



NetApp Verified Architecture

SAP HANA on NetApp All SAN Array

Modern SAN, Data Protection, and Disaster Recovery

Nils Bauer, Roland Wartenberg, Darryl Clinkscales, Daniel Hohman, Marco Schöen, Steve Botkin, Michael Peppers, Vidula Aiyer, Steve Collins, Pavan Jhamnani, Lee Dorrier, NetApp
Jim Zuccherro, Naem Saafein, Ph.D., Broadcom Brocade

June 2020 | NVA-1147-DESIGN | Version 1.0

Abstract

This NetApp® Verified Architecture covers modernizing SAP systems and operations for SAP HANA on NetApp All SAN Array (ASA) storage systems with Brocade FC SAN Fabric. It includes backup and recovery, disaster recovery, and data protection. The solution leverages NetApp SnapCenter® to automate SAP HANA backup, restore and recovery, as well as cloning workflows. Disaster recovery configuration, testing, and failover scenarios are described using synchronous NetApp SnapMirror® data replication software. Additionally, SAP Data Protection with CommVault is outlined.

In partnership with



TABLE OF CONTENTS

1	Forward – A Message from Broadcom/Brocade	5
2	Executive Summary	5
3	Program Summary	7
4	Solution Overview	8
4.1	Backup and Recovery	9
4.2	Disaster Recovery	11
4.3	SAP Lifecycle Management	12
4.4	Target Audience	12
5	SAP Data Protection Overview	13
5.1	Business Application Requirements	13
5.2	Backups	13
5.3	Synchronous or Asynchronous Data Replication	13
5.4	SAP Data Protection: NetApp + Commvault	14
5.5	Recover in Minutes Versus Hours	14
5.6	Automation and Optimization	14
5.7	Hybrid-Cloud: On-Premises or in the Cloud	14
6	Solution Components and Use Cases	14
6.1	Solution Components	15
6.2	Test Lab Setup	16
6.3	All SAN Array High Availability	16
7	Solution Verification	20
7.1	SAP HANA Backup	20
7.2	SAP HANA Restore and Recovery	27
7.3	SAP System Refresh	30
7.4	SAP HANA Disaster Recovery	38
8	Technology Requirements	51
8.1	Hardware Requirements	51
8.2	Software Requirements	52
9	Conclusion	52
	Where to Find Additional Information	53
	Version History	54

LIST OF TABLES

Table 1) Hardware requirements	51
Table 2) Software requirements	52

LIST OF FIGURES

Figure 1) Solution overview	8
Figure 2) Runtime comparison – file-based backup versus Snapshot copy backup.....	10
Figure 3) Solution and use case overview	15
Figure 4) Test lab setup.....	16
Figure 5) Unified ONTAP paths (asymmetric active-active)	17
Figure 6) NetApp AFF A700 ASA symmetric active-active pathing	18
Figure 7) ASA planned failover test results.	19
Figure 8) ASA unplanned failover test result	19
Figure 9) SAP HANA backup.	20
Figure 10) SAP HANA resource details.....	21
Figure 11) Storage footprint configuration.	21
Figure 12) Resource protection configuration.....	22
Figure 13) Snapshot Backup Operation.	22
Figure 14) SnapCenter topology view.	23
Figure 15) SAP HANA backup catalog – SYSTEMDB.	23
Figure 16) SAP HANA backup catalog – tenant database.	24
Figure 17) SAP HANA plug-in deployment.....	24
Figure 18) SAP HANA plug-in hosts.....	25
Figure 19) SnapCenter resource view.....	25
Figure 20) HDB user store key input.	26
Figure 21) HANA resource details after discovery.....	26
Figure 22) Resource protection configuration.....	27
Figure 23) SnapCenter topology view.	28
Figure 24) SnapCenter restore scope.	28
Figure 25) SnapCenter recovery scope.....	29
Figure 26) SnapCenter restore operation summary.	29
Figure 27) SnapCenter restore operation job details.....	30
Figure 28) SAP system refresh using restore operation.....	31
Figure 29) SAP system refresh with SnapCenter.....	31
Figure 30) SAP System refresh operation steps.	32
Figure 31) Automation script.	32
Figure 32) SAP system refresh lab setup.....	33

Figure 33) Clone from backup.	33
Figure 34) SnapCenter clone target.	34
Figure 35) SnapCenter mount and post clone commands.	34
Figure 36) SnapCenter job details of clone create operation.....	35
Figure 37) SnapCenter clone delete.....	36
Figure 38) SnapCenter pre and unmount script.	37
Figure 39) Job details of clone delete operation.....	37
Figure 40) SAP HANA disaster recovery.....	39
Figure 41) Latency with synchronous SnapMirror.	40
Figure 42) Cluster peers, source, and target.....	40
Figure 43) Volume protection with Synchronous SnapMirror.	41
Figure 44) Volume relationships.....	42
Figure 45) Creating FlexClone volumes.	44
Figure 46) List of FlexClone volumes.	45
Figure 47) Mapping LUN to target host.	45
Figure 48) SnapMirror quiesce operation.	48
Figure 49) SnapMirror quiesce operation.	48
Figure 50) List of volume relationships.	49
Figure 51) SnapMirror break operation.	49
Figure 52) SnapMirror break operation.	50
Figure 53) List of volume relationships.	50
Figure 54) Map LUN to disaster recovery server.....	51

1 Forward – A Message from Broadcom/Brocade

SAP software solutions play a crucial role in many organizations' daily operations and future expansions. To maximize the value of the enterprise SAP platform, customers need a secure, robust, and proven infrastructure to match business-critical SAP high-availability and disaster recovery requirements.

With its high-performance, highly available, and easy-to-deploy all-flash storage platforms, NetApp has evolved as a leader in the SAN storage industry. NetApp leads the storage world with many valued innovations, including an industry-leading role in NVMe over Fibre Channel (NVMe/FC), ground-breaking performance, data protection, resiliency, and storage efficiency capabilities. As a result, their flash-based SAN storage portfolio offers the necessary requirements to lead the transformation for high-value workloads such as SAP, now and in the future.

In this NetApp Verified Architecture, NetApp, with its Brocade-empowered fabrics, demonstrates how to improve SAP HANA enterprise data protection while providing superior performance, low latency, return on investment (ROI), and infrastructure availability. These benefits are necessary for enterprise-grade OLTP applications powered by SAP HANA. High performance and persistent storage are a must for large in-memory database solutions, especially for mission-critical tier one applications and workloads such as SAP HANA.

NetApp and Broadcom (Brocade BU) have a deep engineering partnership. They have even developed and deployed the industry's first end-to-end enterprise NVMe/FC SAN architecture. This solution, which is based on NetApp and Broadcom's industry-leading Fibre Channel SAN Gen6 architecture with next-generation low-latency, high-bandwidth, and high-performance NVMe-based storage and fabrics, is designed to provide a future-proof strategy to efficiently and effectively provision SAP test systems and simplify operations for single and multiple SAN HANA environments starting with traditional proven SAP Fibre Channel SAN solution today as outlined in this NVA, while also providing a path to the future with NVMe/FC.

NetApp and Brocade have had a nearly two-decade partnership that has consistently delivered industry-leading FC SAN-based storage solutions for mission-critical enterprise SAN applications.

Dr. Naem Saafein

Director, Technical Development

2 Executive Summary

Now more than ever, customers are transforming businesses with new revenue capabilities for new markets, customer-focused products and services, flexible development, and selling to ensure efficient and fast time-to-market models. SAP offers S/4 HANA as an intelligent and integrated next-generation enterprise resource planning (ERP) system with proven tools and services for seamless transitions. SAP HANA is an in-memory database for high performance transactional and analytical processing; therefore, customers have various deployment options, including cloud, on-premises, and hybrid. Using a proven reference architecture for these deployments offers the following benefits:

- Lowest business risk
- Business agility and accelerated time to value
- Future proof investment
- Highest levels of performance
- Superior data protection and 60x–100x faster backups
- Exceptional ROI/TCO
- Perfect end-user experience
- Proven and verified

A NetApp Verified Architecture describes proven best-in class-systems and solutions that are designed, tested, and documented to facilitate, accelerate, and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that NetApp has developed to help meet demanding enterprise business needs.

This NetApp Verified Architecture provides a solution that combines SAP HANA certified high-performance modern SAN storage with SAP HANA integrated data protection software, giving your organization a highly-available storage solution for your mission-critical SAP HANA workloads.

This document discusses the following topics:

- A typical setup of single and multiple-host HANA systems
- Execution of backup operations in a matter of minutes instead of hours
- Rapid and fully automated restore and recovery operations
- Automated workflows for SAP HANA test system provisioning
- Synchronous mirroring for disaster recovery to mitigate business risk
- Disaster recovery testing without impacting RTO and RPO
- Dramatic TCO savings in your data center operations through consolidation to this all-flash platform

ROI and TCO cost-benefit analysis of this modern SAN for SAP HANA solution:

- 212% ROI
- Seven-month payback period
- 80%–90%+ savings
 - In data center floor space
 - Associated power and cooling costs
- 75%+ savings in labor costs

The financial advantages of a combined NetApp, Brocade, and CommVault solution for an SAP HANA platform offers a strong business value and maximizes investment returns.

These benefits span three key areas, each of which contribute to a compelling TCO story for the joint solution:

- **Infrastructure consolidation.** Modernizing and consolidating legacy HDD–based storage and SAN fabric with this solution results in significantly improved TCO and lower data center operating expenses and labor costs.
 - Note:** It is crucial to understand the full impact of NetApp storage efficiency technologies to be able to fully evaluate the TCO of this enterprise SAN solution.
- **Modern brocade SAN fabric.** Inclusion of Brocade sixth generation fabric provides another compelling component to produce a lower cost of ownership. Brocade Fabric Vision technology with IO Insight, an extension of Gen 6 Fibre Channel, can enable organizations to:
 - Eliminate nearly 50% of maintenance costs through automated testing and diagnostic tools.
 - Save large environments millions of dollars on CapEx costs by eliminating the need for expensive third-party tools through integrated network sensors, monitoring, and diagnostics.
 - Tune device configurations with integrated I/O metrics to optimize storage performance and increase ROI.
- **Additional soft benefits.** These benefits are critical to understanding the true TCO of a storage solution and often result in an even higher (ROI) not quantified in this analysis, for example:
 - Elimination of downtime and business risk
 - Superior performance and platform responsiveness

- Lower overall hardware costs through NetApp data protection technology
- Improved operational efficiency with automation and optimization

Table 1) Cost benefit analysis of refreshing legacy SAN with this NetApp modern SAN solution.

Value	Analysis Results
ROI	212%
Net present value (NPV)	\$1,463,827
Payback period (months)	Seven months
Cost reduction	Approximately \$1.5 million saved over a three-year analysis period compared to the legacy SAN storage system
Savings on power and space	\$263,340
Administration costs savings	\$692,016

3 Program Summary

This report is part of the Modern SAN Best Practices Program, which provides test and validated design and configuration recommendations for an SAP HANA deployment in an FC environment with NetApp All SAN Array (ASA) storage and Broadcom SAN fabrics.

This program is a collaboration between NetApp and Broadcom’s Brocade and Emulex divisions, which together developed the industry’s first end-to-end enterprise NVMe architecture. NVMe powered fabrics are the next frontier in accelerating critical enterprise applications and will be applicable for SAP deployments as SAP phases in support for this future ready technology.

This report describes the system and solution that were designed, tested, and documented to facilitate resilient high-performance modern SAN deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that NetApp has developed to meet the business needs of our most demanding customers.

This report also describes the design choices and the best practices for this shared infrastructure platform. These design considerations and recommendations are not limited to the specific components that are described in this document—they also apply to other component versions.

Table 2) Comparison of legacy SAN and NetApp modern SAN.

	Legacy SAN	NetApp Modern SAN
Host connectivity	FC	FC, NVMe/FC
Future NVMe/FC next-generation support	No	Yes
Unified storage	No	Yes
Staff to manage	2 FTE	½ FTE
Bandwidth	8Gb avg. (max 16G FC)	32Gb
Data migrations	Required	No
Data center footprint	Large	Small

4 Solution Overview

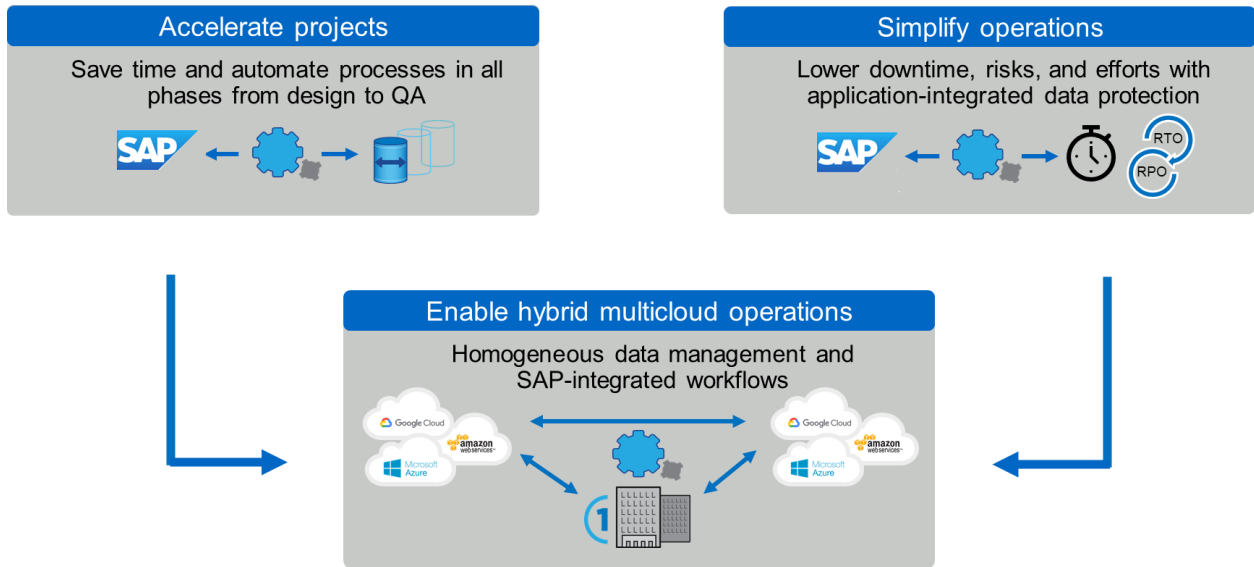
NetApp solutions for SAP HANA are based on tight software integration into SAP providing end-to-end automated workflows for SAP relevant use cases. NetApp provides solutions for SAP enabling consumption of the NetApp unique data management features with an SAP centric view. These solutions include SAP system provision tasks as well as SAP-integrated data protection for backup and disaster recovery.

The solutions and the value proposition can be broken down into three main areas:

- Accelerate projects
- Simplify operations
- Enable hybrid multicloud operations

Figure 1 provides a high-level solution overview.

Figure 1) Solution overview.



Accelerate Projects

Accelerating projects is based on the ability to provision SAP test systems rapidly in an automated fashion. The obvious use case is the SAP System Refresh operation, where data from the production system needs to be loaded into a test system. Combining storage cloning with application-integrated workflows accelerates and highly simplifies these operations. But there are also other use cases, like the handling of logical corruption, where you just spin-up a clone of your production system by using any prior NetApp Snapshot™ backup in a matter of minutes. If you use this clone with data before the corruption occurred, you can export the data that was accidentally deleted and import the data into your production system. And finally, you can use the same process to test your disaster recovery workflow, ensuring that you can recover from a disaster if it really happens.

Simplify Operations

Simplifying operations is based on the ability to execute backup and restore operations rapidly and efficiently. This solution is not only relevant for compliance with SLAs for the production systems during normal operation, it is also critical for any upgrade project or test cycle, where backup and restore operations are part of the workflow. Instead of waiting multiple hours until the backup is finished, these projects can be either accelerated, or the time saved can be used for further testing to reduce risk.

Enable Hybrid Multicloud Operations

The third area is the enablement of hybrid multicloud operations. NetApp wants customers to be able to choose where to run their SAP landscape: on-premises, in the cloud, or distributed between on-premises and different cloud providers. Regardless of the customer's choice, the goal is to provide homogenous data management with SAP-integrated workflows.

4.1 Backup and Recovery

Today, organizations need continuous, uninterrupted availability of SAP applications. Backing up SAP databases is a critical task and can significantly affect the performance production SAP systems. The time that it takes to restore and recover these systems is also a concern.

Traditional SAP HANA backup and restore operations, which are based on streaming data from the database host to the backup target, have the following challenges:


- Long backup operations with performance degradation on production SAP systems
- Unacceptable system downtime due to long restore and recovery operations
- Shrinking backup windows due to increased criticality of applications
- The need for a flexible solution to mitigate logical corruption

Figure 2 shows a comparison of the runtime of file-based and storage Snapshot-based backups for different database sizes. For file-based backups, a throughput of 500MBps is used as a basis for the calculation. The runtime of file-based backups are of course dependent on the size of the database. The graph in Figure 2 shows that it takes approximately two hours to back up a 4TB database and more than 18 hours to back up a 32TB database.

In contrast, the runtime of a Snapshot-based backup operation is independent from the size of the database. Customer data has shown an average of approximately two minutes to complete a Snapshot-based backup, the value which is used in Figure 2. The graph also shows how much faster the Snapshot backup is, compared to file-based backup. For example, for a 32TB database, the Snapshot backup operation is more than 500 times faster.

The same data is valid for restore operations. The NetApp SnapRestore® software feature allows you to restore from any Snapshot copy in a matter of seconds, independent on the size of the dataset. Restoring from a file-based backup would take the same time as the backup operation.

Figure 2) Runtime comparison – file-based backup versus Snapshot copy backup.



The Solution: Backup and Recovery Operations in Minutes Instead of Hours

With NetApp storage solutions that run NetApp ONTAP® data management software, in combination with NetApp SnapCenter® data protection software, you can meet all those challenges. And with the NetApp Snapshot technology that is included in ONTAP software, you can create backups or execute restore operations of any size dataset in a matter of seconds. SAP HANA supports the use of storage-based Snapshot copies as a valid backup operation with documented interfaces.

Backup Operations

NetApp SnapCenter and the plug-in for SAP HANA use ONTAP Snapshot technology and the SAP HANA SQL backup interface to give you an SAP-integrated backup solution. SnapCenter gives you automated workflows for backup operations, including retention management for data backups, for log backups, and for the SAP HANA backup catalog. And for long-term retention, SnapCenter also manages the optional replication of application-consistent backups to an off-site secondary location. Your off-site backup storage can be either a physical storage system on the premises or a NetApp Cloud Volumes ONTAP instance that runs in Amazon Web Services (AWS), Google Cloud Platform, or in Microsoft Azure.

60 to 100 Times Faster Backup Operations

Evaluation of customer data has shown that for SAP HANA, the average backup time with Snapshot copies is in the range of a few minutes. Feedback from one of our large SAP customers indicated that Snapshot copy-based backup operations are 60 to 100 times faster than traditional full database streaming backups. The largest contributor to the overall backup duration is the time that SAP HANA needs to write the synchronized backup savepoint. The amount of time that is required to write the savepoint is a function of the memory of the SAP HANA system and the activity on the system. The NetApp storage Snapshot operation is performed in a matter of seconds, independent of the size of the database.

Reduced System Downtime

Because NetApp Snapshot copy-based backup operations are rapid and do not affect system performance, you can schedule multiple Snapshot copy backups daily instead of creating a single daily backup as with traditional streaming backup technology. When a restore and recovery operation is necessary, your system downtime is significantly reduced by two key features. By using NetApp SnapRestore data recovery technology on the storage layer, the restore operation is executed in mere seconds. And because a higher backup frequency results in fewer database logs that need to be applied, the forward recovery is also accelerated.

4.2 Disaster Recovery

Business continuity is essential in IT organizations. They must be able to provide high availability services for the mission-critical applications that their customers require to run their businesses. Otherwise, their customers will face productivity decrease, and eCommerce organizations could face a direct impact on their revenue.

Therefore, every IT organization needs a disaster recovery plan to increase the resiliency of services in order to meet recovery point objective (RPO) and recovery time objective (RTO). Building disaster recovery plans can be cumbersome. IT organizations face the following problems when defining their disaster recovery plans:

- Lack of a future-proof disaster recovery solution that can respond to the rapidly changing needs of the business
- Inflexible disaster recovery solutions without SAP HANA integration
- Difficulty of testing the disaster recovery plan without affecting the production system
- No seamless integration into the cloud

NetApp Disaster Recovery Solution

NetApp has developed a full portfolio of technologies and tools to help IT organizations build or adapt their disaster recovery plans to respond to all business demands. These NetApp technologies constitute an extraordinarily versatile disaster recovery solution for SAP HANA on the market.

The solution includes NetApp SnapMirror® replication, NetApp MetroCluster™ software, and NetApp FlexClone® thin-cloning technology. The solution supports:

- Asynchronous and synchronous storage replication
- Replication of nondatabase data, such as application server binaries
- Use of disaster recovery resources for development and testing
- Use of replicated data to refresh development and testing systems
- Disaster recovery testing based on cloning

Storage Replication

Storage replication is suitable for low-to-medium RTO requirements, where it is acceptable for the SAP HANA database to be started and for data to be loaded into memory after a disaster recovery failover. Storage replication is also used to replicate nondatabase data, such as SAP application server binaries. NetApp SnapMirror data replication software provides synchronous and asynchronous replication. The replication is configured on the storage volume level.

Disaster Recovery Resources for Development and Testing

With storage replication, the servers at the disaster recovery site can be used for development and testing during normal operation. When you use a SnapMirror-based solution, the disaster recovery site

can be either on premises or in the cloud, or a remote disaster recovery site and the replicated data can be used for performing a development and testing system refresh.

