

Technical Report

SAP Applications on Microsoft Azure

Using Azure NetApp Files

Bernd Herth, Tobias Brandl, Nils Bauer, NetApp

Will Bratton, Geert van Teylingen, Amish Patel, Juergen Thomas, Ralf Klahr, Microsoft

March 2020 | TR-4746

In partnership with



Abstract

This document provides best practices for leveraging Azure NetApp® Files for SAP Applications and SAP HANA deployments. It also details the different use cases from SAP shared file systems to full SAP HANA deployments and specific performance, data protection, and migration considerations with SAP HANA on Azure NetApp Files.

TABLE OF CONTENTS

1	SAP on Microsoft Azure Overview.....	5
1.1	NetApp Values and Solutions on Microsoft Azure.....	5
1.2	Microsoft Certified Solutions for SAP	8
1.3	Enterprise Storage Requirements.....	9
1.4	SAP Use Cases	11
2	Shared Files Requirements	12
2.1	Shared File System for SAP NetWeaver Application Server	12
2.2	Shared Files Solutions for the Cloud.....	13
2.3	Providing Shared Files using Azure NetApp Files.....	13
3	SAP HANA	16
3.1	Certified Azure Native Solutions	16
3.2	Storage Requirements for SAP HANA.....	17
3.3	SAP HANA Azure NetApp Files Sizing (for Capacity/Performance)	19
4	SAP HANA Data Protection, System Copies and Migration.....	22
4.1	Snapshot Data Protection in Azure with Azure NetApp Files.....	22
4.2	Protecting Snapshot Copies Against Local and Regional Disasters	26
4.3	SAP HANA Backup in Seconds Using Snapshot Copies.....	27
4.4	SAP HANA Restore in Seconds Using Snapshot Copies	32
4.5	SAP HANA System Copies Using Snapshot Copies (Thick Clones)	49
5	Example: Installation of an SAP NetWeaver ABAP Stack.....	51
5.1	Installation Planning.....	52
5.2	Prepare Storage on Azure NetApp Files.....	54
5.3	Prepare the Operating System and Mount the Volumes.....	58
6	Example: Cost Optimized Deployments.....	62
6.1	Know Your Limits.....	62
6.2	Optimization for Database Files.....	63
6.3	Optimization for Multiple SIDs.....	64
7	Example: Installation of SAP HANA Scale-Out with Host Auto Failover	64
8	SAP BW HANA Landscape with Azure NetApp Files.....	65
8.1	Overview of SAP BW HANA Landscape with Azure NetApp Files.....	66

Where to Find Additional Information	67
---	-----------

Version History	67
------------------------------	-----------

LIST OF TABLES

Table 1) NetApp ONTAP offerings side-by-side.....	7
Table 2) Azure SAP certified VMs	8
Table 3) SBX NFS file systems.	52
Table 4) Volume and folder structure.	52
Table 5) SAP HANA basic file systems.	53
Table 6) Volume and folder structure for SAP HANA.	53
Table 7) Backup volumes.	54

LIST OF FIGURES

Figure 1) Cloud Volumes ONTAP.	6
Figure 2) Azure NetApp Files.	7
Figure 3) SAP storage performance requirements and certification.	10
Figure 4) Shared files requirements.	11
Figure 5) Shared files for SAP.....	12
Figure 6) Linux cluster NFS server.....	13
Figure 7) Azure NetApp Files.	14
Figure 8) Performance Tiers.....	15
Figure 9) Capacity pool and volumes.	16
Figure 10) SAP HANA single-host system.	17
Figure 11) SAP HANA multiple-host system.	18
Figure 12) Tailored datacenter integration (TDI) performance criteria.....	20
Figure 13) Static versus dynamic storage (performance) provisioning.....	21
Figure 14) Storage tuning & cost optimization.....	21
Figure 15) Traditional backup.....	23
Figure 16) Snapshot copy-based backup.....	23
Figure 17) Snapshot copy creation.....	24
Figure 18) Snapshot restore operation.....	25
Figure 19) SAP HANA local Snapshot copies.	26
Figure 20) Cloud Sync – offload backups to Azure Blob.	27
Figure 21) HANA backup using storage Snapshot copies.....	28
Figure 22) HANA backup using Python script.	32
Figure 23) SAP HANA restore workflow.	33
Figure 24) SAP HANA restore options.	35
Figure 25) Typical SAP landscape.	49
Figure 26) SAP system refresh/copy workflow.	51

Figure 27) SAP NetWeaver on SAP HANA.....52
Figure 28) Cost optimization for shared files – single volume directories to client mountpoints mapping.63
Figure 29) Cost optimization for scale-out – volume directories to client mountpoints mapping.....63
Figure 30) Typical SAP HANA scale-out with host auto failover setup65
Figure 31) SAP BW HANA landscape with Azure NetApp files.66

1 SAP on Microsoft Azure Overview

Today, many customers use Microsoft Azure to accelerate their SAP deployments in order to reduce cost and provide increased agility for their business processes. All of these benefits are important to SAP IT leaders who use a cloud-first strategy. Moreover, moving the SAP estate to Azure and integrating SAP with Azure's vast array of platform as a service (PaaS) features such as Azure Data Factory, Azure IoT Hub, and Azure Machine Learning creates business value to support digitalization ambitions.

Many large enterprises choose Azure as the cloud platform of choice for their enterprise applications, including the SAP Business Suite and S/4HANA. Many customers are embracing the Dev/Ops paradigm by first moving their development and test SAP systems; however, more customers are now choosing to migrate their complete SAP infrastructure, including production, into the cloud.

Azure's vast SAP offering ranges from small virtual machines (VMs) for SAP application servers up to tailored SAP HANA on Azure (large instances). These instances can scale to 24TB single host and 120TB multiple-host configurations. In 2018, Microsoft introduced the Azure M-Series VMs with up to 4TB of memory; VMs with 12TB of memory have been certified recently. These colossal VMs are targeted at specific workloads such as SAP HANA.

To get started with your SAP on Azure journey, see the Microsoft Azure article: [Using Azure for Hosting and Running SAP Workload Scenarios](#).

1.1 NetApp Values and Solutions on Microsoft Azure

NetApp storage and data management solutions, based on the NetApp ONTAP® data management software, have been the foundation for many customers' enterprise workloads such as SAP. NetApp ONTAP systems and NFS services have been used in many of the largest SAP deployments for more than 15 years. These technologies provide a secure and stable operation and simplify the data management, which help speed up projects and reduce risk.

As a global SAP technology partner, NetApp has a long history of excellent solutions and products with a deep integration into SAP applications, enabling customers to use NetApp Snapshot™ technology for fast, storage-efficient, and reliable backup and recovery, as well as fast and storage-efficient cloning for faster time to market while improving quality. The fully supported products help SAP customers to not only automate a comprehensive backup and disaster recovery strategy, but to also integrate other important workflows focusing on the complete SAP application lifecycle management using this Snapshot-based SAP system copies and cloning operations.

Many SAP customers who want to move their SAP systems to the cloud don't want to relinquish the many NetApp benefits for their SAP projects and operations. Customers do not want to give up on the performance, reliability, and enterprise data management capabilities when they move these enterprise file-based workloads to the cloud. Not every cloud can offer a highly available, enterprise-grade, fast, reliable, feature-rich, but simple-to-manage shared file service based around NFS, as it is required for all SAP environments.

On Azure, customers can now benefit from two distinct ONTAP based offerings on which to build their SAP systems. The following sections provide an overview of both solutions, NetApp Cloud Volumes ONTAP and Azure NetApp Files; however, the remainder of this document focuses on Azure NetApp Files only.

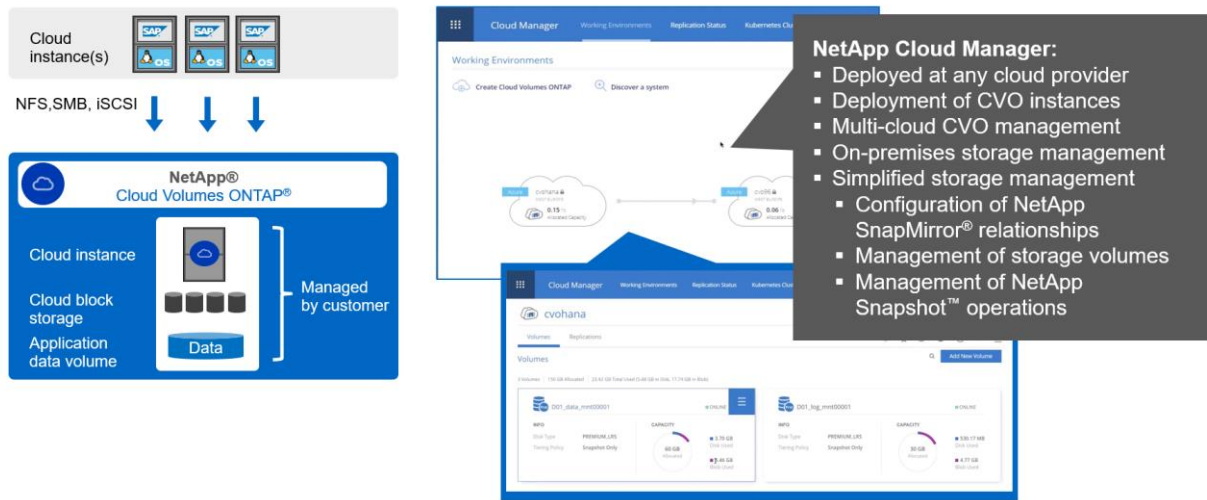
For more information about NetApp Cloud Volumes ONTAP, see [TR-4757 – SAP Applications on Microsoft Azure Using Cloud Volumes ONTAP](#).

Cloud Volumes ONTAP on Azure

[NetApp Cloud Volumes ONTAP](#) extends the trusted enterprise data management capabilities of ONTAP to leading cloud platforms such as Microsoft Azure. In Azure, Cloud Volumes ONTAP provides CIFS-, NFS-, and iSCSI-based services to host SAP workloads. By leveraging the underlying Azure storage and

compute resources, Cloud Volumes ONTAP adds storage-efficiency features such as thin provisioning, deduplication, compression, and now tiered storage to Azure Blob storage.

Figure 1) Cloud Volumes ONTAP.



Cloud Volumes ONTAP is a NetApp proven data management software running in a cloud instance using Cloud Block Storage. For the initial provisioning of Cloud Volumes ONTAP and a cloud-like, simplified, but efficient, management of all resources, the customer must deploy a small appliance NetApp OnCommand® Cloud Manager, as shown Figure 1) Cloud Volumes ONTAP. With OnCommand Cloud Manager customers can deploy and manage multiple Cloud Volume ONTAP instances either as a single node or a highly available dual-node configuration. In addition to managing Cloud Volume ONTAP in Azure, OnCommand Cloud Manager can manage on-premises ONTAP systems as well as Cloud Volumes ONTAP instances at other data centers or even other cloud providers. When provisioning Cloud Volumes ONTAP, customers can select from different system classes and license types. This configuration defines the maximum storage capacity and the possible throughput and performance.

Customers can provision data volumes and shared files to the cloud instance to run their SAP application and databases. New NetApp customers can use OnCommand Cloud Manager for this provisioning, while existing NetApp customers can use their existing NetApp tools and workflows to manage Cloud Volumes ONTAP as well.

Azure NetApp Files

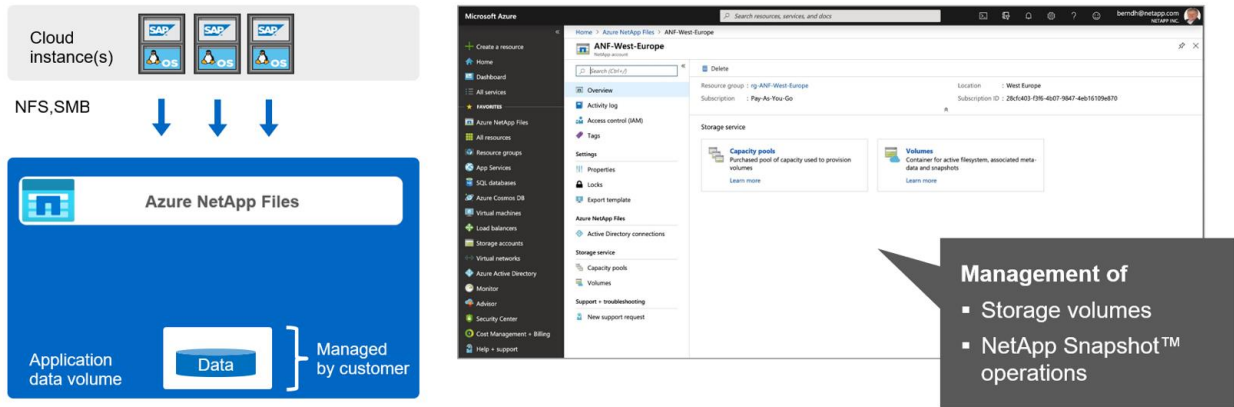
Since May 2019, Microsoft delivers an Azure native, first-party portal service for enterprise NFS and SMB file services based on NetApp ONTAP technology. This development is driven by a strategic partnership between Microsoft and NetApp and further extends the reach of ONTAP's world-class data services to Azure.

This Azure cloud-native data service delivers high performance with low-latency, reliability, and enterprise data management and security for customers who are moving enterprise NFS and SMB workloads to Azure.

Azure NetApp Files is completely integrated into the Azure data center and portal. Customers can use the same comfortable graphical interface and API to create and manage shared files as with any other Azure service, as shown in Figure 2. Azure NetApp Files provides NetApp enterprise class storage and delivers many of the data management capabilities such as the ability to easily create and resize volumes, adapt capacity and performance without downtime, and create space-efficient storage Snapshot copies and clones in seconds, which are very valuable to use to optimize SAP operations.

Compared to Cloud Volumes ONTAP, Azure NetApp Files is built on enterprise-grade ONTAP and All Flash storage systems in the Azure data center, and directly operated and maintained by Microsoft. This configuration results in high storage performance combined with low latency I/O.

Figure 2) Azure NetApp Files.



Comparison

Table 1 shows the different management levels for the various ONTAP-based systems, ranging from on-premises (FAS/AFF ONTAP) to infrastructure as a service (IaaS with Cloud Volumes ONTAP) to PaaS (Azure NetApp Files). Clearly, Azure NetApp Files provides the best on-demand cloud service experience while providing on-premises-like performance. For cloud-based storage with a higher level of management control as well as access to all ONTAP features, Cloud Volumes ONTAP might provide the best experience.

Table 1) NetApp ONTAP offerings side-by-side.

	On-Premises ONTAP AFF Full Control and Ownership	Cloud Volumes ONTAP Full Control and SW Ownership	Azure NetApp Files Service Consumption
Mount volumes	Customer	Customer	Customer
Provision volumes	Customer	Customer	Customer
Performance	High	Low, medium	High
Space, time-efficient Snapshot copies	Yes	Yes	Yes
Space, time-efficient cloning	Yes	Yes	Yes
Deduplication and compression	Yes	Yes	No
Protocols	NFS, SMB, iSCSI, and so on	NFS, SMB, and iSCSI	NFS and SMB
Endpoint replication/migration	SnapMirror, Cloud Sync	SnapMirror, Cloud Sync	Cloud Sync
Encryption	Yes (customer-managed)	Yes (customer-managed)	Yes (Azure-managed)

	On-Premises ONTAP AFF Full Control and Ownership	Cloud Volumes ONTAP Full Control and SW Ownership	Azure NetApp Files Service Consumption
Purchasing commitment	Typically, 3–5 years	Paygo (hourly) BYOL (yearly)	Paygo (hourly)
Disk capacity planning	Customer	Azure	Azure
ONTAP upgrades	Customer	Customer	Azure
ONTAP service deployment	Customer	Customer on Azure IaaS	Azure
Hardware deployment/refresh	Customer	Azure	Azure

1.2 Microsoft Certified Solutions for SAP

For the latest information related to SAP on Azure certifications, see [SAP Note 1928533](#) and the [SAP Certified and Supported SAP HANA Hardware Directory](#).

As of Q1CY19, the Azure VMs and SAP HANA on Azure listed in Table 2 are certified to run SAP workloads (Table 2).

Table 2) Azure SAP certified VMs

VM Series	VM Type	VM Size	SAPs	Supported SAP HANA Scenarios	SAP Certification
DS v2	Various DS1x_v2	2 vCPU; 14GB	3,530		Certified (Any DB, App)
	DS12_v2	4 vCPU; 28GB	6,680		
	DS13_v2	8 vCPU; 56GB	12,300		
	DS14_v2	16 vCPU; 112GB	24,180		
	DS15_v2	20 vCPU; 140GB	30,430		
GS	GS4	16 vCPU; 224GB	22,680		Certified (Any DB, App) GS5 certified for OLAP & controlled availability for S4
	GS5	32 vCPU; 448GB	41,670	OLAP and S4	
DS v3	D2s_v3	2 vCPU; 8GB	2,178		Certified (Any DB, App)
	D4s_v3	4 vCPU; 16GB	4,355		
	D8s_v3	8 vCPU; 32GB	8,710		
	D16s_v3	16 vCPU; 64GB	17,420		
	D32s_v3	32 vCPU; 128GB	34,840		
	D64s_v3	64 vCPU; 256GB	69,680		

VM Series	VM Type	VM Size	SAPs	Supported SAP HANA Scenarios	SAP Certification
ES v3	E2s_v3	2 vCPU; 16GB	2,178		Certified (Any DB, App)
	E4s_v3	4 vCPU; 32GB	4,355		
	E8s_v3	8 vCPU; 64GB	8,710		
	E16s_v3	16 vCPU; 128GB	17,420	OLTP/OLAP	Certified (Any DB, App)
	E32s_v3	32 vCPU; 256GB	34,840	OLTP/OLAP	
	E64s_v3	64 vCPU; 432GB	70,050	OLTP/OLAP	
M	M32ts	32 vCPU; 192GB	33,670	OLTP	SAP HANA, App, Any DB certified
	M32ls	32 vCPU; 256GB	33,300	OLTP	
	M64ls	64 vCPU; 512GB	66,600	OLTP	
	M64s	64 vCPU; 1024GB	67,315	OLTP/OLAP	
	M64ms	64 vCPU; 1792GB	68,930	OLTP	
	M128s	128 vCPU; 2048GB	134,630	OLTP/OLAP	
	M128ms	128 vCPU; 3800GB	134,630	OLTP	
Mv2	M208s_v2	208 vCPU; 2.85 TB	259,950	OLAP/OLTP	SAP HANA, App, Any DB certified
	M208ms_v2	208 vCPU; 5.7 TB	259,950	OLAP/OLTP	
	M416s_v2	416 vCPU; 5.7 TB	488,230	OLAP/OLTP	
	M416ms_v2	416 vCPI; 11.4 TB	488,230	OLAP/OLTP	

1.3 Enterprise Storage Requirements

In the cloud as well as on-premises, enterprises must be able to scale storage capacity and performance with guaranteed availability. They must also ensure compatibility with host clients' data formats and operating systems. This section looks at these challenges and requirements individually.

Extreme File Service Performance

File services are directly dependent on the high level of performance provided by the storage environment in use. Some workloads might require a high-level SLA in order to maintain peak operability and IOPS levels, with consistent low latencies.

Host Client Data Compatibility

An enterprise's file services must be compatible with all their host clients' data formats, operating systems, and access protocols.

Guaranteed Business Continuity

When it comes to running an enterprise file service, any disruption in normal operation can negatively impact the business. Whether an outage is caused by a disaster or through an update process, it is essential to allow the ongoing availability of the file share, with zero downtime and no data loss.

Data Protection

Providing storage for an enterprise require companies to comply with industry specific data security, data protection, and disaster recovery SLAs and regulations. ONTAP's Snapshot technology is key to meeting these requirements, because Snapshot copies can be used to rapidly restore systems in the event of an issue, be that logical/physical corruptions, user errors, malware or even ransomware attacks. And because application data is often the most important part of a business' operation, these copies must be able to be automated in order to meet the most stringent recovery point objective (RPO), recovery time objective (RTO), and backup requirements.

Data Cloning

Further, many organizations need multiple copies of data, mostly for test and development purposes. SAP landscapes are especially littered with various system copies for a variety of uses, so creating and refreshing those are typically cumbersome. This process requires additional capacity as well as an excessive amount of time to create. ONTAP's space and time-efficient cloning drastically improves this situation because creating clones takes zero capacity and near zero time, which improves the time to market and the quality while lowering the cost.

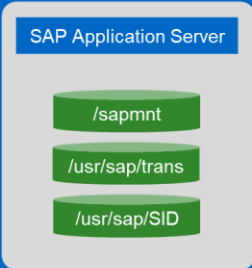
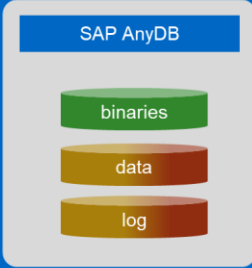
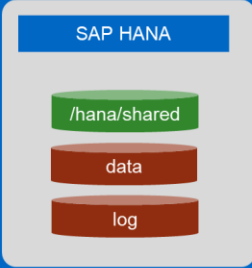
Data Security

With shared storage, data security becomes a major concern. It's important to make sure that access to the file system is in the user's control at all times. Ensuring data security at all levels is crucial to preventing losses before they happen. Role-based access control (RBAC), file/directory access controls and proper key management within organizations can help make data more secure.

SAP Storage Classification

To understand the specific storage requirements and different use cases for SAP, we need to take a closer look at the different types of SAP applications, as shown in Figure 3.

Figure 3) SAP storage performance requirements and certification.

	SAP Application Server	SAP AnyDB	SAP HANA
			
Certification	None	None	Required (for production)
Performance KPIs	None	Based on sizing	Defined by SAP (HCOT)

Performance requirements: ■ low ■ medium ■ medium/high ■ high

Figure 3 shows typical file systems that are used by an SAP application server, an SAP/AnyDB database server (such as Oracle, SQL, DB2, ASE, and MAX DB), and a SAP HANA server. The colors indicate the typical performance requirement for these type of file systems. Showing a low performance requirement for standard file systems that hold binaries, log and trace files, or configuration files using a green color. In contrast, SAP HANA data and log volumes require higher storage performance; therefore, they are displayed in red when compared against traditional databases.

