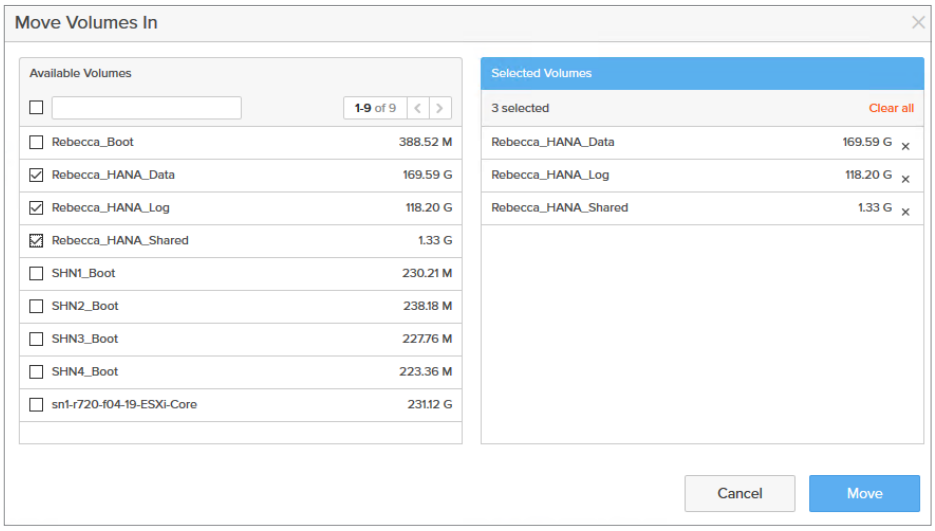

To stretch the pod to another array, select the "+" icon in the top right-hand corner.

Select the array to add to the pod from the drop-down menu and then select add.

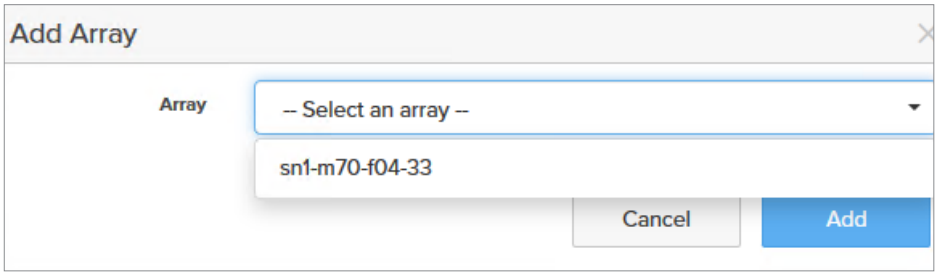

The array will show up in the pod's array view, with the status "resyncing".

When the pod has completed stretching (where all of the volume data has been copied to the other array) the Status will be shown to be online.

Step 3. Set Failover preference

In the event that a specific failover preference is required for the arrays in a pod, this can be changed in the details section of the pod view. The arrays for failover preference can then be appropriately set.

In the pod management view select the three vertical ellipses and select "Add Arrays to Failover Preference...".

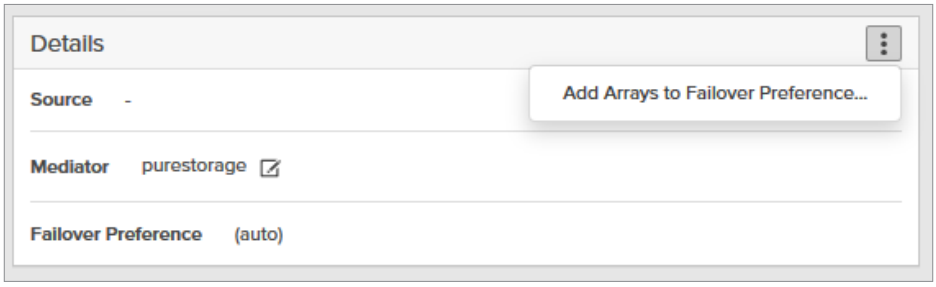

In the prompt displayed select the available array for failover preference in the left-hand window and it will appear in the Select Arrays window. Select Add to confirm this selection.

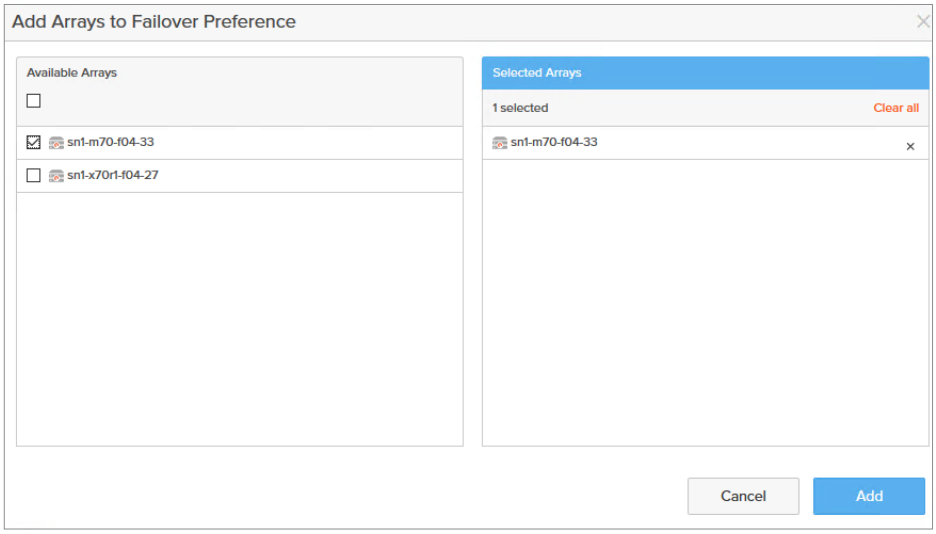

Step 4. Set Preferred array access

To set preferred array access, on the local array the host entry can be set to use which array is local to it. This is appropriate when active cluster is deployed using a multisite scenario where the speed and responsiveness of connectivity will vary. This is a way of ensuring storage requests are routed to a desired array.

In the storage view navigate to the hosts tab and identify a host which the volumes in an ActiveCluster configuration are a part of.

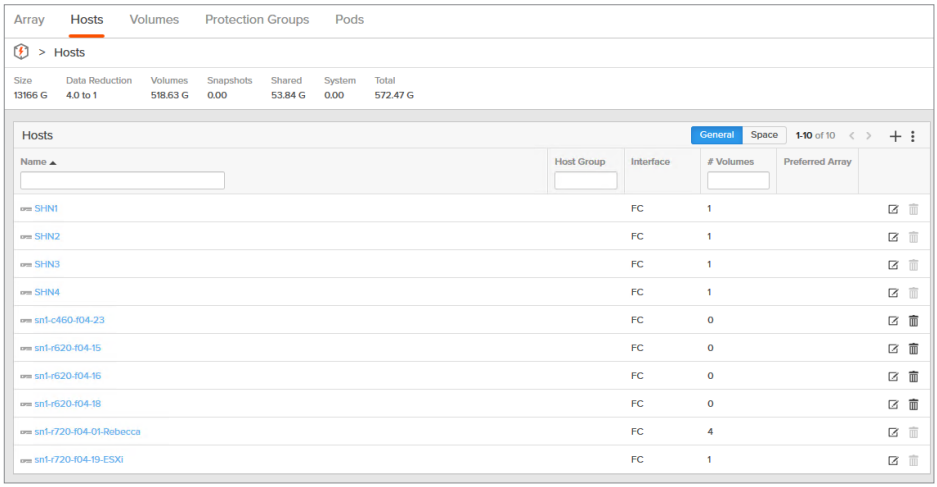

Select the relevant host and navigate to its management view.

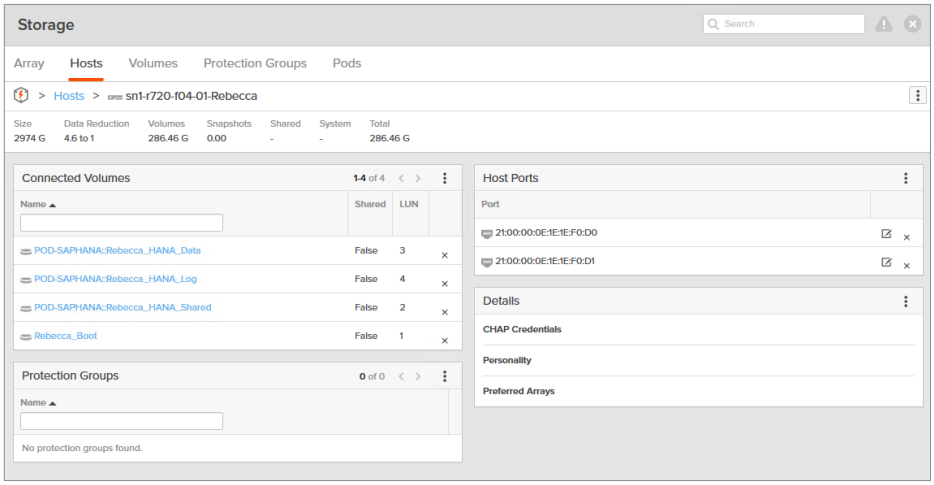

Identify the details section and select the three vertical ellipses and then select "Add Preferred Arrays".

In the prompt select the preferred array, the array local to the system, and then select add.

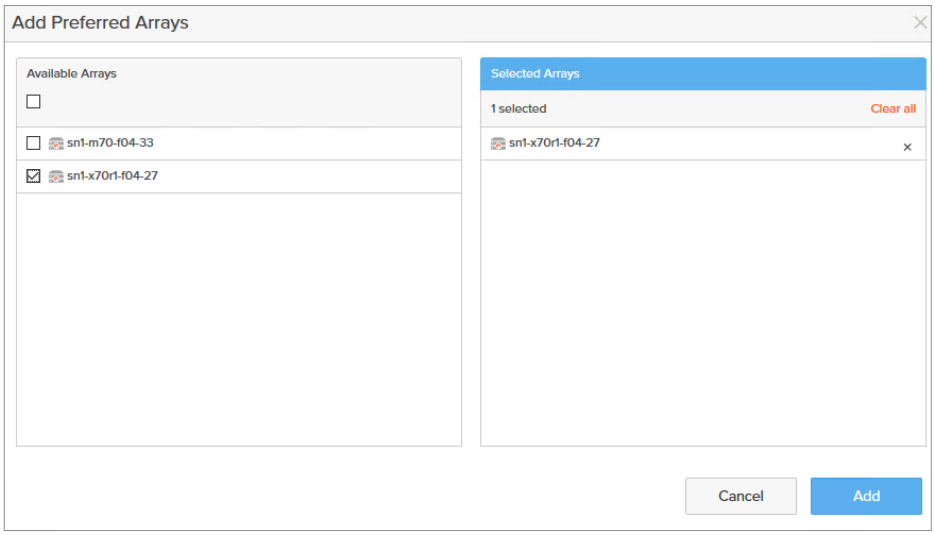

Removing ActiveCluster Configuration

In the event that replication needs to be terminated, pod volume membership changed or replication needs to be temporarily suspended the POD can be “unstretched” to remove operations from the other arrays. If a pod is stretched to two or more arrays and only one array is needed, the FlashArray no longer required is removed from the pod. Before removing a FlashArray from a pod the following must be done:

- Disconnect the pod’s volumes from the host that is connected to the FlashArray that is being removed from the pod. Purity will not allow un-stretching a pod if even one of its volumes is connected to a host or host groups.
- Ensure that any hosts using those volumes have access to the FlashArray that you intend to keep in the pod. If hosts only have access to the pod you plan to remove through the FlashArray that is being removed, they will lose access to their volumes once the FlashArray is removed from the pod, which may result in an application error based on the setup.
- Ensure you are removing the FlashArray you intend to – double-check the array before removing it.