Disaster Recovery Failover Testing

Every organization must test its disaster recovery plan. This testing shows whether the system reacts as stipulated in the disaster recovery plan and documentation. With NetApp FlexClone technology, you can execute a disaster recovery failover test without influencing or interrupting the ongoing replication to the disaster recovery site. In this way, FlexClone lets you run a test without influencing the RTO or RPO.

4.3 SAP Lifecycle Management

Today, enterprises are looking at ways to increase competitiveness by accelerating projects and speeding time to market. To reach this goal, they need to improve the lifecycle of their enterprise applications by automating tasks and simplifying processes. Traditional SAP lifecycle management approaches to development and test-system provisioning are primarily based on manual processes. These manual processes are often error prone and time consuming, delaying innovation and the ability to respond to business requirements.

The main challenges IT organizations face today are:

- Slow implementation of new features
- Lack of automation
- Lost productivity through lack of integration among orchestration tools

The Solution: Automated Provisioning of SAP Development and Test Systems

NetApp is addressing these challenges by providing a lifecycle management solution that is fully integrated. NetApp SnapCenter software and the SAP HANA plug-in brings Snapshot copies that provide application consistency to the solution and can be used to automate the SAP development and test-system provisioning, including the required steps on the HANA database layer.

Fast and Space-Efficient SAP System Provisioning Using Storage Cloning Technology

NetApp FlexClone is the key ingredient of the solution enabling fast system copies. Traditional copies can take many hours to make. With FlexClone thin-cloning technology, even the largest volumes can be cloned in a matter of seconds. This innovative technology also makes sure that a clone uses a small amount of space for metadata and then consumes additional space only as data is changed or added. These clones can be created from either the production, disaster recovery, or backup storage system.

4.4 Target Audience

The target audience for this NetApp Verified Architecture report includes the following groups:

- CIOs and business information officers who seek new and better ways to serve line-of-business owners with benefits from modern technologies and proven solutions
- Architects, administrators, solutions engineers, and business consultants who are responsible for designing and deploying infrastructure for enterprise mission-critical applications
- Database administrators who require new data management capabilities and performance to serve evolving data requirements
- Data scientists who seek to leverage SAP HANA in memory database technology to accelerate and transform data into usable business information
- Application owners who need to accelerate business application projects, improve business outcomes, and provide superior customer experience

- Service delivery managers who must meet SLAs and service-level objectives (SLOs) that require IT infrastructure and solutions to promote consistent and predictable results

5 SAP Data Protection Overview

SAP HANA powers the next generation of tier-one, mission-critical, applications including ERP, customer relationship management (CRM), and supplier relationship manager (SRM) platforms that require a comprehensive data protection and disaster recovery solution.

Studies have shown that business application downtime has a significant negative mission and business impact for enterprise organizations. Such downtime not only has a significant financial impact, but it can also affect the organization's reputation, staff morale, and customer loyalty. Avoiding downtime and its subsequent impact is paramount.

Surprisingly, not all organizations have a comprehensive disaster recovery policy, creating a significant operational and continuity risk. Mitigating that risk through a comprehensive disaster recovery policy involves understanding the mission and business application requirements along with technical capabilities needed for effective data protection and disaster recovery. Many tier-one enterprises who do have a disaster recovery plan actually lack the capability to test their plan nondisruptively and without risk. Fortunately, with a NetApp modern SAN architecture, those risks and gaps are eliminated.

5.1 Business Application Requirements

There are two key availability metrics for business applications:

- The RPO, or the maximum tolerable data loss
- The RTO, or the maximum tolerable business application downtime

5.2 Backups

Backups are created to enable restore and recovery from different point-in-time datasets. Typically, these backups are kept for a couple of days to a few weeks.

The RTO for restore and recovery is defined by the needed restore time, the recovery time (including database start), and the loading of data into memory. For large databases and traditional backup approaches, the RTO can easily be several hours, which might not be acceptable. To achieve very low RTO values, a backup must be combined with a hot-standby solution, which includes preloading data into memory.

In contrast, a backup solution must address logical corruption, because data replication solutions cannot cover all kinds of logical corruption. For more information, see the section 4.1, "Backup and Recovery.

5.3 Synchronous or Asynchronous Data Replication

The RPO primarily determines which data replication method you should use. If the RPO must be zero, even when the primary and backup storage is lost, the data must be replicated synchronously. However, there are technical limitations for synchronous replication such as the distance between the two data centers. In most cases, synchronous replication is not appropriate for distances larger than 100km. Indeed, synchronous replication over a large distance places significant demands on the network infrastructure between the two data centers and therefore can be very expensive.

If a larger RPO is acceptable, asynchronous replication can be used over large distances. The RPO in this case is defined by the replication frequency.

5.4 SAP Data Protection: NetApp + Commvault

While this NetApp Verified Architecture primarily focuses on native NetApp data protection solutions, leveraging integrated data protection software such as Commvault is an alternative option for protecting critical SAP application data on NetApp storage systems.

NetApp and Commvault have a long-standing partnership for more than 10 years. From the outset, solutions for SAP have played an important part. NetApp provides SAP solutions with the ultimate in performance and high availability, with built-in data management and replication. Commvault Complete Backup and Recovery adds policy-based snapshot lifecycle management and automation, alongside management of NetApp Snapshot and NetApp SnapMirror replication.

The NetApp and Commvault technologies used to keep your SAP system available and provide rapid recovery use a range of APIs to interface with your SAP applications and databases. This integration means you can trust that SAP is performing well and is protected in a consistent manner with Commvault. The same tools can also accelerate SAP System refreshes.

5.5 Recover in Minutes Versus Hours

Through tight integration with NetApp Snapshot technology, Commvault Complete Backup and Recovery enables rapid recovery of your most critical SAP applications, allowing you to meet even the most stringent availability requirements. It also works regardless of your SAP HANA database size or whether you're running SAP HANA scale-up or HANA scale-out, so meeting your RTO is one less thing you have to worry about.

The flexibility of being able to leverage both streaming and snap-assisted SAP backups help you meet a variety of SLA requirements.

5.6 Automation and Optimization

Commvault software includes a robust API set and enterprise-grade automation, enabling you to orchestrate operations across your SAP landscapes. It uses machine-learning to deliver AI-driven reporting and data management, which reliably predicts whether you will meet your recovery SLAs. It also removes the need for manual optimization processes over time.

5.7 Hybrid-Cloud: On-Premises or in the Cloud

In hybrid-cloud architectures for SAP on NetApp environments, Commvault can also protect your SAP HANA on Azure (large instances) environments with full snapshot and replication support. Commvault software also provides the ability to migrate data and workloads to and from the cloud for a variety of use cases, providing SAP or any other application owners with a high degree of flexibility.

6 Solution Components and Use Cases

This document describes data protection, SAP lifecycle management, and disaster recovery solutions for SAP HANA running on NetApp ASA storage systems. NetApp SnapCenter software and the HANA plug-in are used to automate the required SAP use case specific workflows.

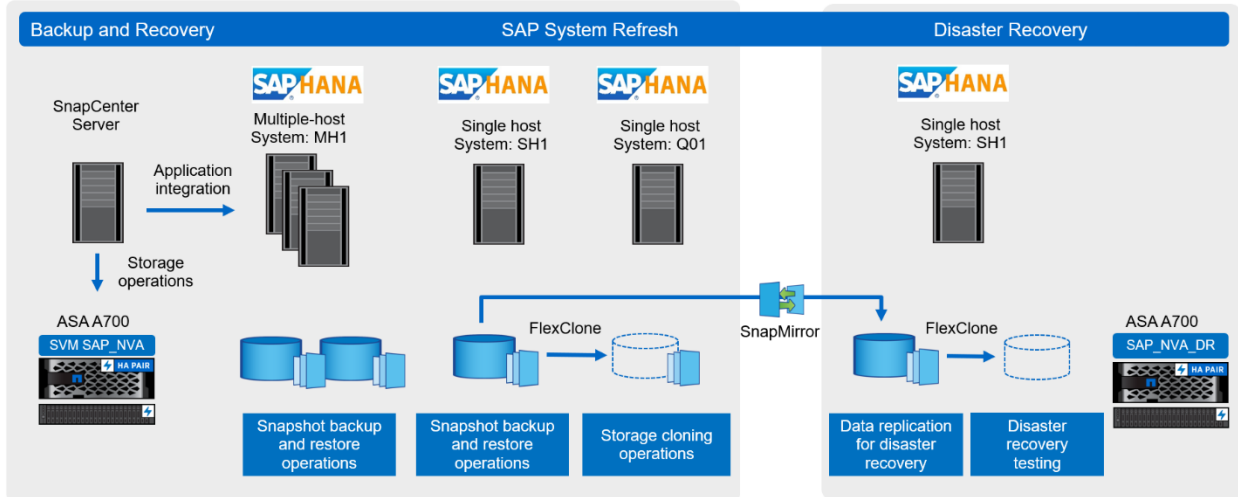
The solution described in this document covers the following SAP use cases:

- SAP HANA data protection:
 - Snapshot copy-based backup operations
- SAP HANA restore and recovery:
 - SnapCenter and Snapshot backups
- SAP System refresh (based on storage cloning technology)
- Disaster recovery:

- Synchronous SnapMirror
- Failover and disaster recovery testing

Figure 3 shows an overview of the solution and use cases.

Figure 3) Solution and use case overview.



6.1 Solution Components

SAP HANA Systems

The SAP HANA systems are installed on bare-metal servers and are connected to the NetApp ASA storage systems using FC SAN. The lab setup includes an SAP HANA multiple-host and an SAP HANA single-host system. Both SAP HANA systems are multitenant database container (MDC) single tenant configurations.

The SAP HANA systems have been installed according to NetApp best practice guide for SAP HANA, [TR-4436: SAP HANA on NetApp AFF Systems with Fibre Channel Protocol Configuration Guide](#).

NetApp All SAN Array

NetApp ASA delivers a simplified and dedicated SAN experience that provides continuous data availability for your organization's mission-critical databases and other SAN workloads. Like all NetApp AFF systems, ASA offers market-leading performance, even with inline storage efficiency, encryption, and active data protection. As with other NetApp storage systems, the ASA systems are certified for SAP HANA and are listed on [SAP's HANA certification web site](#).

Brocade Solution Component

Broadcom's Brocade has been the leading provider of storage networking solutions worldwide for more than 20 years, supporting the mission-critical systems and business-critical applications of most large enterprises. Brocade networking solutions help organizations achieve their critical business initiatives as they transition to a world where applications and information can reside anywhere. Today, Brocade is extending its proven data center expertise across the entire network with open, application-optimized, and efficient solutions that are built for consolidation and unmatched business agility.

The sixth generation of FC is aimed at satisfying the needs of growing deployments of flash storage, hyperscale virtualization, and new high-speed data center architectures such as NVMe. Brocade Gen 6 Fibre Channel platforms shatter application performance barriers with up to 100 million IOPS and 32Gb/128Gb FC performance to meet the demands of flash-based storage workloads.

SnapMirror Data Replication

A second ASA storage system is configured to implement a disaster recovery solution for the SAP HANA systems. The SAP HANA systems data is replicated synchronously from the primary to the disaster recovery ASA system. The replication is done by using synchronous SnapMirror.

6.2 Test Lab Setup

The lab setup consists of four bare-metal servers used for the SAP HANA databases:

- MH1 – SAP HANA multiple hosts system in a 2+1 configuration
- SH1 – SAP HANA single host system

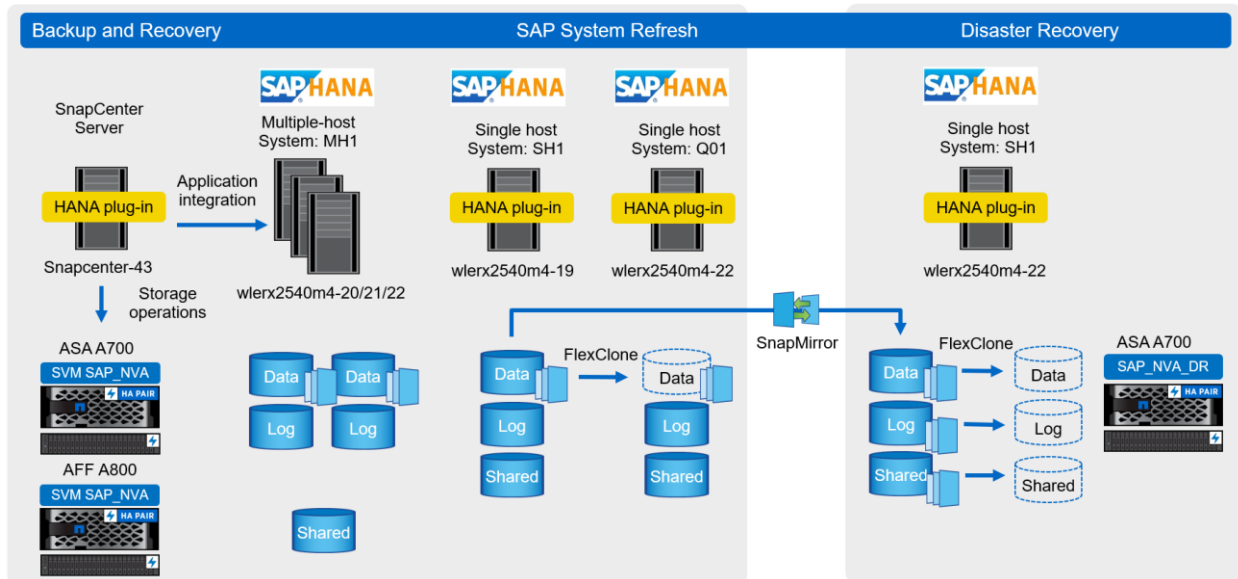
One server of the SAP HANA multiple host system is used to run additional SAP HANA systems. The SAP HANA system Q01 is used as the target system for the SAP system refresh operation. The same server is also configured for the disaster recovery failover tests.

An additional VM is used for the SnapCenter server. The SAP HANA systems are configured in SnapCenter for Snapshot backup and cloning operations.

An ASA A700 storage system is used to host the data and log LUNs of the HANA systems. The `/hana/shared` file system of the HANA multiple host system requires a shared file system, which is accessible by all hosts of the SAP HANA system. For the purposes of this test, the file system is put on an NFS share on a AFF A800 system, but entry-level AFF or FAS systems would be equally sufficient as well.

An additional ASA A700 storage system is used as a disaster recovery storage. This system is the target for the synchronous SnapMirror replication. Figure 4 shows the test lab setup.

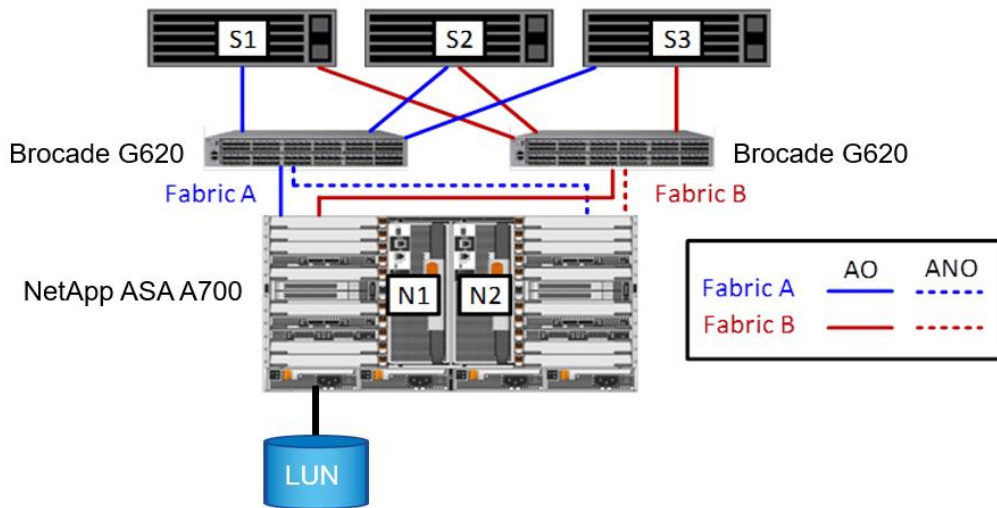
Figure 4) Test lab setup.



6.3 All SAN Array High Availability

Before the release of NetApp ONTAP 9.7, all ONTAP controllers featured the architecture that is shown in Figure 5. This architecture advertised routes directly to the controller that hosted the LUN as active-optimized (AO) paths, with all other paths (indirect paths) advertised as active-non-optimized (ANO) paths. Active nonoptimized paths are not preferred and are not used unless no active optimized paths exist.

Figure 5) Unified ONTAP paths (asymmetric active-active).



With ONTAP 9.7, NetApp introduced AFF ASA systems, which feature symmetric active-active topology, as shown in Figure 6. The ASA supports SAN (block protocols) only and is built on a single HA pair. It currently supports FC and iSCSI protocols, and support for NVMe protocols and larger clusters are expected in later releases.

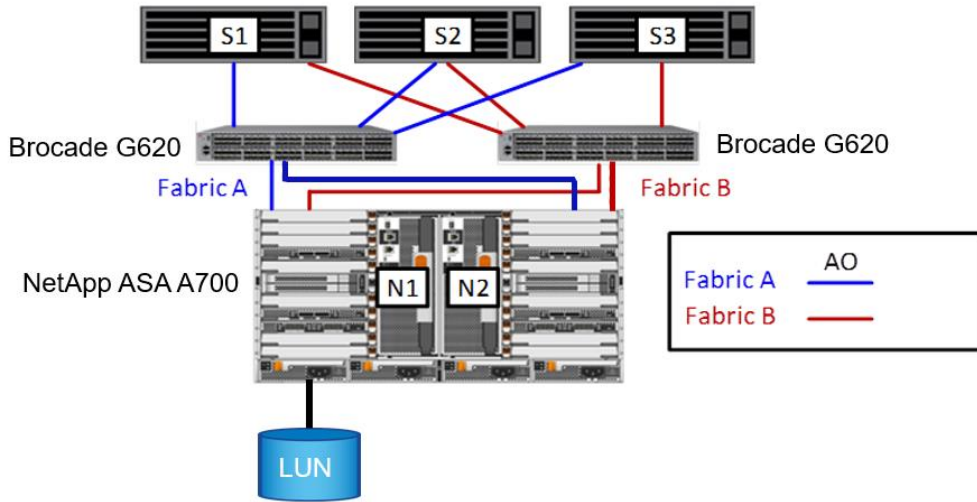
The defining features of ASA systems include:

- Symmetric active-active operations, which means that all paths are active “preferred” paths to all LUNs. ASA advertises all paths as AO, which means that there are always active paths to all LUNs, even if a storage failover (SFO, also called a takeover or giveback) occurs. The practical effect is that hosts always have active paths and don’t need to query for new paths if an SFO occurs. This means that even during a failover, because hosts have access to all of their data, hosts don’t suffer from having no active paths and therefore don’t log “no paths to storage” errors. I/O continues uninterrupted between hosts and storage. This feature minimizes the impact of an SFO so that the host experience resembles the uninterrupted experience seen with frame-style arrays. Unified clusters advertise both AO and ANO paths.

Note: Hosts that connect to a unified cluster see both AO paths (preferred) and ANO paths (not preferred). If the host loses all AO paths and doesn’t receive updates that advertise new AO paths, it changes the ANO paths that it still has to a LUN to AO or preferred paths. However, this process can take some small amount of time for the host to make those adjustments to its storage map because the host typically waits until its link down timer expires.

- The ASA introduced with ONTAP 9.7 takes full advantage of the complete rewrite of ONTAP System Manager (formerly OnCommand® System Manager) GUI that occurred between ONTAP 9.6 and 9.7. The primary objective was to greatly streamline and simplify the System Manager GUI, to be more intuitive and to make more best practice decisions for administrators. All aspects of provisioning, configuring, and managing of ONTAP SANs have been significantly simplified.
- The ASA also offers additional simplicity by removing any NAS protocols and features which reduces configuration options. This feature reduces the skillset that you need to configure, to provision, and to manage the ASA.

Figure 6) NetApp AFF A700 ASA symmetric active-active pathing.



Failover Tests

The SAP HANA Performance Test tool was used to generate load on the NetApp ASA system while performing planned and unplanned controller failover of the ASA HA pair.

The result of the planned failover test is shown in Figure 7. The takeover of one ASA controller was initiated by its high-availability partner. Immediately after the takeover was initiated, throughput was reduced because half of the active paths (for example, the paths through the node being taken over) were lost. Additionally, for the first few seconds after a takeover, there was takeover-related processing that reduced throughput. As you can see in Figure 7, throughput stabilized rapidly to pretakeover levels. Planned takeovers can be used for system maintenance operations such as nondisruptive upgrades of ONTAP.

Note: Hosts always have access to their data and there is active I/O at all times during a planned takeover.

Figure 7) ASA planned failover test results.