When you are running an SAP system in production, an important aspect to keep in mind is the SAP HANA certification or a formal support declaration for AnyDB databases, such as Oracle. In contrast, file systems used for SAP application servers do not require a certification. SAP HANA certification includes the host and storage on which SAP HANA is running in production. Additionally, the certification includes performance tests with stringent key performance indicators (KPIs) that a system must pass in order to be listed in the [Certified and Supported SAP HANA Hardware Directory](#). As of November 2019, [Azure NetApp Files is listed in the SAP HANA Hardware Directory](#) and is therefore fully certified and supported for running SAP HANA production workloads in Azure.

1.4 SAP Use Cases

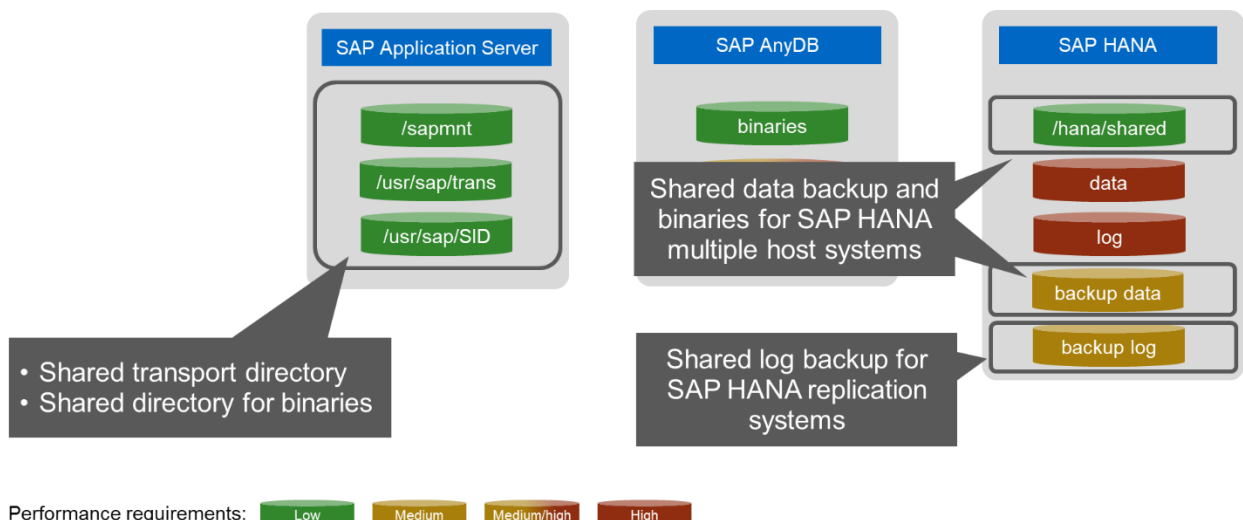
Based on the classification described in the previous section, there are five different use cases where Azure NetApp Files could help improve customers' SAP cloud experience:

- SAP shared files
- SAP databases (including SAP HANA data and log)
- SAP snapshot and file-based backups
- SAP Cross Regional Disaster Recovery
- SAP System Copies for Test/Dev and QA

As shown in Figure 4, shared files are needed for almost any SAP system landscape. Starting with the typical candidates such as `/usr/sap/trans`, SAP HANA shared volume, and SAP HANA log and data backup volumes in SAP HANA multiple hosts and system replication scenarios. The exact performance requirements can vary based on a customer's setup; however, you can find shared file systems in any customer environment, mostly using NFS as the protocol.

In order to benefit fully from the advanced data management capabilities for Azure NetApp Files, and to support SAP HANA multiple hosts scenarios with one or more standby nodes, the SAP HANA data and log volumes also need to be placed on Azure NetApp Files.

Figure 4) Shared files requirements.



2 Shared Files Requirements

2.1 Shared File System for SAP NetWeaver Application Server

Almost every SAP landscape requires a shared file system either provided through the NFS protocol on UNIX and the Linux operating systems, or the SMB protocol in the case of an SAP system running on Windows operating system.

This section discusses the shared file requirements for an example setup using an SAP HANA multiple-host system with SAP HANA system replication and an SAP application server using clustered SAP ABAP Central Services (ASCS) and an SAP Enqueue Replication Server (ERS). This is a very common setup to achieve high availability for the SAP HANA database using HANA system replication, as well as for the application server by implementing clustered ASCS and SAP ERS instances.

Figure 5) Shared files for SAP.

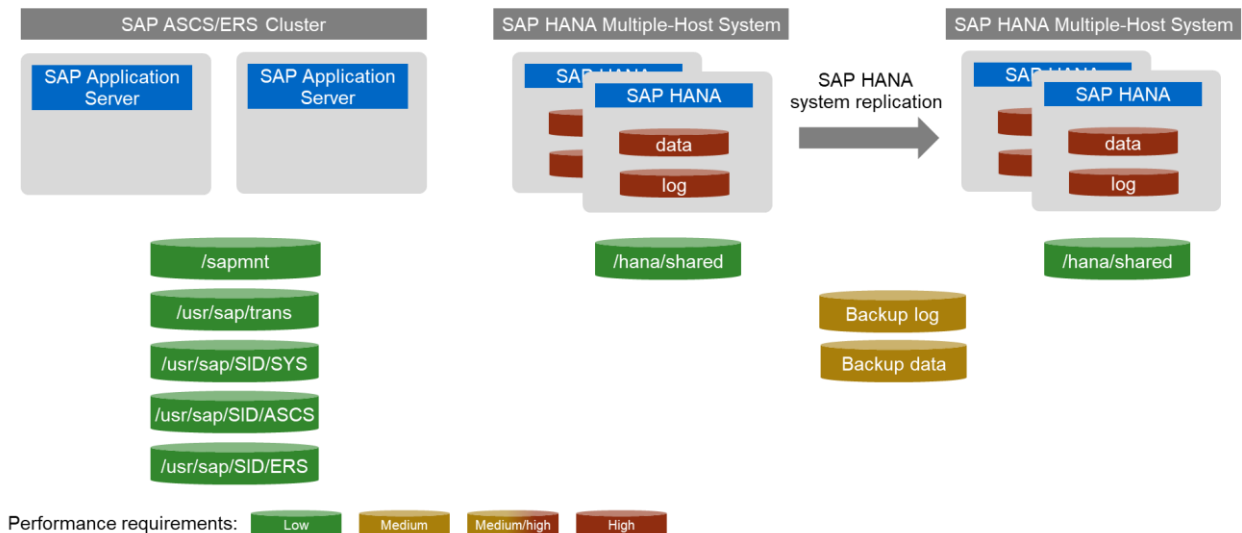


Figure 5 shows the following shared file systems that are required for the system landscape:

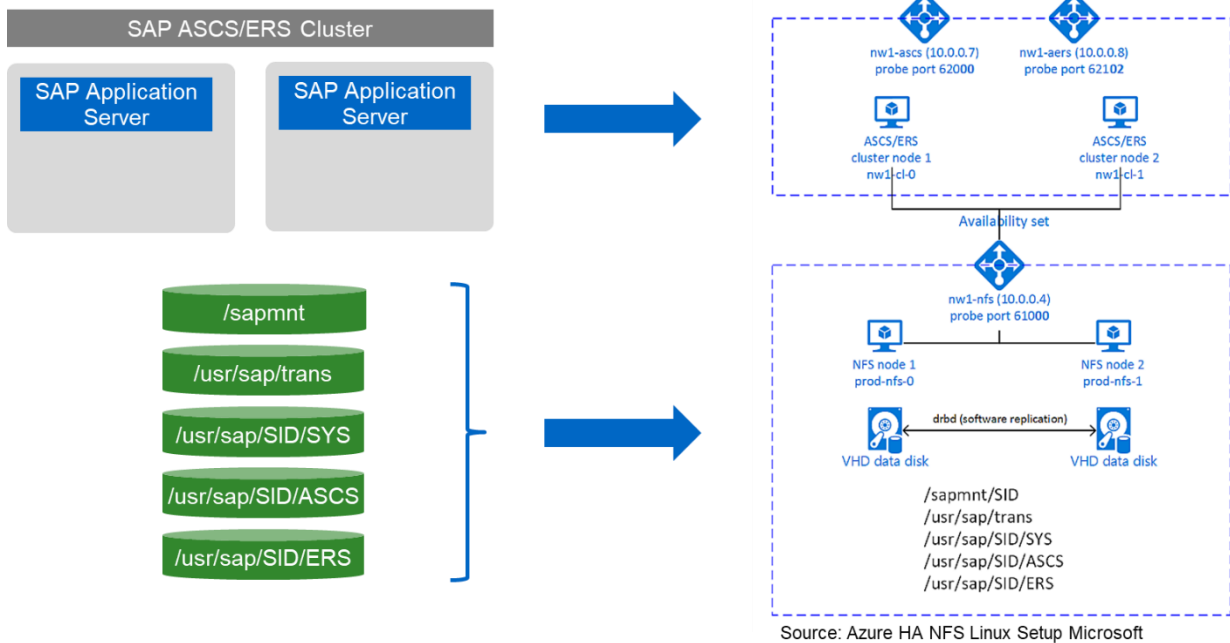
- **/sapmnt.** If you have more than one application instance, use the `/sapmnt` file system to store a common set of binaries and configuration files. The I/O pattern is reading the binaries and configuration files and writing a few logs. In Figure 5, green indicates a lower performance requirement.
- **/usr/sap/trans:** This is a common file system that is used to share (or transport) customer developments or other transports between systems in a single SAP landscape.
- **/usr/sap/<SID>/SYS, /usr/sap/<SID>/ASCS, and /usr/sap/<SID>/ERS.** These file systems are used for the SAP application server instances. The performance requirements are rather low; however, for a high-available setup with the ERS, it is mandatory that the underlying file system is a high-available system as well so that the ERS locking table is preserved in case of an instance failover.
- **/hana/shared.** For a multiple-host SAP HANA system, `/hana/shared` must be an NFS shared file system.
- **Backup data.** For file-based backups in a multiple-host environment, all SAP HANA servers should have access to the backups, which requires an NFS share. File shares used for file-based backups require a significantly higher throughput than the previously discussed file systems to allow the backup of the SAP HANA database to finish as quickly as possible. This requirement can be partly mitigated by the use of Snapshot copies.

- **Backup log.** The automatic SAP HANA log backup is written to this shared location. For an SAP HANA system replication setup, the location must be a shared location between both SAP HANA systems to allow for failover. Even in an SAP HANA single-host setup, this location should still be highly available. Also, for the log backup, a medium performance requirement is recommended.

2.2 Shared Files Solutions for the Cloud

A Linux cluster using a block replication device (as shown in Figure 6) is a commonly used solution within the cloud to provide a highly available NFS service.

Figure 6) Linux cluster NFS server.



At a first glance, this solution looks appealing because the initial setup of the Linux cluster can be completely automated by using Azure ARM templates. However, this type of solution can quickly present the following issues:

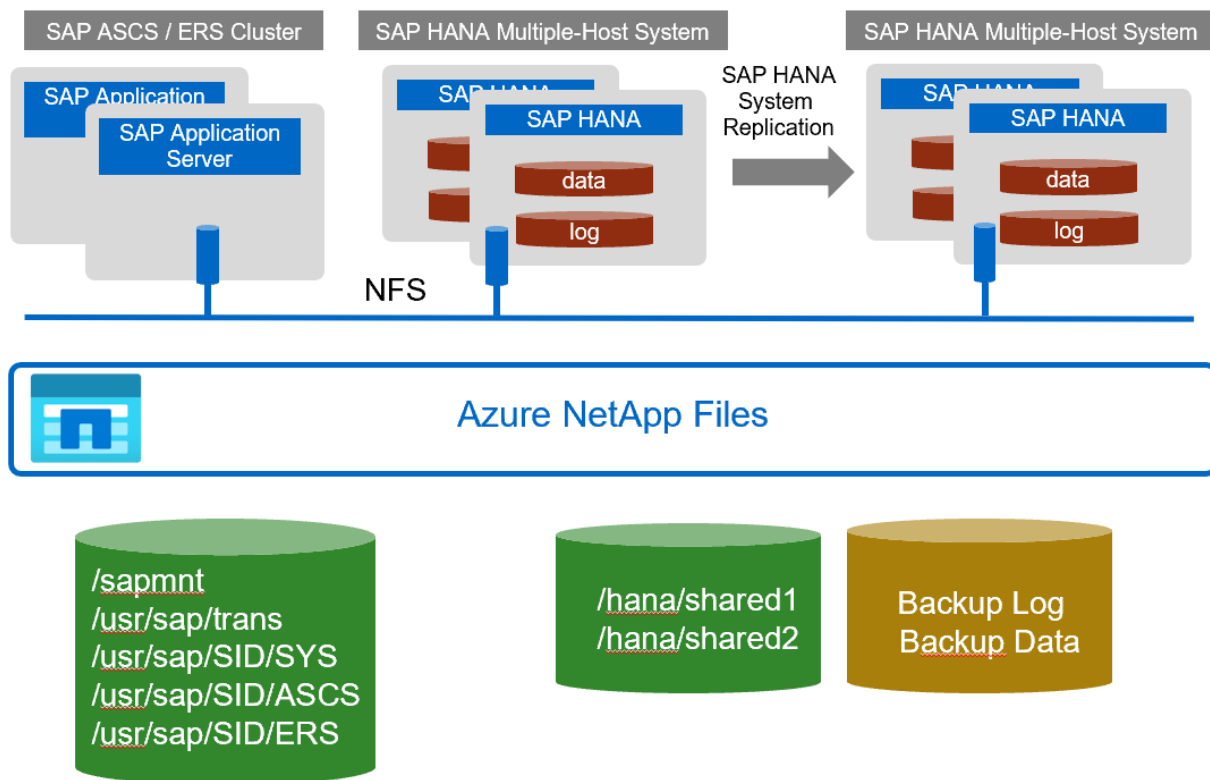
- The first issue is the level of manual administration that these deployments require. For example, allocating a new file system requires allocating new storage, mounting it to the compute hosts that will serve out the data, and potentially initializing the new share with existing data. If the file system needs to grow, this growth must be handled manually. If the performance of the underlying disks needs to be upgraded, the allocation of the new storage and migration of existing files need to be taken care of while still trying to minimize downtime.
- The second issue is the complexity of managing the storage over time as the deployment grows or shrinks. Storage administrators working with production file shares need to maintain uninterrupted access to the files, provide backup or snapshot facilities, allow test copies of the data to be created, and much more. Not all storage administrators have the skills to maintain and administer a Linux cluster. Providing robust support for this kind of functionality requires a high level of technical expertise.

2.3 Providing Shared Files using Azure NetApp Files

This section explains how to use Azure NetApp Files to address the shared files requirements of the example provided in Figure 7.

Note: Cloud Volumes ONTAP also offers a very good solution for shared files. For more information about using Cloud Volumes ONTAP, see the NetApp technical report, "SAP Applications on Microsoft Azure Using Cloud Volumes ONTAP".

Figure 7) Azure NetApp Files.



To provision the shared file systems that are listed in section 2.1, "Shared File System for SAP NetWeaver Application Server", the individual file systems are grouped based on both the protection and performance requirements of each system. It is obvious that managing many smaller volumes for each of the required shared file systems increases the management overhead. To simplify this process further, the following three volumes are created:

- One volume to hold all SAP application server shared files
- One volume for the `/hana/shared` file system of each of the SAP HANA databases
- One volume to store the log backups and the optional file-based backup files

Performance and Cost Optimization

To better understand why it makes sense to optimize usage in relation to performance and costs, we need to look more closely at how Azure NetApp Files volumes are provisioned.

Azure NetApp Files volumes are allocated from capacity pools that are provisioned in an Azure NetApp Files storage account. Each capacity pool is assigned:

- A service level ("tier") that defines the performance as a factor of the provisioned storage capacity
- The provisioned storage capacity for the capacity pool

The performance of a volume is based on the capacity pool service level, in combination with the capacity of that volume. Users can dynamically grow and shrink the volumes manage capacity, performance, or both. Billing occurs for the provisioned amount in the capacity pool (on an hourly basis).

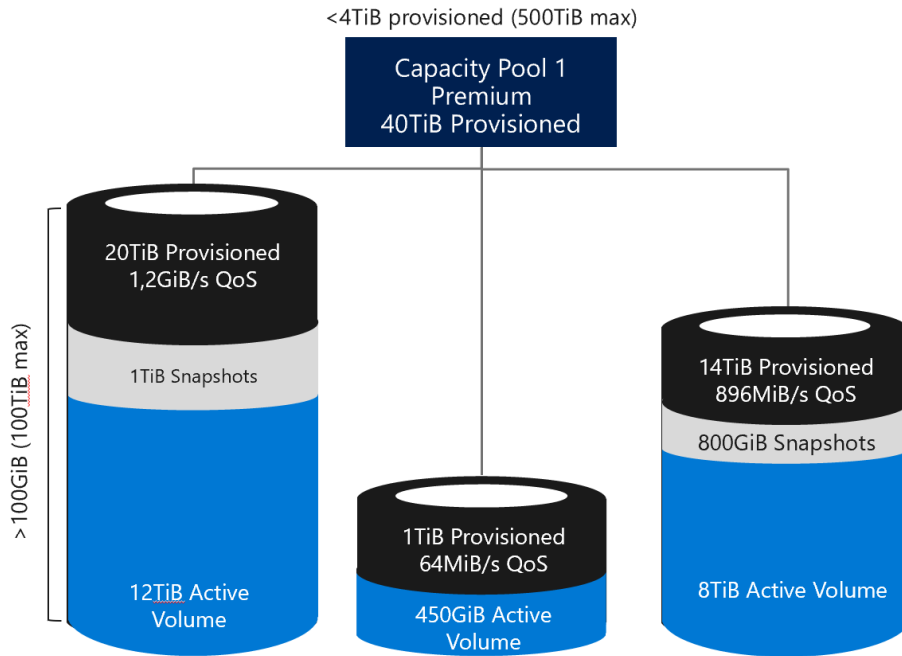
Figure 8) Performance Tiers.

Standard	Premium	Ultra
Throughput Up to 16MiB / s per 1TiB	Throughput Up to 64MiB / s per 1TiB	Throughput Up to 128MiB / s per 1TiB
IOPS Up to 1,000 (16K) per 1TiB	IOPS Up to 4,000 (16K) per 1TiB	IOPS Up to 8,000 (16K) per 1TiB
Suitable for <ul style="list-style-type: none"> • Static web content • File shares • Database backups 	Suitable for <ul style="list-style-type: none"> • Databases • Ent apps incl SAP • Analytics • Engineering apps • Messaging queues 	Suitable for <ul style="list-style-type: none"> • Extreme performance / throughput intensive apps • High Performance Computing
<i>Throughput comparable to mainstream HDD</i>	<i>Throughput comparable to mainstream SSD</i>	<i>Throughput comparable to Flash Arrays</i>

Each of the service levels listed in Figure 8 has an associated cost per provisioned capacity and includes a quality of service (QoS) level that defines the overall maximum throughput or IOPs per provisioned space. For example, a 4TB provisioned single capacity pool with premium service level (“tier”) provides an overall available throughput for all volumes in this capacity pool of 4 x 64MiB/s, which totals 256MiB/s.

Within a capacity pool each volume is provisioned with a specific size ranging from 100GB up to the size of the capacity pool, with a maximum of 100TB. The volume size (times the tier SLA) therefore defines the maximum throughput for the volume, which explained in detail in the [Azure NetApp Files documentation’s performance section](#). Figure 9 illustrates the capacity pool and volumes performance and consumption. For more information see [Azure NetApp Files documentation](#).

Figure 9) Capacity pool and volumes.



For each provisioned volume, the active file system and the space occupied by the Snapshot copies consumes the space within the volume (used space).

Note: Space is consumed from the provisioned capacity pool, with active file system data and snapshot data stored inside the volume. There is no separate storage volume required to store snapshots.

Note: Azure NetApp Files will not reject write I/O when the space consumption reaches the maximum size of the volume or capacity pool. By design the volume is thinly provisioned on the capacity pool, which will auto grow. However, the performance is always defined by the user configured size of the volume. If the performance needs to increase as well, the volume must be manually increased in size. For more information about this behavior, see the [Azure NetApp Files documentation](#).

As previously mentioned, the provisioned capacity of volumes and the provisioned capacity and service level of the capacity pool can be easily and dynamically changed with an immediate effect on the active file systems – both upwards and downwards. This allows the storage performance and capacity demands (up or down) to be instantly adapted without any downtime.

3 SAP HANA

3.1 Certified Azure Native Solutions

As part of Microsoft Azure's offerings for SAP HANA, Microsoft continuously certifies several IaaS platforms. For an official and current list of these platforms, see the [Certified and Supported SAP HANA Hardware Directory](#).

The certified Azure offerings for SAP HANA workloads on VMs range from 112GB (DS14v2) up to 12TB (M416ms_v2). For SAP HANA large instances on Azure (bare-metal servers), the host memory can be up to 20TB for SAP HANA single host systems (S960m) and up to 60TB for SAP HANA multiple host systems(15 x S384).

Note: With [TDIv5](#), a 24TB single host and 120TB multiple hosts are possible.

Note: SAP HANA IaaS platform certification includes compute, network and storage.

The certified VMs for SAP HANA currently offer three different storage options:

- Premium Disk with Write Accelerator for the HANA log volume
- Ultra-Disk
- Azure NetApp Files with NFS v4.1

Note: Since November 2019 Azure NetApp Files is fully supported and certified for SAP HANA database workloads. All KPIs and certifications have been successfully passed. In order to meet these KPIs, it is important to understand the performance sizing aspect discussed in this chapter.