Navigate to the pod management view for a specific pod in the storage view under Pods.

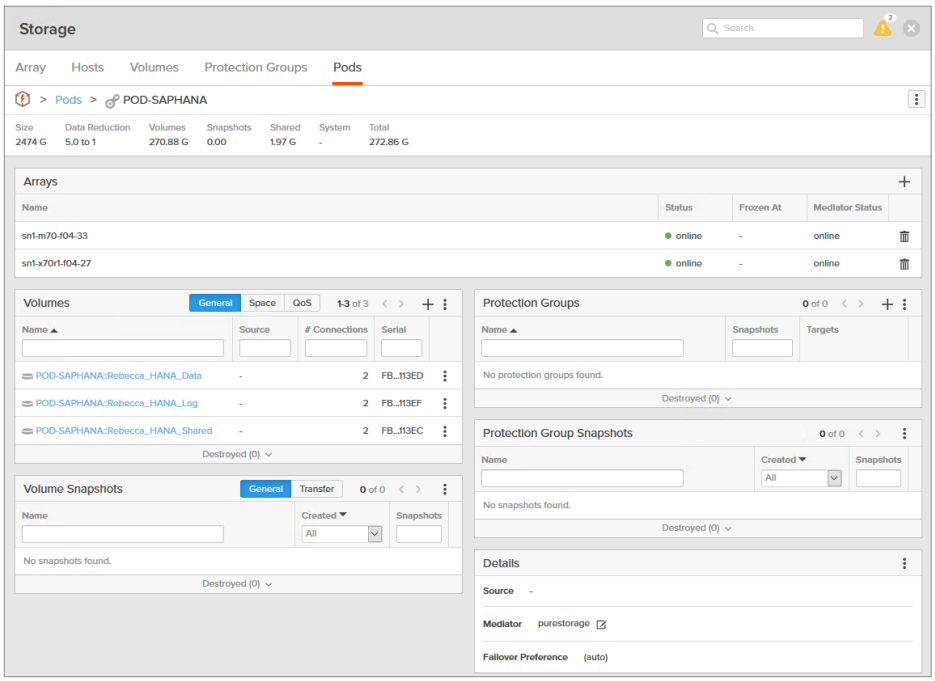

Identify the section for Arrays. To remove an array, select the “bin” icon for the relevant array to remove.

A prompt will show, ensure the correct array is being removed and select remove. This will cause the pod to be "un-stretched".

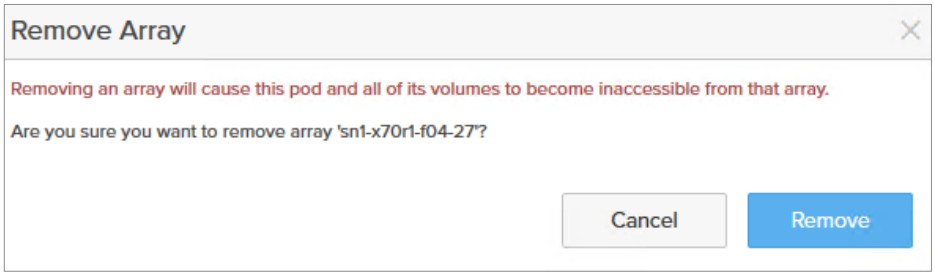

Multi-Site Disaster Recovery

Overview

Multi-Site Disaster recovery is delivered through FlashRecover Replication – a snapshot-based asynchronous replication solution leveraging space efficient snapshots to replicate point-in-time consistent copies of one of more block storage volumes. This solution combines snapshots and replication into a single concept allowing the creation of multiple recovery points on a replication target. The replication target can be an on-premise FlashArray or Cloud Block Store instance, allowing for organizations to adapt this solution towards the required business need, be it a low recovery point objective or multiple site protection.

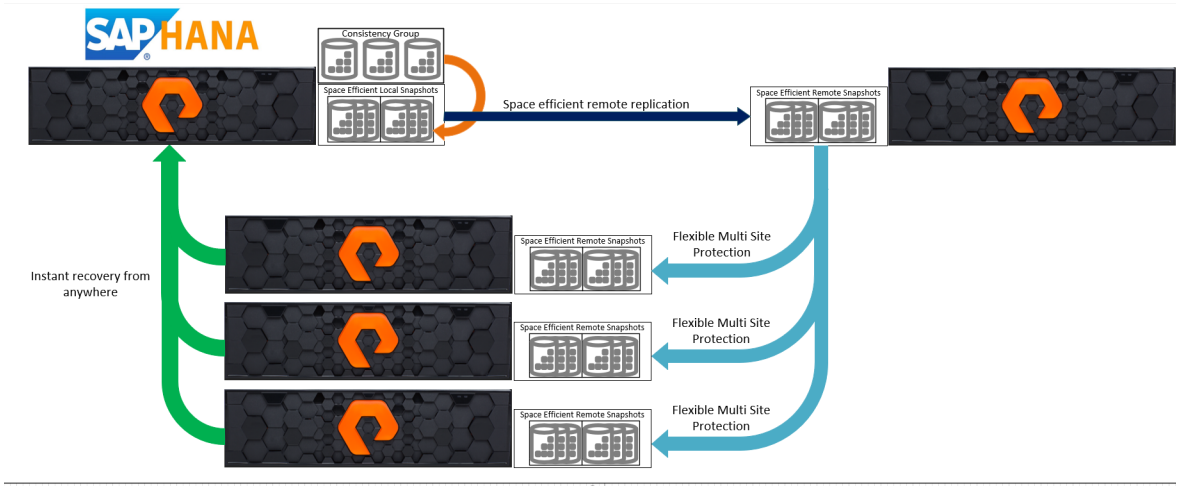

Fig. 16 Multi-Site Disaster Recovery achieved through FlashRecover replication.

Configuration and operation

In order to configure FlashRecover replication the following requirements must be met:

— The source and target arrays must be connected to ethernet switches.

— Each array must have a replbond interface or interfaces configured for use with the replication service.

In the FlashArray web user interface, in the Storage view navigate towards the Network tab to identify and configure the network settings. To configure a new interface or an existing one, use either the + in the top right-hand corner or the edit clipboard alongside each individual interface.

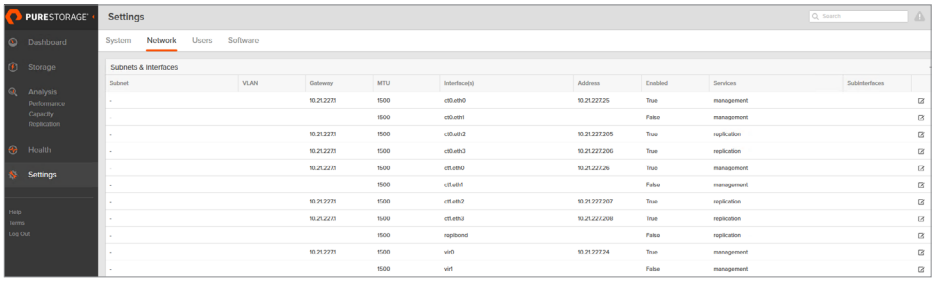

— Any replication interface must be configured with an IP address, netmask and (optionally) gateway while being set to enabled on all arrays.

— Connecting arrays to one another requires access on port 443 and 8117. In the event of cross site replication through firewalls these ports must be set to allow transfers.

Step 1. Connect to the remote storage array(s) and create an asynchronous connection

Connecting to another FlashArray requires that a connection key is acquired from the corresponding array attempting to be connected to. This key is used as a secure way of ensuring only authorised individuals can connect to the other array. The other array can be connected to using the virtual IP address of fully qualified domain name of the remote array.

In the FlashArray web user interface for both the local and remote system, in the Storage section navigate to the Array tab.

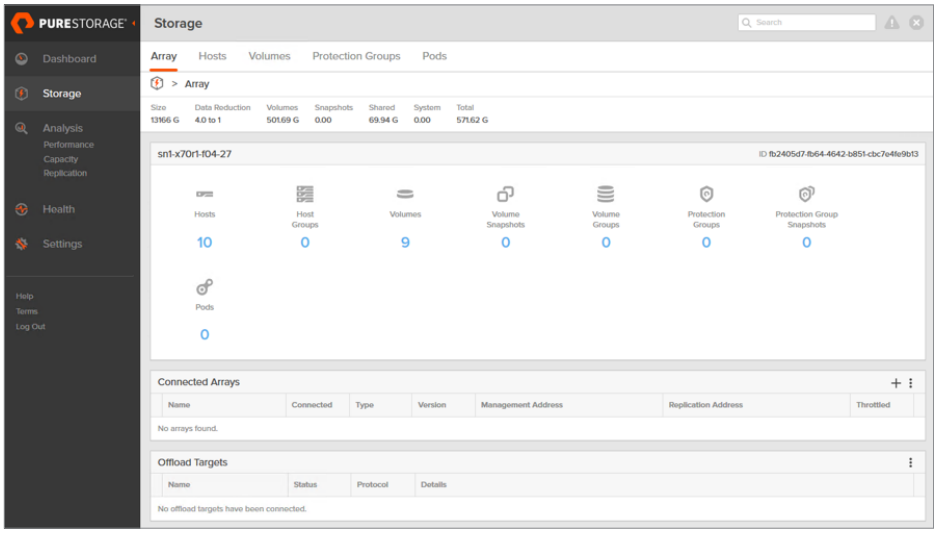

In the interface for the array to be connected to identify the section for connected arrays and select the 3 vertical ellipses and select "Get Connection Key".

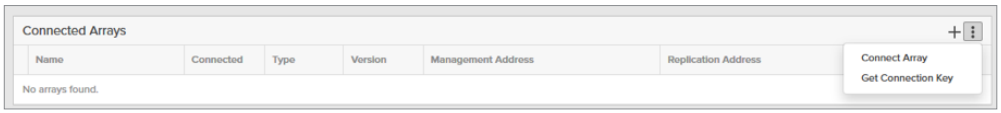

A prompt will show up with the connection key for that array, Select Copy and then navigate to the array to connect from.