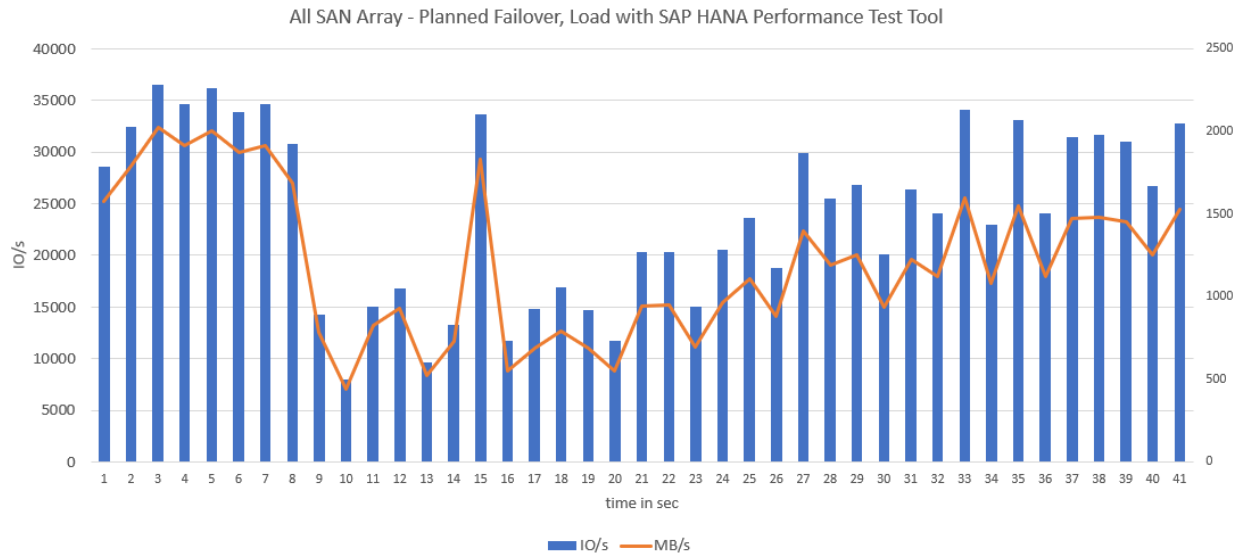
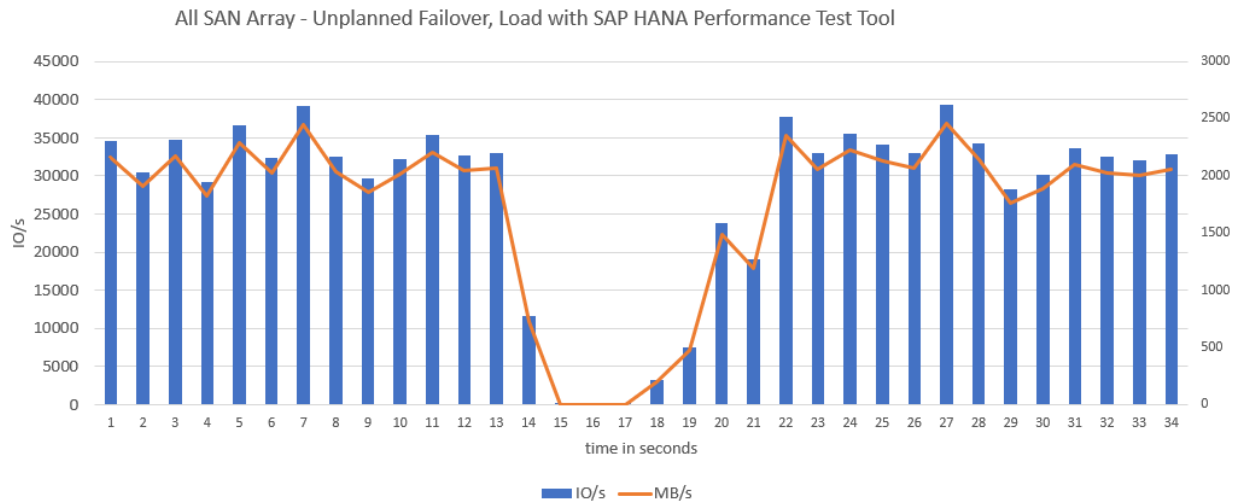


Figure 8 shows the test result of the unplanned failover test. In this example, we simulated an unplanned failover by inducing a power outage. After the unplanned failover, there was a slightly longer impact and a more pronounced impact, as there were no I/Os for a very short time while ONTAP was performing the takeover, moving storage ownership, and updating path advertisements. There was also a follow-on short interval after I/O resumed at a reduced rate, while takeover processing completed. Within several seconds of the failover, throughput was restored to pretakeover levels. From a host or application perspective, this will appear as a spike in latency. Remember, with the ONTAP documented six 9s uptime (<99.9999), which translates to 31.56 seconds of unplanned downtime, unplanned failovers are rare.

Figure 8) ASA unplanned failover test result.



7 Solution Verification

7.1 SAP HANA Backup

NetApp SnapCenter and the SAP HANA plug-in provides an application-integrated backup solution with end-to-end automated workflows. SnapCenter supports all SAP HANA architectures, single host and multiple host systems, with single container or MDC configurations.

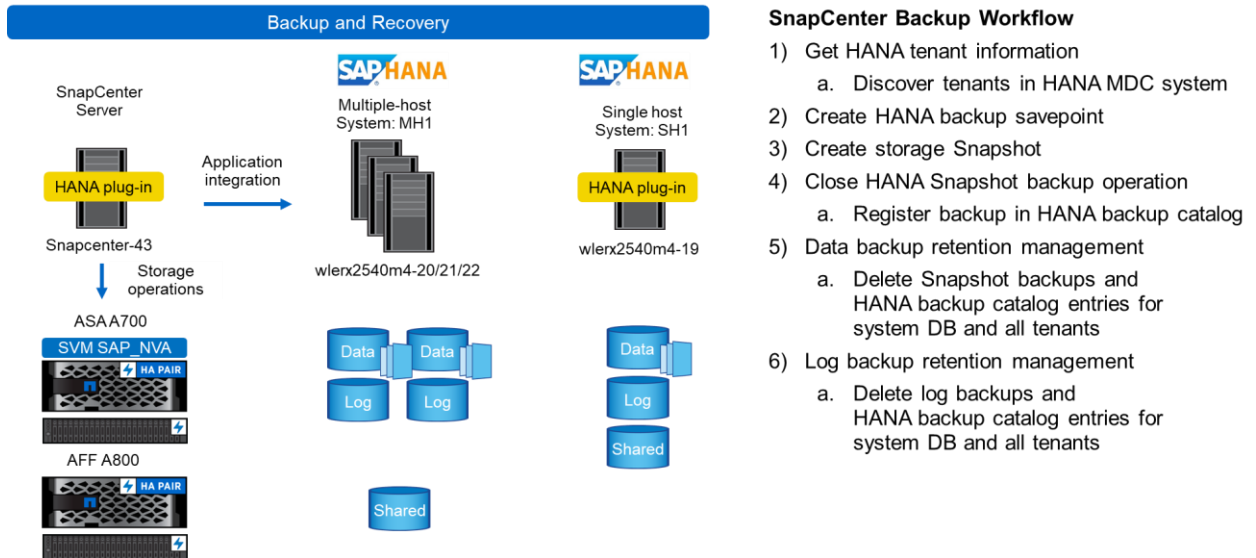
SnapCenter handles all required steps for a backup solution based on storage Snapshot technology:

- SAP HANA quiesce and unquiesce database operations
- Storage Snapshot copy operations
- SAP HANA backup catalog management
- Retention management of data and log backups
- SAP HANA database block integrity check operations
- Optional replication to offsite backup or disaster recovery storage

The SAP HANA systems in SnapCenter are configured according to the best practice guide [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#).

Figure 9 shows the lab setup with the two SAP HANA systems used for this document. The SAP HANA multiple host system is managed with a central SAP HANA plug-in on the SnapCenter server. For the SAP HANA single host system, the SAP HANA plug-in is installed on the database host.

Figure 9) SAP HANA backup.



Backup Operation for SAP HANA Multiple Host System – MH1

The images in this section provide a high-level overview of the backup configuration and operation for the SAP HANA multiple host system MH1. SAP HANA multiple host systems are configured manually in SnapCenter, while single host systems are auto-discovered by SnapCenter.

The first configuration (Figure 10) is the resource details configuration. Important parameters are the plug-in host (in this example, the SnapCenter server itself) and the HDB user store keys. For a multiple host system, configure a key for each host to support backup operations in case of a host failure or failover.

Figure 10) SAP HANA resource details.

The screenshot shows the 'Add SAP HANA Database' wizard at the 'Provide Resource Details' step. The left sidebar has three steps: '1 Name', '2 Storage Footprint', and '3 Summary'. The main area contains the following fields:

Resource Type	Multitenant Database Container
HANA System Name	MH1 - Multiple hosts system
SID	MH1
Plug-in Host	SnapCenter-43.sapcc.stl.netapp.com
HDB Secure User Store Keys	MH1KEYHOST1,MH1KEYHOST2,MH1KEYHOST3
HDBSQL OS User	SYSTEM

Information icons (i) are present to the right of the SID, Plug-in Host, HDB Secure User Store Keys, and HDBSQL OS User fields. A blue box highlights the Plug-in Host and HDB Secure User Store Keys fields. At the bottom right, there are 'Previous' and 'Next' buttons.

The second configuration (Figure 11) is the storage footprint configuration. In this step, select the storage volumes and LUNs. In this example, the 2+1 HANA multiple host system MH1 has two data volumes for the two worker hosts.

Figure 11) Storage footprint configuration.

The screenshot shows the 'Add SAP HANA Database' wizard at the 'Provide Storage Footprint Details' step. The left sidebar has three steps: '1 Name', '2 Storage Footprint', and '3 Summary'. The main area contains the following fields:

Storage Systems for storage footprint

SAP_NVA

Modify SAP_NVA

Select one or more volumes and if required their associated Qtrees and LUNs

Volume name	LUNs or Qtrees
S_20_SAP_Data_vol	MH1_data_rmt00001
S_21_SAP_Data_vol	MH1_data_rmt00002

A 'Save' button is located at the bottom right of the 'Modify SAP_NVA' dialog. At the bottom right of the main wizard area, there are 'Previous' and 'Next' buttons.

The third configuration (Figure 12) is the resource protection configuration. The most important part is the policy configuration, where you select backup policies and combine them with schedule configurations. In

this example, local Snapshot backup operations at the primary storage and block integrity check operations are configured.

Figure 12) Resource protection configuration.

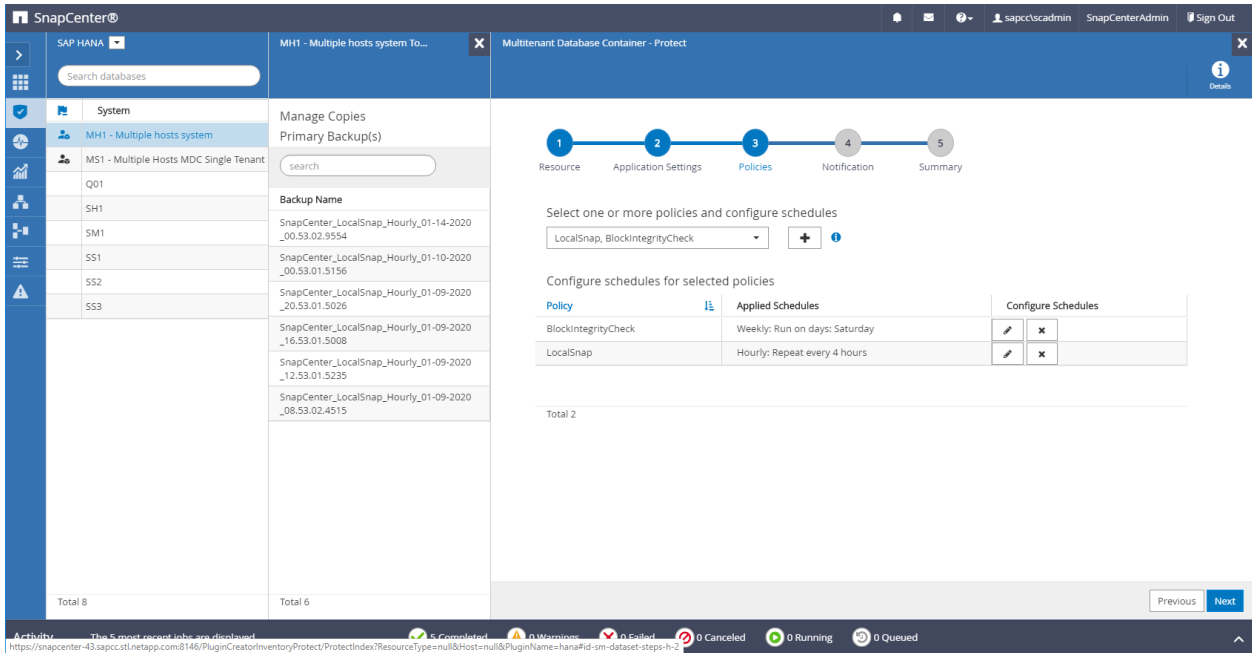
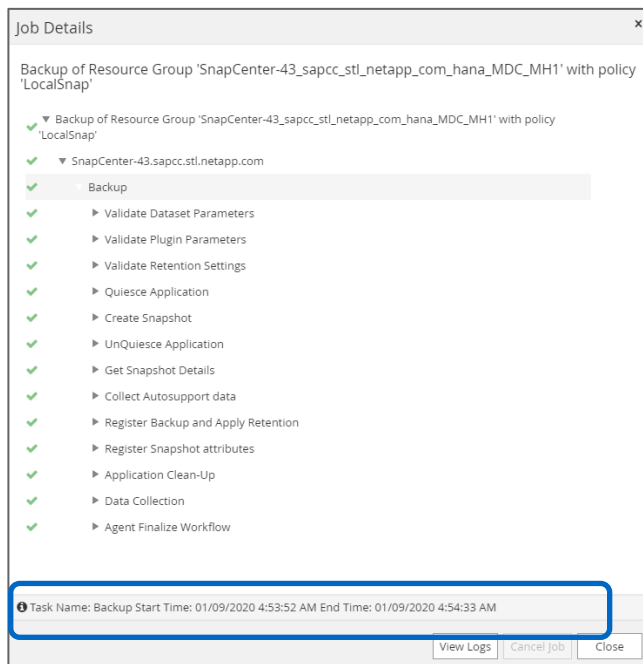


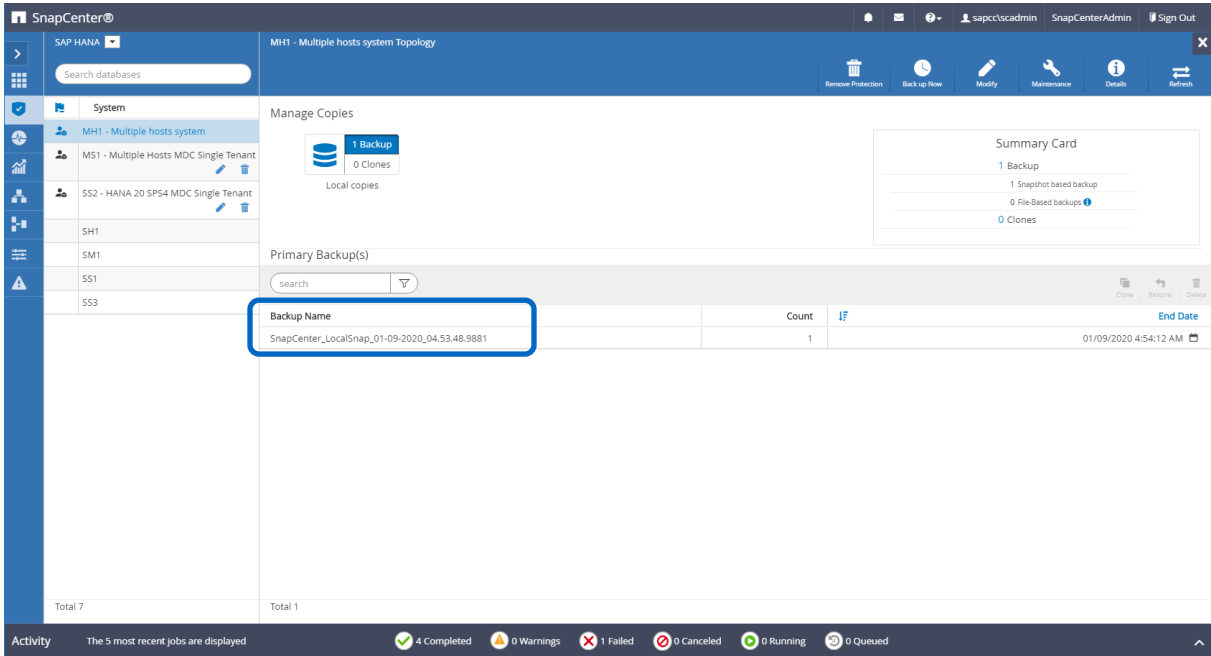
Figure 13 shows the job details of a backup operation. Using start and end time a total runtime of 41 seconds can be calculated. As discussed in section 4.1, “Backup and Recovery,” the runtime of an SAP HANA backup using Snapshot backup operations is independent of the size of the database. Therefore, for even a very large SAP HANA database, the backup operation is in a range of a few minutes.

Figure 13) Snapshot Backup Operation.



The newly created backup is listed in the SnapCenter topology view of the SAP HANA resource MH1 (Figure 14). The backup name used in SnapCenter is also used as the Snapshot name on the storage layer and is also visible in the HANA backup catalog.

Figure 14) SnapCenter topology view.



The same backup name used in SnapCenter is also written as the comment field and the external backup ID in the HANA backup catalog for the SystemDB and the tenant database (Figure 15 and Figure 16).

Figure 15) SAP HANA backup catalog – SYSTEMDB.

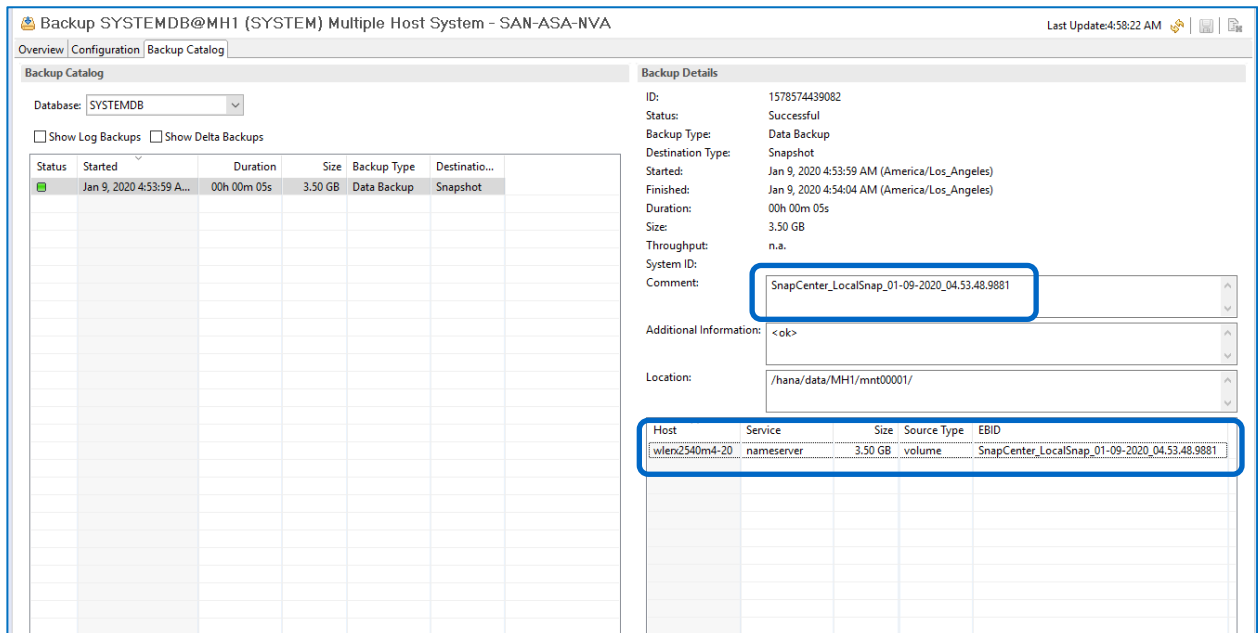
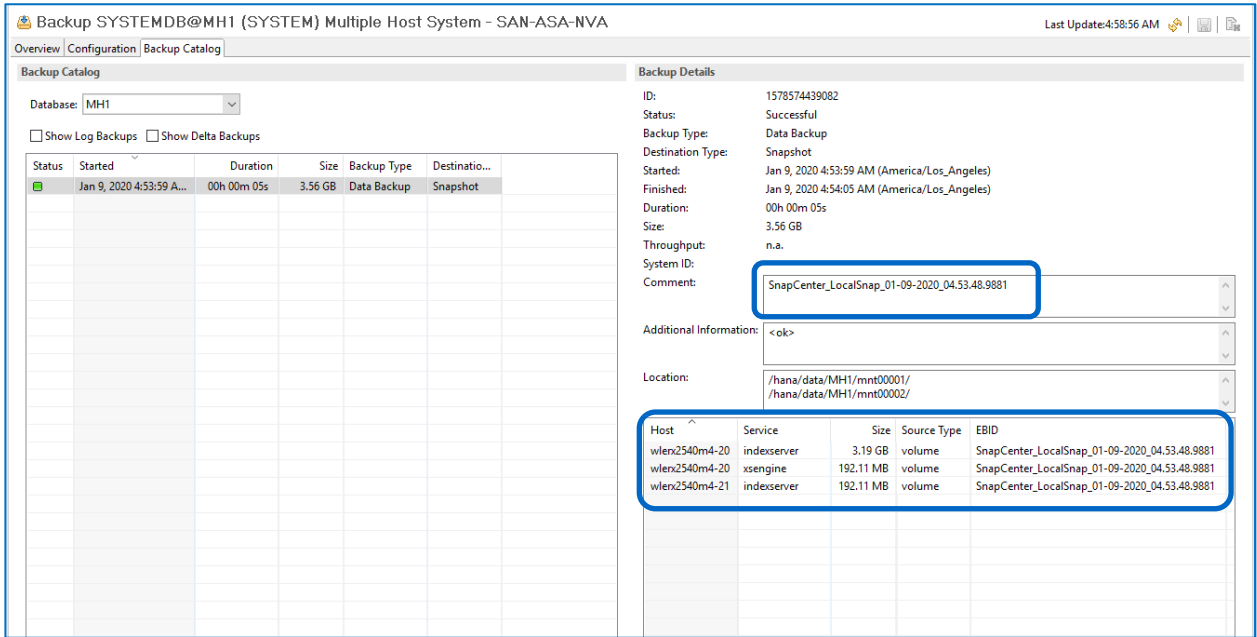


Figure 16) SAP HANA backup catalog – tenant database.