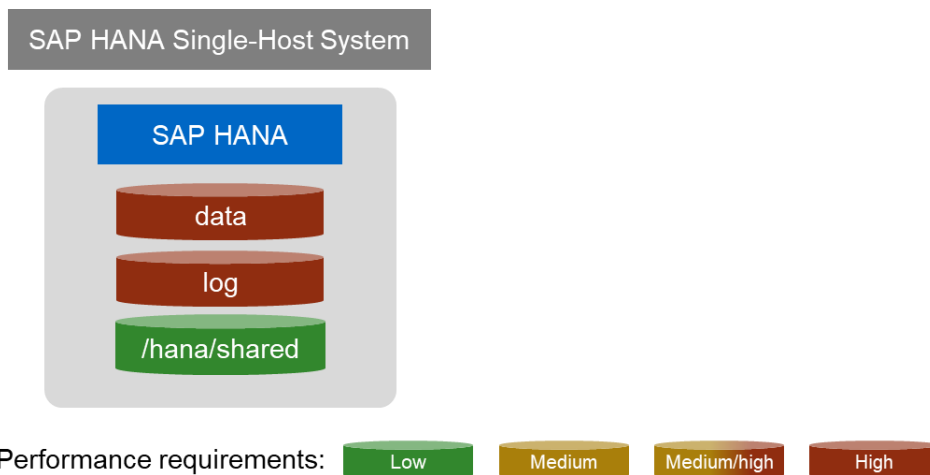
3.2 Storage Requirements for SAP HANA

This section describes the storage requirements for SAP HANA.

Single-Host SAP HANA

The most basic SAP HANA configuration requires three different storage volumes, as shown in Figure 10.

Figure 10) SAP HANA single-host system.



The following storage volumes are required for an SAP HANA single-host system:

- **/hana/shared.** This storage volume is used for shared files such as binary, logs, and configuration. For more information, see section 2.1, "Shared File System for SAP NetWeaver Application Server." There are no specific storage requirements for the shared files file system for a single-host SAP HANA system. It can be either direct attached storage or NFS mounted storage. For NFS both protocol versions NFSv3 as well as NFSv4.1 are supported.
- **Data volume.** The SAP HANA data volume must have at least the size of the memory size allocated for the SAP HANA database, such as the memory of the SAP HANA VM in most of the setups. For the data and the log volume, SAP defines performance criteria and specific certification KPIs. For a single-host system, the data volume can be direct attached storage or on NFS mounted storage. The SAP certification with Azure NetApp Files required the use of NFS protocol v4.1.
- **Log volume.** An SAP HANA log volume is required to store the most recent redo logs. SAP recommends that the size be 50% of the HANA memory, with a maximum size of 0.5TB. For a single-host system, the log volume can be direct attached storage or NFS mounted storage. The SAP certification with Azure NetApp Files required the use of NFS protocol v4.1.

Note: For more information about SAP HANA storage requirements, see the SAP whitepaper [SAP HANA TDI-Storage Requirements](#).

Multiple-Host SAP HANA

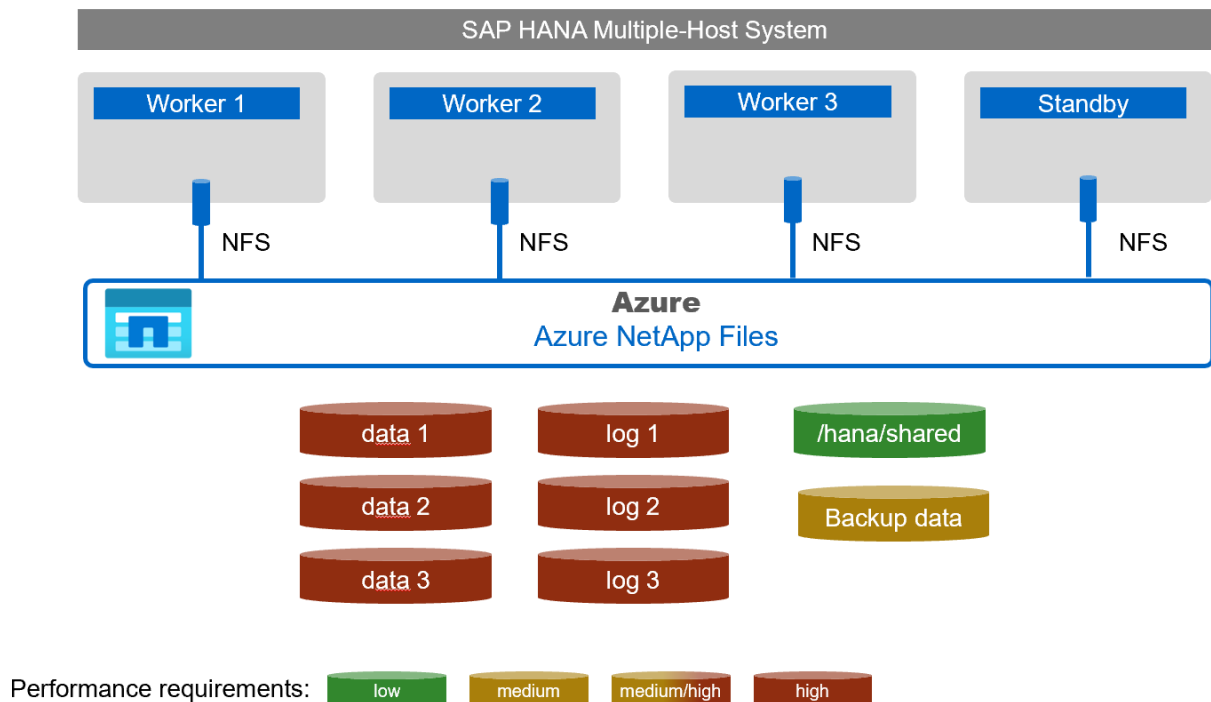
When the memory size of a single server is not sufficient, SAP HANA allows you to combine the memory of multiple servers to run SAP HANA in a multiple-host configuration. In this setup, SAP HANA allows you to configure a standby host that takes over the role of the failed server if a host failure of one of the worker hosts occurs, as shown in Figure 11. This SAP HANA cluster mechanism requires that the `/hana/shared` file system is available on all SAP HANA nodes, which in turn requires NFS.

For a multiple-host setup that supports the SAP HANA cluster with active standby nodes, the data and log files need to be on shared file systems. In public cloud environments, a shared NFS volume is the only viable option. With Azure NetApp Files using NFSv4.1 for the data and log volume, all data and log volumes need to be mounted on all SAP HANA hosts.

In addition, an SAP HANA multiple-host setup requires the following common backup volumes are shared between all SAP HANA hosts for a file-based backup:

- **Backup volume–data backup.** Depending on how many backup sets must be kept, SAP recommends that, for each set, the size of the SAP HANA data area should be allocated at the volume level.
- **Backup volume–log backup.** SAP HANA automatically backs up the log segments from the log volume to a shared file system, which all hosts must have access to. Many customers use the same volume for storing the file-based data backups as well as the automated log backups. The sizing depends on the change rate within the SAP HANA database.

Figure 11) SAP HANA multiple-host system.



3.3 SAP HANA Azure NetApp Files Sizing (for Capacity/Performance)

SAP HANA KPIs

To gain production support for SAP HANA, SAP requires that all the underlying infrastructure provides sufficient performance. For on-premises infrastructures, SAP created a Hardware Configuration Check Tool (HWCCT) and its successor the Hardware & Cloud Optimization Tools (HCOT) and a set of performance metrics for the data and log volume that customers can use to test whether their setup fulfills the required storage performance.

These performance values are guidelines for cloud environments, but tests are required to certify a cloud setup.

Note: SAP requires that customers use certified configurations to get production support. For test or development SAP HANA installations, no specific requirements need to be fulfilled.

Note: Specific KPIs are only required for the data and log volume and are independent from the SAP HANA database size or the application usage on top of HANA.

Capacity & Performance

There are two requirements that need to be considered in order to properly size Azure NetApp Files storage for the data and log volumes that meet SAP HANA production requirements.

- **Performance**

For a production SAP HANA database SAP specifies the different throughput values for different block sizes. The following numbers show the highest values from these KPIs

- Data volume: 400MBps read / 250MBps write
- Log volume: 250MBps read, write @ 1000 µsec maximum (<1ms) latency

- **Capacity**

SAP describes the capacity requirements for the data and log volume in the [SAP HANA Storage Requirements White Paper](#) with the following rules for:

- Data volume: The minimum size for the data volume = MEM of the SAP HANA database
- Log volume: The minimum size for the log volume = minimum (½ memory | 0.5TB)

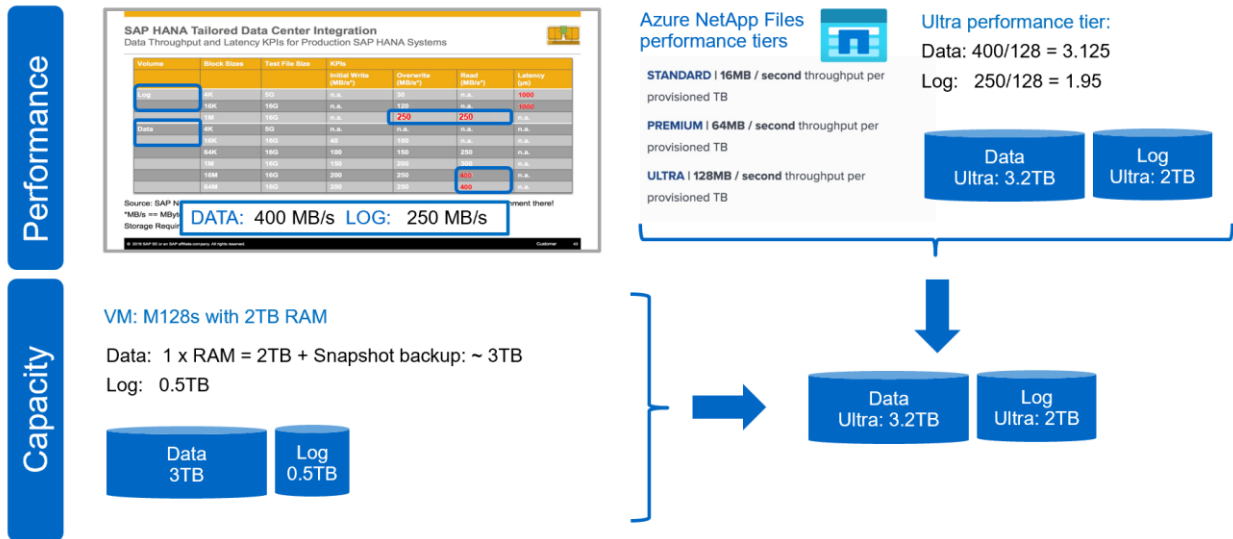
For the data volume there is another factor that influences the required capacity, which is the potential use of storage snapshot-based data protection. Whenever SAP HANA changes data on the storage layer – when snapshots are being used – Azure NetApp Files will store the changed data in new, additional blocks of the data volume – while preserving all blocks that belong to a snapshot from being overwritten. This may require (additional) capacity on the data volume depending on the change rate of the individual HANA database and the number of snapshots that need to be kept (retention).

As described in the previous section, there are two knobs to tune Azure NetApp Files to meet both requirements:

- Capacity pool service level. The service level defines the throughput per provisioned terabyte that a capacity pool can achieve.
- Volume size. Defines the maximum throughput a volume gets assigned, as well as the amount of data the volume can store.

As shown in Figure 12, the right service level and volume sizes need to be selected to meet the performance as well as the capacity requirement. In the following example this results in 'overprovisioning' of the log volume (from a capacity standpoint) to meet the SAP performance KPIs, while the capacity and performance requirement for the data volume lead to an almost identical volume size.

Figure 12) Tailored datacenter integration (TDI) performance criteria.



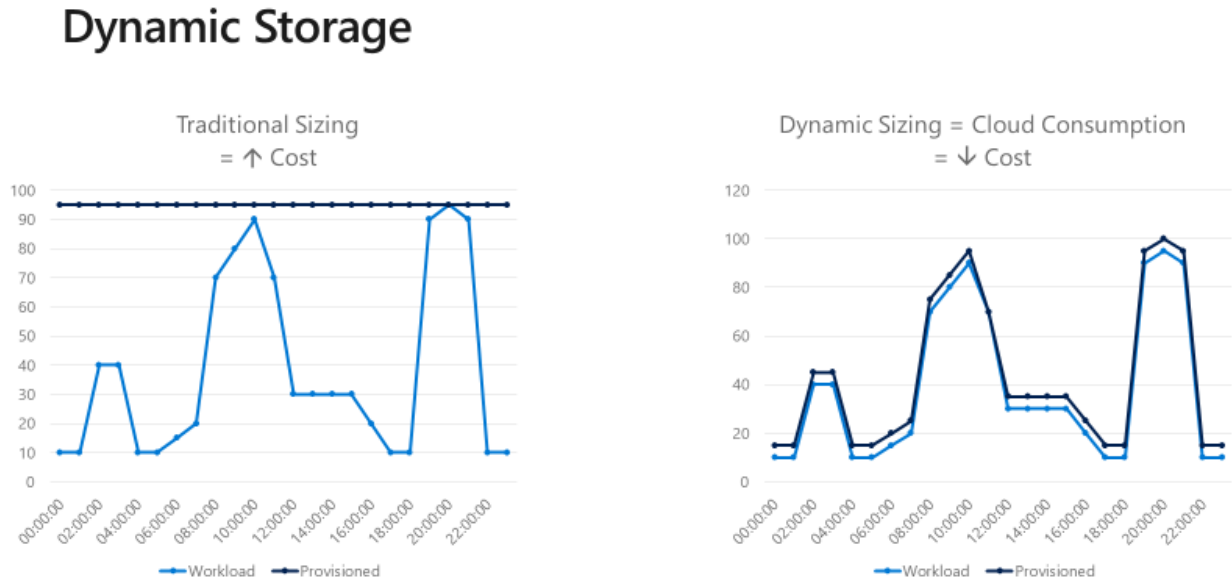
Note: Depending on SAP HANA database/instance sizes and the associated performance requirements, some (capacity) 'overprovisioning' may occur. However, since often time snapshots will be used for advanced data protection and recovery, this capacity will be used to retain snapshot data – basically resulting in (first-line-of-defense) backup/recovery capacity 'free of charge'.

Online Tuning Options

One of the most important features of Azure NetApp Files is that the volume and capacity pool sizes – and therefore performance – can be changed at runtime without any reconfiguration at the OS, NFS filesystem or database level and that these changes will have immediate influence on the SAP HANA database performance.

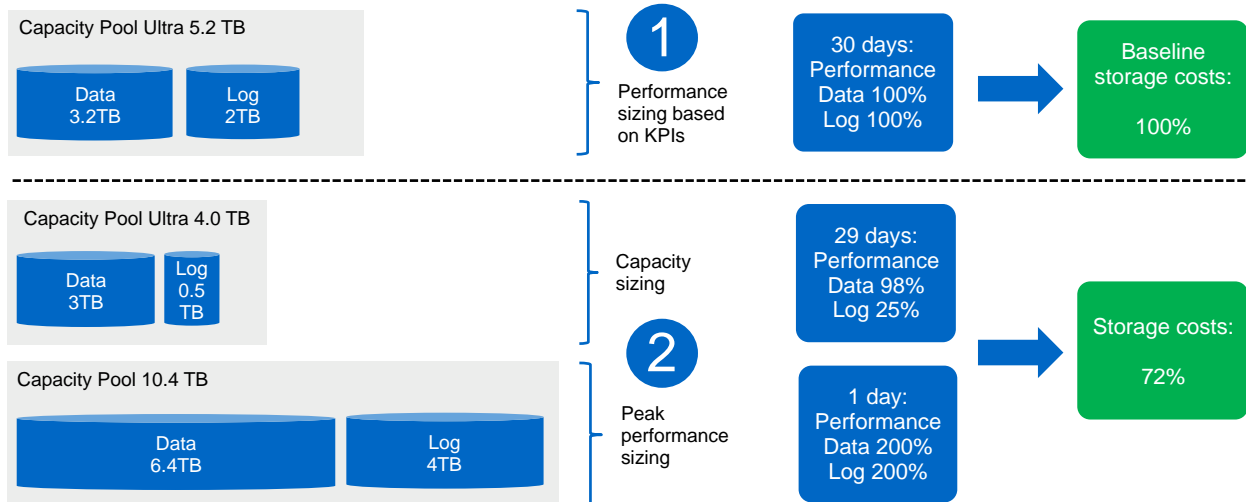
Often times SAP systems are sized based on **peak performance requirements** such as month-end closing, data load or other peak workloads – wasting valuable resource for the remaining periods. While there is almost no other choice in on-premises setups, and even with most standard cloud deployments – while being flexible in assigning compute resources – the typical (fixed-block) storage offerings usually do not offer the same flexibility (Figure 13).

Figure 13) Static versus dynamic storage (performance) provisioning.



This is where Azure NetApp Files makes a difference. Volume/pool sizes can easily be scaled down to the used capacity level in periods when the throughput is not needed, while dynamically scaling up to be prepared for the database for peak performance requirements.

Figure 14) Storage tuning & cost optimization.



In the example of Figure 14 an 28% cost reduction can be seen compared to the initial KPI based sizing, even when adjusting the storage for a 1 day peak performance **that provides twice the throughput** while for the rest of the month the size is based on capacity requirements only.

4 SAP HANA Data Protection, System Copies and Migration

In addition to performance, ease of management, and flexibility, SAP customers also require enterprise-grade data protection, not only for their shared file systems but also for their databases, as well as a way to migrate from their on-premises ONTAP environment to Azure. By using Azure NetApp Files, customers can implement comprehensive data protection for their SAP shared file systems by using Azure NetApp Files' ONTAP storage-based Snapshot copies.

In addition, Cloud Volumes ONTAP can be leveraged as a landing zone for test and development of SAP HANA landscapes in Azure migrated from on-prem, before moving the environments to Azure NetApp Files for full production purposes.

This chapter covers a variety of data protection, system copy, and migration use cases.

4.1 Snapshot Data Protection in Azure with Azure NetApp Files

The Challenges of SAP HANA Backups into the Cloud

The typical challenge with medium-to-large-sized database environments such as SAP HANA is the rapid creation and restoration of backup data. Most HANA implementations are immensely mission critical and don't allow for low SLAs with very high RPO and RTO characteristics. Typical legacy approaches to backup using data dump to disk or tape lead to unacceptable RPO and RTO SLAs that are measured in hours, and sometimes even days. As enterprise customers move mission-critical SAP HANA estates into Azure, more enterprise-grade solutions are required. This is where application-consistent Snapshot copies come into play.

Solving the Challenges with SAP HANA Backups Using Snapshot Copies

The key benefits of the Azure NetApp Files storage-based Snapshot copies include:

- Snapshot copies can be created and kept on the service (for example, inside the original volume) with no performance impact on the storage service.
- Because these copies are created on the storage system, they don't consume database resources and no data needs to be copied out.
- Recovery from a data Snapshot copy is much faster than recovery from a data backup, resulting in aggressively short RTOs (down to minutes).
- Snapshot copies and restores are near-instantaneous, allowing frequent creation, which means even shorter RTO because only a few log backups need to be applied after restoring a recently created Snapshot copy.

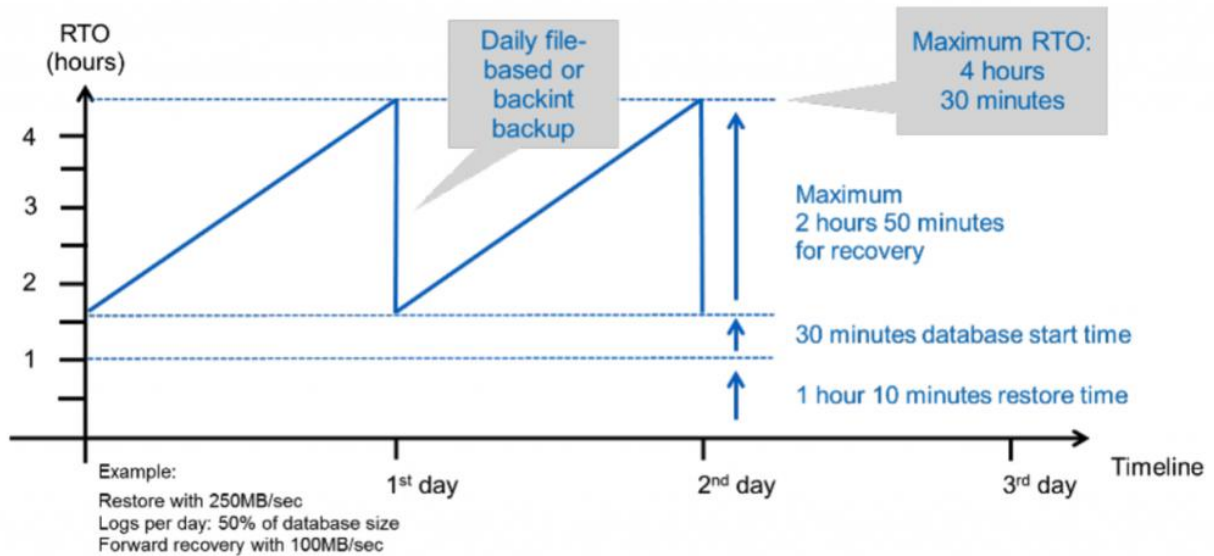
Note: Snapshot copies serve as a first-line-of-defense and can account for the vast majority of the required restore operations of any given organization. However, Snapshot copies are not to be considered a complete backup. To cover all BU/R requirements typically external snapshot replicas and/or other backup copies must be created in a remote location. This is further described in paragraph 4.2 "Protecting Snapshot Copies Against Local and Regional Disasters."

The following sections illustrate these benefits.