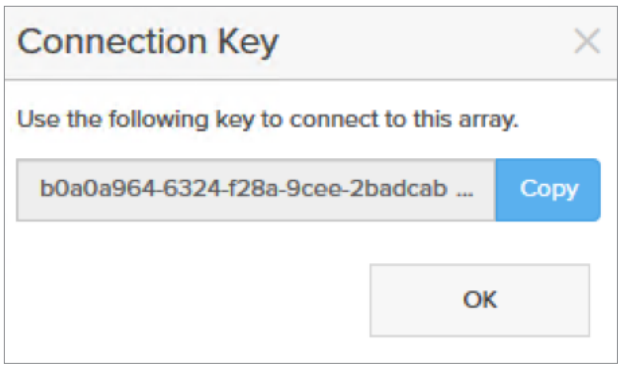

In the user interface for the FlashArray to connect from, return to the Connected Arrays section, select the three vertical ellipses and select "Connect Array".

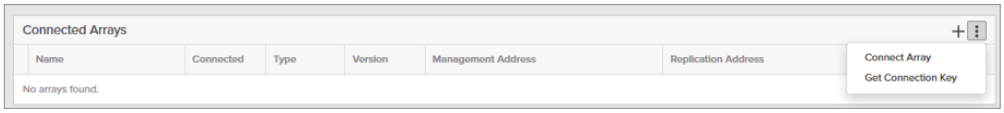

A prompt will show, enter the relevant information for Management Address, Type (Asynchronous replication), the connection key and any replication addresses required. Select Connect when all of the relevant information has been entered.

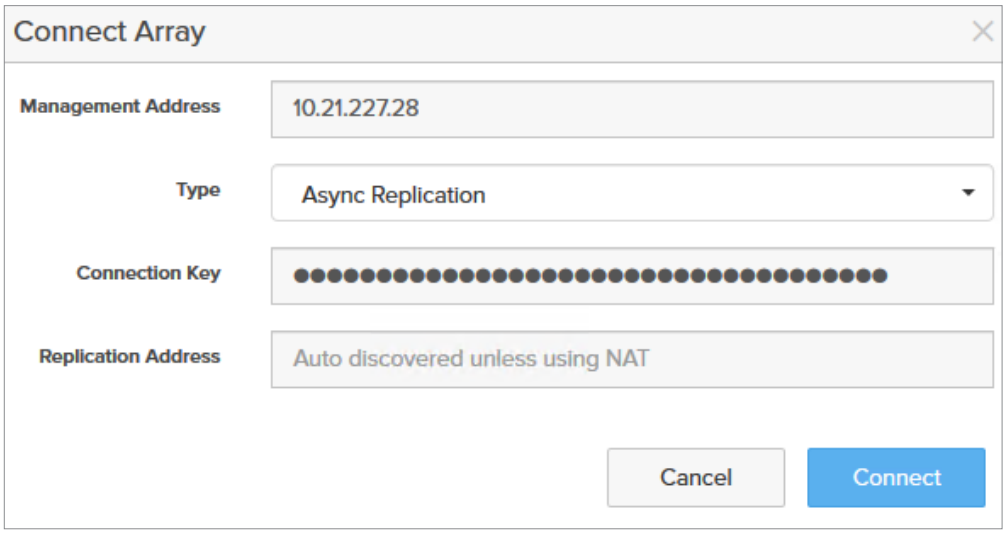

Once connected, the connected array will show under the relevant section.

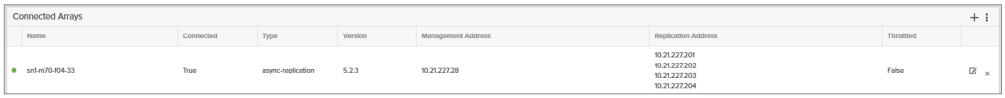

Step 1a. Set throttling for asynchronous replication target

When organizations need to limit the amount of bandwidth used by a service, to ensure that other services are still capable of communication without a delay due to saturation of the network link, a network throttling policy is usually applied. With asynchronous replication a throttling policy is applied per connected array. Throttling policies can be applied to bandwidth limits at all times or during a time range.

For the array which needs to be throttled, select the "Edit Bandwidth Throttle" icon on the left-hand side.

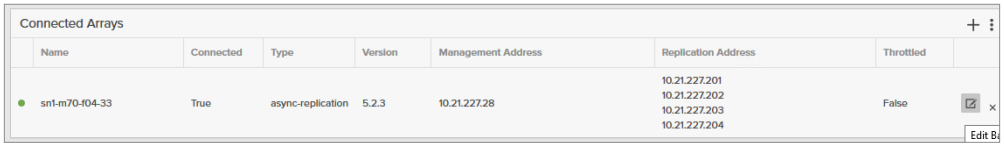

In the prompt which appears a Default throttle can be applied or a window throttle.

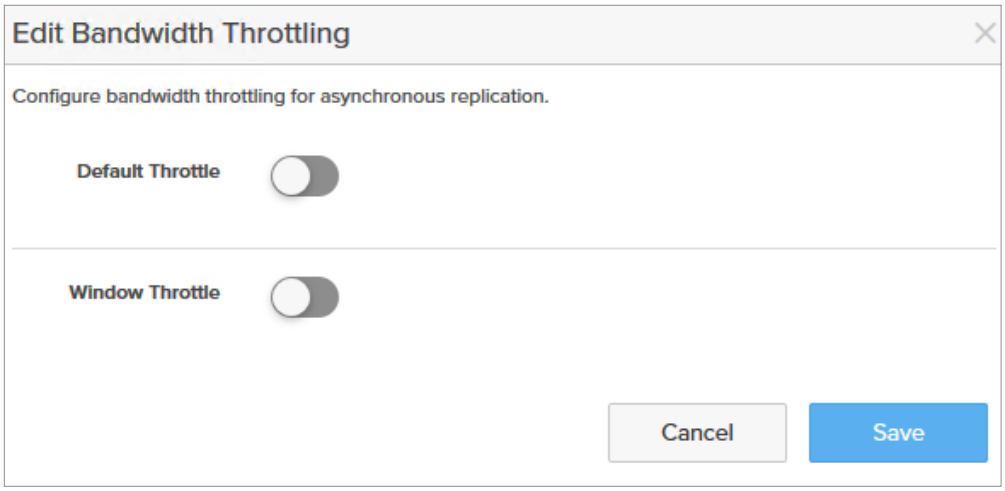

For the default throttling which will be on at all times , set the desired bandwidth limit.

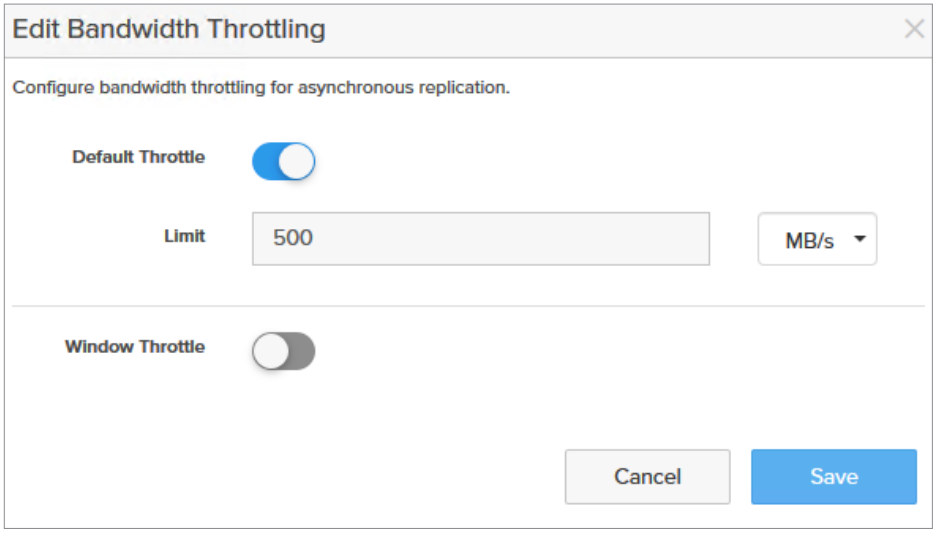

For throttling to be applied during a period of time, set the time range and the bandwidth limit.

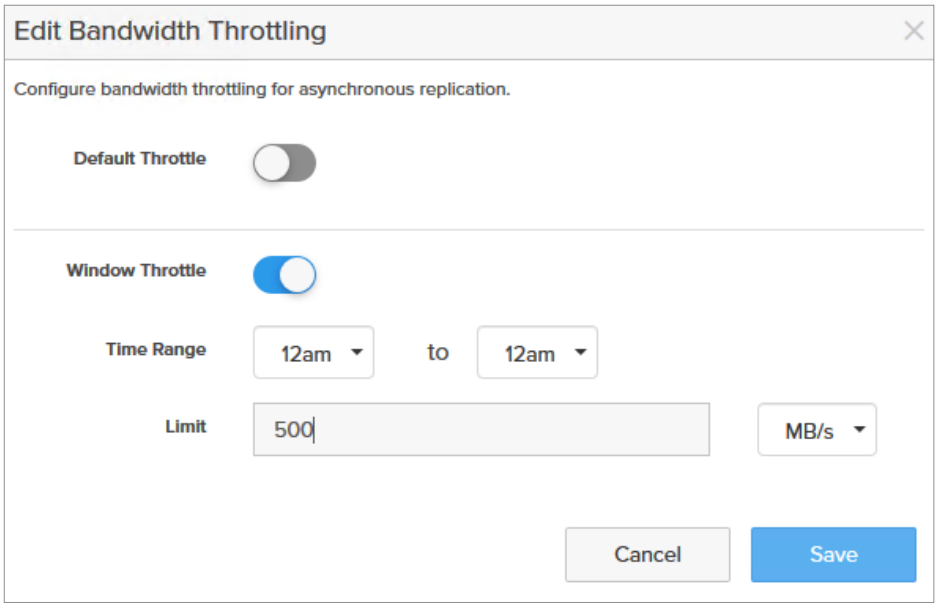

The array will now show us as throttled.

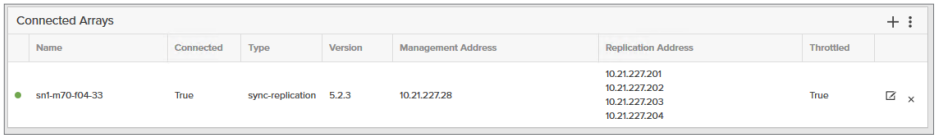

Step 2. Create a Protection Group

FlashRecover makes use of a single interface to manage the creation and/or replicate of snapshots. This interface is located in the Storage View under Protection Groups. A unique name must be given to the Protection Group.

In the Storage view, the Protection Groups tab is used to manage consistency groups of volumes, hosts and host groups.