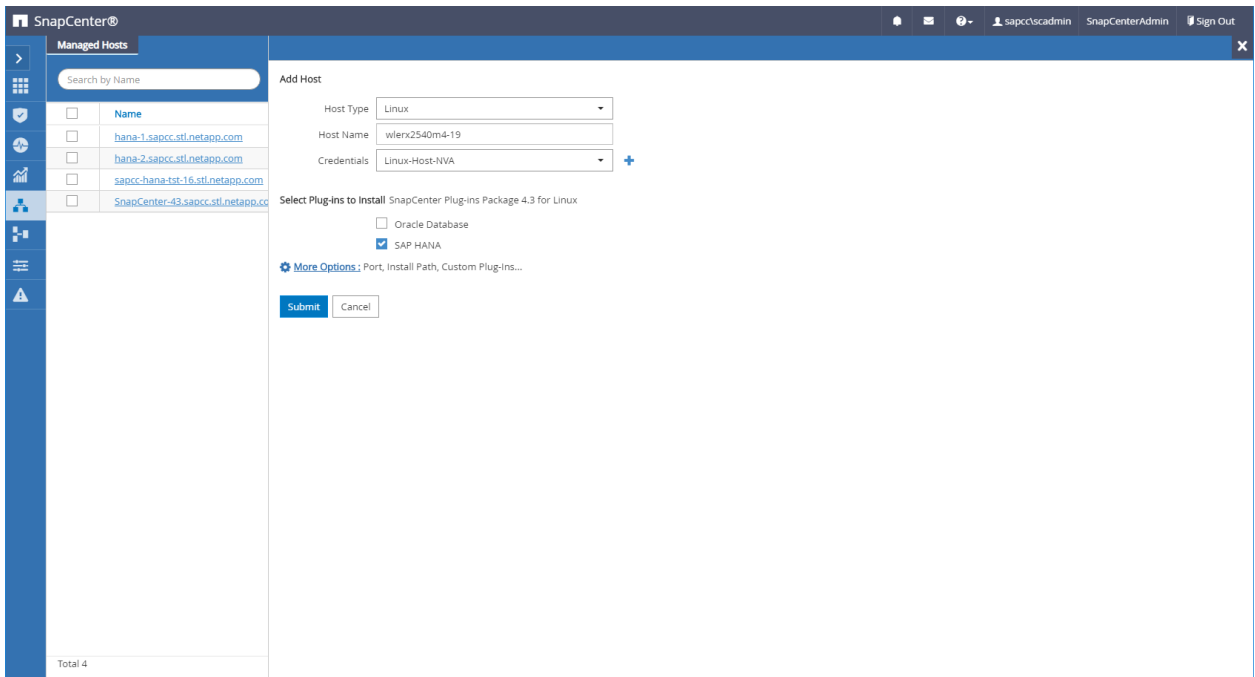


Backup Operation for SAP HANA Single Host System – SH1

The images in the section provide a high-level overview of the backup configuration and operation for the SAP HANA single host system SH1. SAP HANA single host system can be automatically discovered by SnapCenter.

The first configuration (Figure 17) is the deployment of the SnapCenter HANA plug-in on the database host. This is done by adding the SAP HANA host to SnapCenter.

Figure 17) SAP HANA plug-in deployment.



The SAP HANA plug-in is now deployed on the SAP HANA host.

Figure 18) SAP HANA plug-in hosts.

Name	Type	System	Plug-in	Version	Overall Status
hana-1.sapcc.stl.netapp.com	Linux	Stand-alone	UNIX, SAP HANA	4.3	Running
hana-2.sapcc.stl.netapp.com	Linux	Stand-alone	UNIX, SAP HANA	4.3	Running
sapcc-hana-tst-16.stl.netapp.com	Linux	Stand-alone	UNIX, SAP HANA	4.3	Running
SnapCenter-43.sapcc.stl.netapp.com	Windows	Stand-alone	Microsoft Windows Server, SAP HANA	4.3	Running
wlerx2540m4-19.stl.netapp.com	Linux	Stand-alone			Installing plug-in


With the plug-in deployment, the first phase of the auto discovery starts. This includes the SAP HANA architecture (single container, MDC system) as well as the SID and tenant name. A new SAP HANA resource is automatically added during the discovery process. By clicking on the resource, the second level discovery starts.

Figure 19) SnapCenter resource view.

System	System ID (SID)	Tenant Database	Plug-in Host	Resource Groups	Policies	Last backup	Overall Status
MS1 - Multiple Hosts MDC Single Tenant	MS1		SnapCenter-43.sapcc.stl.netapp.com		BlockIntegrityCheck LocalSnap	01/08/2020 11:15:23 PM	Backup succeeded
SH1	SH1	SH1	wlerx2540m4-19.stl.netapp.com				Not protected
SS2 - HANA 20 SPS4 MDC Single Tenant	SS2		SnapCenter-43.sapcc.stl.netapp.com		BlockIntegrityCheck LocalSnap	01/08/2020 10:05:23 PM	Backup succeeded
SM1	SM1	SM1 TENANT2	hana-2.sapcc.stl.netapp.com		BlockIntegrityCheck LocalSnap	01/08/2020 10:28:56 PM	Backup succeeded
SS1	SS1	SS1	hana-1.sapcc.stl.netapp.com		BlockIntegrityCheck LocalSnap LocalSnapAndSnapVault	01/08/2020 10:30:56 PM	Backup succeeded
SS3	SS3	SS3	sapcc-hana-tst-16.stl.netapp.com		BlockIntegrityCheck LocalSnap	01/08/2020 10:36:29 PM	Backup succeeded

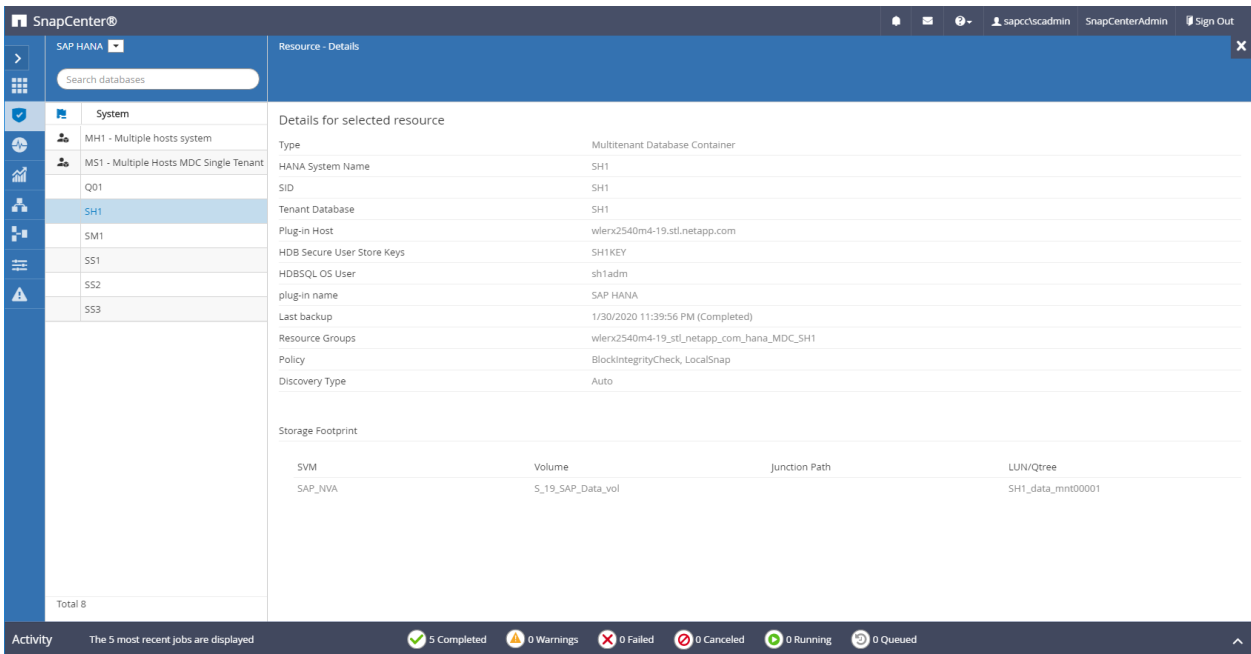
You must provide an HDB user store key.

Figure 20) HDB user store key input.



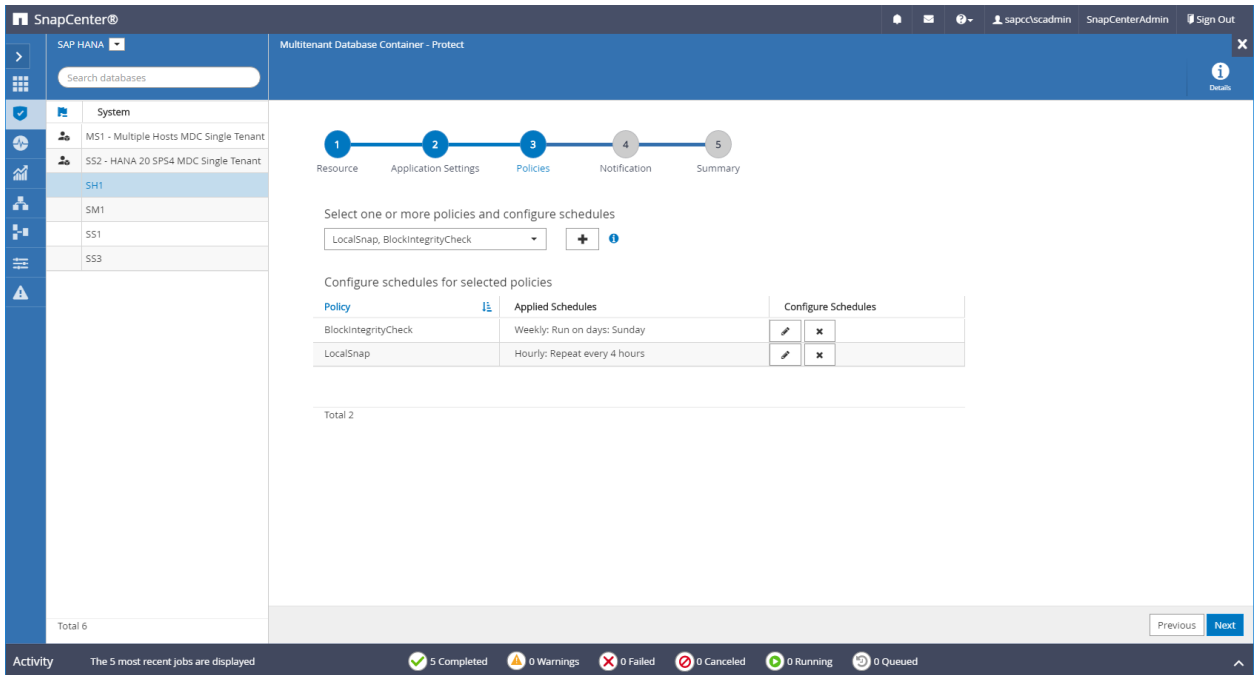
During the second-level discovery, additional database information as well as the storage footprint of the SAP HANA database is discovered.

Figure 21) HANA resource details after discovery.



The last configuration (Figure 22) is the resource protection configuration. The most important part of this configuration is the policy configuration, where you select backup policies and combine them with schedule configurations. In this example, local Snapshot backup operations at the primary storage and block integrity check operations are configured.

Figure 22) Resource protection configuration.



The backup operation of a SAP HANA single host system is identical to what was already described for the SAP HANA multiple host system in section titled, “Backup Operation for SAP HANA Multiple Host System – MH1”.

7.2 SAP HANA Restore and Recovery

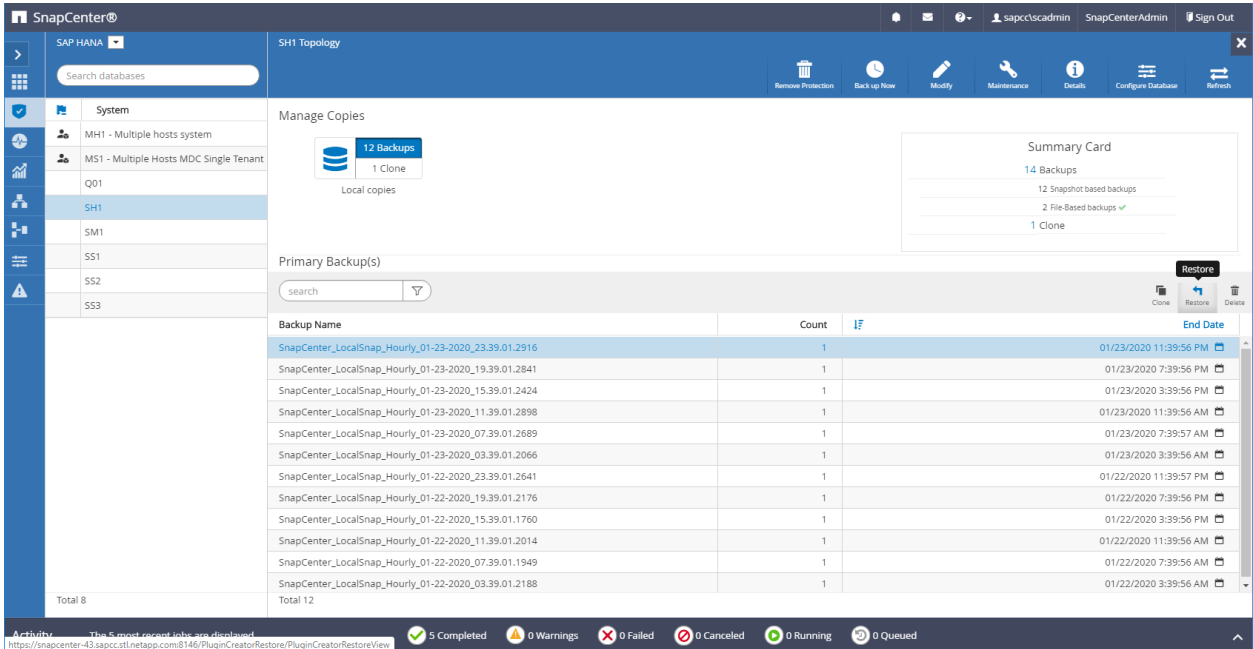
SnapCenter is used to execute a fully automated restore and recovery operation. This includes all the required operations on the operating system and the storage and database layer.

SnapCenter performs the following operations:

1. Stops the SAP HANA database
2. Restores the database:
 - a. Unmounts the LUNs
 - b. Restores the LUNs on the storage layer
 - c. Discovers and mounts the LUNs
3. Recovers the database:
 - a. Recovers the system database
 - b. Recovers the tenant database
 - a. Starts the HANA database

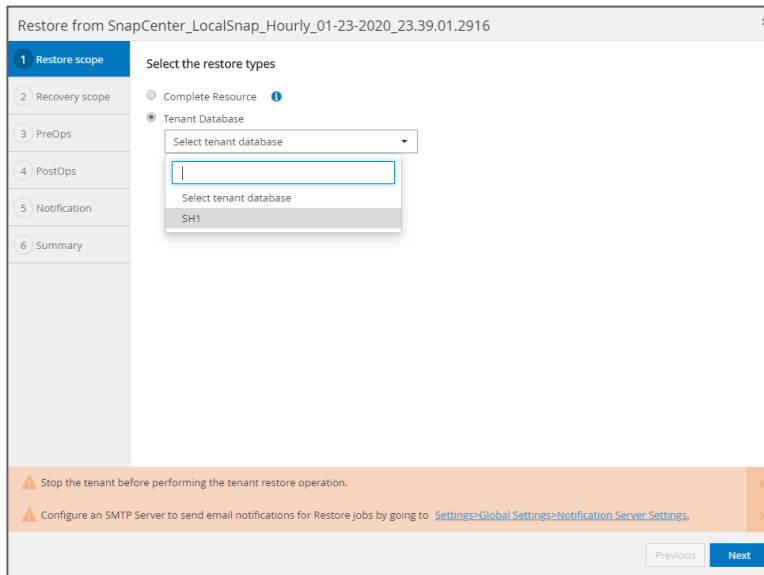
Any of the available Snapshot backups can be selected for the restore operation, as shown in Figure 23.

Figure 23) SnapCenter topology view.



In the restore scope, you can select a complete restore or a tenant restore. In this example, a complete restore is selected, which means that the System DB and the tenant database are restored and recovered.

Figure 24) SnapCenter restore scope.



In the recovery scope, you can select different recovery options including Most Recent State, Point in Time, or To Specific Backup.

Note: If the recovery should be done manually, you can select No Recovery. In this example, SnapCenter will only execute the restore operation.

Figure 25) SnapCenter recovery scope.

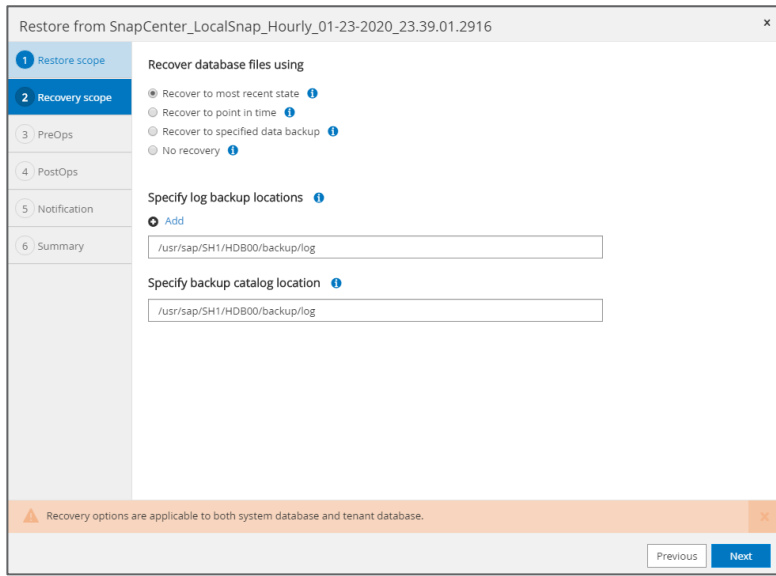
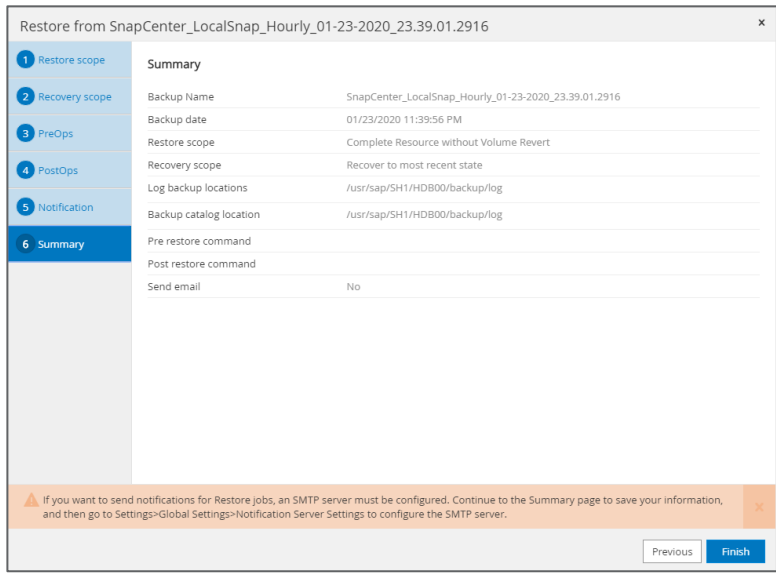
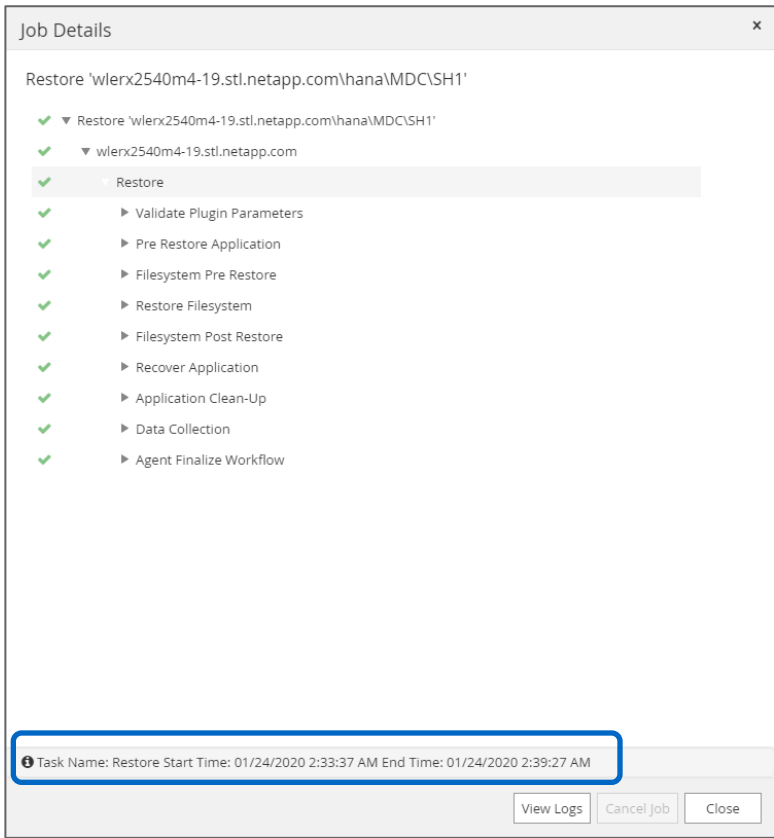


Figure 26) SnapCenter restore operation summary.