Traditional Backup

In traditional backups, the complete database payload is streamed to the backup media. Given the large (and with larger and larger SAP HANA systems increasing) amount of data (and typically limited bandwidth), creation of the backups typically takes hours. As a result, often times it is possible to only complete one backup (or two backups at best) per day. In addition, restores usually take longer than the backup itself, so a lot of compute and network resources are being consumed, as shown in Figure 15.

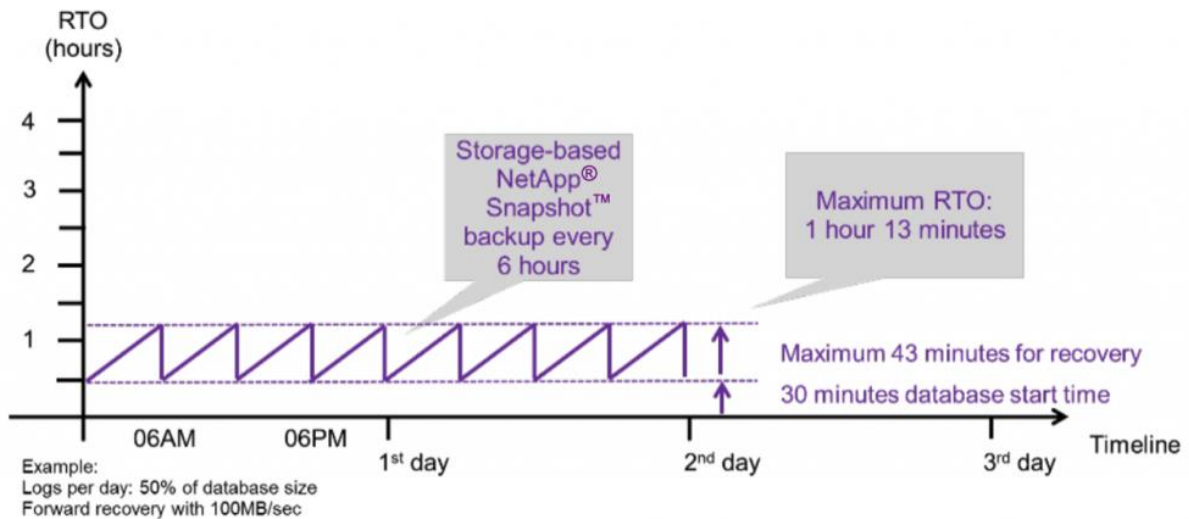
Figure 15) Traditional backup.



Snapshot Copy-Based Backup

With Snapshot copy-based backups, more backups can be created during the day. Data does not need to be copied on the storage service, therefore, creating a Snapshot copy takes only seconds. Regardless of data size, recovery from an ONTAP Snapshot copy is almost instantaneous and can be done from a fairly recent copy. Therefore, recovery time is shorter. Also, because of the more frequent creation of the Snapshot backups, fewer redo logs need to be applied after the Snapshot copy is restored.

Figure 16) Snapshot copy-based backup.



Regardless of the volume size, number of Snapshot copies kept, and the frequency of creation, Azure NetApp Files Snapshot copies based on ONTAP are fast (in seconds), space-efficient, and without performance impact. Instead of copying data, ONTAP marks the blocks on the active file system (volume) to be part of the new Snapshot copy and ensures that whenever a block is changed, the block is written to another empty location, preserving the snapped data block and avoiding any additional I/O. In other

words, Snapshot copies are pointers to data blocks that allow restore operations to be fast as well, since only pointers are changed, and no data will be copied.

A Snapshot copy is a point-in-time file system image. Low-overhead Snapshot copies are made possible by the unique features of the WAFL® (Write Anywhere File Layout) storage virtualization technology that is part of Azure NetApp Files' Data ONTAP®. Like a database, WAFL uses pointers to the actual data blocks on disk, but, unlike a database, WAFL does not rewrite existing blocks; it writes updated data to a new block and changes the pointer. An Azure NetApp Files Snapshot copy simply manipulates block pointers, creating a “frozen” read-only view of a WAFL volume that lets applications access older versions of files, directory hierarchies without special programming.

Because actual data blocks aren't copied, Snapshot copies are extremely efficient both in the time needed to create them and in storage space. An Azure NetApp Files Snapshot copy takes only a few seconds to create, regardless of the size of the volume or the level of activity on the Azure NetApp Files storage volume.

Meanwhile, the Snapshot copy of the data remains completely stable. An Azure NetApp Files Snapshot copy incurs no performance overhead; users can comfortably store up to 255 Snapshot copies per volume, all of which are accessible as read-only and online versions of the data (Figure 16).

Data in Azure NetApp Files Snapshot copies can be restored in three different ways:

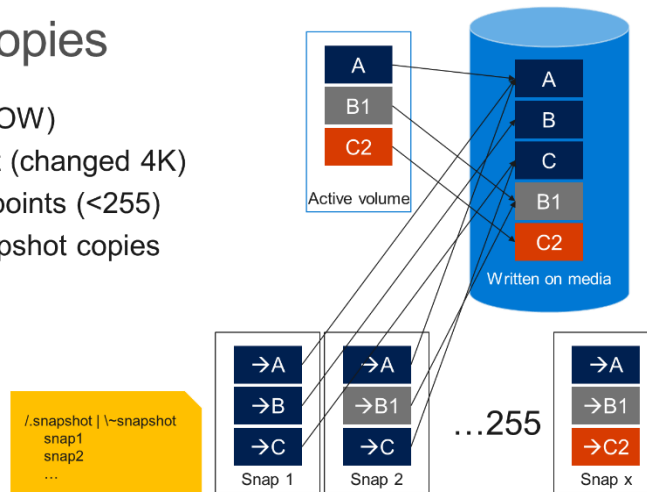
- By copying files and directories from the read-only Snapshot copy folders in the `/.snapshot` directory of a volume
- By restoring a volume Snapshot copy to a new volume (thick clone)
- By reverting a volume from a Snapshot copy of a volume (Snapshot copy restore)¹

Snapshot technology forms the basis of a unique ecosystem of high-availability, disaster-tolerant, and data protection solutions, as shown in Figure 17 and Figure 18.

Figure 17) Snapshot copy creation.

Snapshot copies

- Performant (ROW)
- Space efficient (changed 4K)
- Many restore points (<255)
- Readable Snapshot copies

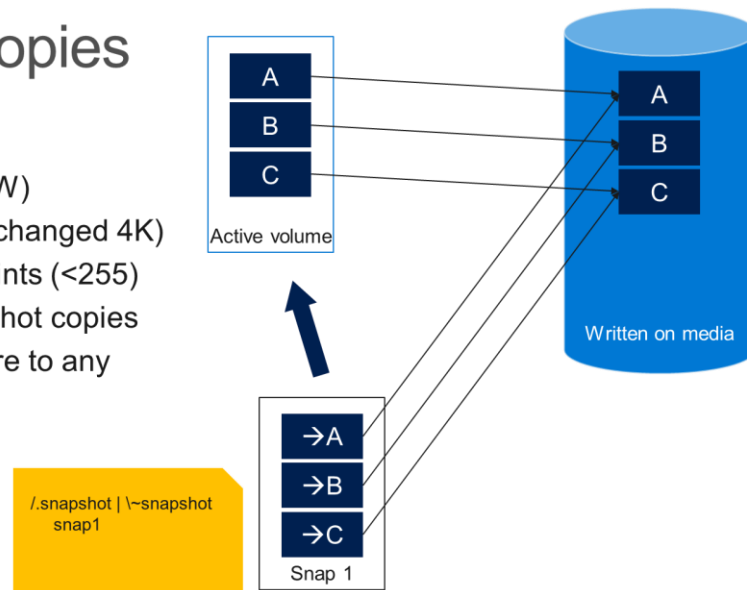


¹ At the time of writing, the in-place Snapshot restore (revert) feature was in Public Preview.

Figure 18) Snapshot restore operation.

Snapshot copies

- Performant (ROW)
- Space efficient (changed 4K)
- Many restore points (<255)
- Readable Snapshot copies
- Immediate restore to any Snapshot copy



Note: For more information about NetApp Snapshot copies, see: [NetApp Snapshot Technology Datasheet](#).

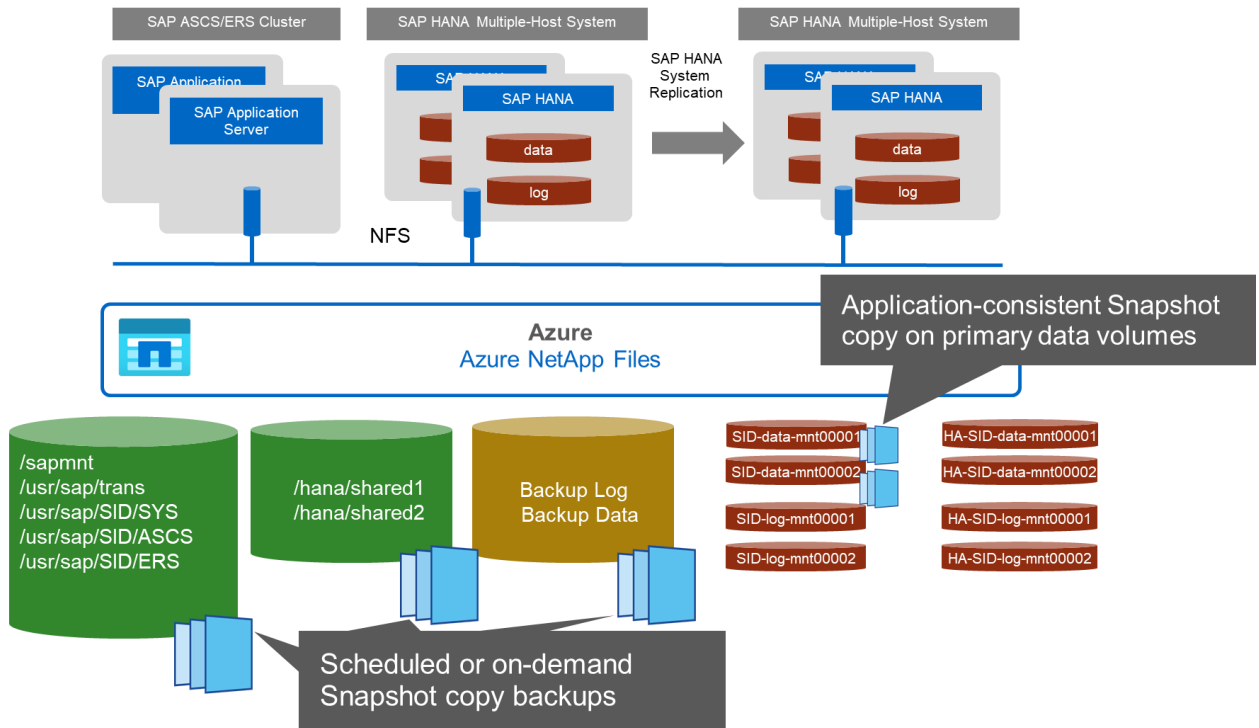
SAP HANA Snapshot Backups

Azure NetApp Files fully supports the fast and reliable ONTAP Snapshot capabilities, using either the standard Azure portal or a programmatic Azure API approach. Either way, an application-consistent Snapshot backup can be orchestrated and completed in seconds, regardless of the size of the volumes, with no impact on the storage performance, similar to an on-premises situation. In the past, many SAP HANA customers were looking for fast and reliable snapshot-based data protection in the cloud. With Azure NetApp Files, this protection is now available.

The Azure NetApp Files service allows customers to easily create storage Snapshot copies by using the Azure console or integrate Snapshot functionality into their own scripts by using REST API or PowerShell. This allows for full end-to-end orchestration of application-consistent Snapshot copies in SAP HANA environments. This is described in detail in the following paragraphs.

These Snapshot copies are stored in the Azure NetApp Files volumes and can be used to quickly restore volumes or create (thick) clone copies of the volumes, as shown in Figure 19.

Figure 19) SAP HANA local Snapshot copies.



4.2 Protecting Snapshot Copies Against Local and Regional Disasters

Many customers want to offload backups from the primary storage onto a different storage system or into a different region. One feature that makes the access of Snapshot copy-based backups very easy, is that Azure NetApp Files allow you access to all files of each of the Snapshot copies in a hidden directory called `.snapshot`. This is explained in more detail in section 4.4, “SAP HANA Restore in Seconds Using Snapshot Copies.”

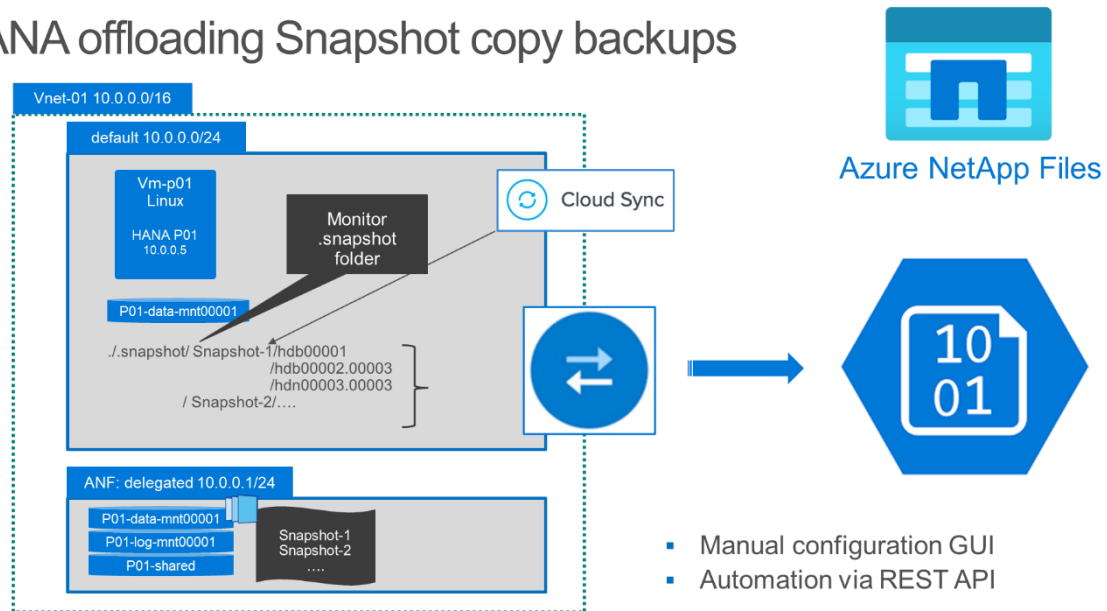
In order to offload these backups, customers can use various tools such as AzCopy to copy the data onto a different storage account, for example by using Azure blob storage.

NetApp Cloud Sync

One option that can automate this copy is NetApp [Cloud Sync](#), a fast cloud-based secure data synchronization service. Cloud Sync can be configured to automatically monitor the `.snapshot` folder so that whenever a Snapshot copy has been created, Cloud Sync starts offloading all files that are in the backup from Azure NetApp Files to Azure Blob storage, as shown in Figure 20.

Figure 20) Cloud Sync – offload backups to Azure Blob.

HANA offloading Snapshot copy backups



With Cloud Sync, not only can customers transfer the backups to Azure Blob for offloading or archiving, but they can build true hybrid and multicloud synchronization scenarios, supporting various storage architectures and cloud providers as synchronization target. This method incurs some additional load on the storage service, which might require some performance optimizations. The volume and pool resizing feature of Azure NetApp Files, as described in section 3.3, can be used to increase storage performance for the duration of the backup creation and restore processes.

Third-Party Backup

In addition to using Azure tools or services such as AzCopy or Cloud Sync, oftentimes, customers want to continue using third-party backup software, such as Commvault or Veeam, to create backups or offload Snapshot copies to Azure Blob storage. For more information, see the third-party vendors' solutions. These methods incur a significant load on the storage service, which can require some performance optimizations. The volume and pool resizing feature of Azure NetApp Files, as described in section 3.3, can be used to increase storage performance for the duration of the backup creation and restore processes.

SAP File-Based Backups

For backup and restore and database integrity check purposes, SAP recommends creating and using file-based backups on a regular basis (for example, weekly) for an integrity check. In order to achieve good backup streaming performance, SAP recommends that you provision separate volumes for data and log file backups on a standard capacity pool and with pool/volume sizes that accommodate the required backup storage capacity. The volume and pool resizing feature of Azure NetApp Files, as described in section 3.3, can then be used to increase storage performance for the duration of the backup creation, integrity check, and restore processes. For more information, see [SAP Administration on how to use file-based backups](#).

4.3 SAP HANA Backup in Seconds Using Snapshot Copies

SAP HANA 2.0 SPS1 and later supports the creation of Snapshot copy-based backups for systems with a single tenant. Starting with SAP HANA 2.0 SPS4, also for databases with more than one tenant. Also,

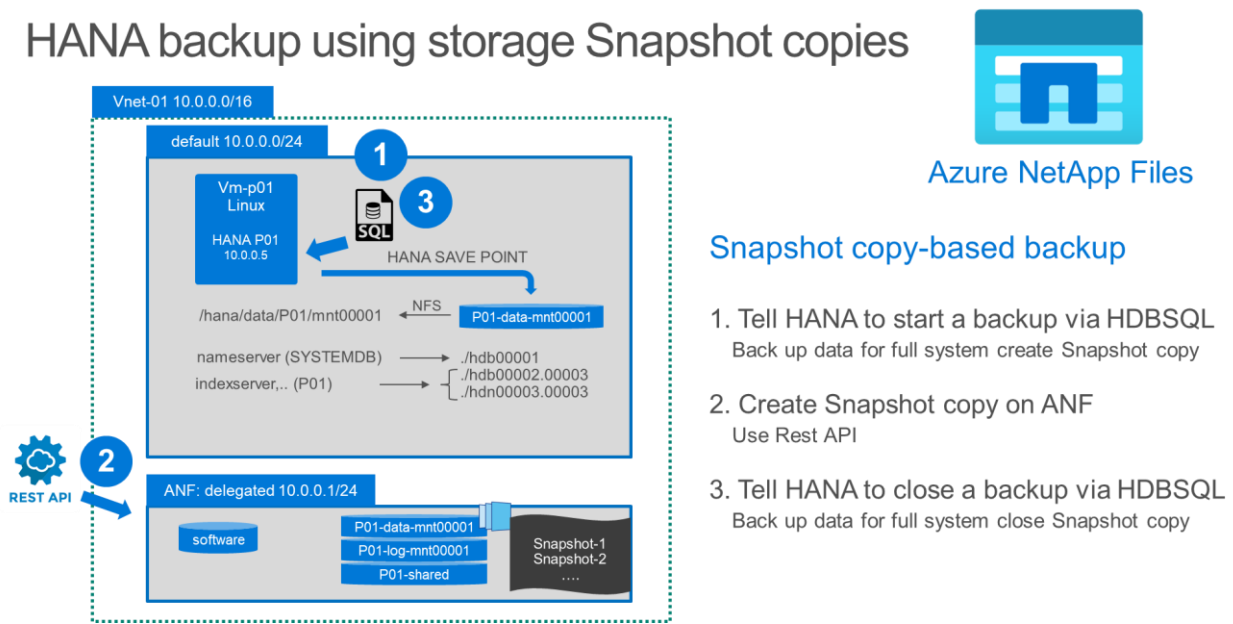
customers who are still using HANA 1.0 SPS12, the single-container version, can use Snapshot-based backups, as shown in Figure 21.

Creating storage-based Snapshot backups is a simple three-step procedure, as shown in Figure 21, using SAP’s SQL API. Using SAP HANA Studio and Azure Portal, the same result can be achieved by using a manual GUI-driven process:

1. To ensure the data on the persistence layer (the SAP HANA data volume) is consistent, SAP provides a SQL command to create an internal save point that is written to the data volume.
2. At that point, the storage Snapshot copy of the data volumes must be created. With Azure NetApp Files, this is done through the API or GUI, and no additional file system freeze is required, as it would be for other storage systems.
3. After a few seconds, when the Snapshot copy creation is complete, the SAP HANA database backup must be closed. With this command, a backup entry is added to the SAP HANA backup catalog.

Figure 21) HANA backup using storage Snapshot copies.

HANA backup using storage Snapshot copies



Snapshot copy-based backup

1. Tell HANA to start a backup via HDBSQL
Back up data for full system create Snapshot copy
2. Create Snapshot copy on ANF
Use Rest API
3. Tell HANA to close a backup via HDBSQL
Back up data for full system close Snapshot copy

Manual Workflow Example: Using Storage-Based Snapshot Backups

The example in this section illustrates how to implement the three-step procedure to back up an SAP HANA 2.0 Multi Database Container (MDC) system. Starting with SAP HANA 2.0 SPS1, SAP supports a single SQL command to back up the whole system, SYSTEMDB, and all the tenant databases with a single storage Snapshot copy.

Use the SQL commands provided by SAP to create a backup to set a comment as well as an external backup ID to identify and map the backup to the external storage Snapshot copy. It is a NetApp best practice to use the Snapshot copy name as the external backup ID and comment. That way, you can always identify the Snapshot copy that belongs to the backup from within the SAP HANA backup catalog.

Preparation

There are many different ways to execute the SQL command against the SAP HANA system database. In this example, we use the SQL client on the database host as user <SID>adm. This example demonstrates how to include this CLI command into your own workflow. To avoid the use of passwords

on the command line, NetApp recommends storing the login information in an encrypted keystore by running the SAP `hdbuserstore` command.

- In a SAP HANA MDC single tenant setup, port `3<instanceNo>13` is the standard port for SQL access to the system database and must be used in the `hdbuserstore` configuration.
- For a SAP HANA single container setup, port `3<instanceNo>15` is the standard port for SQL access to the index server and must be used in the `hdbuserstore` configuration.

In this example, we use the SAP System with SID P01 and instance number 00 installed on host `vm-p01` (`ip.10.0.0.7`). The system database index server is listening on port `30013` (for system number 00).