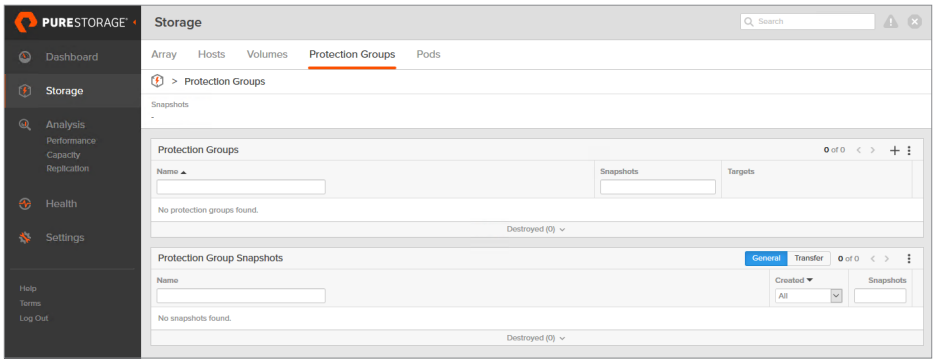

In the Protection Groups tab, identify the section for Protection Groups where all existing Protection Groups will be located. To create a new Protection Group, select the 3 vertical ellipses in the top right-hand corner and select "Create… ".

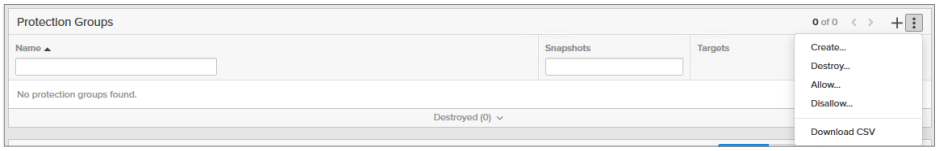

When the prompt is shown give the protection group an appropriate name and then select "Create".

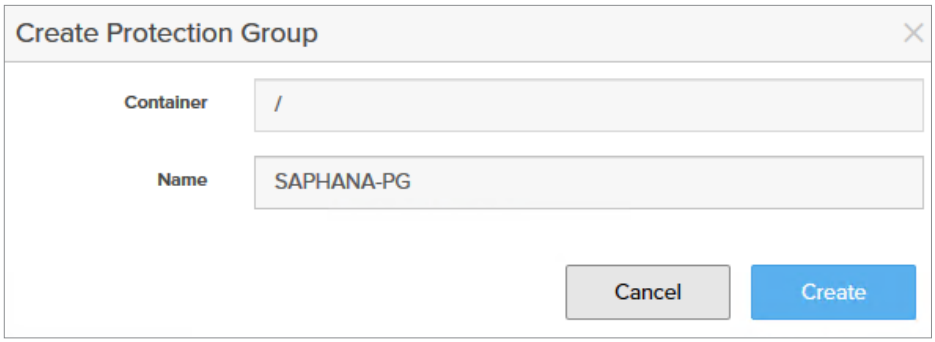

The new Protection Group will now be shown.

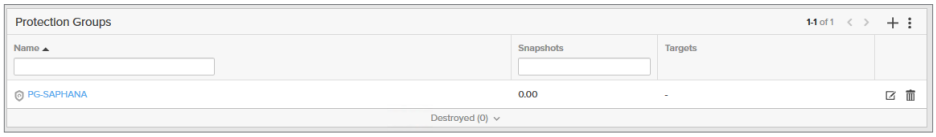

Selecting the Protection Group will navigate towards the management view for it.

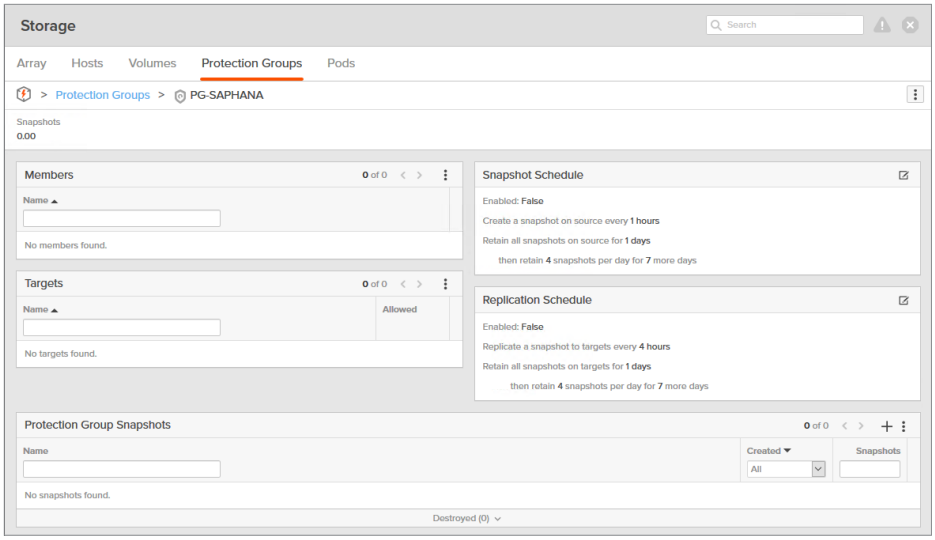

Step 3. Add Asynchronous replication target(s)

In the management view for the specific protection group, the area for "Targets" allows for multiple FlashArray or Pure Cloud Block Store instances to add. Each system added as a target in the Protection Group will be a target for any replication policies set up for it.

In the management view for the Protection Group, identify the section for "Targets". Select the three vertical ellipses and select "Add...".

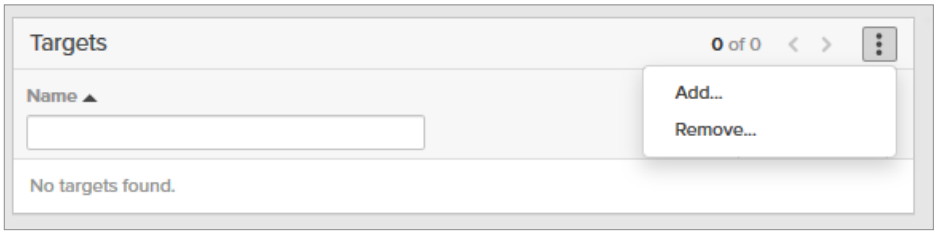

Identify any available targets, select them and then select "Add" to complete the process.

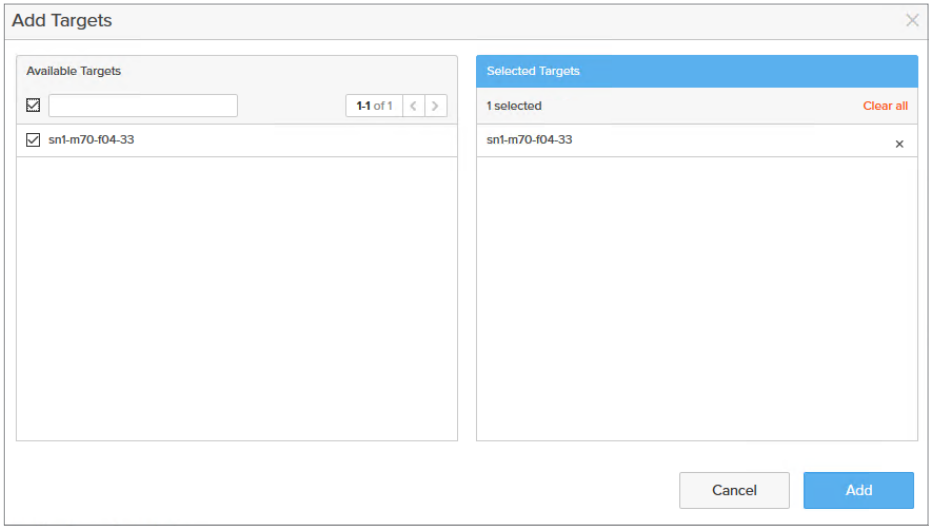

The target will show up in the Protection Group management view once it has been added.

Step 4. Add Host, Host Group or Volumes to Protection Group

In order to asynchronously replicate volume data, the volumes required to be replicated must be added to the Protection Group. Volumes can be added by making the host or host group to which those volumes are attached, or just adding the volumes themselves, a member of the Protection Group. Any volumes added to a Protection group are considered a consistency group, ensuring that all of the data on each volume is synchronised to the same point in time when a snapshot is created.

In the Protection Group Management view, navigate to the Members section and select the three vertical ellipses. Then select the type of member to add.

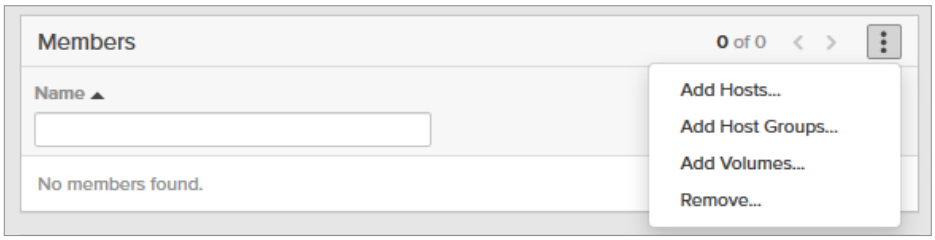

To add the volumes attached to a host as a member, select the host and then select Add.

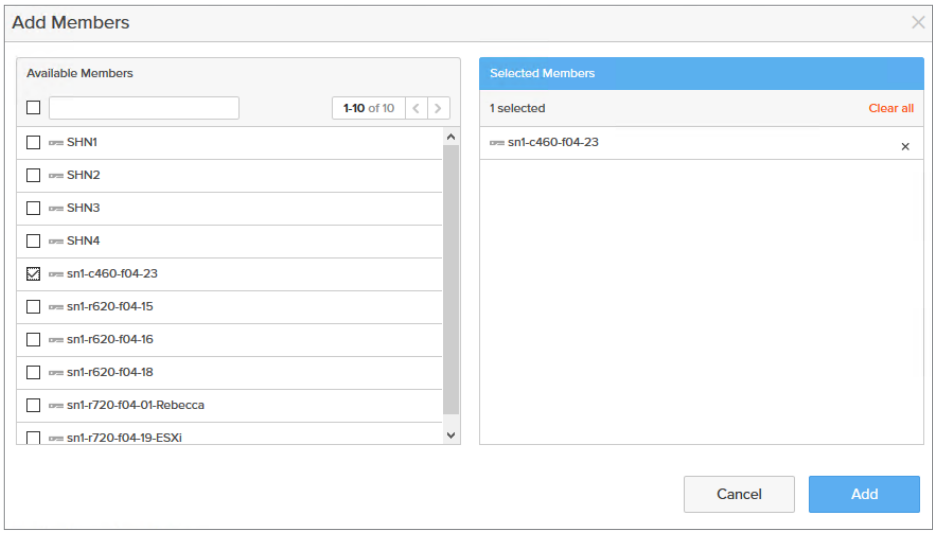

To add volume(s) as a member select the volumes and then select Add.