SnapCenter executes all required restore operations on the operating system, storage, and database levels. The restore and recovery operation is completed in approximately six minutes. The restore operation itself, as discussed in the section 4.1, “Backup and Recovery,” only takes a few seconds. Most of the time is spend for the SAP HANA recovery operation. Also keep in mind that the runtime of an SAP HANA restore operation based on Snapshot backup operations is independent of the size of the database. So even for very large HANA database, the restore operation will be in a range of a few minutes.

Figure 27) SnapCenter restore operation job details.

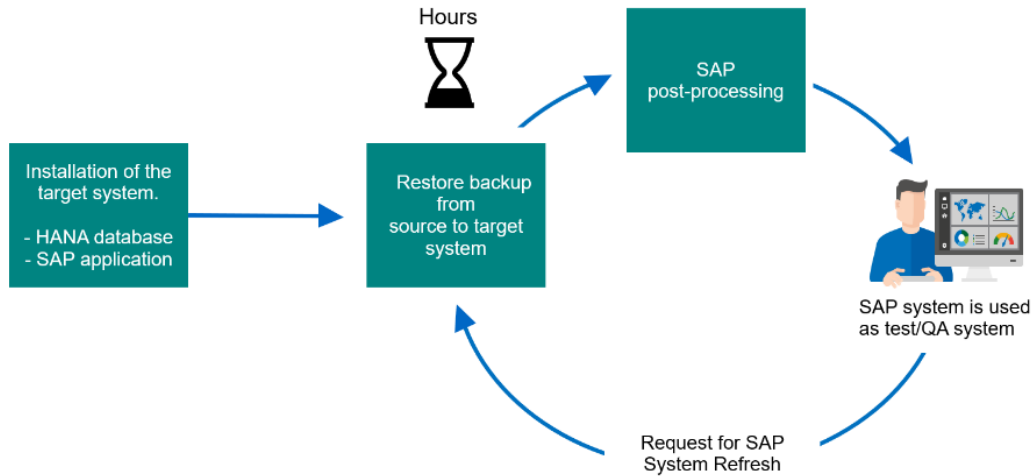


7.3 SAP System Refresh

An SAP system refresh is a refresh of an existing target SAP system with data from a source SAP system. The target system is typically part of an SAP transport landscape; for example, a quality assurance system that is refreshed with data from the production system. The host name, instance number, and SID are different for the source and target systems.

A common approach is to restore a file-based backup from the source system into the target system, as shown in Figure 28. The challenge with this approach can be the long run time of the restore operation, specifically with large databases. For example, restoring a 4TB database takes more than two hours with a restore throughput of 500MBps. For more information, see section 4.1, “Backup and Recovery.”

Figure 28) SAP system refresh using restore operation.



As an alternative, Snapshot backups from the source system can be used to refresh the target system using storage cloning technology. The advantage is the speed of the cloning process, which is in the range of a couple of seconds independent of the size of the database. This allows businesses to increase competitiveness by accelerating projects and speeding time to market. SnapCenter provides two workflows, clone create and clone delete, which are used for the SAP system refresh operation.

Figure 29) SAP system refresh with SnapCenter.

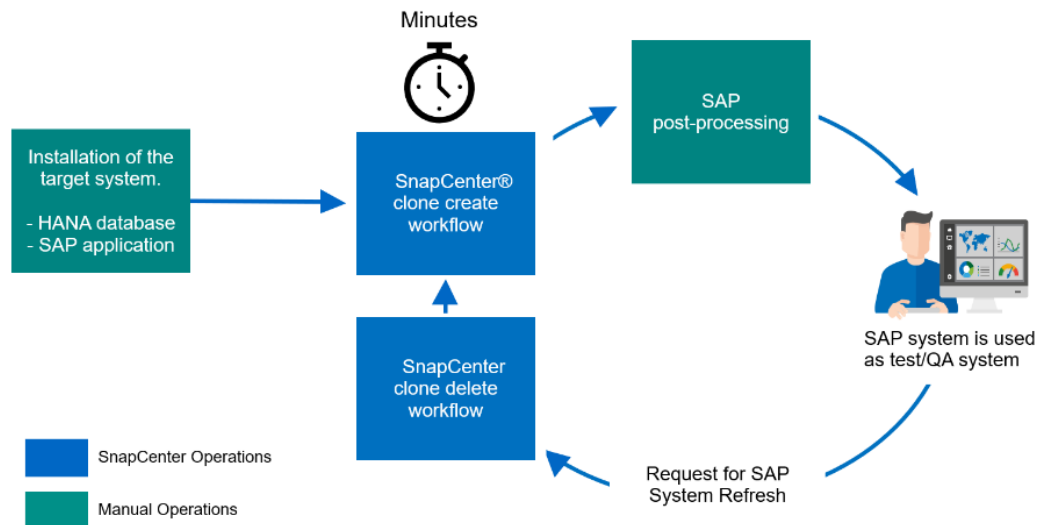
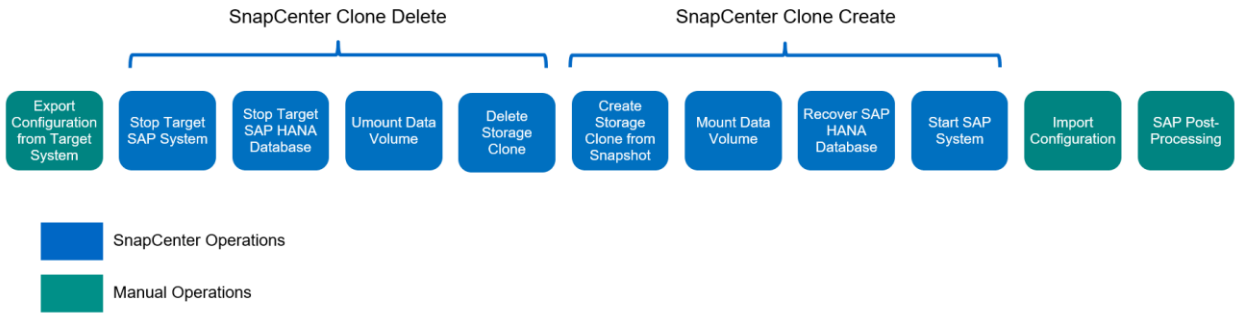


Figure 30 shows the different steps of an SAP system refresh operation. SnapCenter automates a significant portion of the workflow (blue boxes) while the SAP application specific steps (green boxes) must be executed manually. With the SnapCenter clone delete workflow, the SAP HANA database is stopped, the file system is unmounted and the storage clone is deleted. With the clone create workflow, a FlexClone volume, based on the selected Snapshot backup is created, the file system is mounted and the SAP HANA database is recovered.

Figure 30) SAP System refresh operation steps.



The SnapCenter clone create and clone delete workflows are extended using external scripts, which are executed at a specific workflow step. In our setup, we are using a script for mounting, unmounting, shutdown, and recovery. Figure 31 shows the interaction of the external scripts with the different SnapCenter workflow steps.

Figure 31) Automation script.

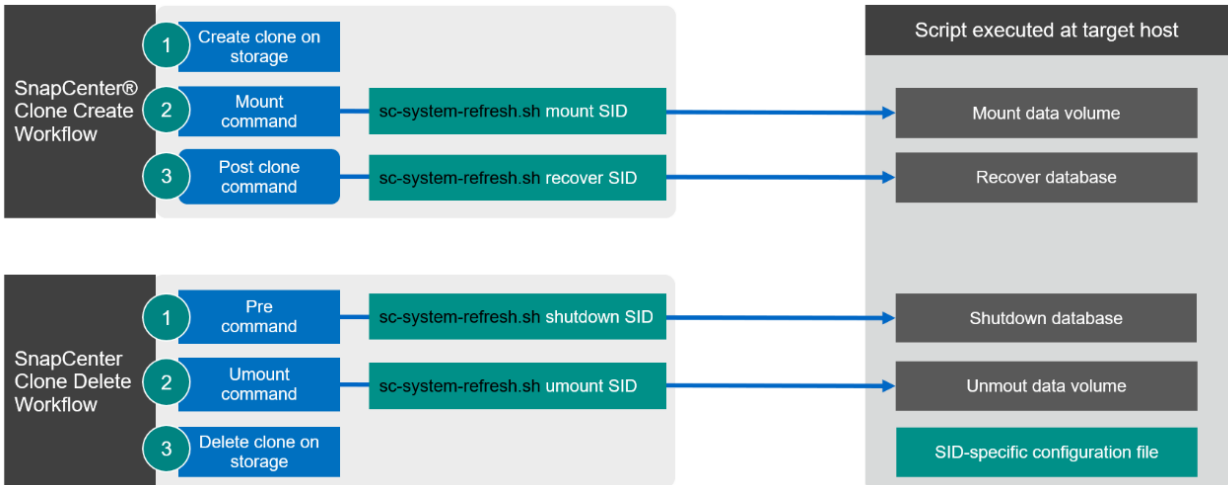
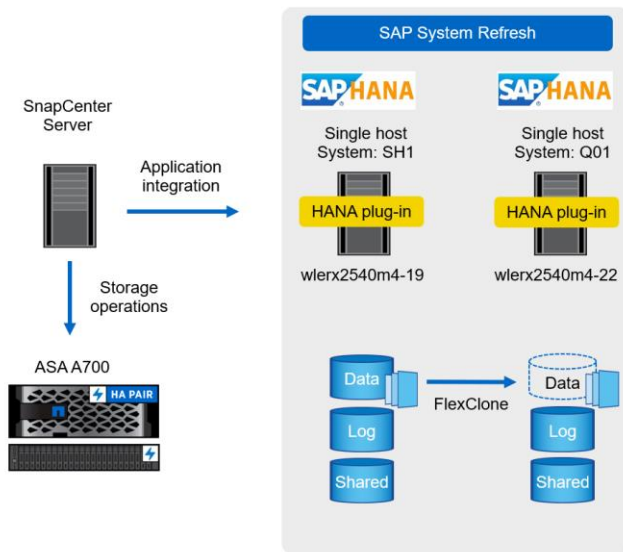


Figure 32 shows the lab setup, which is used for the SAP system refresh operation.

Figure 32) SAP system refresh lab setup.



SnapCenter clone create workflow

- 1) User selects Snapshot backup of source system
- 2) SnapCenter creates FlexClone volume
 - a. LUNs are mapped to target host
- 3) Mount script
 - a. Discover LUN
 - b. Mount file system
- 4) Recover script
 - a. Recover system database
 - b. Recover tenant database

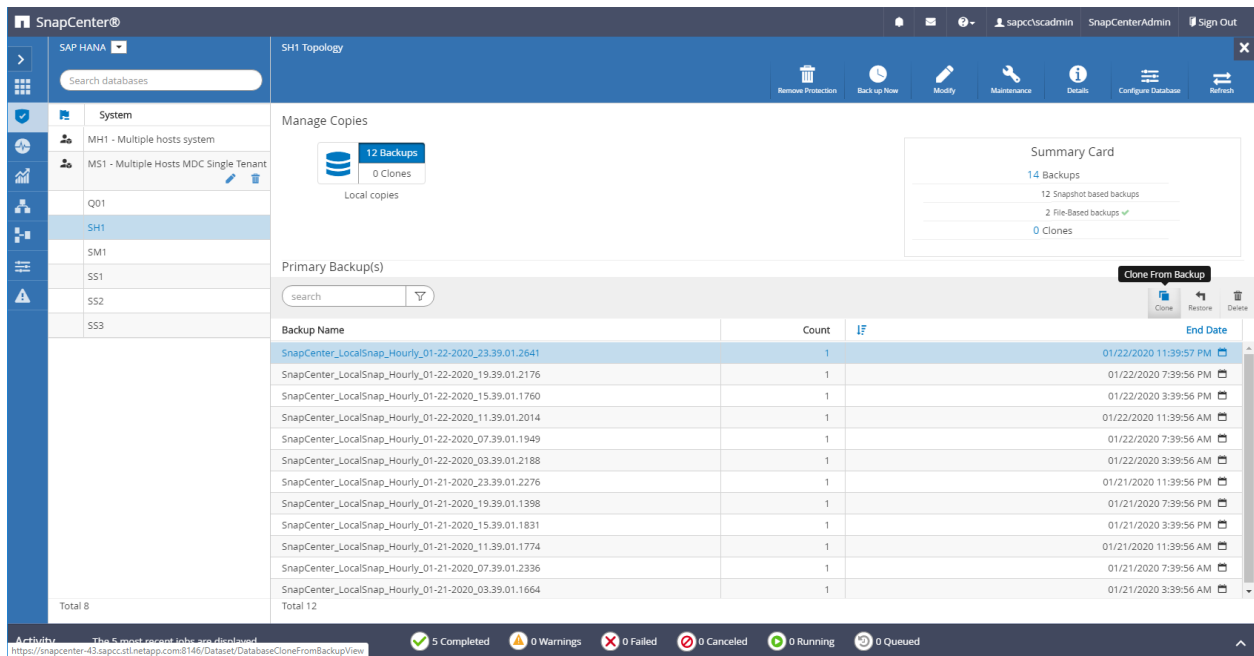
SnapCenter clone delete workflow

- 1) Pre script
 - a. Shutdown database
- 2) Umount script
 - a. Unmount file system
 - b. Clean-up multipathing, device files
- 3) SnapCenter deletes FlexClone volume

SnapCenter Clone Create Workflow

The clone create workflow is started from the topology view of the source system by selecting a Snapshot backup, which should be used for the SAP system refresh operation.

Figure 33) Clone from backup.



The target server must be selected. In an FC setup, SnapCenter maps the cloned LUN to the target server.

Figure 34) SnapCenter clone target.

Clone From Backup

1 Location Select the host to create the clone

2 Settings Clone server wlerx2540m4-22.stl.netapp.com

3 Scripts NFS Export IP Address

4 Notification

5 Summary

Configure an SMTP Server to send email notifications for Clone jobs by going to [Settings>Global Settings>Notification Server Settings](#).

Previous Next

The `pre clone`, `mount`, and `post clone` commands can now be provided. In this example, we used the `sc-system-refresh.sh` script for the mount and the recovery operation.

Figure 35) SnapCenter mount and post clone commands.

Clone From Backup

1 Location Enter optional commands to run before performing a clone operation

2 Settings Pre clone command

3 Scripts Enter optional commands to mount a file system to a host

4 Notification Mount command /snapcenter/sc-system-refresh.sh mount Q01

5 Summary Enter optional commands to run after performing a clone operation

Post clone command /snapcenter/sc-system-refresh.sh recover Q01

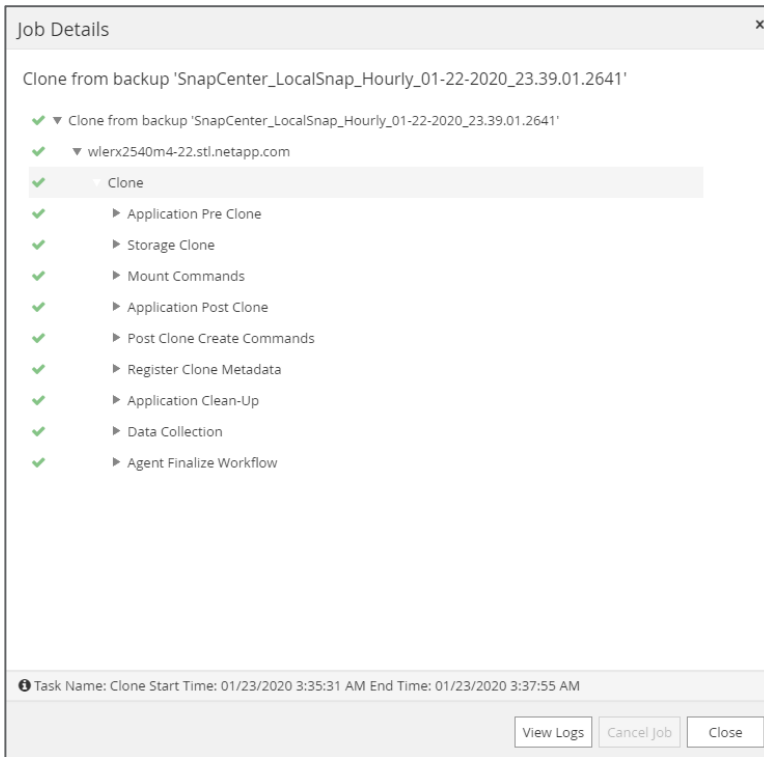
Configure an SMTP Server to send email notifications for Clone jobs by going to [Settings>Global Settings>Notification Server Settings](#).

Previous Next

The Job Details screen shows the different workflow steps as well as the start and end times. The complete workflow took two minutes and 20 seconds.

Note: The runtime is independent of the size of the SAP HANA database. Even for very large SAP HANA systems the storage cloning operation is finished in less than a minute.

Figure 36) SnapCenter job details of clone create operation.



In the log file of the `sc-system-refresh.sh` script, the different steps for the mount and recovery operation can be checked.

```

20200123063541###wlerx2540m4-22###sc-system-refresh.sh: Discover LUN and get UUID.
20200123063541###wlerx2540m4-22###sc-system-refresh.sh: Get source volume name and SnapCenter job
ID from environment.
20200123063541###wlerx2540m4-22###sc-system-refresh.sh: Source volume: S_19_SAP_Data_vol,
SnapCenter job ID: 5366
20200123063541###wlerx2540m4-22###sc-system-refresh.sh: Get volume and LUN name from SnapCenter
log file.
20200123063541###wlerx2540m4-22###sc-system-refresh.sh:
/vol/S_19_SAP_Data_vol01232003353819926/SH1_data_mnt00001
20200123063541###wlerx2540m4-22###sc-system-refresh.sh: Rescan SCSI bus.
20200123063608###wlerx2540m4-22###sc-system-refresh.sh: Get device name.
20200123063608###wlerx2540m4-22###sc-system-refresh.sh: Get UUID of LUN.
20200123063608###wlerx2540m4-22###sc-system-refresh.sh: Device: /dev/sdcr
20200123063608###wlerx2540m4-22###sc-system-refresh.sh: LUN UUID:
3600a0980383041334a3f4f6133735175
20200123063608###wlerx2540m4-22###sc-system-refresh.sh: Adding entry in /etc/fstab.
20200123063608###wlerx2540m4-22###sc-system-refresh.sh:
/dev/mapper/3600a0980383041334a3f4f6133735175 /hana/data/Q01/mnt00001 xfs relatime,inode64 0 0
20200123063608###wlerx2540m4-22###sc-system-refresh.sh: Mounting data volume: mount
/hana/data/Q01/mnt00001.
20200123063608###wlerx2540m4-22###sc-system-refresh.sh: Data volume mounted successfully.
20200123063608###wlerx2540m4-22###sc-system-refresh.sh: Change ownership to q0ladm.
20200123063616###wlerx2540m4-22###sc-system-refresh.sh: Recover system database.
20200123063616###wlerx2540m4-22###sc-system-refresh.sh: /usr/sap/Q01/HDB06/exe/Python/bin/python
/usr/sap/Q01/HDB06/exe/python_support/recoverSys.py --command "RECOVER DATA USING SNAPSHOT CLEAR
LOG"
20200123063647###wlerx2540m4-22###sc-system-refresh.sh: Wait until SAP HANA database is started
....
20200123063647###wlerx2540m4-22###sc-system-refresh.sh: Status: GRAY
20200123063657###wlerx2540m4-22###sc-system-refresh.sh: Status: GRAY
20200123063707###wlerx2540m4-22###sc-system-refresh.sh: Status: GREEN
20200123063707###wlerx2540m4-22###sc-system-refresh.sh: SAP HANA database is started.