1. To store the connection information in the keystore, run the following command as `<SID>adm`:

```
p01adm@vm-p01:/usr/sap/P01/HDB00> hdbuserstore set BACKUP vm-p01:30013 System password
p01adm@vm-p01:/usr/sap/P01/HDB00> hdbuserstore LIST
DATA FILE      : /usr/sap/P01/home/.hdb/vm-p01/SSFS_HDB.DAT
KEY FILE       : /usr/sap/P01/home/.hdb/vm-p01/SSFS_HDB.KEY

KEY BACKUP
  ENV : vm-p01:30013
  USER: System
...
```

2. Open a backup:

- a. Instruct SAP HANA to start a backup (open backup) to ensure that the data volume is in an application-consistent state. To issue the SQL command, we need the SAP CLI tool `HDBSQL` as `<SID>adm` by using the previously created user store key `BACKUP`.

```
p01adm@vm-p01:/mnt/software/script> hdbsql -U BACKUP

Welcome to the SAP HANA Database interactive terminal.

Type:  \h for help with commands
       \q to quit

hdbsql SYSTEMDB=> BACKUP DATA FOR FULL SYSTEM CREATE SNAPSHOT COMMENT 'TEST003';
0 rows affected (overall time 132.236 msec; server time 129.906 msec)
```

- b. To find the internal backup ID that SAP created, run the following command:

```
hdbsql SYSTEMDB=> SELECT BACKUP_ID FROM M_BACKUP_CATALOG WHERE ENTRY_TYPE_NAME = 'data snapshot'
AND STATE_NAME = 'prepared';
1 row selected (overall time 9665.268 msec; server time 255 usec)
```

- c. A screen holding the internal backup ID is displayed.

```
BACKUP_ID
1576767153528
lines 1-2/2 (END)
```

Note: Creating a Snapshot-based backup for an MDC system with multiple tenants requires that all the tenants are online.

3. Create a Snapshot copy:

- a. You can use the REST API, a language SDK such as for Python, or the Azure portal to create a Snapshot copy. In this example, the Snapshot is created by using the Azure Portal. We need to select the data volume `P01-data-mnt00001` and select the Snapshot copies. On the Snapshot page, click `Add Snapshot` to create a new Snapshot copy.

P01-data-mnt00001 (ANF-West-Europe/ANF-WE-01/P01-data-mnt00001) - Snapsh

Volume

Search (Cmd+/) << + Add snapshot Refresh

Name *
TEST003 ✓

Search snapshots

Name	Location
initial-offline	West Europe
2019-04-15T09	West Europe
2019-04-15T100531.540469	West Europe
TESTEBID	West Europe
TEST001	West Europe
TESTID	West Europe
vm-sc-server_hana_P01_P01_vm...	West Europe
vm-sc-server_hana_P01_P01_vm...	West Europe
vm-sc-server_hana_P01_P01_vm...	West Europe
vm-sc-server_hana_P01_P01_vm...	West Europe
vm-sc-server_hana_P01_P01_vm...	West Europe
vm-sc-server_hana_P01_P01_vm...	West Europe
vm-sc-server_hana_P01_P01_vm...	West Europe
vm-sc-server_hana_P01_P01_vm...	West Europe
Offline	West Europe
vm-sc-server_hana_P01_P01_vm...	West Europe
2019-08-30T070801.100767	West Europe
TECHED001	West Europe
btobias01	West Europe
2019-11-05T141137.437862	West Europe
TEST002	West Europe

OK

b. After a few seconds, the Snapshot copy is created and included in the list of Snapshot copies.

P01-data-mnt00001 (ANF-West-Europe/ANF-WE-01/P01-data-mnt00001) - Snapshots

Volume

Search (Cmd+/) << + Add snapshot Refresh

Search snapshots

Name	Location	Created
initial-offline	West Europe	04/11/2019, 02:25:50 PM
2019-04-15T09	West Europe	04/15/2019, 11:18:13 AM
2019-04-15T100531.540469	West Europe	04/15/2019, 12:05:33 PM
TESTEBID	West Europe	04/15/2019, 12:05:46 PM
TEST001	West Europe	04/15/2019, 12:12:53 PM
TESTID	West Europe	04/15/2019, 12:23:54 PM
vm-sc-server_hana_P01_P01_vm...	West Europe	08/01/2019, 09:41:10 PM
vm-sc-server_hana_P01_P01_vm...	West Europe	08/01/2019, 11:23:30 PM
vm-sc-server_hana_P01_P01_vm...	West Europe	08/02/2019, 06:15:05 AM
vm-sc-server_hana_P01_P01_vm...	West Europe	08/02/2019, 06:33:39 AM
vm-sc-server_hana_P01_P01_vm...	West Europe	08/02/2019, 07:03:59 AM
Offline	West Europe	08/20/2019, 05:02:46 PM
vm-sc-server_hana_P01_P01_vm...	West Europe	08/26/2019, 04:39:17 PM
2019-08-30T070801.100767	West Europe	08/30/2019, 09:08:07 AM
TECHED001	West Europe	09/23/2019, 05:24:55 PM
btobias01	West Europe	10/04/2019, 12:11:53 PM
2019-11-05T141137.437862	West Europe	11/05/2019, 03:11:40 PM
TEST002	West Europe	12/19/2019, 03:44:09 PM
TEST003	West Europe	12/19/2019, 03:53:28 PM

4. Close the backup:

- a. As a final step, we need to tell SAP HANA to close the backup. For this command, you need to know the internal backup number from step 1 and the Snapshot copy name.

```
hdbsql SYSTEMDB=> BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID 1576767153528 SUCCESSFUL
'TEST003';
0 rows affected (overall time 361.822 msec; server time 359.357 msec)
```

- b. Use the SAP HANA studio to display the new entries in the SAP HANA backup catalogue.

The screenshot shows the SAP HANA Backup Catalogue for SYSTEMDB@P01. The 'Backup Catalog' tab is active, displaying a table of backup entries. The most recent backup is highlighted in blue.

Status	Started	Duration	Size	Backup Type	Destination...
Success	Dec 19, 2019 2:52:33 ...	00h 01m 48s	1.64 GB	Data Backup	Snapshot
Success	Dec 19, 2019 2:42:39 ...	00h 07m 06s	1.70 GB	Data Backup	Snapshot
Success	Nov 5, 2019 2:11:38 ...	00h 00m 03s	1.72 GB	Data Backup	Snapshot
Success	Oct 4, 2019 10:11:37 ...	00h 00m 08s	1.73 GB	Data Backup	Snapshot
Success	Sep 23, 2019 3:24:42 ...	00h 00m 04s	1.69 GB	Data Backup	Snapshot
Success	Aug 30, 2019 7:08:01...	00h 00m 07s	1.56 GB	Data Backup	Snapshot
Success	Aug 30, 2019 6:44:37...	00h 00m 13s	1.56 GB	Data Backup	Snapshot
Success	Aug 26, 2019 2:39:07...	00h 00m 13s	1.52 GB	Data Backup	Snapshot
Success	Aug 2, 2019 5:03:57 ...	00h 00m 08s	1.50 GB	Data Backup	Snapshot
Success	Aug 2, 2019 4:33:37 ...	00h 00m 09s	1.50 GB	Data Backup	Snapshot
Success	Aug 2, 2019 4:15:03 ...	00h 00m 08s	1.50 GB	Data Backup	Snapshot
Success	Aug 1, 2019 9:23:26 ...	00h 00m 09s	1.50 GB	Data Backup	Snapshot
Success	Aug 1, 2019 7:41:05 ...	00h 00m 11s	1.50 GB	Data Backup	Snapshot
Success	Apr 15, 2019 10:23:1...	00h 00m 47s	1.47 GB	Data Backup	Snapshot
Success	Apr 15, 2019 10:15:3...	00h 01m 04s	1.47 GB	Data Backup	File

The 'Backup Details' panel on the right shows the following information for the selected backup:

- ID: 1576767153528
- Status: Successful
- Backup Type: Data Backup
- Destination Type: Snapshot
- Started: Dec 19, 2019 2:52:33 PM (UTC)
- Finished: Dec 19, 2019 2:54:22 PM (UTC)
- Duration: 00h 01m 48s
- Size: 1.64 GB
- Throughput: n.a.
- System ID:
- Comment: TEST003
- Additional Information: <ok>
- Location: /hana/data/P01/mnt00001/

The 'Host' table below the details panel shows:

Host	Service	Size	Name	Source Type	EBID
vm-p01	nameserver	1.64 GB	hdb00001	volume	TEST003

- c. Verify that the comment and the external backup ID (EBID) are set to the Snapshot copy name. SAP HANA creates this entry for all tenant databases.

The screenshot shows the SAP HANA Backup Catalogue for P01. The 'Backup Catalog' tab is active, displaying a table of backup entries. The most recent backup is highlighted in blue.

Status	Started	Duration	Size	Backup Type	Destination...
Success	Dec 19, 2019 2:52:33 ...	00h 01m 48s	1.84 GB	Data Backup	Snapshot
Success	Dec 19, 2019 2:42:39 ...	00h 07m 06s	2.00 GB	Data Backup	Snapshot
Success	Nov 5, 2019 2:11:38 ...	00h 00m 03s	1.94 GB	Data Backup	Snapshot
Success	Oct 4, 2019 10:11:37 ...	00h 00m 08s	1.94 GB	Data Backup	Snapshot
Success	Sep 23, 2019 3:24:42 ...	00h 00m 04s	1.89 GB	Data Backup	Snapshot
Success	Aug 30, 2019 7:08:01...	00h 00m 07s	1.77 GB	Data Backup	Snapshot
Success	Aug 30, 2019 6:44:37...	00h 00m 13s	1.77 GB	Data Backup	Snapshot
Success	Aug 26, 2019 2:39:07...	00h 00m 13s	1.73 GB	Data Backup	Snapshot
Success	Aug 2, 2019 5:03:57 ...	00h 00m 08s	1.70 GB	Data Backup	Snapshot
Success	Aug 2, 2019 4:33:37 ...	00h 00m 09s	1.70 GB	Data Backup	Snapshot
Success	Aug 2, 2019 4:15:03 ...	00h 00m 08s	1.70 GB	Data Backup	Snapshot
Success	Aug 1, 2019 9:23:26 ...	00h 00m 09s	1.70 GB	Data Backup	Snapshot
Success	Aug 1, 2019 7:41:05 ...	00h 00m 11s	1.70 GB	Data Backup	Snapshot
Success	Apr 15, 2019 10:16:5...	00h 01m 09s	1.67 GB	Data Backup	File

The 'Backup Details' panel on the right shows the following information for the selected backup:

- ID: 1576767153528
- Status: Successful
- Backup Type: Data Backup
- Destination Type: Snapshot
- Started: Dec 19, 2019 2:52:33 PM (UTC)
- Finished: Dec 19, 2019 2:54:22 PM (UTC)
- Duration: 00h 01m 48s
- Size: 1.84 GB
- Throughput: n.a.
- System ID:
- Comment: TEST003
- Additional Information: <ok>
- Location: /hana/data/P01/mnt00001/

The 'Host' table below the details panel shows:

Host	Service	Size	Name	Source Type	EBID
vm-p01	xsengine	80.00 MB	hdb00002.00003	volume	TEST003
vm-p01	indexserver	1.77 GB	hdb00003.00003	volume	TEST003

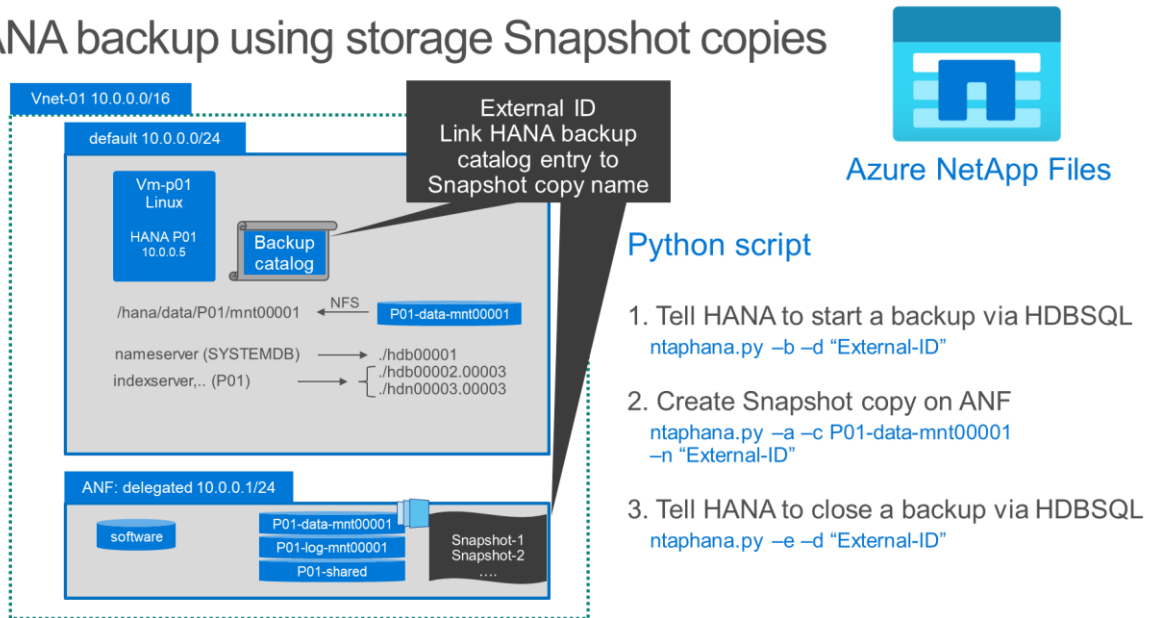
Note: It is important to close the Snapshot copy; otherwise, SAP HANA will reject any new backup.

Using Azure NetApp Files API Calls to Automate the Backup

To automate the Snapshot copy backup process, use scripting to orchestrate the start backup > snapshot > close backup sequence. Azure NetApp Files provides a variety of [REST API options](#) as well as [SDKs to list and create snapshots](#). This allows the three-step process described above to be automated following the example process shown in Figure 22.

Figure 22) HANA backup using Python script.

HANA backup using storage Snapshot copies



For example, to map the backup entry in the SAP HANA backup catalog to the corresponding Snapshot copy on the Azure NetApp Files volume, you can use the CLI parameter External-ID. With this parameter, the Snapshot backup can be easily identified from the SAP HANA backup catalog for later restore operations.

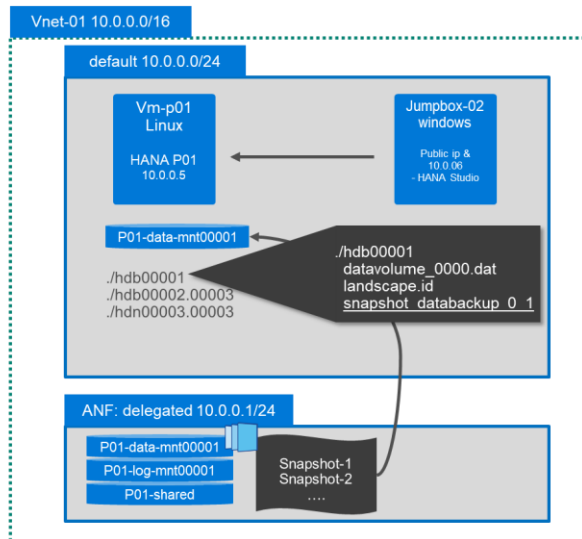
Note: This information is meant for example purposes only. The actual Python scripts are not provided by NetApp or Microsoft and must be created by the customer.

4.4 SAP HANA Restore in Seconds Using Snapshot Copies

Restoring a SAP HANA MDC database is best explained with an illustration, as shown in Figure 23.

Figure 23) SAP HANA restore workflow.

HANA restore workflow



Azure NetApp Files

Restore and recovery workflow

1. Start SYSTEM DB recovery using HANA studio or HANA cockpit
SAP HANA will be stopped
2. Restore the volumes content to its snapshot stage
Several options how to restore
3. Recover the SYSTEM DB using HANA studio, HANA cockpit, or SAP script
4. Start recovery of the Tenant database using HANA studio, HANA cockpit or SAP script

A SAP HANA MDC database can be installed on a single host (scale up) or in a multiple-host setup (scale out). It contains a SYSTEM database and one or more tenant databases. For a full restore and recovery, the SAP HANA database must first be stopped. After the file structure (or volume) has been restored, the SYSTEM database needs to be recovered first, followed by recovery of each existing tenant.

The SAP HANA Persistence Layer

To understand the different options to restore the SAP HANA persistence layer, we need to look at the file system structure of a SAP HANA database. This example has a single-host SAP HANA system installed that has one tenant database. At the storage level, each host has a single mount point where an Azure NetApp Files volume for each host is mounted:

```
vm-p01:/hana/data/P01/mnt00001 # df -h .
Filesystem                Size  Used Avail Use% Mounted on
10.0.1.4:/P01-data-mnt00001 101T  7.8G 100T   1% /hana/data/P01/mnt00001
```

For the SAP HANA host, a numbered mount point (mnt00001 for a single-host system) is created in the folder /hana/data/P01, where P01 is the system identifier (SID) of the HANA database. On this mount point, a single Azure NetApp Files volume is mounted. Following a standard naming convention, the volume name and export path are P01-data-mnt00001, where P01 is the SAP HANA SID and mnt00001 is the first and only mount point of the single-host database.

Inside the mounted volume, the following typical folder structure is present:

```
vm-p01:/hana/data/P01/mnt00001 # ls -als .
total 20
drwxr-x- 5 p01adm sapsys 4096 May  8 17:56 .
drwxr-x- 3 p01adm sapsys  22 Apr 10 12:13 ..
drwxrwxrwx 14 root    4096 Aug 20 15:02 .snapshot
drwxr-x- 2 p01adm sapsys 4096 Aug  2 05:04 hdb00001
drwxr-x-x 2 p01adm sapsys 4096 Apr 11 08:45 hdb00002.00003
drwxr-x-x 2 p01adm sapsys 4096 Aug  2 05:04 hdb00003.00003
-rw-r-r- 1 p01adm sapsys  17 Aug  5 15:09 nameserver.lck
```

The SYSTEM database and each tenant have one or more folders, again following an SAP HANA internal naming and numbering schema. Each of the folders contains a single data file, `datavolume_0000.dat`, which contains the persistence data of the SAP HANA database:

```
hdb00002.00003:
total 329088
__DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
datavolume_0000.dat
```

There is also a file called `__DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__`, indicating that users are not supposed to touch any of the files in this structure.

At the time that a storage-based Snapshot backup is initiated, SAP HANA ensures that the data volume is in a consistent state, and it creates an additional file, `snapshot_databackup_0_1`, in this structure, which is later used to identify the backup in case a SAP HANA database recovery is necessary.

Snapshots are Always Visible

When using Azure NetApp Files for a SAP HANA database, an additional `.snapshot` directory is visible at the top of the mount point. Under this directory, all created Snapshot copies are available—giving read-only access to the stored backup data.

```
vm-p01:/hana/data/P01/mnt00001/.snapshot # ls -als
total 48
drwxr-x- 5 p01adm sapsys 4096 May  8 17:56 Offline
drwxr-x- 5 p01adm sapsys 4096 Apr 11 08:45 TEST001
drwxr-x- 5 p01adm sapsys 4096 Apr 11 08:45 TESTEBID
drwxr-x- 5 p01adm sapsys 4096 Apr 11 08:45 TESTID
```

The folder is given the Snapshot copy name and contains the complete file structure at the time the Snapshot copy was created:

```
vm-p01:/hana/data/P01/mnt00001/.snapshot/TESTID # ls -als *
-rw-r-r- 1 p01adm sapsys 17 Apr 15 10:23 nameserver.lck
hdb00001:
total 3376564
-rw-r-- 1 p01adm sapsys 0 Apr 11 08:44 __DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
-rw-r-- 1 p01adm sapsys 3690987520 Apr 15 10:23 datavolume_0000.dat
-rw-r-- 1 p01adm sapsys 36 Apr 15 06:51 landscape.id
-rw-r-- 1 p01adm sapsys 155648 Apr 15 10:23 snapshot_databackup_0_1
hdb00002.00003:
total 67996
-rw-r-r- 1 p01adm sapsys 0 Apr 11 08:45 __DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
-rw-r-r- 1 p01adm sapsys 335659008 Apr 15 10:23 datavolume_0000.dat
hdb00003.00003:
total 3591184
-rw-r-r- 1 p01adm sapsys 0 Apr 11 08:45 __DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
-rw-r-r- 1 p01adm sapsys 3892314112 Apr 15 10:23 datavolume_0000.dat
-rw-r-- 1 p01adm sapsys 155648 Apr 15 10:23 snapshot_databackup_0_1
vm-p01:/hana/data/P01/mnt00001/.snapshot/TESTID #
```

As shown in the listing for the TESTID Snapshot copy, there are several files named `snapshot_databackup_0_1` for each of the databases, indicating that this is a valid application-consistent SAP HANA backup that is also referenced in the SAP HANA backup catalog.

Options to Restore

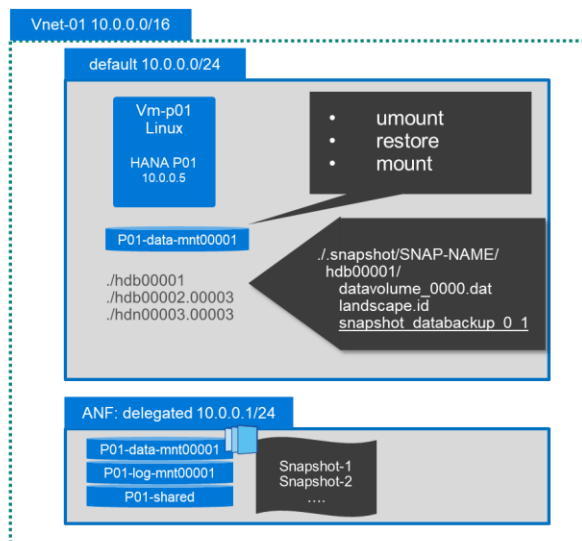
These features of Azure NetApp Files enable three different methods to restore a SAP HANA database:

- Copy files from the `.snapshot` directory. Because all Snapshot copies are available in the `.snapshot` folder, it is possible to copy the files from the Snapshot to the original mount point. This host-based copy might take some time, but it has the advantage of being able to restore and recover a single tenant instead of the whole database (volume), simplifying the recovery process.