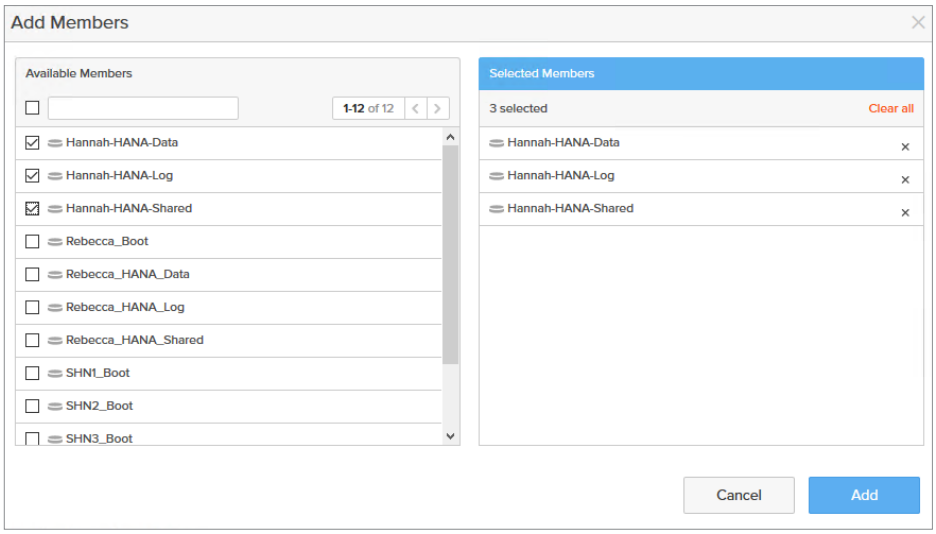

To add any volume(s) attached to all of the hosts in a host group select the host group and then select Add.

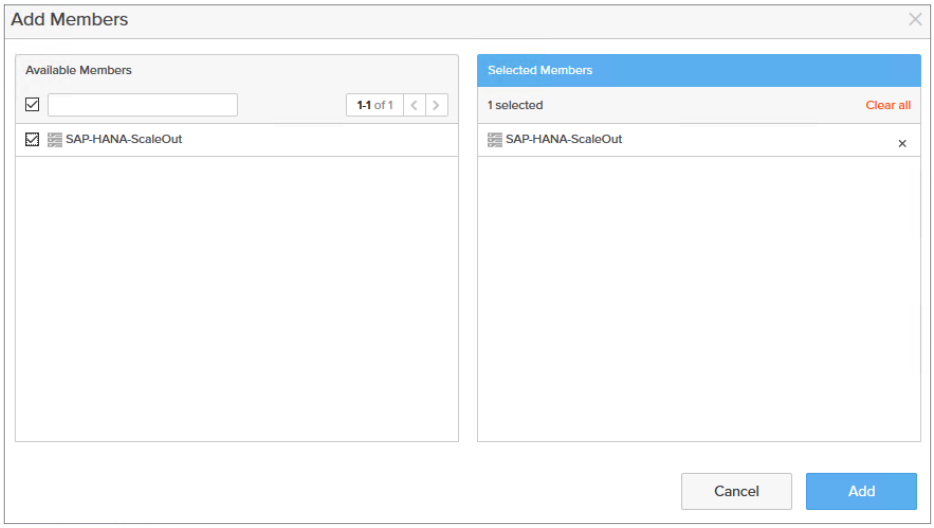

Step 5. Set Snapshot Schedule (Optional)

A snapshot schedule will create a snapshot for any volumes present as a member of the Protection Group, on a set time basis. The retention of snapshots can also be managed here ensuring that effective capacity management is exercised. The recovery point objective (RPO) is set here, for an hour RPO a new snapshot must be created every hour and retained for long enough that it can be utilised as a recovery point. This only creates volume snapshots on the local array.

In the Protection Group management view, identify the section for Snapshot Schedules. To edit the snapshot schedule, select the edit box in the top right-hand corner.

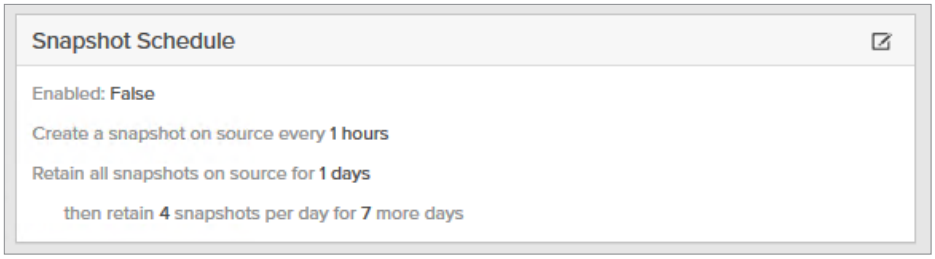

A prompt will show, ensure that snapshot schedule is enabled and the correct values for creation frequency and retention are entered. Select Save when the values are satisfactory.

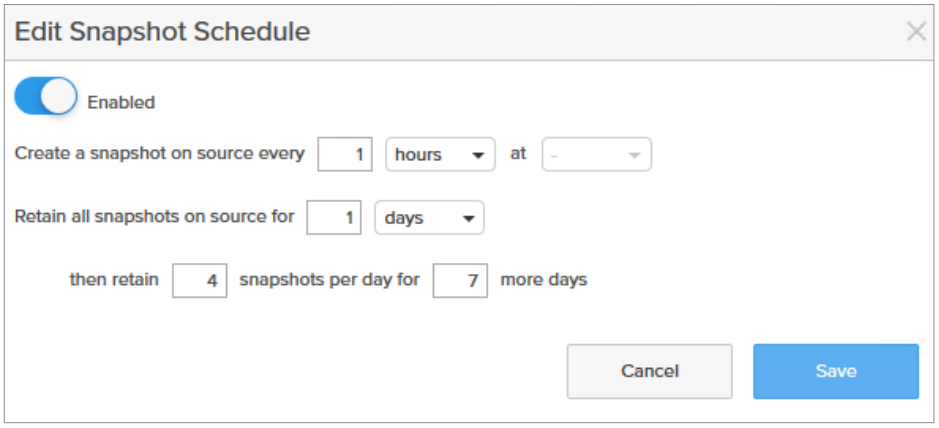

After some time, snapshots created by the policy will be shown in the Protection Group management view.

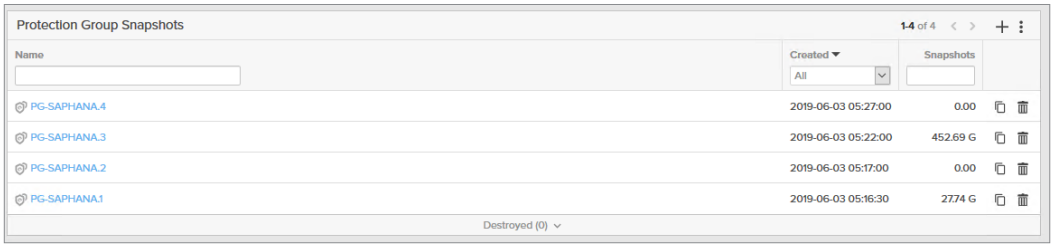

Step 6. Set Replication Schedule

To enable multi-site disaster recovery the snapshots being created on the FlashArray or Cloud Block Store instance must be replicated to another location. This is accomplished by using a replication schedule, also used to manage retention of those snapshots on the target(s) for capacity management. All targets added in step 3 will be replicated to when setting this policy and any volumes added as members of the Protection group will be replicated as a part of this policy.

In the Protection Group management view, identify the section for Replication Schedules. To edit the replication schedule, select the edit box in the top right-hand corner.

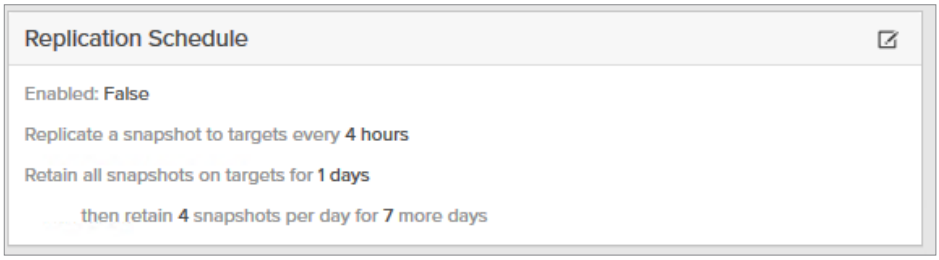

In the prompt set the replication schedule to enabled and then set the relevant values for when to replicate a snapshot, any exclusion times and the relevant retention information.

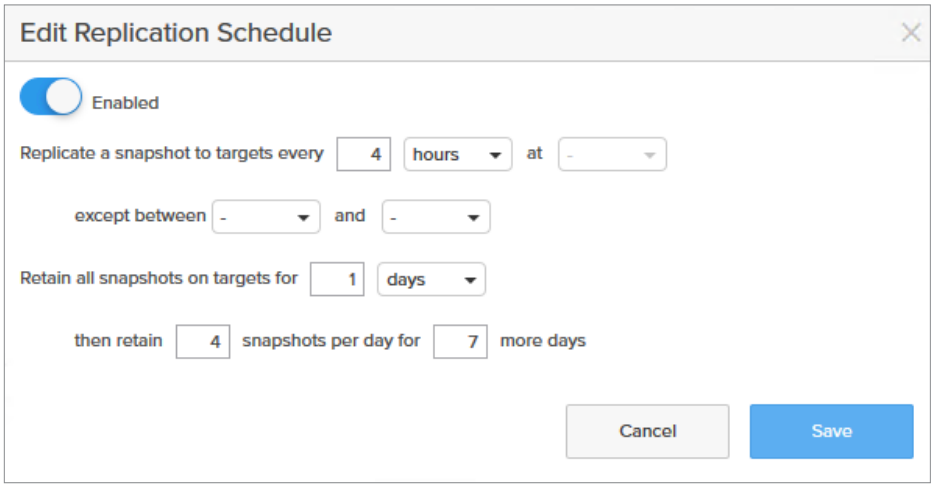

Adding additional replication target(s)

FlashRecover can perform one to many, many to many and many to one replication techniques. In order to add an additional replication target, repeat the process for adding a connected array in Step 1, and then adding that array to the relevant protection group as set out in Step 3.

In the FlashArray web user interface find the connected arrays section under the Array tab in the Storage view. Once the array has been added it will display alongside any other connected arrays.

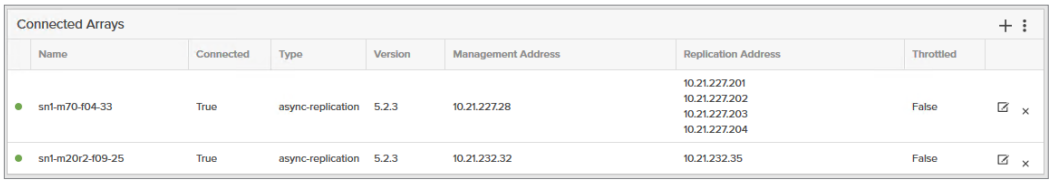

In the Protection Group management view , identify the section for targets and select the three vertical ellipses to add a new target.