```

```

20200123063707###wlerx2540m4-22###sc-system-refresh.sh: Recover tenant database Q01.
20200123063707###wlerx2540m4-22###sc-system-refresh.sh: /usr/sap/Q01/SYS/exe/hdb/hdbsql -U Q01KEY
RECOVER DATA FOR Q01 USING SNAPSHOT CLEAR LOG
0 rows affected (overall time 30.154339 sec; server time 30.152323 sec)

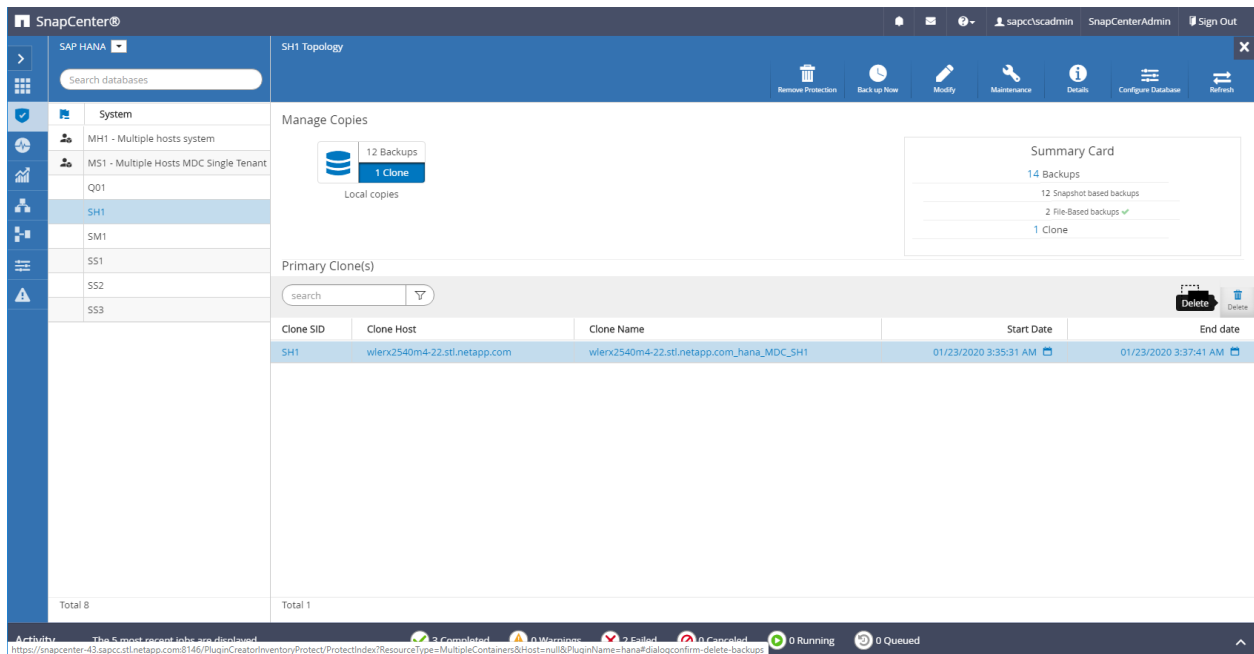
20200123063737###wlerx2540m4-22###sc-system-refresh.sh: Checking availability of Indexserver for
tenant Q01.
20200123063738###wlerx2540m4-22###sc-system-refresh.sh: Recovery of tenant database Q01
succesfully finished.
20200123063738###wlerx2540m4-22###sc-system-refresh.sh: Status: GREEN

```

SnapCenter Clone Delete Workflow

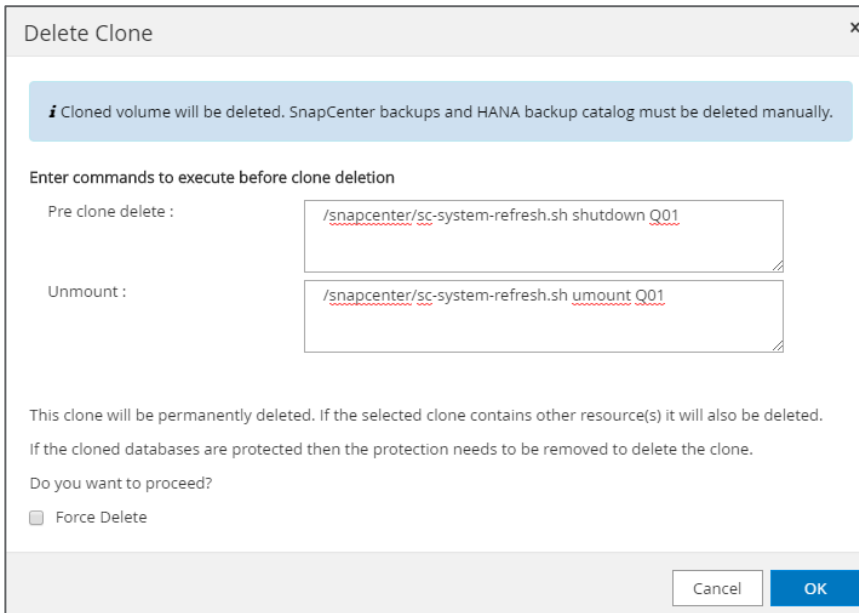
the SnapCenter clone delete workflow starts from the topology view of the source system by selecting the clone, which should be deleted.

Figure 37) SnapCenter clone delete.



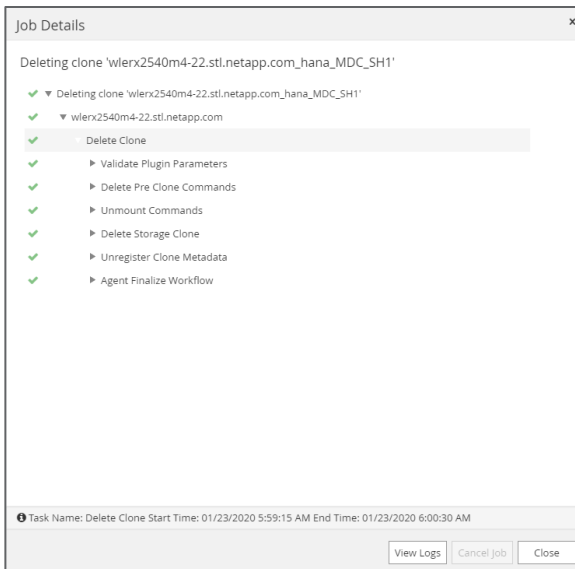
The pre clone and unmount scripts can now be added.

Figure 38) SnapCenter pre and unmount script.



The Job Details page shows the different steps of the workflow.

Figure 39) Job details of clone delete operation.



In the log file of the `sc-system-refresh.sh` script, the different steps for the mount and recovery operation can be checked.

```
20200123085915###wlerx2540m4-22###sc-system-refresh.sh: Stopping HANA database.
20200123085915###wlerx2540m4-22###sc-system-refresh.sh: sapcontrol -nr 06 -function StopSystem
HDB

23.01.2020 08:59:15
StopSystem
OK
20200123085915###wlerx2540m4-22###sc-system-refresh.sh: Wait until SAP HANA database is stopped
....
```

```

20200123085915###wlerx2540m4-22###sc-system-refresh.sh: Status: GREEN
20200123085926###wlerx2540m4-22###sc-system-refresh.sh: Status: GREEN
20200123085936###wlerx2540m4-22###sc-system-refresh.sh: Status: GREEN
20200123085946###wlerx2540m4-22###sc-system-refresh.sh: Status: GREEN
20200123085956###wlerx2540m4-22###sc-system-refresh.sh: Status: GREEN
20200123090006###wlerx2540m4-22###sc-system-refresh.sh: Status: GRAY
20200123090006###wlerx2540m4-22###sc-system-refresh.sh: SAP HANA database is stopped.
20200123090010###wlerx2540m4-22###sc-system-refresh.sh: Unmounting data volume.
20200123090010###wlerx2540m4-22###sc-system-refresh.sh: umount /hana/data/Q01/mnt00001
20200123090010###wlerx2540m4-22###sc-system-refresh.sh: Deleting /etc/fstab entry.
20200123090010###wlerx2540m4-22###sc-system-refresh.sh: Data volume unmounted successfully.
20200123090010###wlerx2540m4-22###sc-system-refresh.sh: Removing multipath map and device files.

```

7.4 SAP HANA Disaster Recovery

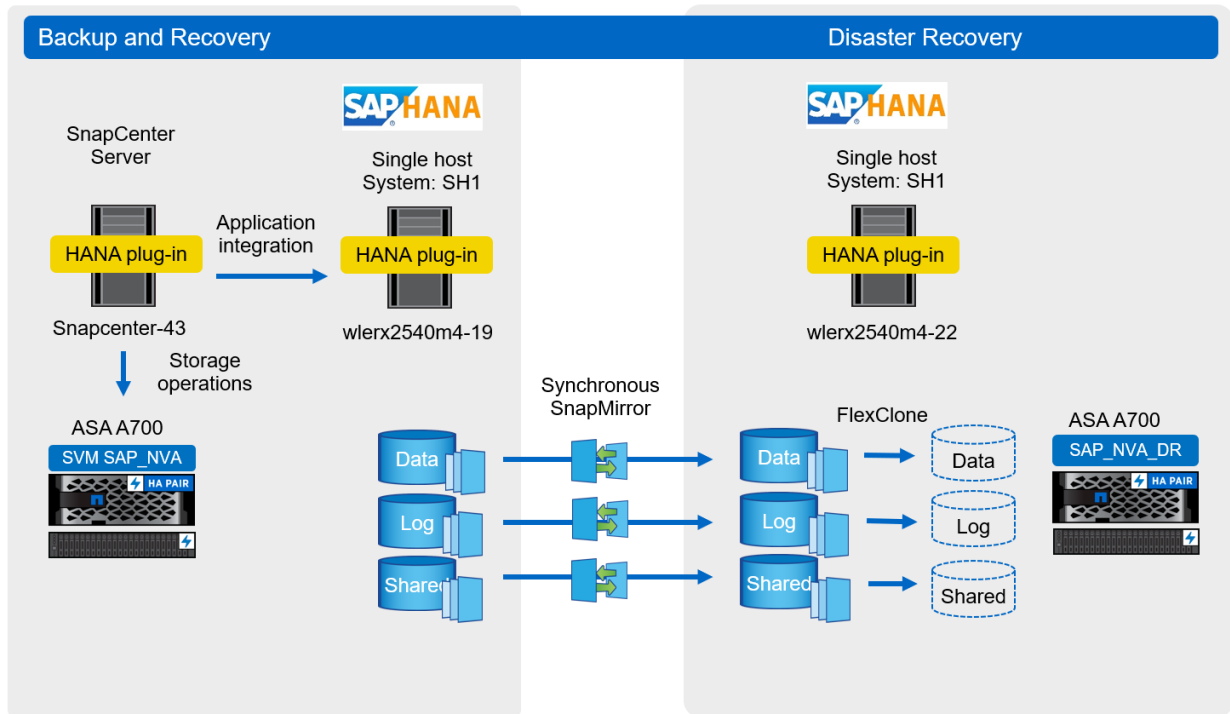
In addition to disaster recovery protection, the NetApp solution covers additional use cases when using SnapMirror storage replication technology. The use cases leverage NetApp Snapshot and cloning technologies to actively use the replicated data at the disaster recovery site to perform disaster recovery testing and for running SAP test systems without interrupting the ongoing replication. SnapMirror storage replication is supported for SAP HANA either with synchronous or with asynchronous replication.

Both replication methods can be used for the following use cases:

- Disaster recovery failover:
 - Replication is disabled
 - Replicated volumes are attached to the disaster recovery server
- Disaster recovery testing:
 - Replication is not interrupted
 - FlexClone volumes are attached to the disaster recovery server
- Running SAP test systems at the disaster recovery site:
 - Replication is not interrupted
 - SAP test system refresh with production data
 - Application consistent Snapshot backups are replicated
 - FlexClone volumes are attached to the SAP test/QA system

In this document, synchronous SnapMirror is used as the replication method. The two use cases, disaster recovery testing and disaster recovery failover, are described in the following sections. Figure 40 shows an overview of the solution setup.

Figure 40) SAP HANA disaster recovery.



Synchronous SnapMirror Impact on Latency

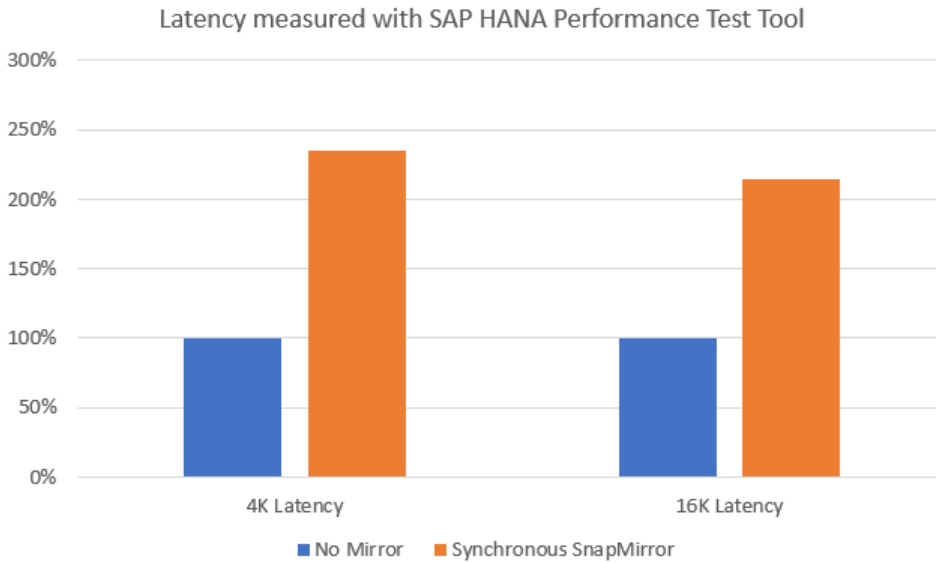
In order to understand the performance and latency impact of synchronous SnapMirror for SAP HANA, we validated the performance characteristics by using SAP HANA Performance Test Tool.

The graph in Figure 41 shows a comparison of a single host test run using a synchronously mirrored volume and a nonmirrored volume. The graph is normalized so that the nonmirrored values, the dark blue bars, are set to 100%. The orange bars show the latency values of the synchronously mirrored volume.

From the latency perspective, we see around 2 to 2.5 higher latency values for the 4K and the 16K block size with synchronous mirroring. But the latency numbers are still far below the required KPIs defined by SAP. Based on these results, data replication with synchronous SnapMirror is a suitable disaster recovery solution for SAP HANA.

Note: The measurements were performed with lab conditions and with the two storage systems located in the same data center. Additional latency, based on speed of light, must be considered, depending on the distance between the primary and disaster recovery storage system.

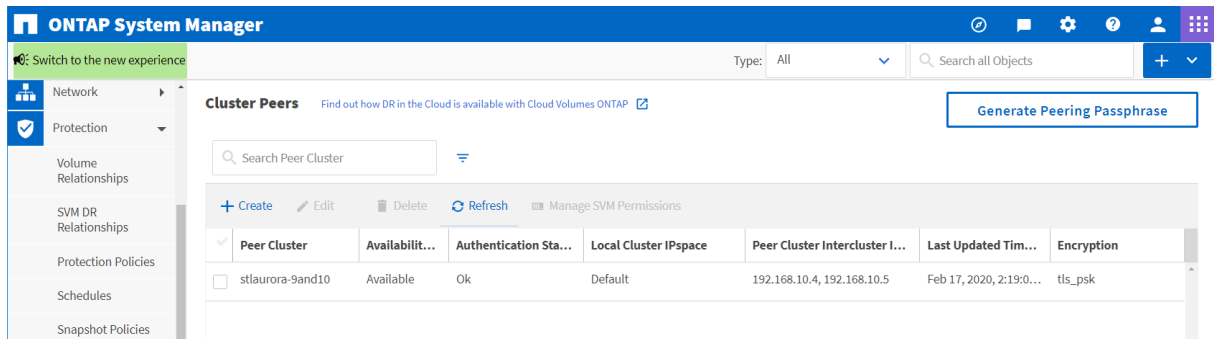
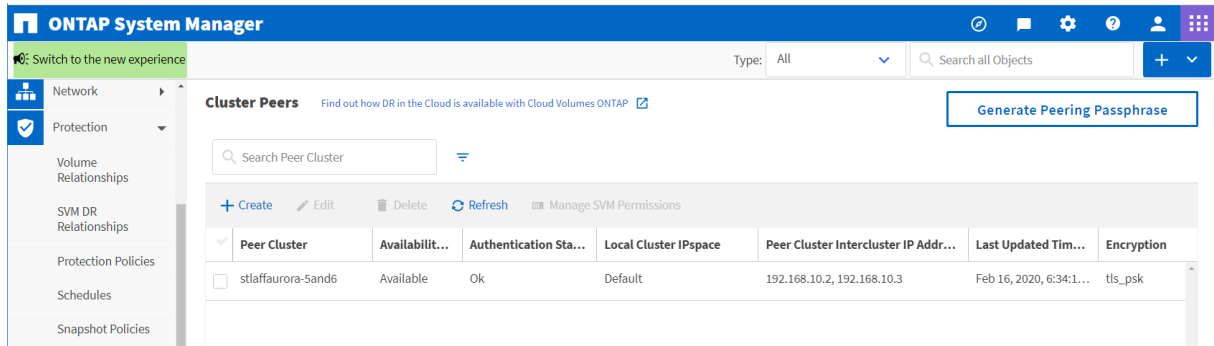
Figure 41 Latency with synchronous SnapMirror.



Synchronous Mirroring Configuration

Before storage volumes can be replicated with SnapMirror, the source and target clusters must be peered. In our lab setup, we configured one network interface on each storage node on each cluster for the replication network. Figure 42 shows the cluster peers.

Figure 42) Cluster peers, source, and target.



Volume replication can be configured by selecting the source volume and selecting Protect. Within the Protection Configuration window, the target storage cluster and the target volume must be selected.

SnapMirror replication could be either asynchronous or synchronous. Synchronous replication can be configured with Sync or StrictSync mode. With StrictSync mode RPO=0 is always guaranteed. If data can't be replicated to the secondary, an I/O error will occur at the primary site application.

Note: The predefined protection policy SnapCenterSync enables the replication of Snapshot copies from the source, which have a specific label. For disaster recovery testing, replication of Snapshot copies is required, therefore, the policy SnapCenterSync is selected.

Figure 43) Volume protection with Synchronous SnapMirror.

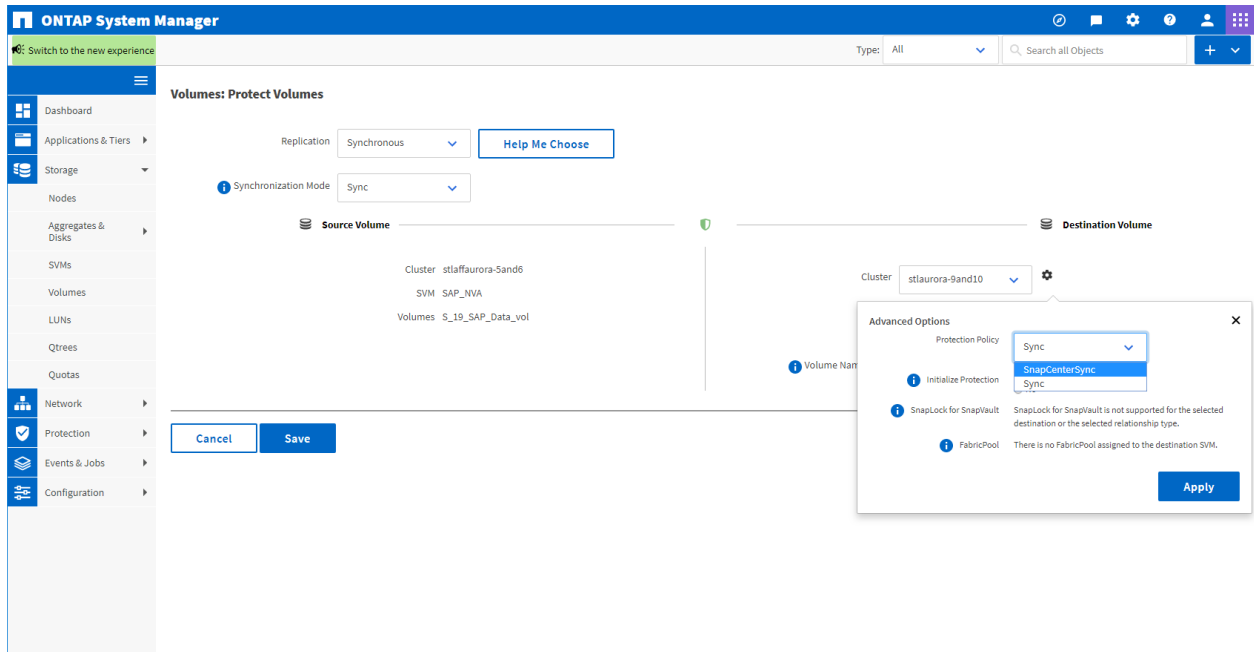
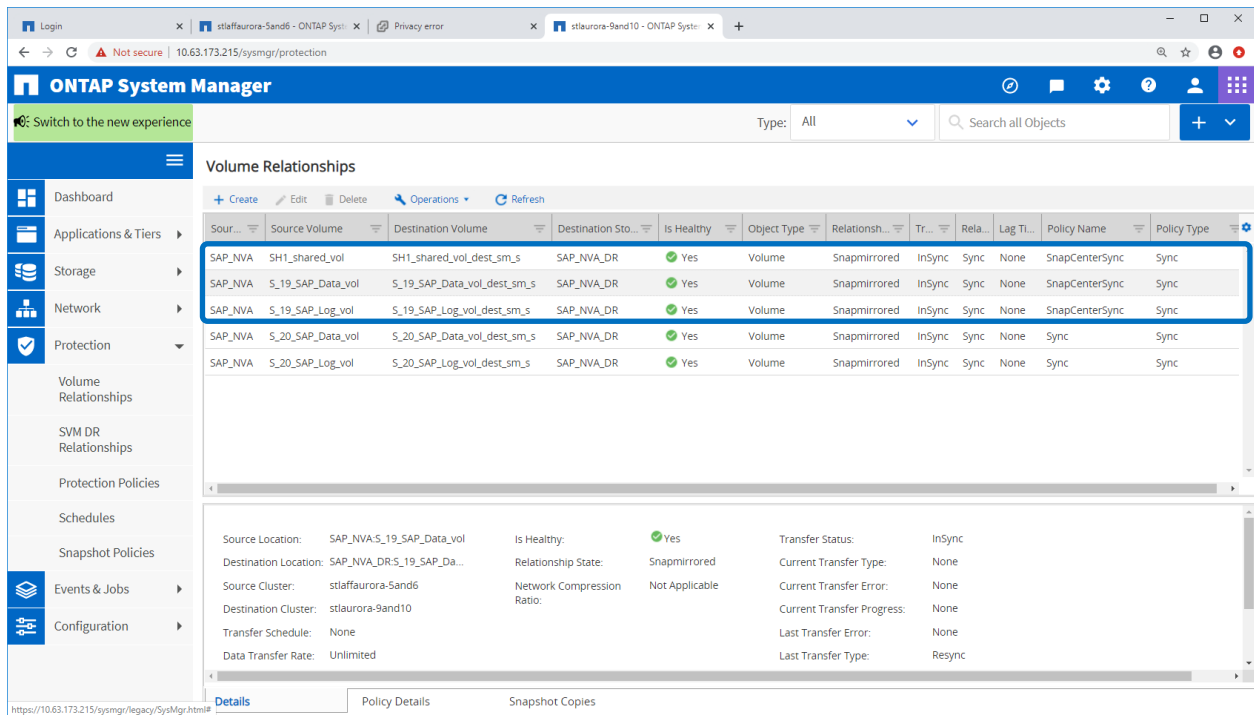


Figure 44 shows the final configuration, where data, log, and shared volumes are synchronously replicated.