- Restore Snapshot copy to a new volume (thick clone). Azure NetApp Files enables the creation of a new volume based on an existing Snapshot copy. This process is extremely fast, just a few seconds, benefiting from ONTAP advanced Snapshot and cloning technology. Instead of the original volume, the new volume is mounted to start the SAP HANA recovery process. This method of restore is immediate, however does require data to be copied, making the restore operation extremely fast but less space efficient. The advantage of this procedure is that it leaves the original volume unchanged, enabling the recovery of the system into a new volume while preserving the original volume (for example, for forensics and root-cause analysis purposes).
- Restore in-place Snapshot copy (revert). At the time of this writing, the in-place Snapshot restore (revert) feature was available in the Azure NetApp Files preview. In-place Snapshot restore is the fastest method to reset a volume back in time (revert) to the moment of the creation of the Snapshot copy. This proven ONTAP technology has been used for many years to restore SAP systems using NetApp storage in on-premises environments. Instead of copying data or creating a new volume, the data is restored in the existing volume. This method of restore is immediate and does not require any data to be copied, making the restore extremely fast and space efficient, regardless of the size of the volume and the amount of changed data in the volume Snapshot copies, as shown in Figure 24.

Figure 24) SAP HANA restore options.

HANA restore options



Azure NetApp Files

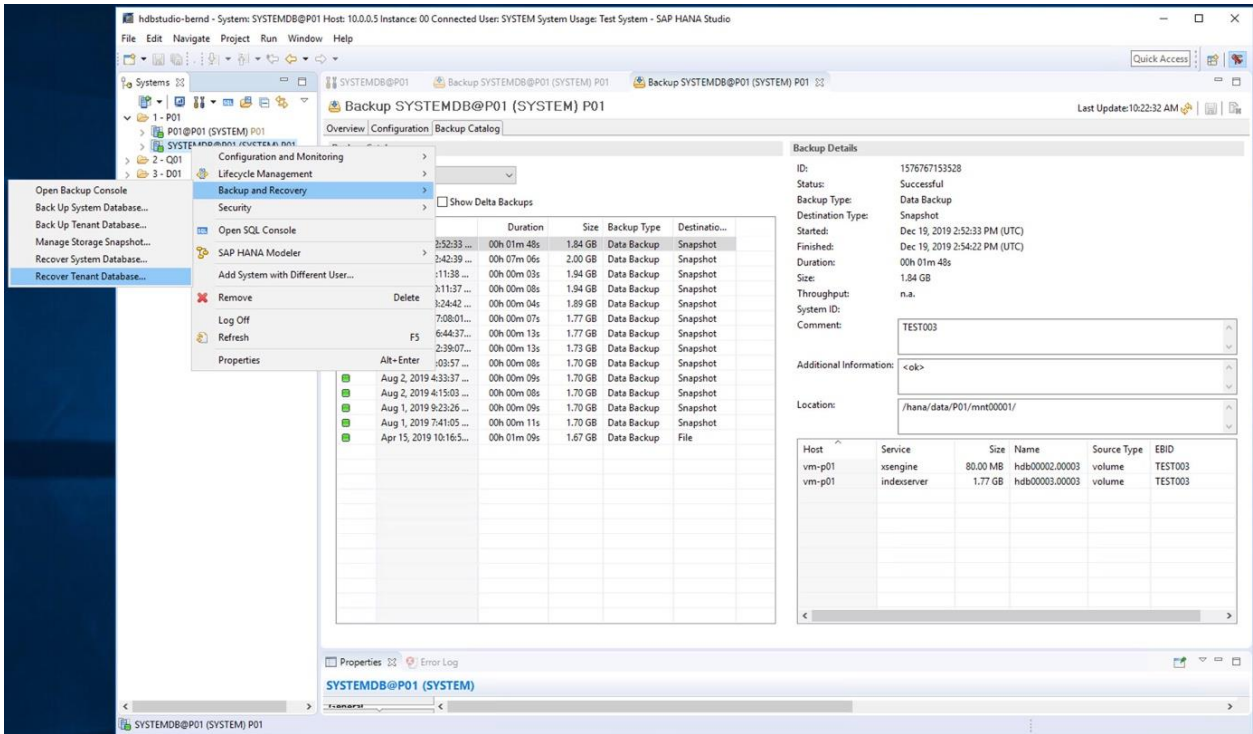
Restore and recovery workflow

1. Copy on client using `.snapshot` folder
 - Granular copy of single tenants possible
 - Preserves all existing Snapshot copies
 - Slow-moving data via client
2. Restore Snapshot copy at storage side via API call
 - Super fast, no data moved
 - Newer Snapshots copie/backups will be lost
 - API call not available in preview
3. Create new volume based on Snapshot copy via API or GUI
 - Requires you to mount a new volume


Example Workflow 1: Restore SAP HANA to New Volume

This example uses the Azure portal to create a new volume based on an existing Snapshot copy to restore the system. It also uses SAP HANA studio to recover the system database as well as the tenant database.

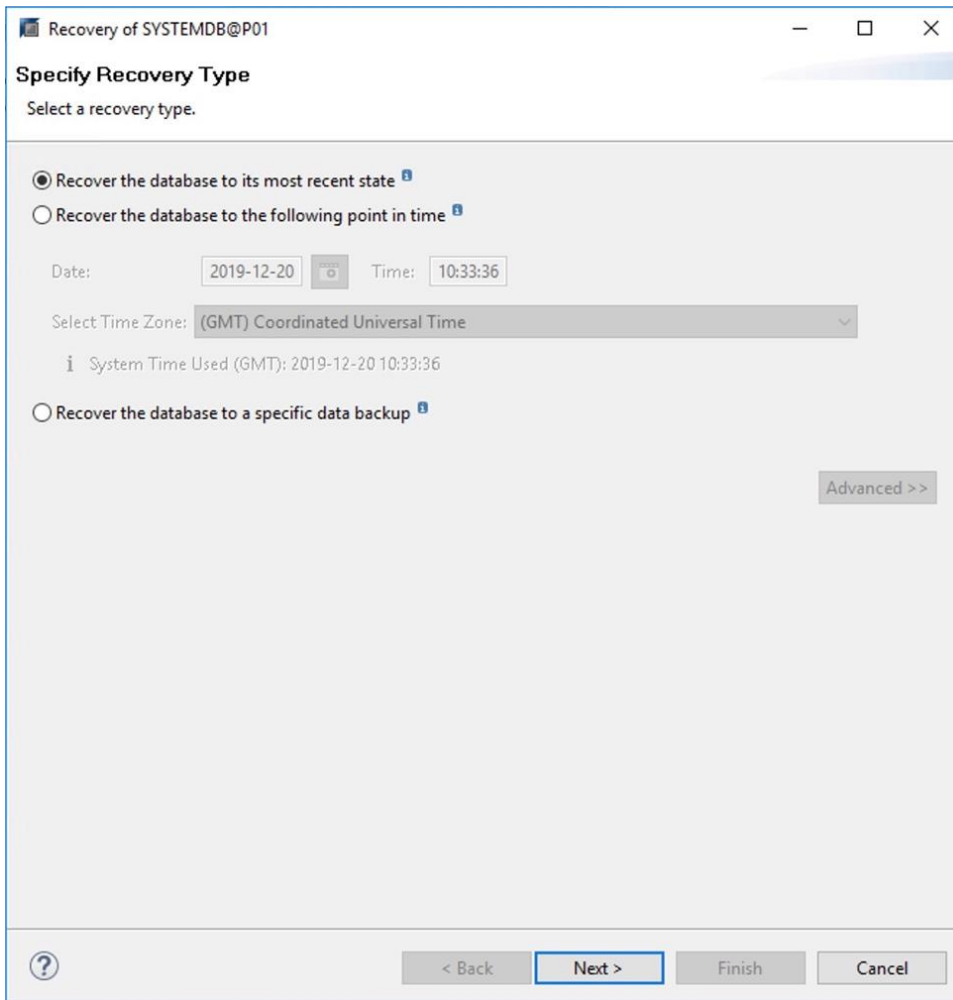
1. To follow the workflow logic described above, you first need to stop the SAP HANA database. Stopping the database is part of the SAP HANA studio recovery workflow, so we'll use the recovery workflow to combine this step. Start HANA studio and select Recover System database.



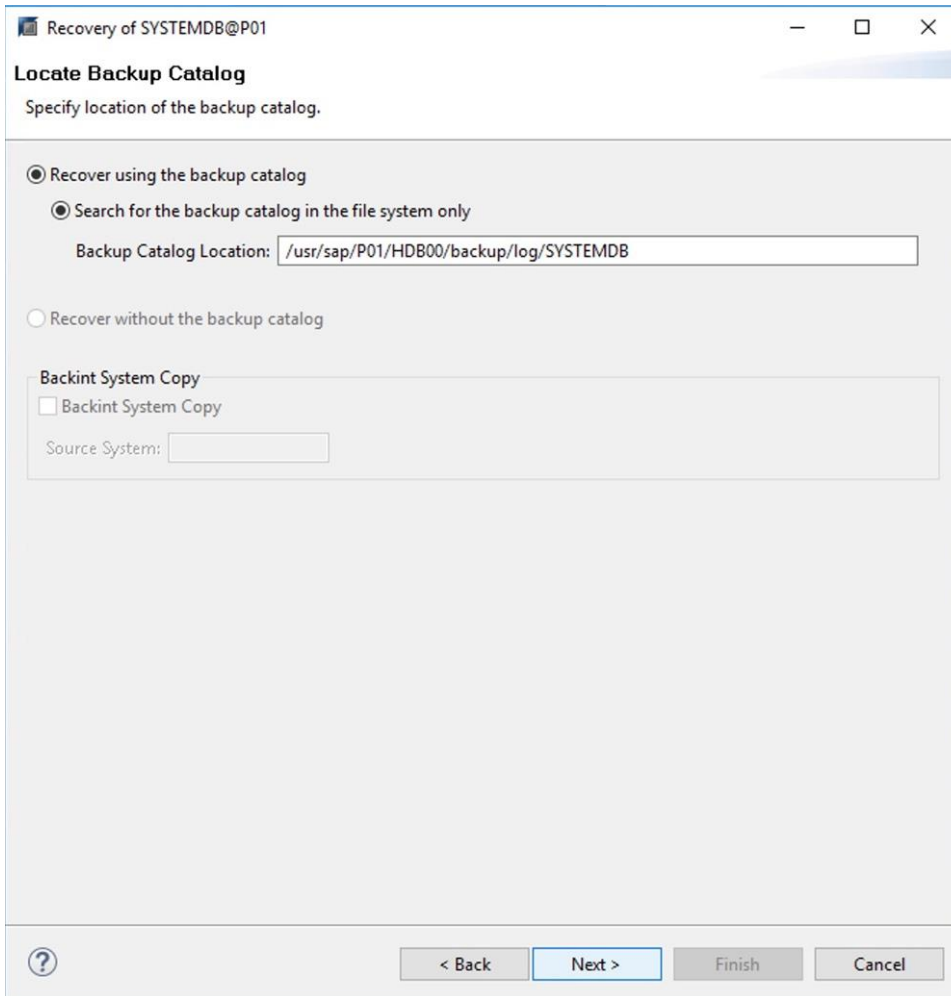
- The recovery wizard starts and asks for confirmation to stop the SAP HANA database. Confirm by clicking OK.



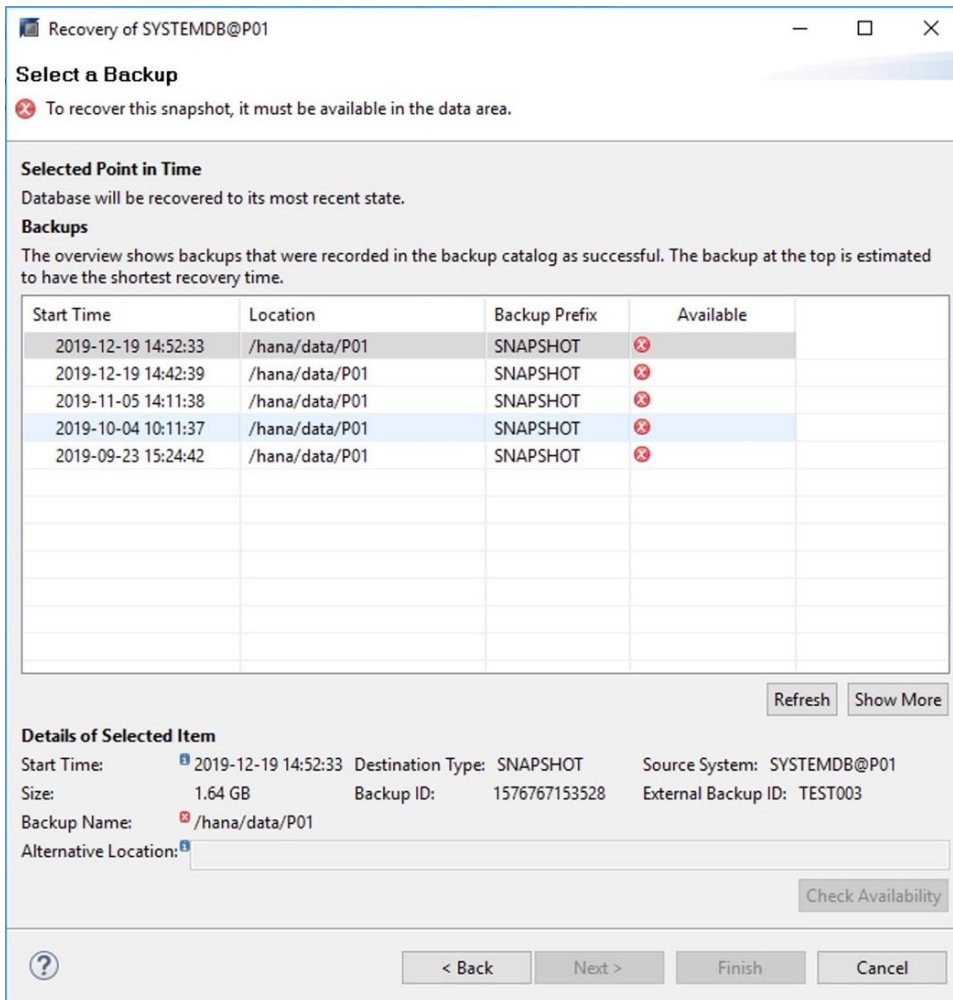
- On the Specify Recovery Type page, select the type of restore. Click Next.



4. Specify where the log backups will be found in case they are required for a recovery. Click Next.



5. The wizard displays the backup entries and highlights how there is no available Snapshot-based backup. This means that the Recovery wizard has not yet found the special file that validates whether the files are part of a Snapshot-based backup. You can use this page to verify which Snapshot copy you need to restore. The external backup ID that was stored in the SAP HANA backup has been set as the Snapshot copy name. In this example, it's TEST003.



- The database is now stopped. Use the Azure portal to create a new volume based on the TEST003 Snapshot copy that you selected. Start the Azure portal and select the P01-data-mnt0001 volume within the volume the Snapshot copy area. All available Snapshot copies are displayed. Select the Snapshot TEST003 and right-click to select Restore to New Volume.

P01-data-mnt00001 (ANF-West-Europe/ANF-WE-01/P01-data-mnt00001) - Snapshots

Volume

Search (Cmd+/) Add snapshot Refresh

Search snapshots

Name	Location	Created	
initial-offline	West Europe	04/11/2019, 02:25:50 PM	...
2019-04-15T09	West Europe	04/15/2019, 11:18:13 AM	...
2019-04-15T100531.540469	West Europe	04/15/2019, 12:05:33 PM	...
TESTEBID	West Europe	04/15/2019, 12:05:46 PM	...
TEST001	West Europe	04/15/2019, 12:12:53 PM	...
TESTID	West Europe	04/15/2019, 12:23:54 PM	...
vm-sc-server_hana_P01_vm-sc-se...	West Europe	08/01/2019, 09:41:10 PM	...
vm-sc-server_hana_P01_vm-sc-se...	West Europe	08/01/2019, 11:23:30 PM	...
vm-sc-server_hana_P01_vm-sc-se...	West Europe	08/02/2019, 06:15:05 AM	...
vm-sc-server_hana_P01_vm-sc-se...	West Europe	08/02/2019, 06:33:39 AM	...
vm-sc-server_hana_P01_vm-sc-se...	West Europe	08/02/2019, 07:03:59 AM	...
Offline	West Europe	08/20/2019, 05:02:46 PM	...
vm-sc-server_hana_P01_vm-sc-se...	West Europe	08/26/2019, 04:39:17 PM	...
2019-08-30T070801.100767	West Europe	08/30/2019, 09:08:07 AM	...
TECHED001	West Europe	09/23/2019, 05:24:55 PM	...
btobias01	West Europe	10/04/2019, 12:11:53 PM	...
2019-11-05T141137.437862	West Europe	11/05/2019, 03:11:40 PM	...
TEST002	West Europe	12/19/2019, 03:44:09 PM	...
TEST003	West Europe	12/19/2019, 03:44:09 PM	...

Restore to new volume
Delete

- On the Create a Volume page, specify the volume name and size for the new volume. NetApp recommends that you follow the naming convention of adding a `-clone` suffix for the new volume `P01-data-mnt00001-clone`.

Create a volume

[Basics](#) [Protocol](#) [Tags](#) [Review + create](#)

This page will help you create an Azure NetApp Files volume in your subscription and enable you to access the volume from within your virtual network. [Learn more about Azure NetApp Files](#)

Volume details

Volume name *	<input type="text" value="P01-data-mnt00001-clone"/>
Restoring from snapshot	<input type="text" value="TEST003"/>
Available quota (GiB)	<input type="text" value="1096"/> 1.07 TiB
Quota (GiB) *	<input type="text" value="1000"/> 1000 GiB
Virtual network	<input type="text" value="vnet-01 (10.0.0.0/16)"/>
Subnet	<input type="text" value="delegated (10.0.1.0/24)"/>

[Review + create](#)

[Next: Protocol »](#)

[Download a template for automation](#)

8. Finish the workflow and create the volume. If you follow the naming convention, you can list all the volumes belonging to P01 by using the filter.

+ Add volume Refresh

Name	Quota	Protocol ty...	Mount path	Service level	Capacity pool
P01-data-mnt00001	100 GiB	NFSv3	10.0.1.4/P01-data-mnt00001	Premium	ANF-WE-01
P01-data-mnt00001-clone	1000 GiB	NFSv3	10.0.1.4/P01-data-mnt00001-...	Premium	ANF-WE-01
P01-log-mnt00001	500 GiB	NFSv3	10.0.1.4/P01-log-mnt00001	Premium	ANF-WE-01
P01-shared	700 GiB	NFSv3	10.0.1.4/P01-shared	Premium	ANF-WE-01

9. Select the new volume to verify the mount path.

P01-data-mnt00001-clone (ANF-West-Europe/ANF-WE-01/P01-data-mnt00001-clone)
Volume

Search (Cmd+ /) Resize Delete

Overview

- Activity log
- Access control (IAM)
- Tags

Settings

- Properties
- Locks
- Export template

Storage service

- Mount instructions
- Export policy
- Snapshots
- Replication

Monitoring

- Metrics

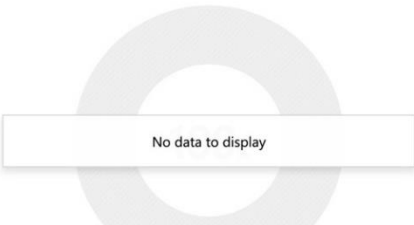
Support + troubleshooting

- New support request

<p>Resource group rg-ANF-West-Europe</p> <p>Mount path 10.0.1.4:/P01-data-mnt00001-clone</p> <p>Subscription Pay-As-You-Go</p> <p>Subscription ID 28cfc403-f3f6-4b07-9847-4eb16109e870</p> <p>Virtual network/subnet vnet-01/delegated</p>	<p>Capacity pool ANF-WE-01</p> <p>Protocol type NFSv3</p> <p>Location West Europe</p> <p>Service level Premium</p> <p>Quota 1000 GiB</p>
--	--

Usage

Current volume logical size



No data to display

10. Switch to the SAP HANA host as root user and unmount the old data file.

```
vm-p01:~ # df
Filesystem                1K-blocks      Used    Available Use% Mounted on
devtmpfs                  28823664         0    28823664   0% /dev
tmpfs                     43260912         0    43260912   0% /dev/shm
tmpfs                     28840608    17344    28823264   1% /run
tmpfs                     28840608         0    28840608   0% /sys/fs/cgroup
/dev/sda2                  30392836    4185224    26207612  14% /
/dev/sda1                  10383336     92656     945680    9% /boot
10.0.1.4:/P01-shared/shared 107456932160 135987840 107320944320 1% /hana/shared
10.0.1.4:/P01-data-mnt00001 107371899840 9070976 107362828864 1% /hana/data/P01/mnt00001
10.0.1.4:/P01-log-mnt00001 107374594240 4148800 107370445440 1% /hana/log/P01/mnt00001
10.0.1.4:/software        107369274560 69447168 107299827392 1% /mnt/software
10.0.1.4:/P01-shared/usr-sap 107456932160 135987840 107320944320 1% /usr/sap/P01
tmpfs                     5768120         0     5768120   0% /run/user/175
tmpfs                     5768120         0     5768120   0% /run/user/1001
/dev/sdb1                  115072060    61468    109122240  1% /mnt/resource
tmpfs                     5768120         0     5768120   0% /run/user/0
vm-p01:~ #
vm-p01:~ # umount /hana/data/P01/mnt00001
vm-p01:~ #
```

11. Edit the /etc/fstab to add the cloned volume so that it is mounted instead of the original data volume Vi /etc/fstab.

```
# Software share
10.0.1.4:/software        /mnt/software    nfs    rw,hard,rsize=65536,wsiz=65536,vers=3,tcp    0
0

# SAP HANA P01 - ANF volumes
#10.0.1.4:/P01-data-mnt00001    /hana/data/P01/mnt00001    nfs
rw,vers=3,hard,timeo=600,rsiz=1048576,wsiz=1048576,intr,noatime,nolock 0 0
10.0.1.4:/P01-data-mnt00001-clone    /hana/data/P01/mnt00001    nfs
rw,vers=3,hard,timeo=600,rsiz=1048576,wsiz=1048576,intr,noatime,nolock 0 0
10.0.1.4:/P01-log-mnt00001    /hana/log/P01/mnt00001    nfs
rw,vers=3,hard,timeo=600,rsiz=1048576,wsiz=1048576,intr,noatime,nolock 0 0
```

```

10.0.1.4:/P01-shared/usr-sap /usr/sap/P01 nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0
10.0.1.4:/P01-shared/shared /hana/shared nfs
rw,vers=3,hard,timeo=600,rsz=1048576,wsz=1048576,intr,noatime,nolock 0 0

```

12. Mount the cloned volume.

```

vm-p01:~ # mount /hana/data/P01/mnt00001
vm-p01:~ # df
Filesystem                1K-blocks      Used    Available  Use% Mounted on
devtmpfs                  28823664         0    28823664   0% /dev
tmpfs                     43260912         0    43260912   0% /dev/shm
tmpfs                     28840608    17344    28823264   1% /run
tmpfs                     28840608         0    28840608   0% /sys/fs/cgroup
/dev/sda2                 30392836    4185484    26207352  14% /
/dev/sda1                 1038336         92656     945680   9% /boot
10.0.1.4:/P01-shared/shared 107456932160 135987904 107320944256 1% /hana/shared
10.0.1.4:/P01-log-mnt00001 107374594240 4148800 107370445440 1% /hana/log/P01/mnt00001
10.0.1.4:/software        107369274560 69447168 107299827392 1% /mnt/software
10.0.1.4:/P01-shared/usr-sap 107456932160 135987904 107320944256 1% /usr/sap/P01
tmpfs                     5768120         0     5768120   0% /run/user/175
tmpfs                     5768120         0     5768120   0% /run/user/1001
/dev/sdb1                 115072060     61468    109122240  1% /mnt/resource
tmpfs                     5768120         0     5768120   0% /run/user/0
10.0.1.4:/P01-data-mnt00001-clone 107377966848 9036800 107368930048 1%
/hana/data/P01/mnt00001

```

13. Back in the Recovery wizard, refresh the page to see if the newly mounted volume is identified as a valid backup. If the green light identifies that the backup is valid, click Next.