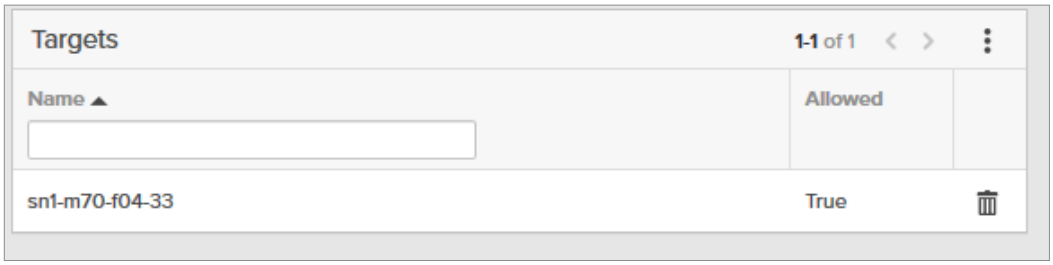

In the prompt which appears select the array to add as an asynchronous replication target and then select Add.

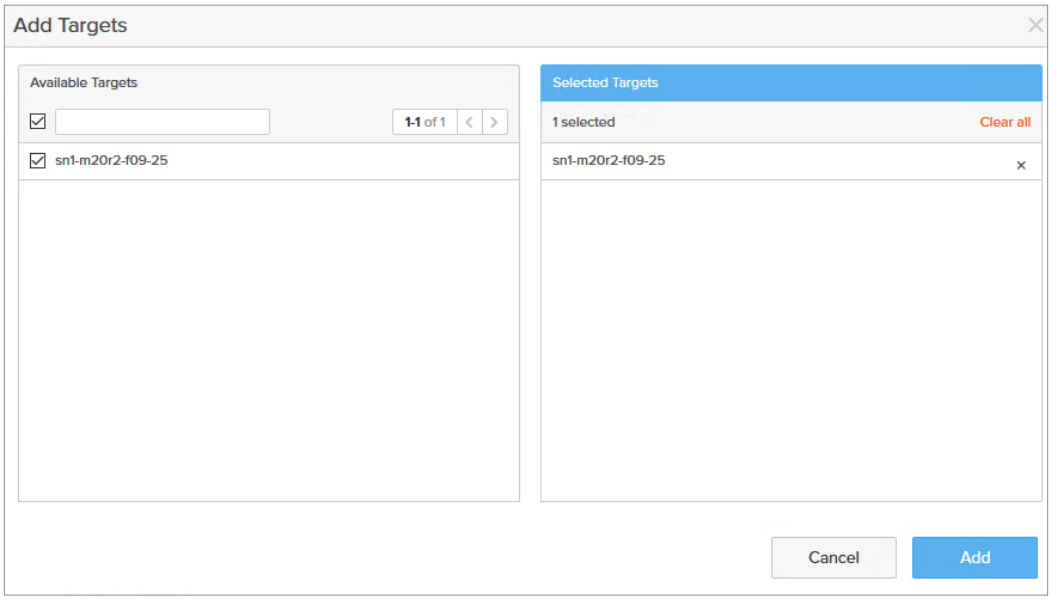

Once added to the Protection Group the additional array(s) will be shown in the targets view.

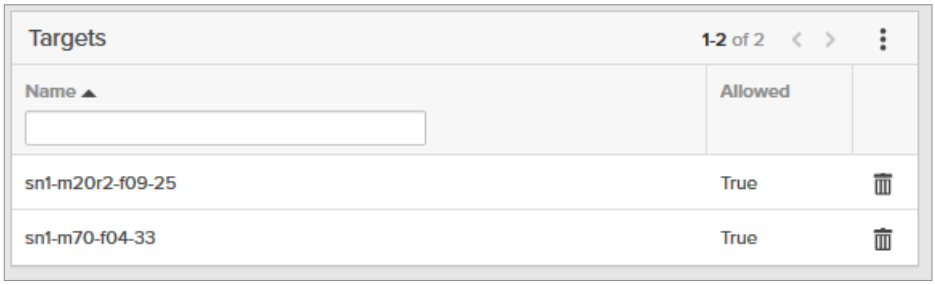

Once replication has occurred, snapshots of the Protection group will be shown in the Protection Group management view created with the source array as the name – appended with the relevant protection group name.

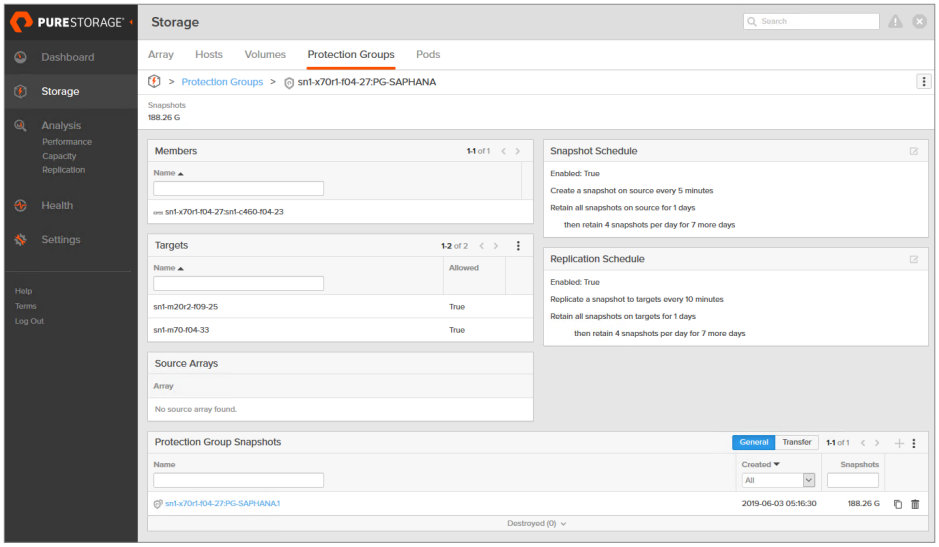

Adding an asynchronous snapshot replication target to an ActiveCluster configuration

Third site Disaster recovery for Active Cluster

Any ActiveCluster configuration can be extended for disaster recovery scenarios by having a third array configured for asynchronous snapshot replication for an existing pod. In the event of the first two arrays failing, data on the third site can be resynchronised back to the first two sites allowing for business processes to resume in a rapid fashion. Any third site array can be used by multiple ActiveCluster configurations.

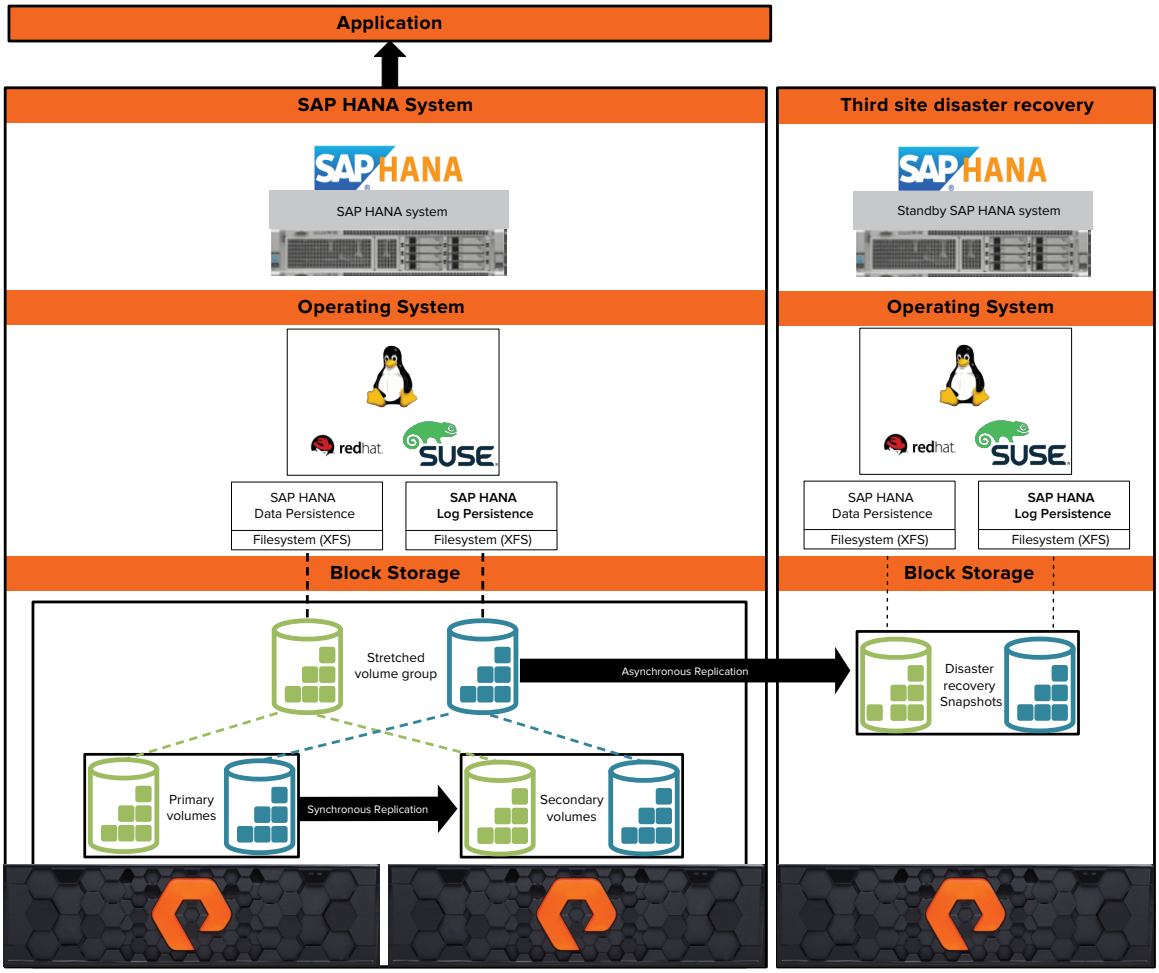

Fig. 17 Extending uniform host access to a third site for DR purposes.