Figure 44) Volume relationships.



Disaster Recovery Testing

A good disaster recovery strategy requires testing of the required workflow. Testing not only demonstrates whether the strategy works and the internal documentation is sufficient, but also allows administrators to train on the required procedures.

The use of NetApp FlexClone technology allows you to execute a disaster recovery failover test without influencing or interrupting the ongoing replication to the disaster recovery site. Therefore, a test can be run without influencing the RTO or the RPO.

The disaster recovery testing workflow can be divided into four main steps:

1. Prepare the target host:
 - a. Install the SAP host agent.
 - b. Add the sapservice entry of the source SAP HANA system.
 - c. Add the <sid>adm user of HANA source system.
 - d. Create the mount points for the SAP HANA source system file systems.
2. Create Snapshot copies at the source storage system:
 - a. Create crash-consistent Snapshot copies of data, log, and shared volumes.
3. Create FlexClone volumes at the target storage system:
 - a. Create FlexClone volumes of crash-consistent Snapshot copies.
 - b. Map LUNs to the target host.
4. Mount file systems and start HANA database at target host:
 - a. Discover LUNs.
 - b. Mount the file systems.
 - c. Start sapservices.

- a. Start the SAP HANA database.

The following sections briefly describe the required steps and operations.

Prepare Target Host

As a first target host preparation step, the SAP host agent software must be installed. The software can be downloaded from SAP's website <https://support.sap.com/en/my-support/software-downloads.html>.

After installing the host agent software, a sapservice for the SAP HANA system must be added to the `/usr/sap/sapservices` file. It can just be copied from the source system.

```
wlerx2540m4-22:/usr/sap # cat /usr/sap/sapservices
#!/bin/sh
limit.descriptors=1048576
LD_LIBRARY_PATH=/usr/sap/SH1/HDB00/exe:$LD_LIBRARY_PATH;export
LD_LIBRARY_PATH;/usr/sap/SH1/HDB00/exe/sapstartsrv
pf=/usr/sap/SH1/SYS/profile/SH1_HDB00_wlerx2540m4-19 -D -u shladm
```

The sh1adm user of the source system must be configured at the target host, if no central user management is used.

The mount points for the source HANA system must be created. In our example:

```
/hana/data/SH1/mnt00001, /hana/log/SH1/mnt00001, /hana/shared.
```

Create Snapshot at the Source Storage System

At the source storage system, Snapshot copies must be created for the data, the log, and the shared volume. The SnapMirror label `app_consistent`, as configured in the Synchronous SnapMirror policy, must be used. All Snapshot copies with this label will be replicated to the target storage system.

Note: It is important that the log volume Snapshot copy is created after the data volume Snapshot copy to ensure that the log data is newer than the data in the data volume.

```
stlaffaurora-5and6::> snap create -vserver SAP_NVA -volume S_19_SAP_Data_vol -snapshot
crash_consistent_snap1 -snapmirror-label app_consistent

stlaffaurora-5and6::> snap create -vserver SAP_NVA -volume S_19_SAP_Log_vol -snapshot
crash_consistent_snap1 -snapmirror-label app_consistent

stlaffaurora-5and6::> snap create -vserver SAP_NVA -volume SH1_shared_vol -snapshot
crash_consistent_snap1 -snapmirror-label app_consistent
```

Snapshot copy list at the source storage system.

```
Stlaffaurora-5and6::> snap show -vserver SAP_NVA -volume S_19_SAP_Data_vol -fields snapmirror-
label
vserver volume snapshot snapmirror-label
-----
SAP_NVA S_19_SAP_Data_vol SnapCenter_LocalSnap_Hourly_01-29-2020_03.39.01.4691 -
SAP_NVA S_19_SAP_Data_vol SnapCenter_LocalSnap_Hourly_01-29-2020_07.39.01.4968 -
SAP_NVA S_19_SAP_Data_vol SnapCenter_LocalSnap_Hourly_01-29-2020_11.39.01.5151 -
SAP_NVA S_19_SAP_Data_vol SnapCenter_LocalSnap_Hourly_01-29-2020_15.39.01.4857 -
SAP_NVA S_19_SAP_Data_vol SnapCenter_LocalSnap_Hourly_01-29-2020_19.39.01.5123 -
SAP_NVA S_19_SAP_Data_vol SnapCenter_LocalSnap_Hourly_01-29-2020_23.39.01.4980 -
SAP_NVA S_19_SAP_Data_vol SnapCenter_LocalSnap_Hourly_01-30-2020_03.39.01.5060 -
SAP_NVA S_19_SAP_Data_vol SnapCenter_LocalSnap_Hourly_01-30-2020_07.39.01.5287 -
SAP_NVA S_19_SAP_Data_vol SnapCenter_LocalSnap_Hourly_01-30-2020_11.39.01.5804 -
SAP_NVA S_19_SAP_Data_vol SnapCenter_LocalSnap_Hourly_01-30-2020_15.39.01.5346 -
SAP_NVA S_19_SAP_Data_vol SnapCenter_LocalSnap_Hourly_01-30-2020_19.39.01.4956 -
SAP_NVA S_19_SAP_Data_vol SnapCenter_LocalSnap_Hourly_01-30-2020_23.39.01.5336 -
SAP_NVA S_19_SAP_Data_vol crash_consistent_snap1 app_consistent
SAP_NVA S_19_SAP_Data_vol snapmirror.917ca9fb-435c-11ea-ad39-00a098af905e_2151923725.2020-02-
13_070500
```

14 entries were displayed.

Snapshot copy list at the target storage system.

```
stlaurora-9and10:> snap list -vserver SAP_NVA_DR -volume S_19_SAP_Data_vol_dest_sm_s -fields
snapmirror-label
vserver      volume                snapshot                snapmirror-label
-----
SAP_NVA_DR S_19_SAP_Data_vol_dest_sm_s crash_consistent_snap1 app_consistent
SAP_NVA_DR S_19_SAP_Data_vol_dest_sm_s snapmirror.917ca9fb-435c-11ea-ad39-
00a098af905e_2151923725.2020-02-13_070500
```

2 entries were displayed.

Create FlexClone Volumes at Target Storage System

At the target storage system, FlexClone volumes will now be created for the data, the log, and the shared volumes, using the replicated crash-consistent Snapshot copies. Figure 45 shows the FlexClone volume creation for the data volume.

Figure 45) Creating FlexClone volumes.

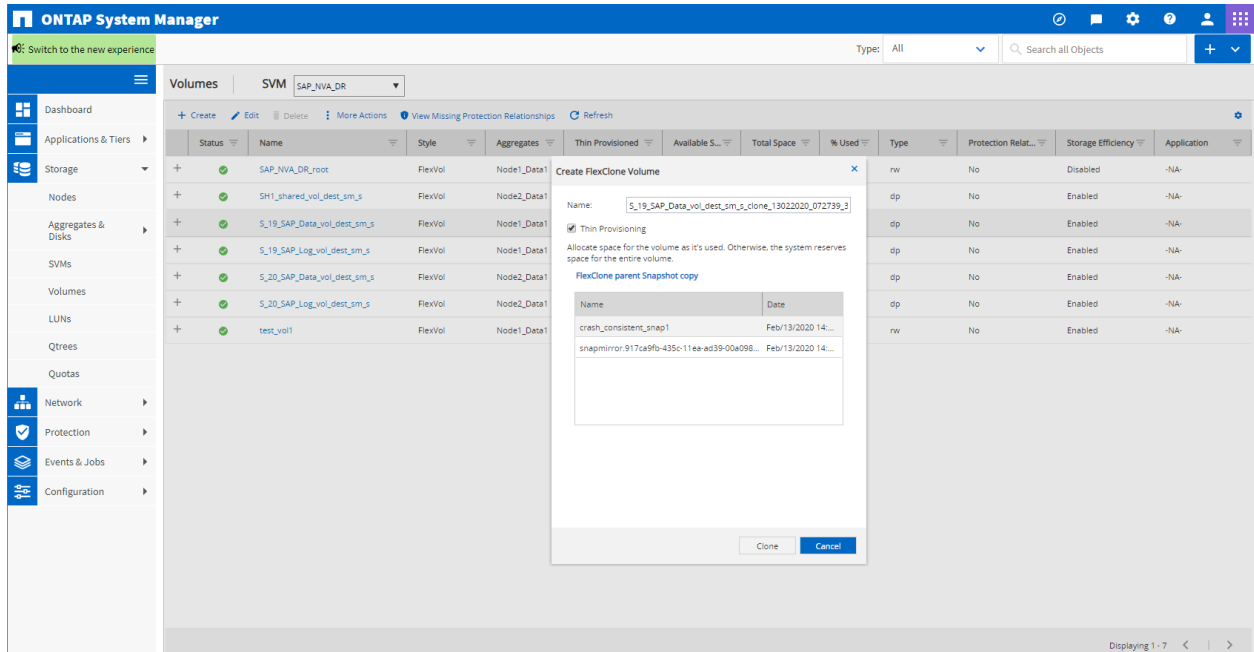


Figure 46 shows all three FlexClone volumes.

Figure 46) List of FlexClone volumes.

Status	Name	Style	Aggrega...	Thin Provid...	Availabl...	Total Sp...	% Used	Type	Protection R...	Storage Effi...	Application
+	SH1_shared_vol_dest_sm_s_clone_13022020_072904_29	FlexVol	Node2_Data1	Yes	5.21 GB	39.75 GB	86	rw	No	Enabled	-NA-
+	S_19_SAP_Log_vol_dest_sm_s_clone_13022020_072841_29	FlexVol	Node1_Data1	Yes	33.14 GB	361.93 GB	90	rw	No	Enabled	-NA-
+	S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31	FlexVol	Node1_Data1	Yes	31.04 GB	363.39 GB	91	rw	No	Enabled	-NA-
+	SAP_NVA_DR_root	FlexVol	Node1_Data1	No	970.81 MB	1 GB	0	rw	No	Disabled	-NA-
+	SH1_shared_vol_dest_sm_s	FlexVol	Node2_Data1	Yes	4.71 GB	37.45 GB	87	dp	No	Enabled	-NA-
+	S_19_SAP_Data_vol_dest_sm_s	FlexVol	Node1_Data1	Yes	33.53 GB	367.47 GB	90	dp	No	Enabled	-NA-
+	S_19_SAP_Log_vol_dest_sm_s	FlexVol	Node1_Data1	Yes	34.83 GB	366.38 GB	90	dp	No	Enabled	-NA-
+	S_20_SAP_Data_vol_dest_sm_s	FlexVol	Node2_Data1	Yes	53.19 GB	411.7 GB	87	dp	No	Enabled	-NA-
+	S_20_SAP_Log_vol_dest_sm_s	FlexVol	Node2_Data1	Yes	30.25 GB	372.07 GB	91	dp	No	Enabled	-NA-
+	test_vol1	FlexVol	Node1_Data1	Yes	46.65 GB	300 GB	84	rw	No	Enabled	-NA-

The LUN in each volume must now be mapped to the target host. Figure 47 shows the mapping for the data volume.

Figure 47) Mapping LUN to target host.

Name	Container Path
SH1_shared	/vol/SH1_shared_vol_dest_sm_s
SH1_shared	/vol/SH1_shared_vol_dest_sm_s_clone_13022020_072904_29
SH1_data_mnt00001	/vol/S_19_SAP_Data_vol_dest_sm_s
SH1_data_mnt00001	/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31
SH1_log_mnt00001	/vol/S_19_SAP_Log_vol_dest_sm_s
SH1_log_mnt00001	/vol/S_19_SAP_Log_vol_dest_sm_s_clone_13022020_072841_29
MH1_data_mnt00001	/vol/S_20_SAP_Data_vol_dest_sm_s
MH1_log_mnt00001	/vol/S_20_SAP_Log_vol_dest_sm_s

Map	Initiator Group Name	Type	LUN ID (Optional)
<input checked="" type="checkbox"/>	wlex2540m4_22	Linux	2

Mount File Systems and Start SAP HANA Database at the Target Host

After the LUNs have been mapped, the discover process at the target host can be started.

```
wlerx2540m4-19:/usr/sap # rescan-scsi-bus.sh -a
```

The device files of the different LUNs can be determined with the NetApp `sanlun` utility. Depending on the multipathing configuration, there will be multiple device files. The following output shows the output for the data volume LUN. The same command must be executed for the log and the shared volume LUNs.

```
wlerx2540m4-19:/usr/sap # sanlun lun show | grep SAP_NVA_DR | grep data
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sddp
host12 FCP 500.1g cDOT
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sddm
host12 FCP 500.1g cDOT
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sddj
host12 FCP 500.1g cDOT
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sddg
host12 FCP 500.1g cDOT
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sddd
host11 FCP 500.1g cDOT
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sdda
host11 FCP 500.1g cDOT
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sdcx
host11 FCP 500.1g cDOT
SAP_NVA_DR
/vol/S_19_SAP_Data_vol_dest_sm_s_clone_13022020_072739_31/SH1_data_mnt00001 /dev/sdcu
host11 FCP 500.1g cDOT
```

One of the device files can now be selected to determine the UUID of the LUN. The same step must be done for the log and shared LUNs.

```
wlerx2540m4-19:/usr/sap # /lib/udev/scsi_id -g -u -d /dev/sddp
3600a098038304132793f4f7050757369
```

This UUIDs are now used to mount the file systems.

```
wlerx2540m4-19:~ # cat /etc/fstab
/dev/system/swap swap swap defaults 0 0
/dev/system/root / btrfs defaults 0 0
/dev/system/root /.snapshots btrfs subvol=@/.snapshots 0 0
/dev/system/root /var btrfs subvol=@/var 0 0
/dev/system/root /usr/local btrfs subvol=@/usr/local 0 0
/dev/system/root /tmp btrfs subvol=@/tmp 0 0
/dev/system/root /srv btrfs subvol=@/srv 0 0
/dev/system/root /root btrfs subvol=@/root 0 0
/dev/system/root /opt btrfs subvol=@/opt 0 0
/dev/system/root /home btrfs subvol=@/home 0 0
/dev/system/root /boot/grub2/x86_64-efi btrfs subvol=@/boot/grub2/x86_64-efi 0 0
/dev/system/root /boot/grub2/i386-pc btrfs subvol=@/boot/grub2/i386-pc 0 0
UUID=AD43-17E1 /boot/efi vfat defaults 0 0

/dev/mapper/3600a098038304132793f4f705075736a /hana/log/SH1/mnt00001 xfs
relatime,inode64,nobarrier,noauto 0 0
/dev/mapper/3600a098038304132793f4f7050757369 /hana/data/SH1/mnt00001 xfs
relatime,inode64,noauto 0 0
/dev/mapper/3600a098038304133535d4f7466746a2f /hana/shared xfs defaults,noauto 0 0
```

```
wlerx2540m4-19:/usr/sap # df -h | grep 3600
/dev/mapper/3600a098038304132793f4f7050757369 500G 7.1G 493G 2% /hana/data/SH1/mnt00001
/dev/mapper/3600a098038304132793f4f705075736a 500G 5.7G 495G 2% /hana/log/SH1/mnt00001
/dev/mapper/3600a098038304133535d4f7466746a2f 500G 67G 434G 14% /hana/shared
```

After mounting the file systems, the sapservices and the SAP HANA database can be started.

```
wlerx2540m4-19:/usr/sap # systemctl start sapinit
```

```
wlerx2540m4-19:~ # su - shladm
shladm@wlerx2540m4-19:/usr/sap/SH1/HDB00> sapcontrol -nr 00 -function StartSystem HDB

14.02.2020 04:40:44
StartSystem
OK

shladm@wlerx2540m4-19:/usr/sap/SH1/HDB00> sapcontrol -nr 00 -function GetSystemInstanceList

14.02.2020 04:42:12
GetSystemInstanceList
OK
hostname, instanceNr, httpPort, httpsPort, startPriority, features, dispstatus
wlerx2540m4-19, 0, 50013, 50014, 0.3, HDB|HDB_WORKER, GREEN
shladm@wlerx2540m4-19:/usr/sap/SH1/HDB00>
```

Disaster Recovery Failover

The disaster recovery failover workflow is very similar to the disaster recovery testing workflow. The preparation of the target host, the mounting of the file systems, and the starting of the SAP HANA database is identical. The difference between the two workflows is the break operation of the replication.

1. Prepare the target host:
 - a. Install the SAP host agent.
 - b. Add the sapservice entry of source the SAP HANA system.
 - c. Add the <sid>adm user of the SAP HANA source system.
 - d. Create mount points for the SAP HANA source system file systems.
2. Break the SnapMirror replication at the target storage system:
 - a. Initiate the SnapMirror quiesce operation.
 - b. Initiate the SnapMirror break operation.
 - c. Map LUNs to disaster recovery SAP HANA host.
3. Mount the file systems and start the SAP HANA database at the target host:
 - a. Discover the LUNs.
 - b. Mount the file systems.
 - c. Start sapservices.
 - d. Start the SAP HANA database.

The following sections describe step 2 of the workflow. Step 1 and step 3 are identical to the steps described in the Disaster Recovery Testing section.

Break the SnapMirror Replication at Target Storage System

The replication must first be quiesced, as shown in Figure 48 and Figure 49.

Figure 48) SnapMirror quiesce operation.

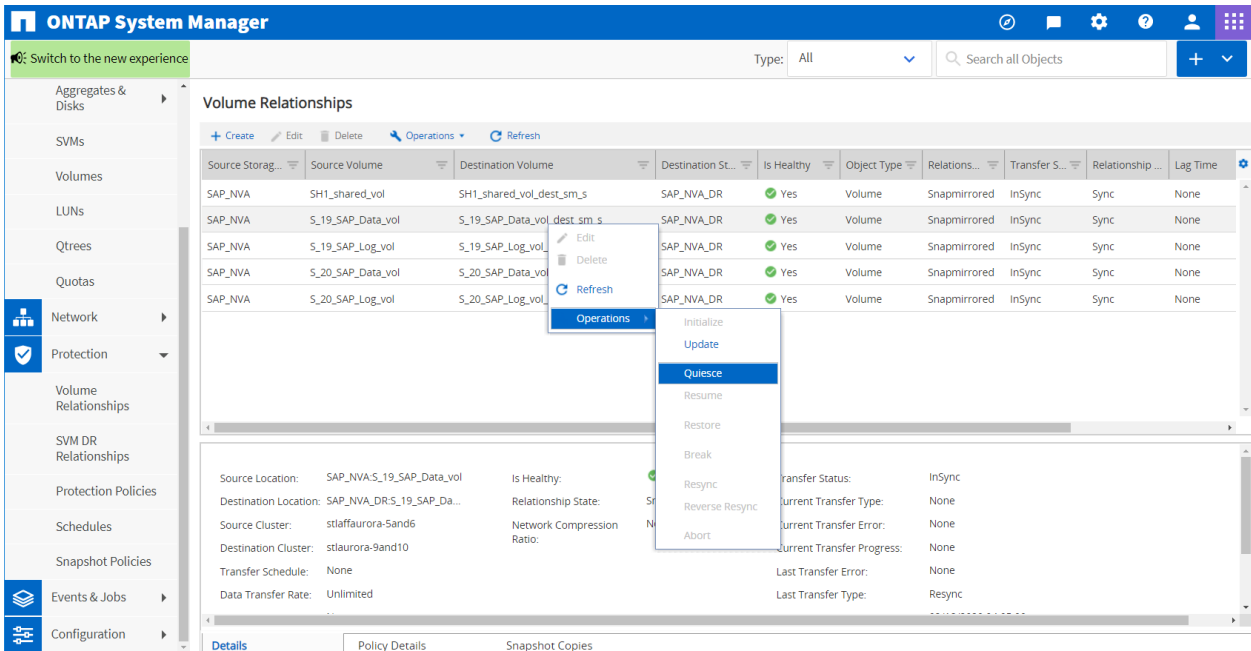


Figure 49) SnapMirror quiesce operation.