Recovery of SYSTEMDB@P01

Select a Backup
Select a backup to recover the SAP HANA database

Selected Point in Time
Database will be recovered to its most recent state.

Backups
The overview shows backups that were recorded in the backup catalog as successful. The backup at the top is estimated to have the shortest recovery time.

Start Time	Location	Backup Prefix	Available
2019-12-19 14:52:33	/hana/data/P01	SNAPSHOT	●
2019-12-19 14:42:39	/hana/data/P01	SNAPSHOT	✘
2019-11-05 14:11:38	/hana/data/P01	SNAPSHOT	✘
2019-10-04 10:11:37	/hana/data/P01	SNAPSHOT	✘
2019-09-23 15:24:42	/hana/data/P01	SNAPSHOT	✘

Refresh Show More

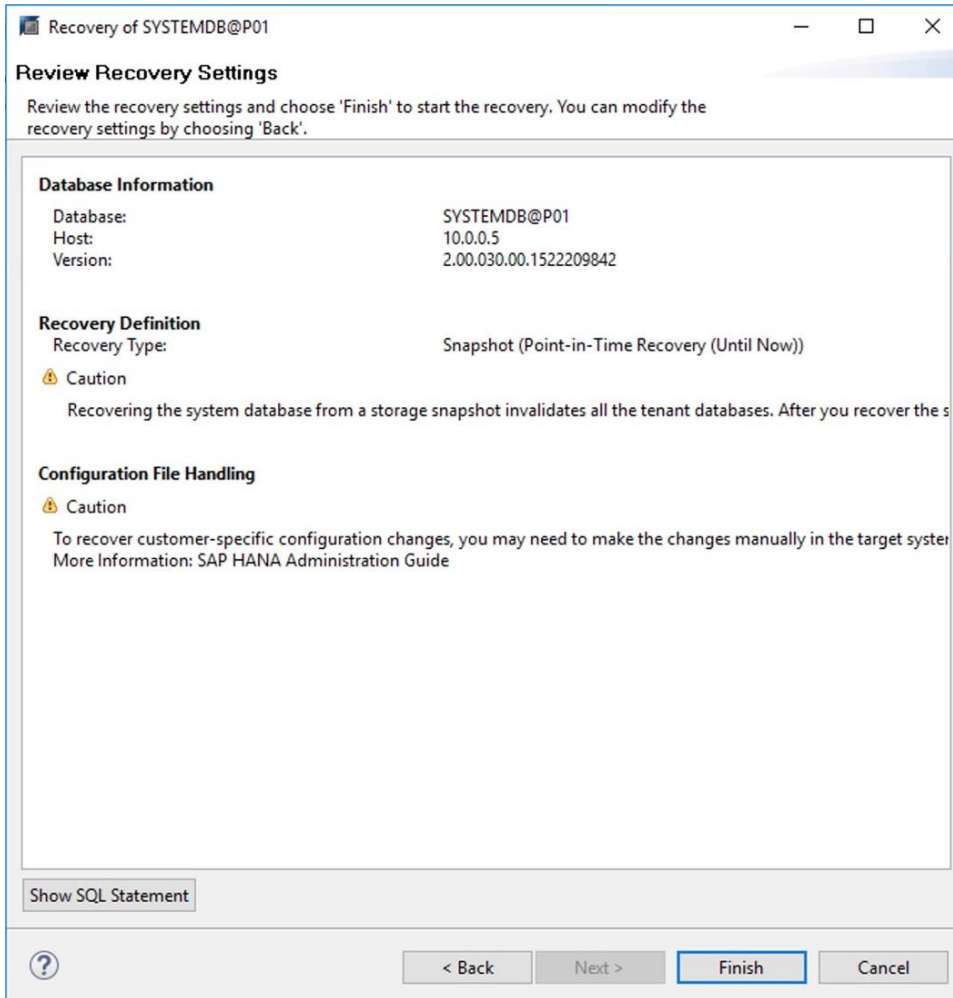
Details of Selected Item

Start Time: 2019-12-19 14:52:33 Destination Type: SNAPSHOT Source System: SYSTEMDB@P01
 Size: 1.64 GB Backup ID: 1576767153528 External Backup ID: TEST003
 Backup Name: /hana/data/P01
 Alternative Location:

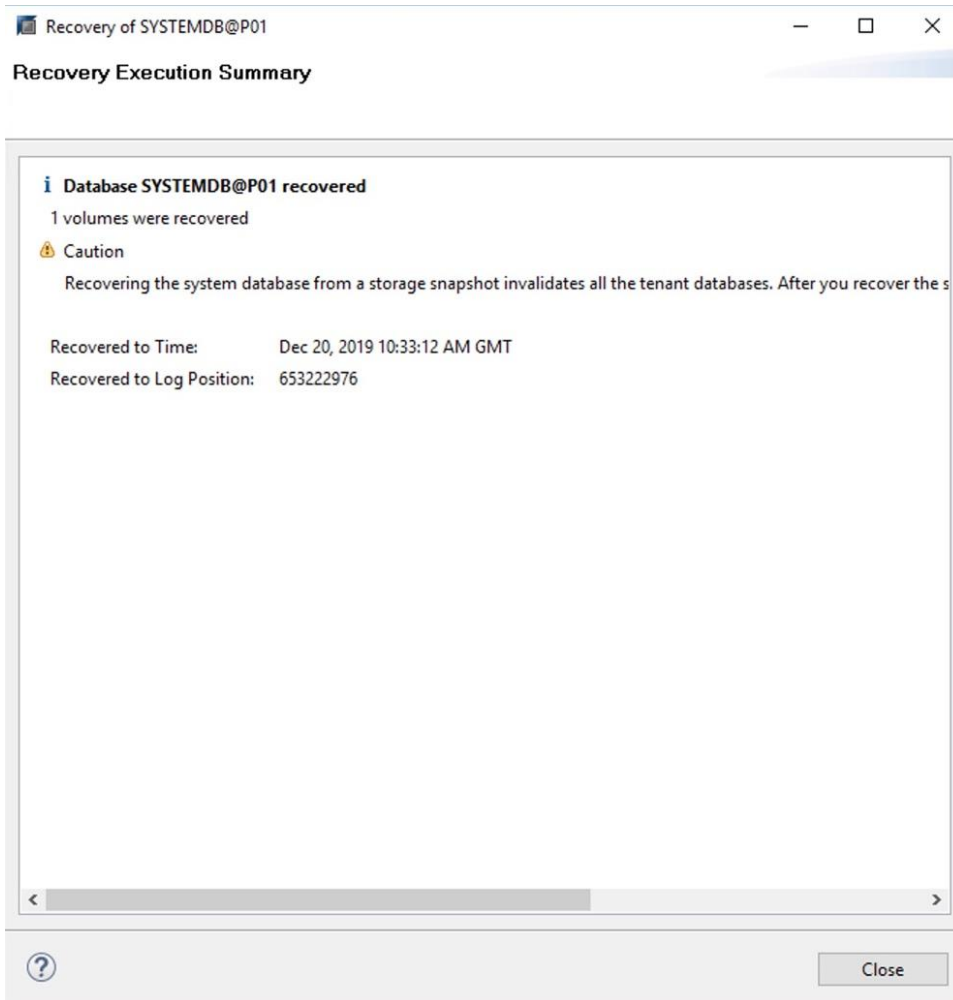
Check Availability

< Back **Next >** Finish Cancel

- Continue the backup workflow until the final confirmation page is displayed. Click Finish to start the recovery process.

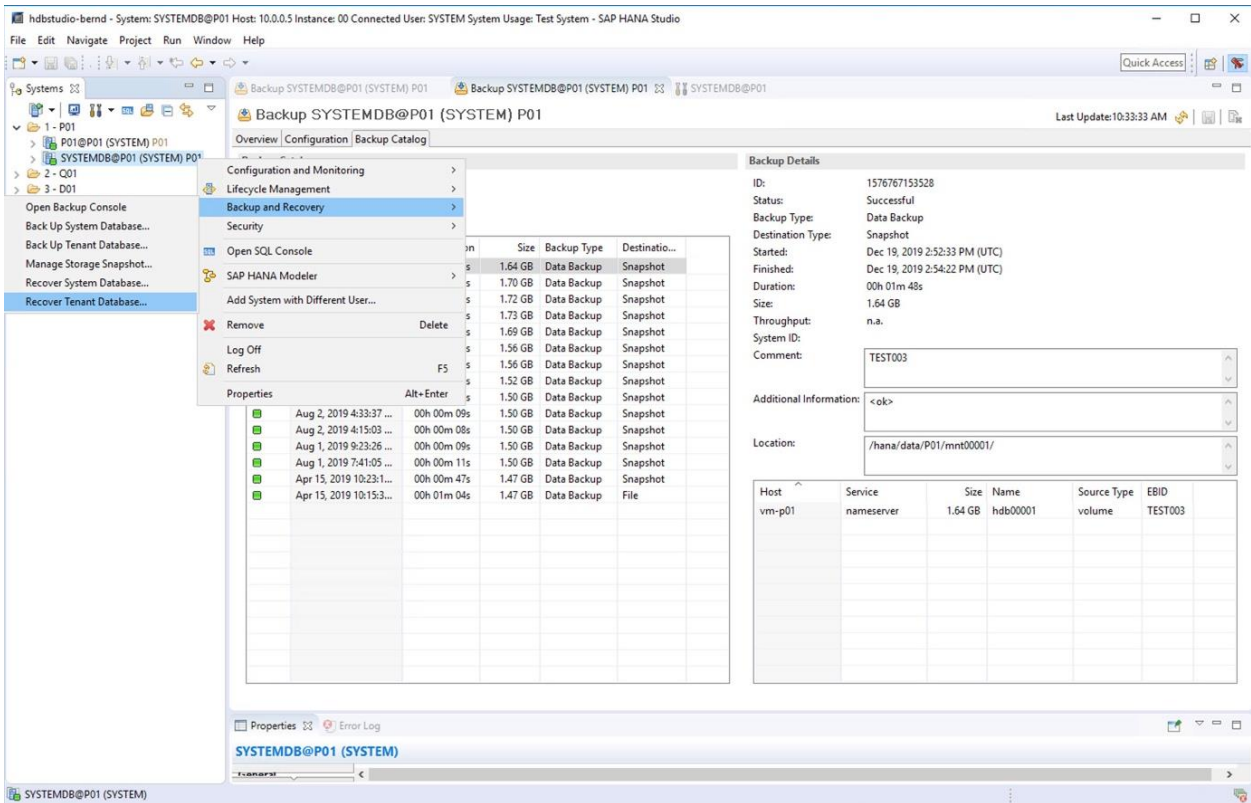


- After the SYSTEMDB recovery is complete, close the workflow.

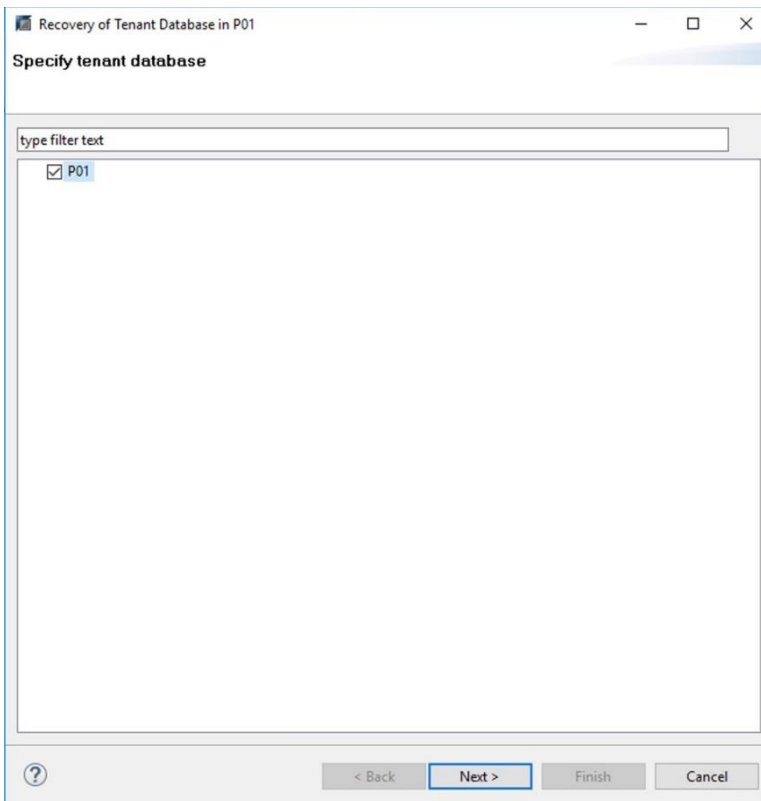


16. Start the recovery workflow for the tenant database.

Note: This menu item is found under the SYSTEMDB.



17. On the start screen, select the tenant you want to recover and click Next.



18. The rest of the workflow is similar to the recovery of the SYSTEMDB, except that this time the wizard already detects the existing Snapshot copy because the tenant files are part of the restored data volume. Continue and wait for the tenant to be recovered.

Note: You can also watch a [YouTube video](#) of this workflow. This video uses a small script that automates the copy from the `.snapshot` directory instead of creating a new volume.

Example Workflow 2: Reverting SAP HANA Previous Snapshot

With Azure NetApp Files, using the new revert feature to restore a volume to its state when the backup (Snapshot copy) was taken is the fastest and most convenient method. Compared to host-based copy or storage-based cloning, the most important advantages of this new revert feature include:

- No data needs to be moved, and no network or host CPU resources are required.
- Fast restoration process (seconds instead of tens of minutes, if not hours)
- No reconfiguration of mount points is required.

Note: Reverting a volume to an older Snapshot copy restores the volume to the exact point-in-time when the Snapshot copy was created. This means that all Snapshot copies created after this 'to be reverted' Snapshot copy will be lost.

For SAP HANA backups, using the revert feature has the following impact:

- When reverting SAP HANA to an older backup, all backups (Snapshot copies) created after this backup are lost; therefore, customers must be sure when to revert to an older-than-most-recent backup (Snapshot copy).
- To recover the database to its most recent state, you must keep all log backups from the time from this reverted Snapshot copy has been created up to the time they need to recover the database.
- You might need to manually clean up the SAP HANA backup catalogue to remove those Snapshot copy-based backups when those more recent Snapshot copies no longer exist.

When considering these implications, it's important to note that volume revert is beneficial when you want to restore SAP HANA to a specific state, for example:

- Taking an initial backup before starting a test cycle that changes the date, so you can easily repeat these tests reverting to the initial state. For example, this process helps to secure system upgrades to give project teams with more time to test compared to scenarios where file-based backups are used.
- For training systems, reverting is a perfect method to freeze a specific state to reset the training system for the next class.
- Immediate restoration of a production environment when there is no need to hold on to the more recent backup (Snapshot copy)², or when storage consumption must be kept to a minimum.

Revert a Volume

Before you start the revert, make sure that the volume is not mounted. For example, when the SAP HANA studio recover workflow for the system database is showing that there is no valid Snapshot backup. At that time, the database is stopped, and you can unmount the data volume.

To revert a volume to a specific Snapshot copy you need to complete the following steps:

²When reverting a volume to any previous Snapshot copy, any evidence of the file system at the moment of data corruption will be lost. This eliminates the possibility for Root Cause Analysis (RCA). If RCA is needed, the better option would be to restore a Snapshot copy to a new volume and change production to that volume. The original, corrupted, data in the volume can then still be used for RCA.

1. Go to the Azure portal and navigate to the data volume. Select Snapshots to see all the Snapshot copies taken for this volume. Select the Snapshot, click on the three dots to the right, and then select Revert Volume. Alternatively, you can use an API call.

The screenshot shows the Azure portal interface for the Snapshots page of volume P01-data-mnt00001. The left sidebar contains navigation options like Overview, Activity log, Access control (IAM), Tags, Settings, Properties, Locks, Export template, Storage service, Mount instructions, Export policy, Snapshots, Replication, Monitoring, Metrics, and Support + troubleshooting. The main area displays a table of snapshots with columns for Name, Location, and Created. The snapshot '2019-11-05T141137.437862' is highlighted, and a context menu is open over it, showing options: Restore to new volume, Revert volume, and Delete.

Name	Location	Created
initial-offline	West Europe	04/11/2019, 2:25:50 PM
2019-04-15T09	West Europe	04/15/2019, 11:18:13 AM
2019-04-15T100531.540469	West Europe	04/15/2019, 12:05:33 PM
TESTEBID	West Europe	04/15/2019, 12:05:46 PM
TEST001	West Europe	04/15/2019, 12:12:53 PM
TESTID	West Europe	04/15/2019, 12:23:54 PM
vm-sc-server_hana_P01_P01_vm-sc-server_08-01-2019_...	West Europe	08/01/2019, 9:41:10 PM
vm-sc-server_hana_P01_P01_vm-sc-server_08-01-2019_...	West Europe	08/01/2019, 11:23:30 PM
vm-sc-server_hana_P01_P01_vm-sc-server_08-02-2019_...	West Europe	08/02/2019, 6:15:05 AM
vm-sc-server_hana_P01_P01_vm-sc-server_08-02-2019_...	West Europe	08/02/2019, 6:33:39 AM
vm-sc-server_hana_P01_P01_vm-sc-server_08-02-2019_...	West Europe	08/02/2019, 7:03:59 AM
Offline	West Europe	08/20/2019, 5:02:46 PM
vm-sc-server_hana_P01_P01_vm-sc-server_08-26-2019_...	West Europe	08/26/2019, 4:39:17 PM
2019-08-30T070801.100767	West Europe	08/30/2019, 9:08:07 AM
TECHED001	West Europe	09/23/2019, 5:24:55 PM
btobias01	West Europe	10/04/2019, 12:11:53 PM
2019-11-05T141137.437862	West Europe	11/05/2019, 3:11:40 PM
TEST002	West Europe	12/19/2019, 3:44:09 PM
TEST003	West Europe	12/19/2019, 3:53:28 PM

2. The portal displays a warning message. Enter the volume name to start the revert workflow.

The screenshot shows the Azure portal interface for the Snapshots page of volume P01-data-mnt00001. The left sidebar is the same as in the previous screenshot. The main area displays the same table of snapshots. A dialog box titled 'Revert volume to snapshot' is open on the right. It contains a warning icon and the text: 'This action is irreversible and it will delete all the volumes snapshots that are newer than 2019-11-05T141137.437862. Please type 'P01-data-mnt00001' to confirm.' Below this, it asks: 'Are you sure you want to revert 'P01-data-mnt00001' to state of 2019-11-05T141137.437862?' with an input field.