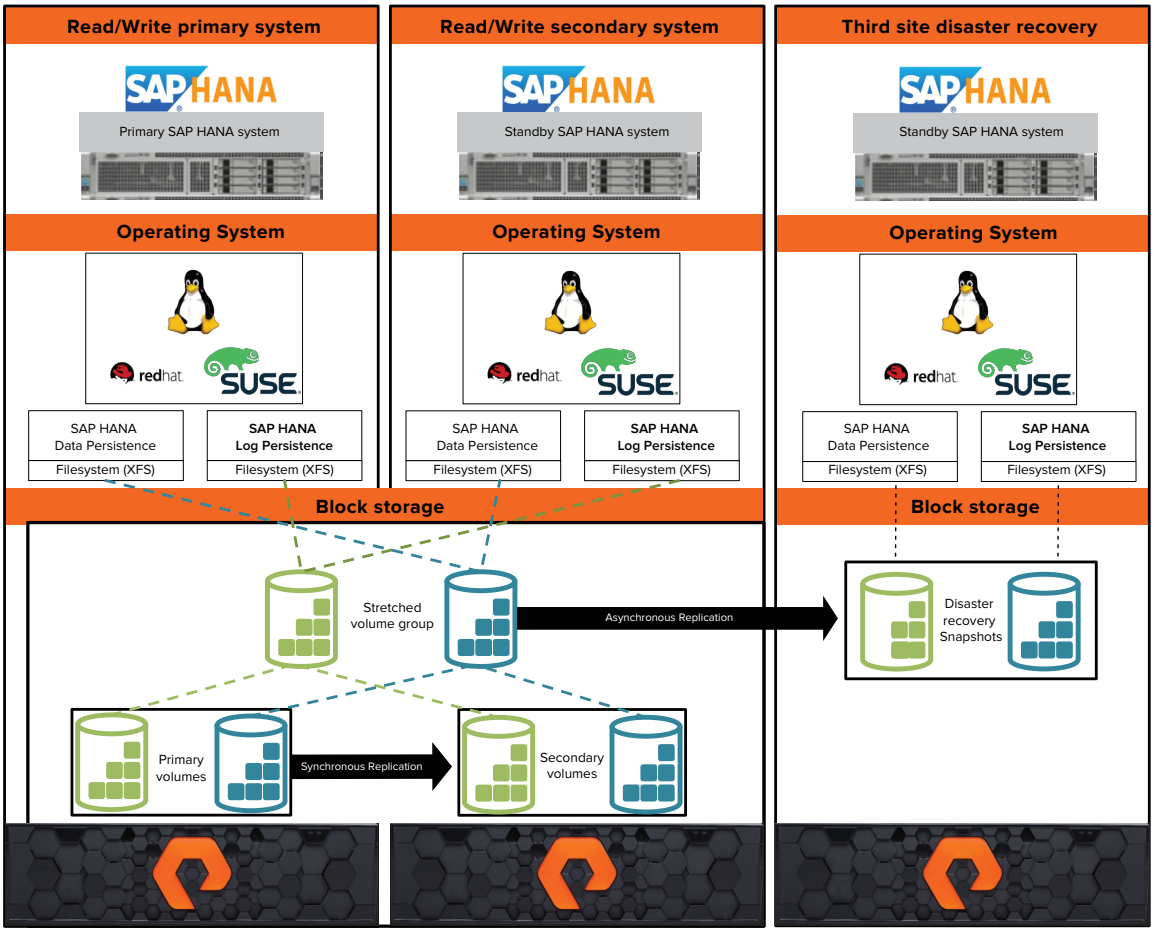

Fig. 18 Extending Symmetric Active/Active host access to a third site for DR purposes.

With the release of Purity 5.2, synchronous (ActiveCluster) and asynchronous (FlashRecover) functionality can be combined to achieve a zero-recovery target and zero recovery point objective and further enhance business continuity solutions. Volumes must be added to a protection group before being added to a pod for active cluster in order to configure this solution.

In the Pure storage web user interface, where an existing active cluster configuration is present find the section for Connected Arrays and select the "+" to add a new array.

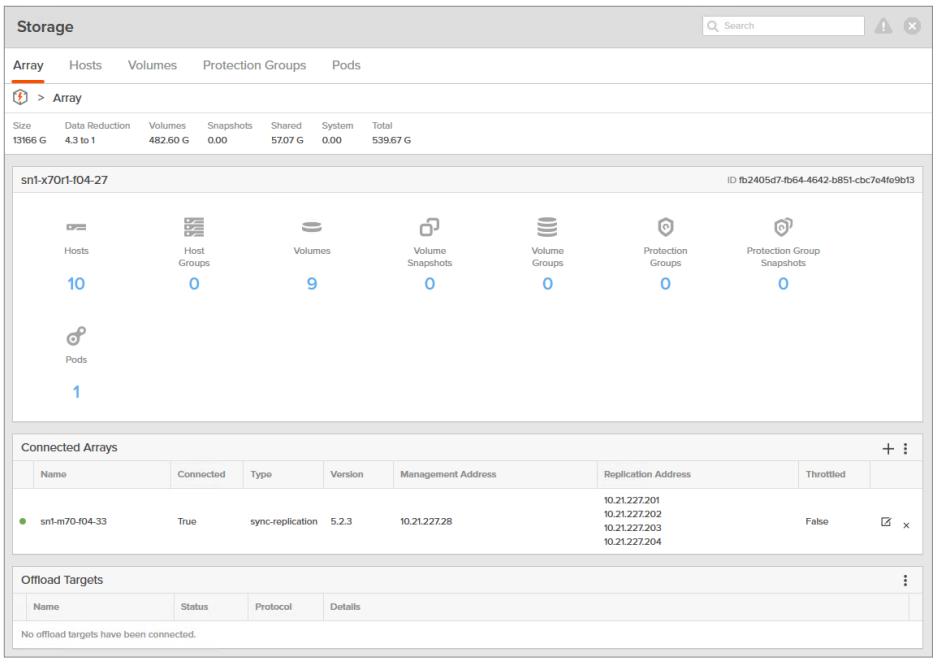

Connect the third array as an asynchronous replication target, ensure the connection key and any relevant replication addresses have been retrieved from it to enter here.

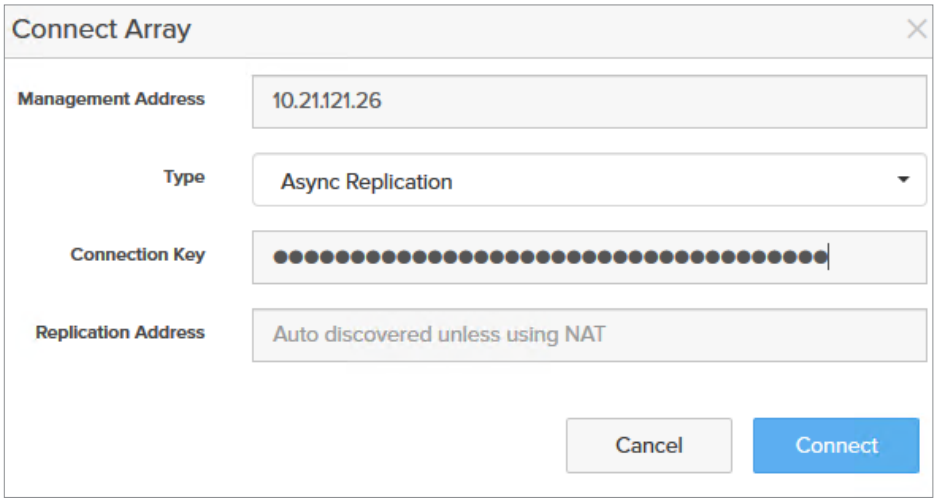

The array will show up as an async-replication target type under connected arrays.

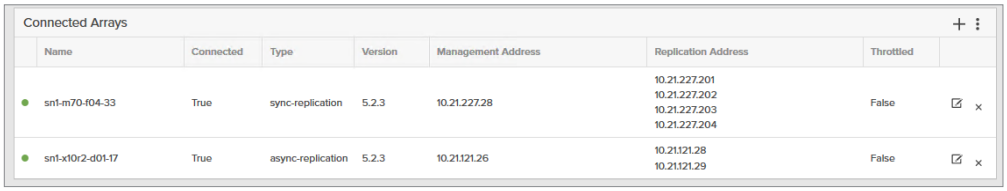

Navigate to the management view for the Pod for the SAP HANA instance.

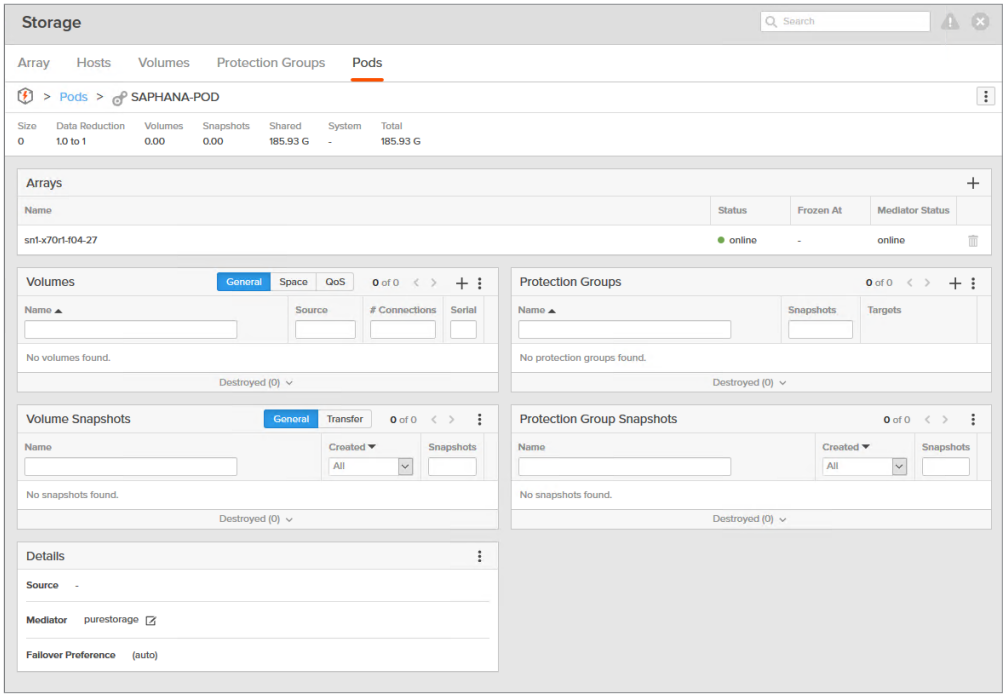

In the pod identify the Protection Groups section and select the three vertical ellipses and then select "Create...".

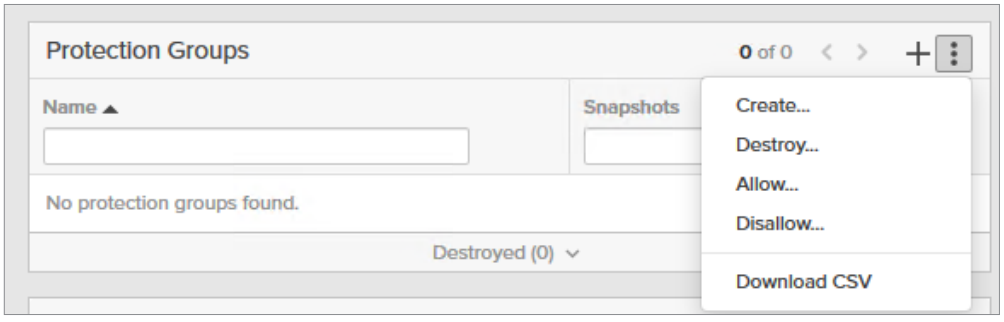

In the prompt ensure the Container is the same name as the pod, and then give the protection group a unique name before selecting Create.