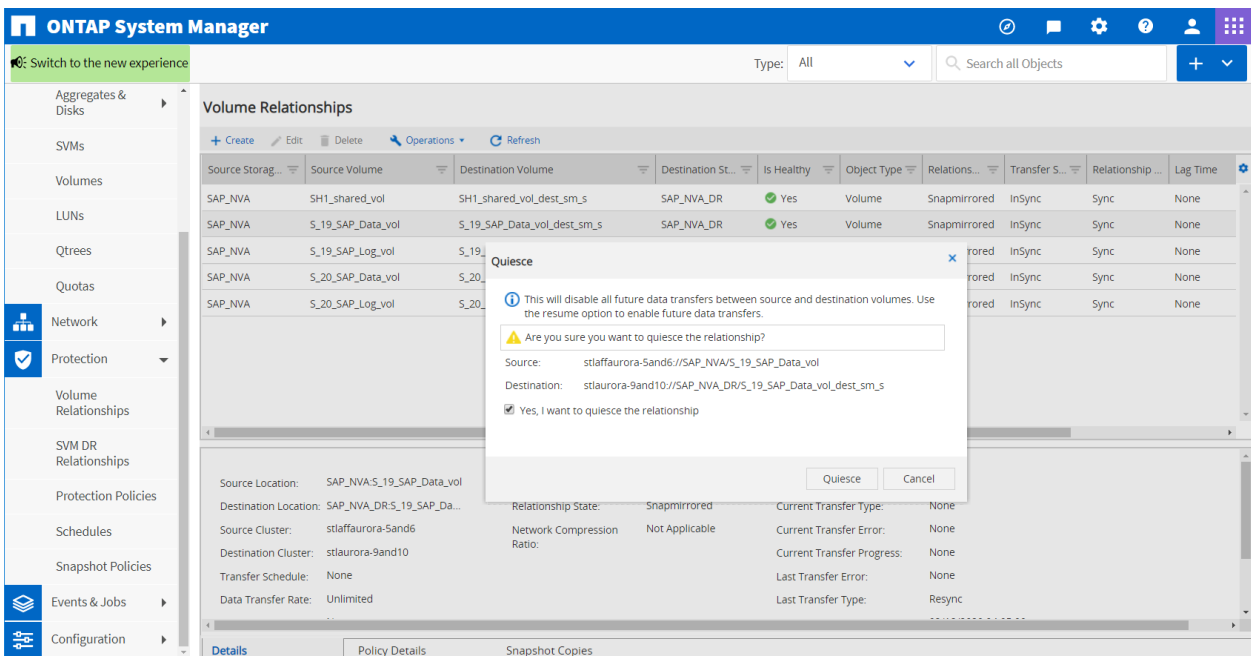
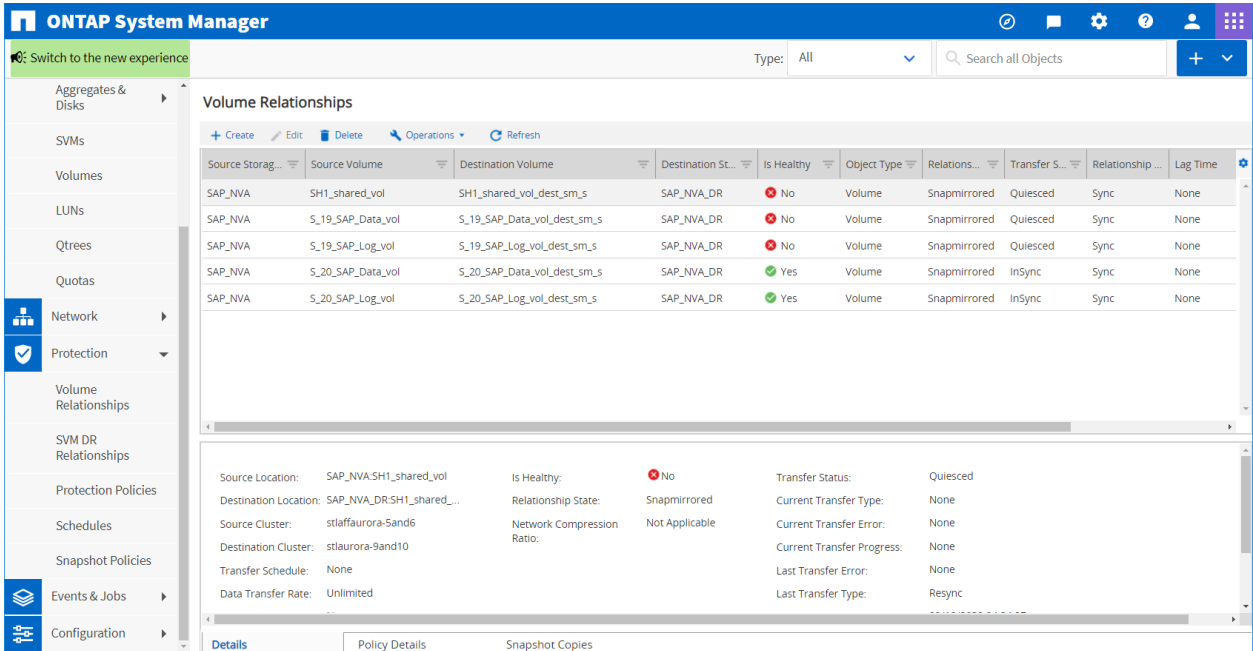


Figure 52 shows the list of volumes that display after the quiesce operation.

Figure 50) List of volume relationships.



The replication must be broken off, as shown in Figure 51 and Figure 52.

Figure 51) SnapMirror break operation.

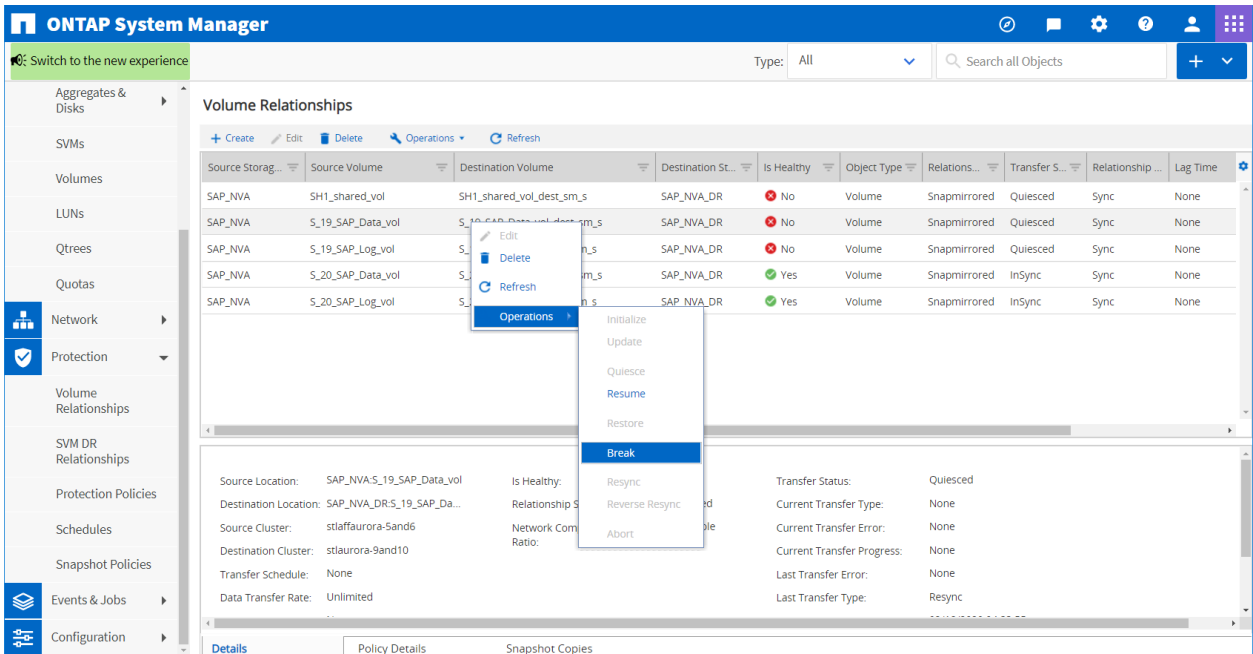


Figure 52) SnapMirror break operation.

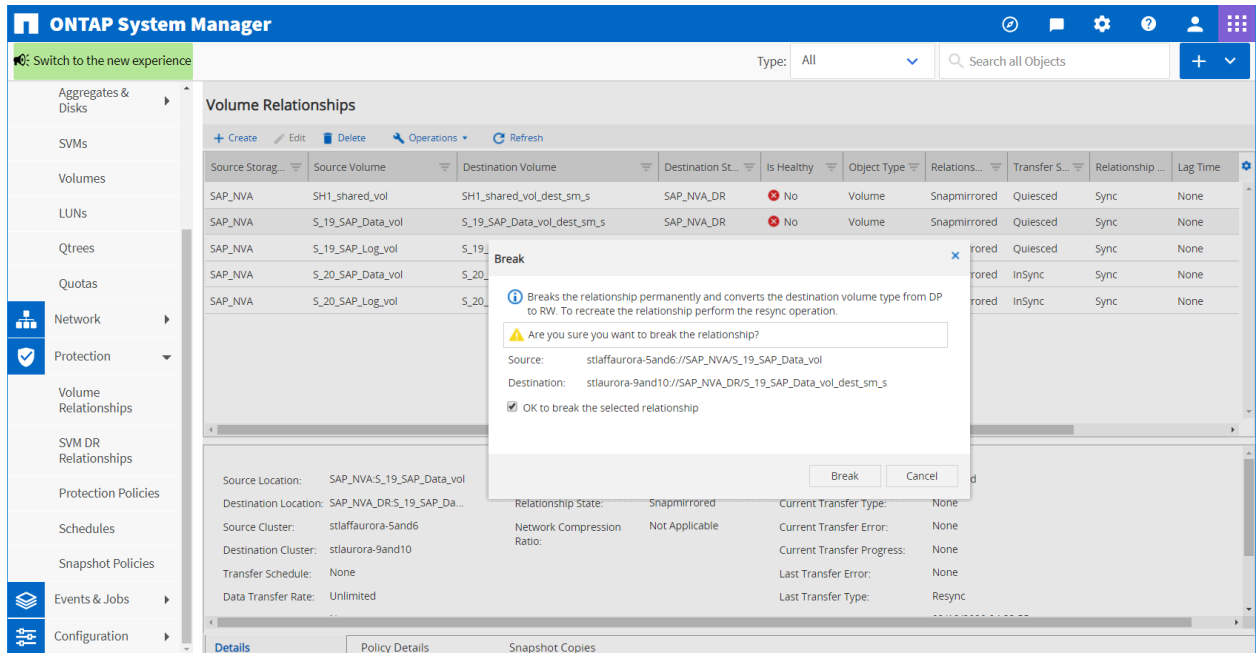
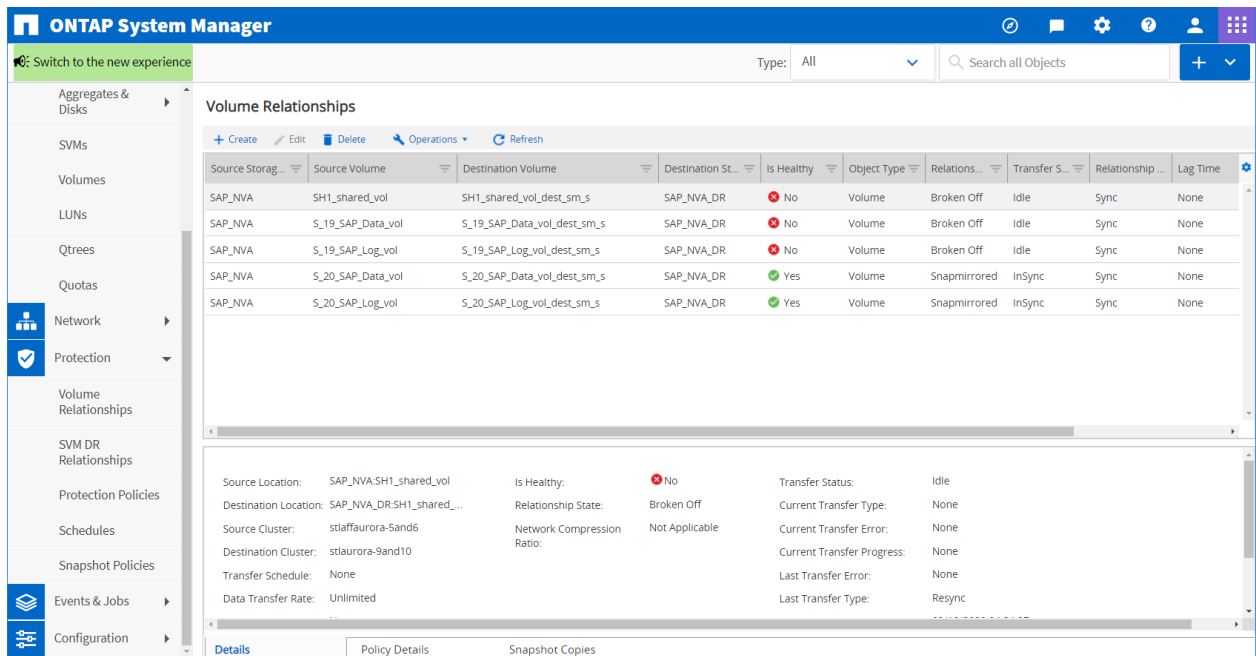


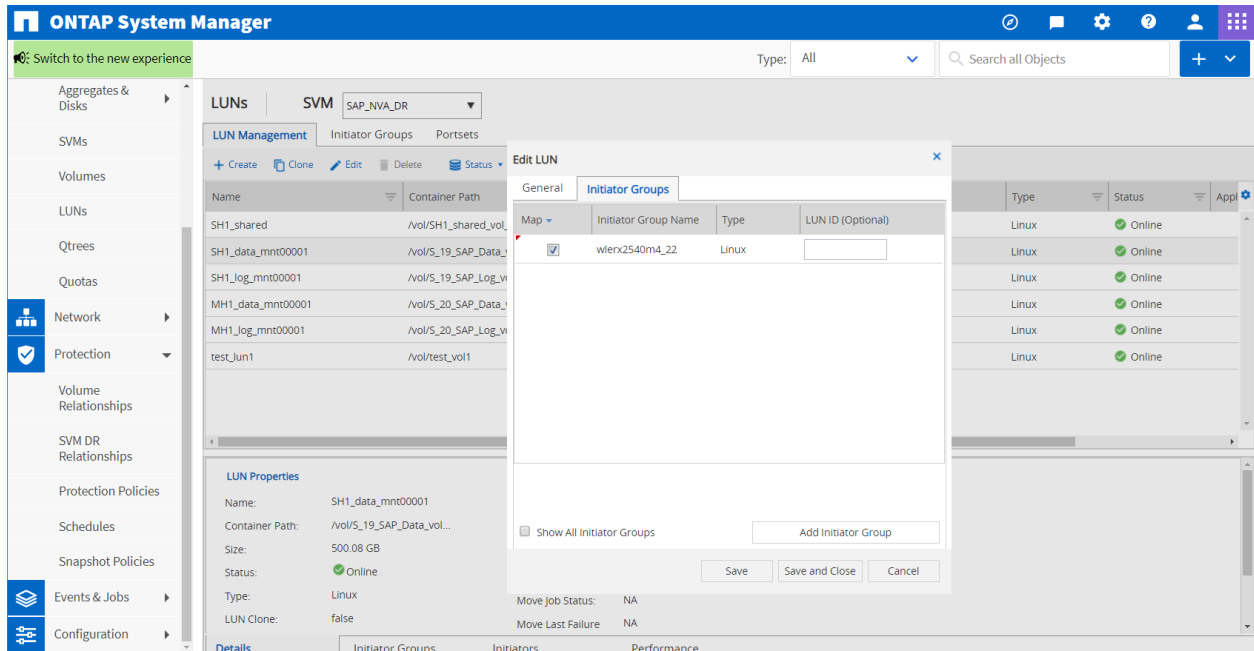
Figure 53 shows the list of volumes that display after the break operation.

Figure 53) List of volume relationships.



All LUNs (data, log, and shared) must be mapped to the disaster recovery server, as shown in Figure 54.

Figure 54) Map LUN to disaster recovery server.



After all the LUNs have been mapped to the disaster recovery server, the mount operation of the LUNs and the start of the SAP HANA database must be executed, as described in the section titled, “Disaster Recovery Testing.”

8 Technology Requirements

8.1 Hardware Requirements

Table 1 lists the hardware components that are required to implement the solution. The hardware components that are used in any particular implementation of the solution might vary based on customer requirements.

Table 1) Hardware requirements.

Hardware	Quantity
ASA A700	<ul style="list-style-type: none"> • 2 dual-node solutions • 1x primary storage system • 1x disaster recovery storage system
High available NFS storage. Note: Only required for /hana/shared on NFS with SAP HANA multiple host systems. NetApp AFF A800 system was used in the lab setup for this purpose, though an entry-level AFF or FAS system is sufficient.	1
Fujitsu PRIMERGY RX2540 M4, with 20 Intel Xeon CPU cores and 256GB of DDR4 memory	4
Emulex Corporation Lancer Gen6 LPe32000 Fibre Channel Host Adapters (32Gb) by Broadcom	<ul style="list-style-type: none"> • 8 • 2x for each server

Hardware	Quantity
FC Switch: Brocade Gen6 G620 running FOS V8.1.0a	2

8.2 Software Requirements

Table 2 lists the software components that are required to implement the solution. The software components that are used in any particular implementation of the solution might vary based on customer requirements.

Table 2) Software requirements.

Software	Version
NetApp SnapCenter	4.3
SAP HANA	2.0 SPS 4 revision 44
SuSE SLES for SAP	15
NetApp ONTAP	9.7P1

9 Conclusion

This report presents a modern NetApp and Broadcom enterprise SAN verified architecture for deploying and using SAP HANA on a modern NetApp ASA with Brocade Gen6 High Performance Fabric.

This solution provides an optimal infrastructure approach for organizations to leverage best-in-class, end-to-end, modern SAN technologies, and deliver business-critical IT services for your SAP solutions today, while preparing for the future—and at a compelling ROI and TCO.

With the increasing need for rapid innovation, faster data delivery and analysis, and business agility that future includes serving high-performance database, analytics, AI/machine learning, and IoT requirements from modern SAN infrastructure.

NetApp and Broadcom have created an architecture framework that is future-ready, usable today, and accessible for organizations to implement within current operational processes and procedures. A primary objective is to enable organizations like yours to quickly and nondisruptively streamline and modernize your traditional SAN infrastructure and the critical IT services like SAP HANA that rely on it.

To meet this objective, these modern platforms must:

- Be high performing to offer more real-time analysis and availability of critical data
- Adopt modern future-facing and disruptive technologies in a nondisruptive manner
- Provide agility, flexibility, and high scalability
- Fit within current operational frameworks
- Align with organizational objectives to consolidate and streamline infrastructure and operations
- Provide compelling ROI and TCO
- Protect and secure your organizations critical asset...it's data!

Exhaustive tests with an SAP HANA workload demonstrate the benefits of a modern SAN architecture, suited for multiple use cases and critical SAN-based workloads. These benefits directly apply to the SAP HANA workload presented in this report.

With the flexibility and scalability of this NetApp Verified Architecture, you can start with a framework to modernize and right-size your organization's SAP HANA infrastructure and can continue to grow and adapt to evolving business requirements.

This modern SAN solution will enable you to accelerate your SAP HANA implementation projects by streamlining your SAP lifecycle management. Development cycles will be shortened by using instantaneous cloning. Important data protection operations such as backup and recovery will be simplified, with the impact on the performance of production systems minimized, and your data protected, wherever it lives with the NetApp fully integrated SAP HANA data protection and CommVault solutions.

Where to Find Additional Information

To learn more about the information that is described in this document, review the following documents and/or websites:

- TR-4436: SAP HANA on NetApp All Flash FAS Systems with Fibre Channel Protocol - Configuration Guide
<http://www.netapp.com/us/media/tr-4436.pdf>
- TR-4384: SAP HANA on NetApp FAS Systems with Fibre Channel Protocol - Configuration Guide
<http://www.netapp.com/us/media/tr-4384.pdf>
- TR-4614: SAP HANA Backup and Recovery with SnapCenter
<https://www.netapp.com/us/media/tr-4614.pdf>
- TR-4667: Automating SAP System Copies Using the SnapCenter
<https://www.netapp.com/us/media/tr-4667.pdf>
- TR-4719: SAP HANA System Replication, Backup and Recovery with SnapCenter
<https://www.netapp.com/us/media/tr-4719.pdf>
- TR-4646: SAP HANA Disaster Recovery with Asynchronous Storage Replication
<https://www.netapp.com/us/media/tr-4646.pdf>
- TR-4018: Integrating NetApp ONTAP Systems with SAP Landscape Management
<https://www.netapp.com/us/media/tr-4719.pdf>
- TR-4080 NetApp Best Practice for Modern SAN and ONTAP 9
<https://www.netapp.com/us/media/tr-4080.pdf>
- TR-4515 ONTAP AFF All SAN Array Systems
<https://www.netapp.com/us/media/tr-4515.pdf>
- NetApp SAN Solutions Website
<https://www.netapp.com/SAN>
- Brocade Fibre Channel Networking Switches
<https://www.broadcom.com/products/fibre-channel-networking/switches/>
- Brocade Fibre Channel Networking Directors
<https://www.broadcom.com/products/fibre-channel-networking/directors/>
- Brocade/NetApp Partner Documents
<https://www.broadcom.com/company/oem-partners/fibre-channel-networking/netapp>
- NetApp SAN Health Program
https://www.netapp.com/us/forms/campaign/amer-us-fy19q3-sss-san-san-health-check-inquiry-form.aspx?ref_source=smc&cid=27476

Version History

Version	Date	Document Version History
Version 1.0	June 2020	Initial version.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2020 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

NVA-1147-DESIGN