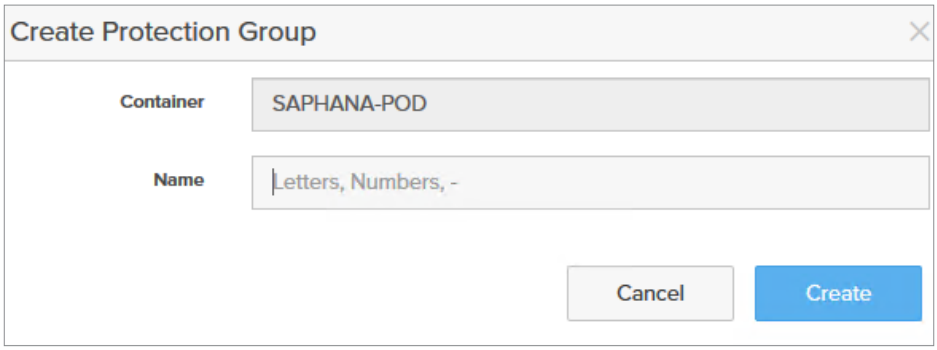

The Protection group will now be shown in the pod.

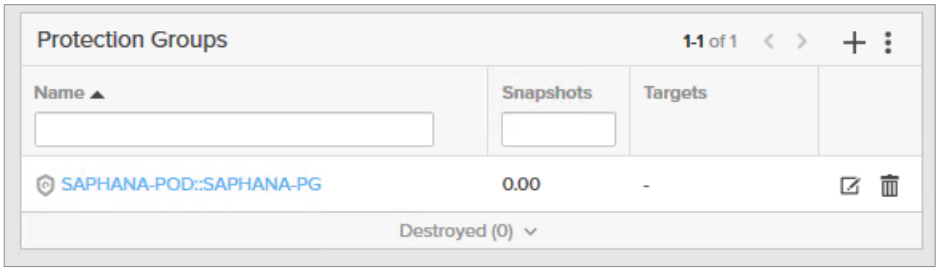

Navigate to the Protection group name which has been created as a part of the pod (the naming convention will be POD NAME::PROTECTION GROUP NAME).

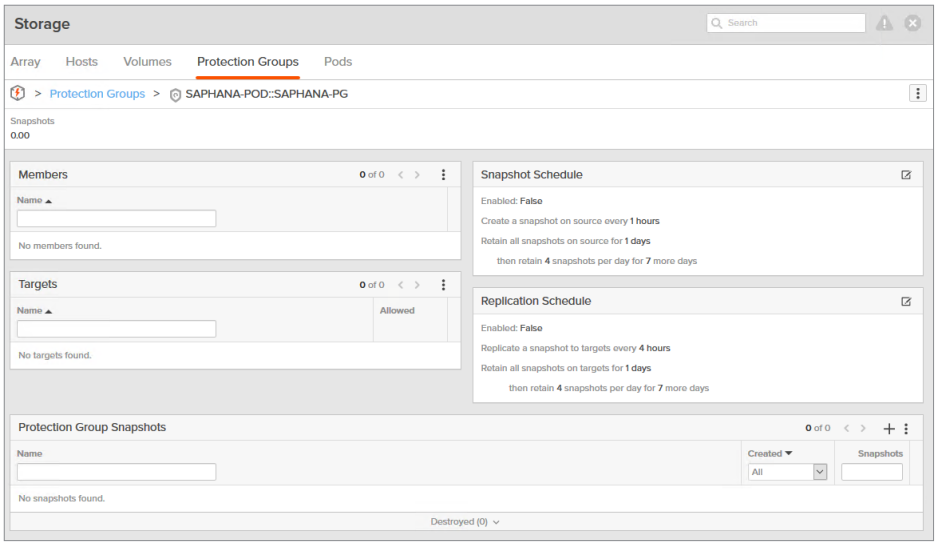

Add the volumes which will be synchronously replicated in the ActiveCluster configuration, to the protection group.

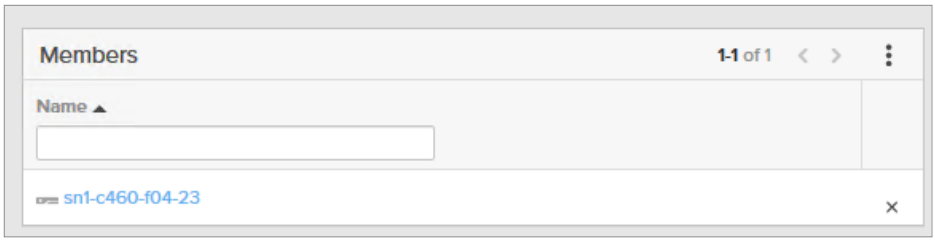

Add the asynchronous replication target to the protection group. Do not add the ActiveCluster synchronous target.

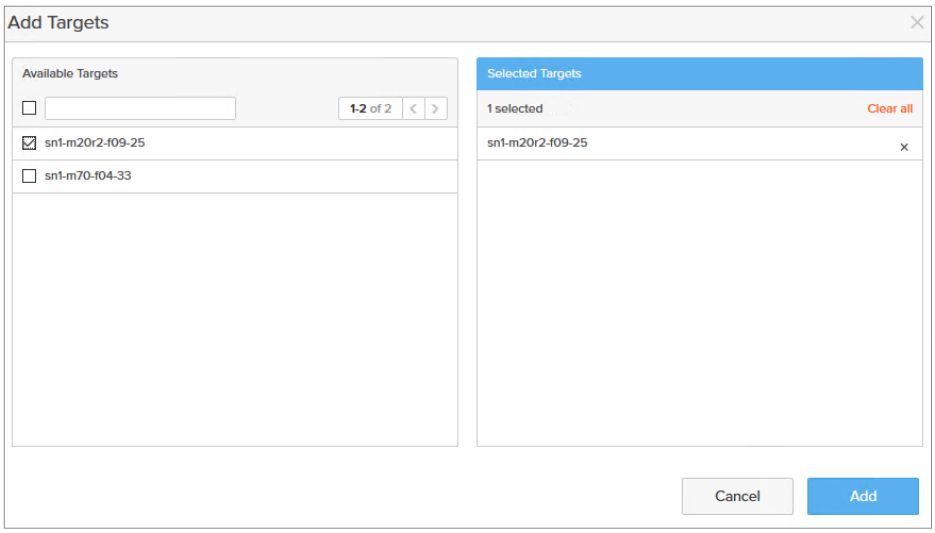

Configure the replication schedule for the protection group.

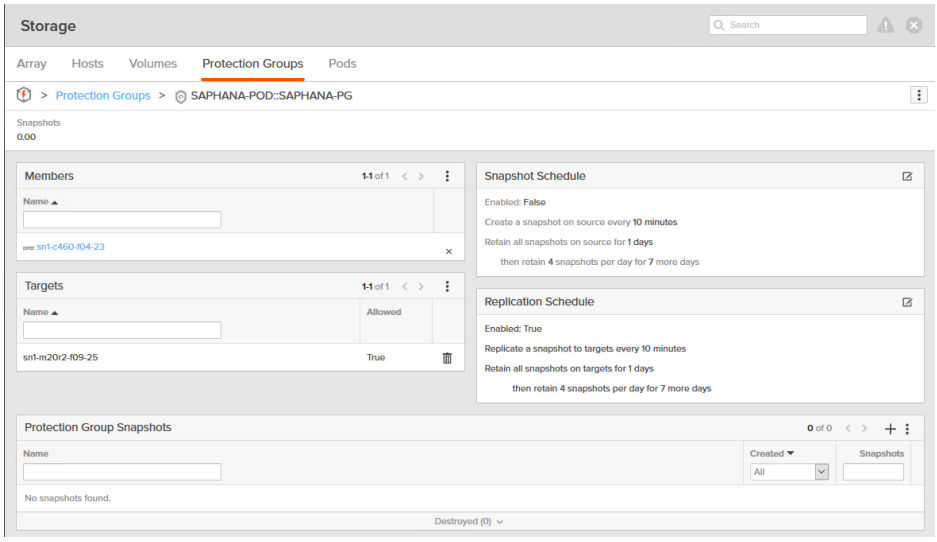

Return to the pod management view and add the array for synchronous replication to it.

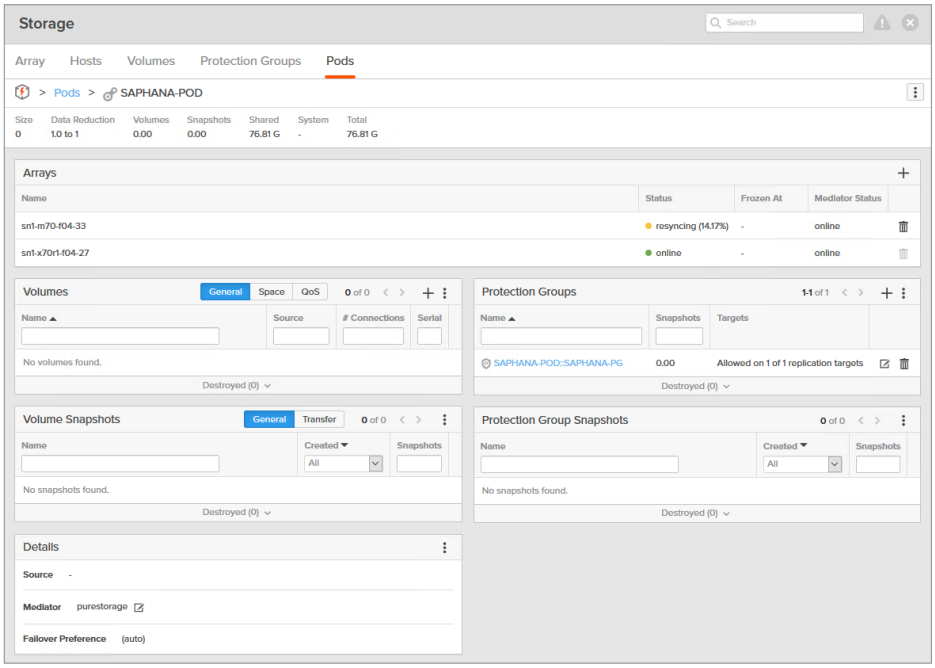

Once the array has completed synchronous replication it will display its status as online.

References

SAP HANA Hardware and Software Requirements

SAP HANA Supported Operating Systems

SAP HANA Revision Strategy for HANA 1.0

//X Product Data Sheet

Purity//FA Data Sheet

Cloud Block Store Data Sheet

Purity ActiveCluster Data Sheet

How to perform system replication for SAP HANA

SAP HANA Backup types

SAP HANA Intel DC Optane Memory

SAP HANA multitenant database

Best practices for using snapshots in the vSphere environment

SAP HANA System replication