3. Remount the data volume and continue the recovery.

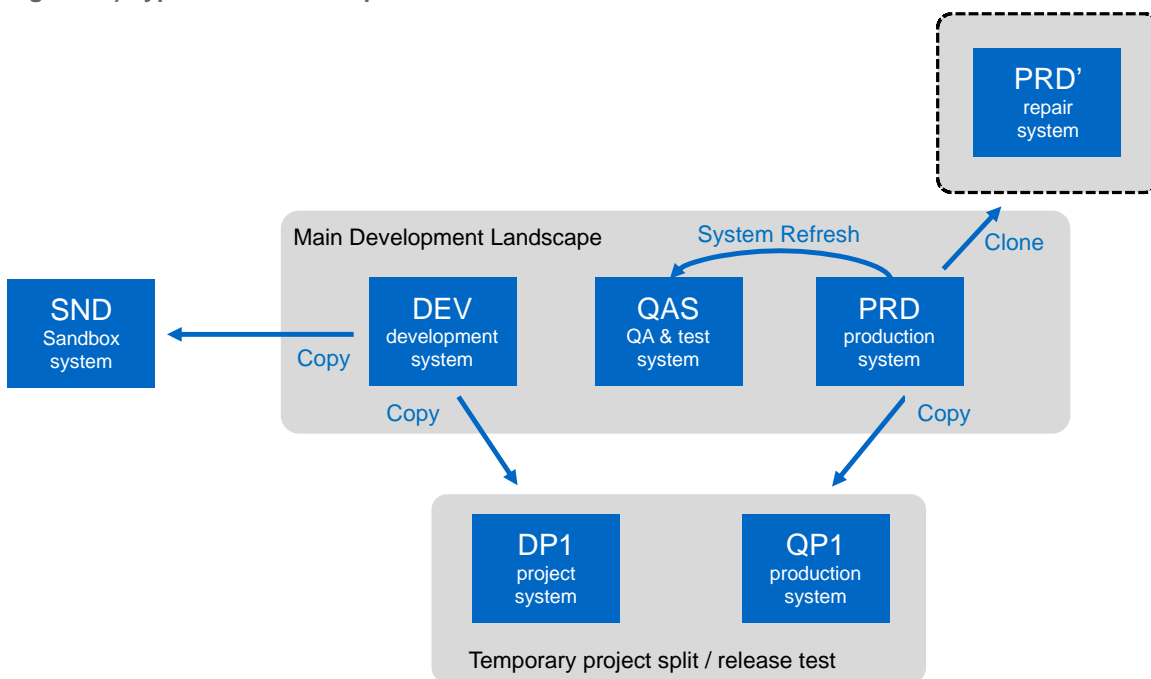
4.5 SAP HANA System Copies Using Snapshot Copies (Thick Clones)

SAP System Copies

To understand the importance and variety of system copies, we need to look at a typical SAP system landscape. Customers need to install the system landscape to ensure that customizations and developments can be made and tested before they are released to production. A few customers might have small landscapes, but most have larger SAP landscapes of 10 or more individual systems with different SAP software running, often dedicated to different organizational units of the company.

As shown in Figure 25, the most basic setup is a three-system landscape with development, QA and test, and production systems (DEV → QAS → PRD). Customers want to test their development in an environment that is as close as possible to the production environment, so the QA and test system should be periodically refreshed with current data from the production system.

Figure 25) Typical SAP landscape.



System Refresh

In the basic system refresh process, you take a backup copy of the production system and restore it to the QA system. After this technical restore, you need to make the following adjustments:

- SAP technical adjustment to change the system ID and all the related activities to convert the production system into a new QA system
- Logical adjustment to (for example) remove or mask some critical data and adjust user rights and privileges

Many customers need to support different project teams to test settings outside of the development system. Some of those settings introduce irreversible changes to the system, and others might never be included in the main development. Therefore, it's advisable to develop and test in a separate "sandbox" system.

Many customers even create a temporary system landscape to test—for example, an upgrade of an SAP landscape or other large changes. For this purpose, the development system is copied. To test these

massive changes with close-to-production data, the production system is used as a source to create a dedicated QA system for this development strand. See the systems DP1 and QP1 in Figure 25.

System Copy

Creating a system copy basically means installing a new system based on an existing one. As part of the system copy process, all the post-copy activities need to be applied, as described in the previous section. System copy and system refresh are closely related. The main difference is that a system refresh updates the target system periodically with data from the source system, while a system copy builds a new system based on an existing system as the source.

But what happens if there is a logical error in production, or if some data is accidentally deleted? There are cases where a restore and recovery of the system is not the right solution, because that would reset the whole system, including all the data that has already been created or changed. In that case, it's necessary to have a clone of the system from before the error happened, to be able to restore the lost data. In such an emergency situation, you don't want to go through the post-copy activities, because the post-copy changes are often not wanted; you need the clone to be as exact as possible to be used as a repair system.

Repair System

The main difference between a repair system and a system copy is that you must not change the system ID or apply any other changes. Therefore, the new system must be isolated or firewalled to ensure that the copy of the production system does not interfere with any connected systems.

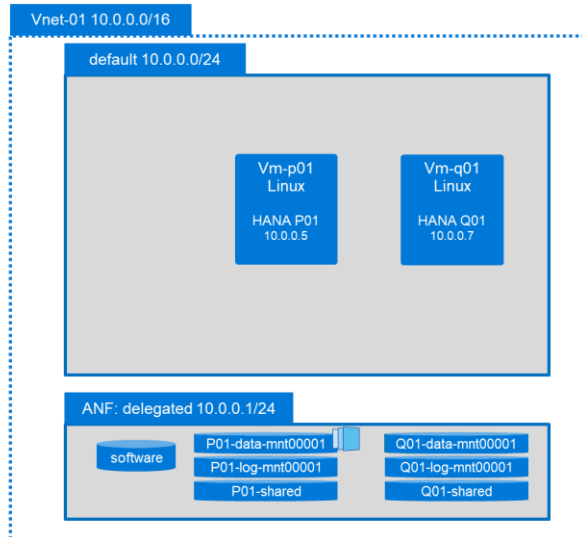
Technical Implementation and Relevance in the Cloud

SAP system copy, clone, refresh, and repair workflows are part of every SAP administrator's worklist. The technical basis for these workflows is an offline or online backup of the underlying database or file systems. Based on the size of the database, the backup and restore process can take up to several hours, and it has an impact on the source system itself, which customers usually want to avoid.

With Azure NetApp Files, application-consistent backups are based on Snapshot copies. They are fast, with almost no impact on capacity of the storage and no impact on performance of the storage service. With a simple API command, or through the Azure portal, customers can create a new volume based on such a Snapshot copy, which replaces a lengthy file-based backup and restore cycle. This storage clone is done in minutes and can be instantaneously used to start the nontechnical post-copy activities; or, in the case of a repair system, to apply the necessary redo logs to move to the point in time before the error happened, as shown in Figure 26.

Figure 26) SAP system refresh/copy workflow.

System copy



Azure NetApp Files

System refresh workflow

1. Shutdown Q01
Unmount or remove Q01-data-mnt00001
2. Create clone from P01 Snapshot copy
Via API or GUI
3. Start recovery of SYSTEMDB Q01
From Snapshot copy/without catalog/clear log
4. Start recovery of Tenant DB

These are the enterprise features that are required to allow enterprise applications such as SAP HANA to benefit most from the elasticity and flexible environment of a cloud deployment. The video [SAP System Refresh using Azure NetApp Files](#) shows the complete workflow.

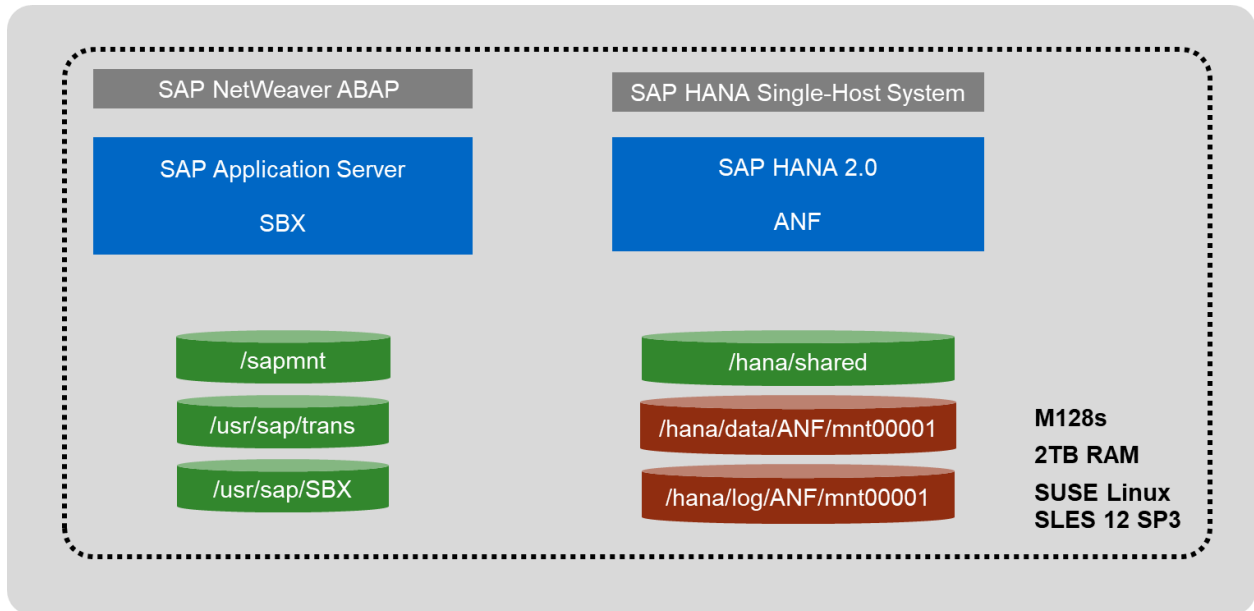
5 Example: Installation of an SAP NetWeaver ABAP Stack

The example described in this section runs a complete system installation of an SAP NetWeaver 7.5 ABAP server with System ID <SID> “SBX” together with a SAP HANA system 2.0 with <SID> “ANF” on a Microsoft Azure M128s instance type. This installation uses Azure NetApp Files for all the shared SAP file systems including the SAP HANA data and log volume.

Note: The use of Azure NetApp Files to provision shared file systems through NFS or SMB for any type of SAP system in a production environment is supported.

Figure 27 shows the installed components and required storage for SAP NetWeaver on SAP HANA.

Figure 27) SAP NetWeaver on SAP HANA.



Note: Most of the for SAP HANA configuration follows the traditional rules that apply for installing SAP HANA with on-premises ONTAP systems, as outlined in [TR-4435: SAP HANA on NetApp AFF Systems with NFS](#).

5.1 Installation Planning

To plan for the storage portion of the installation process, you must know the capacity and performance requirements and decide how many capacity pools need to be created.

File System Sizing and Volume Layout for the NetWeaver Instance

Use the information listed in Table 3 to create the file systems for SBX NFS.

Table 3) SBX NFS file systems.

Path	Size	Comment
/usr/sap/trans	100GB	Transport shared file system
/sapmnt/SBX	128GB	Shared executables for <SID> SBX
/usr/sap/SBX	128GB	Installation directory for <SID> SBX

As described in section 2.3, "Providing Shared Files using Azure NetApp Files," for all file systems with similar protection and performance requirements, NetApp recommends that you combine these file system into a single volume for ease of management and better overall performance and capacity planning, as shown in Table 4 .

In this example, a single volume, SBX-shared (with a size of 1TB and throughput of 64MBps), is created with three folders that function as mount points for the file systems listed in Table 4.

Table 4) Volume and folder structure.

Volume	Service Level	Size	Folder	Mount Point
SBX-shared	Premium	1TB	usr_sap	/usr/sap/SBX

Volume	Service Level	Size	Folder	Mount Point
			trans	/usr/sap/trans
			sapmnt	/sapmnt/SBX

File System Sizing and Storage Layout for the SAP HANA Instance

Use the information listed in Table 5 to create the file systems for the SAP HANA instance.

Table 5) SAP HANA basic file systems.

Path	SAP Recommended Size	Comment
/hana/data/ANF/mnt00001	2048GB	Data file mount point
/hana/log/ANF/mnt00001	512GB	Log file mount point
/hana/shared	1024GB	Hana shared
/usr/sap/ANF	128GB	Installation directory for SID ANF

Use the information in Table 6 to create the volumes and folder structures for the SAP HANA instance. For the data and log files, create an individual volume, where `/hana/shared` and `/usr/sap/ANF` is combined into a single volume. The service level and size chosen for the data and log volumes in this example meet the SAP defined KPIs for production workloads as described in chapter SAP HANA Azure NetApp Files Sizing.

Table 6) Volume and folder structure for SAP HANA.

Volume	Service Level	Size	Folder	Mount Point
ANF-data-mnt00001	Ultra	3.2TB	–	/hana/data/ANF/mnt00001
ANF-log-mnt00001	Ultra	2TB	–	/hana/log/ANF/mnt00001
ANF-shared	Premium	1.2TB	Shared	/hana/shared
			Usr_sap	/usr/sap/ANF

File Based Backup and Log Backup Volume

As described in section 2.1, "Shared File System for SAP NetWeaver Application Server," SAP HANA can benefit from using high performance NFS volumes for the file-based backups and the automated log backup. In a default configuration these backups are stored on:

- `/hana/shared/<SID>/HDB<nn>/backup/data/`
- `/hana/shared/<SID>/HDB<nn>/backup/log/`

Where `<SID>` is the system-ID and `<nn>` is the two-digit system number of the SAP HANA database.

SAP and NetApp highly recommend that you move these file-based backups as well as the log backups from the `/hana/shared` volume onto its own volume.

Note: Staying with the default configuration adds additional capacity requirements to `/hana/shared`.

Note: When you use storage Snapshot copies to back up the `/hana/shared` volume, it is counterproductive to keep the backups on this volume, since the Snapshot copies are Snapshot backups of the file backups as well.

Additional storage requirements depend largely on the customer's backup strategy using file-based backup, storage Snapshot-based backups, or facilitating SAPs BACKINT interface. Table 7 lists the storage requirements for the backup volumes to store two file-based backups and one week's worth of log backups assuming an estimated log rate of 30% of the database size.

Table 7) Backup volumes.

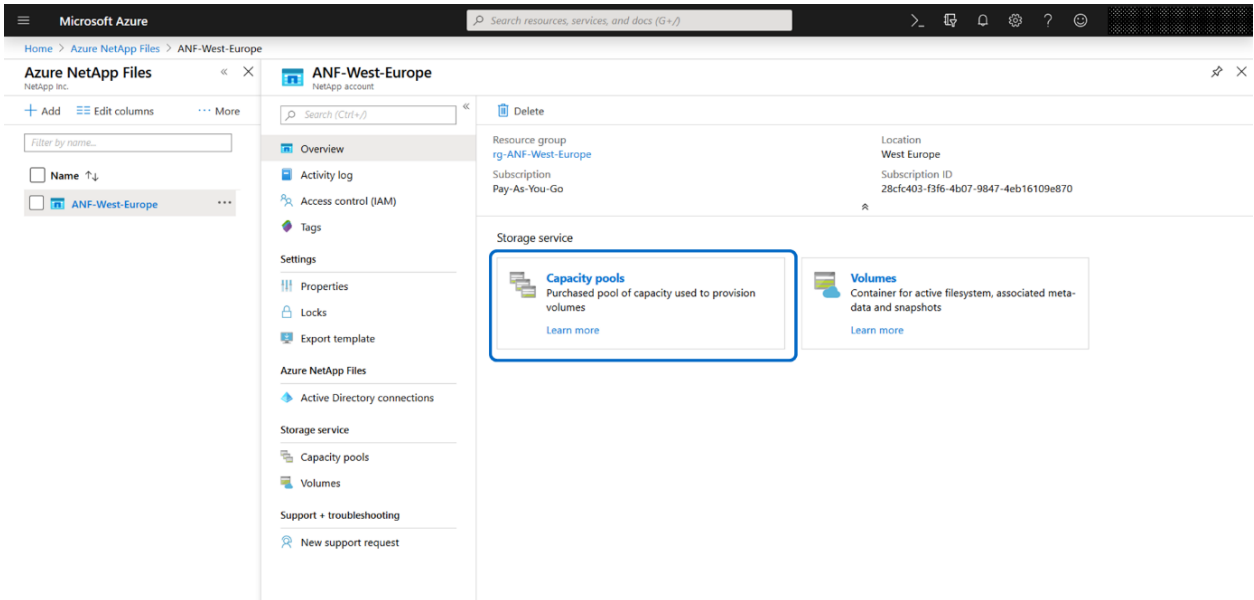
Volume	Service Level	Size	Mount Point
ANF-backup-file	Premium	4TB	e.g. /backup/data
ANF-backup-log	Premium	4TB	e.g. /backup/log

5.2 Prepare Storage on Azure NetApp Files

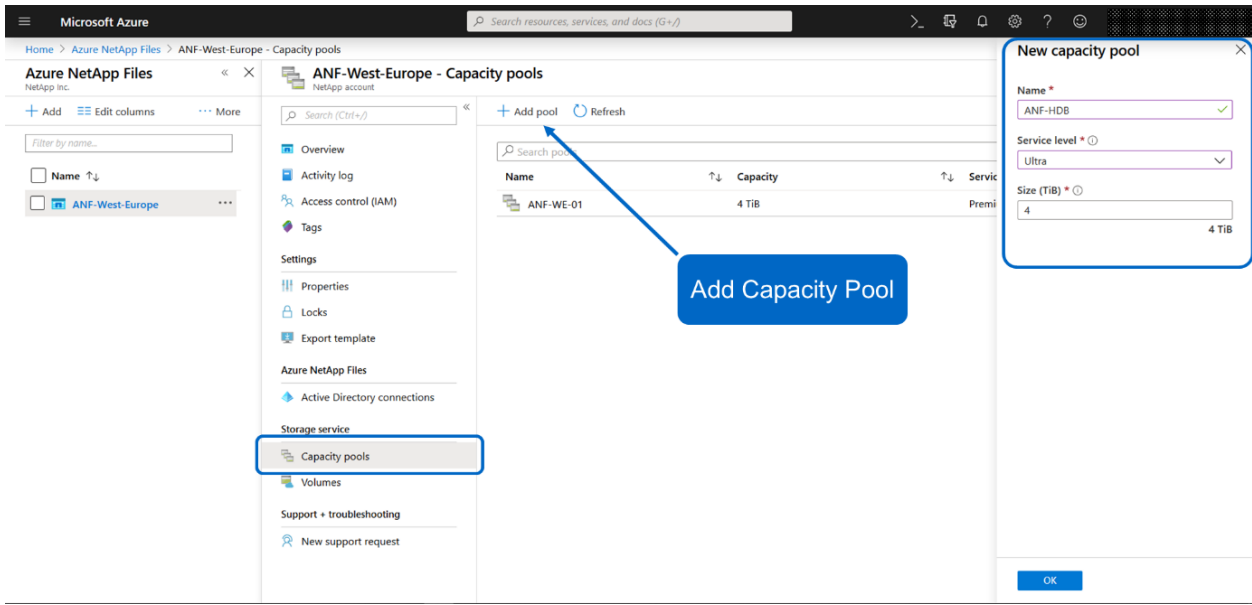
Create Capacity Pools

To create a capacity pool, complete the following steps:

1. From the Microsoft Azure portal, select Azure NetApp Files.
2. If necessary, create a NetApp account in the desired Azure region, for example, West Europe.
3. Click on the NetApp account.
4. On the next screen, click on Capacity pools.



5. Click Add Pool.



6. On the right side of the screen, enter a name for the capacity pool (in this example, ANF-HDB) and the size in terabytes (TB). The minimum capacity is 4TB, as service level we select ULTRA for the SAP HANA database files.
7. Click OK to create the capacity pool.
8. Repeat step 1–7 to create all required capacity pools. We are using one capacity pool per required service level in our example: ANF-NW (Premium), ANF-HDB (Ultra).

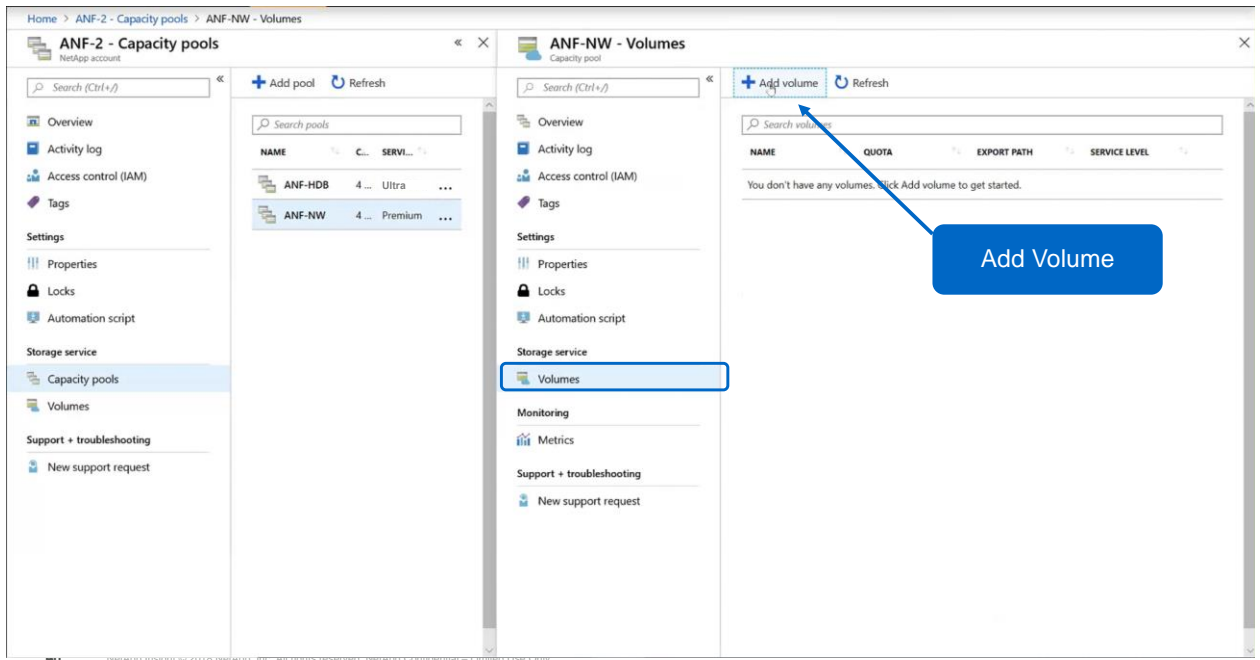
Create Storage Volume for SAP NetWeaver

To create a storage volume for SAP NetWeaver, complete the following steps:

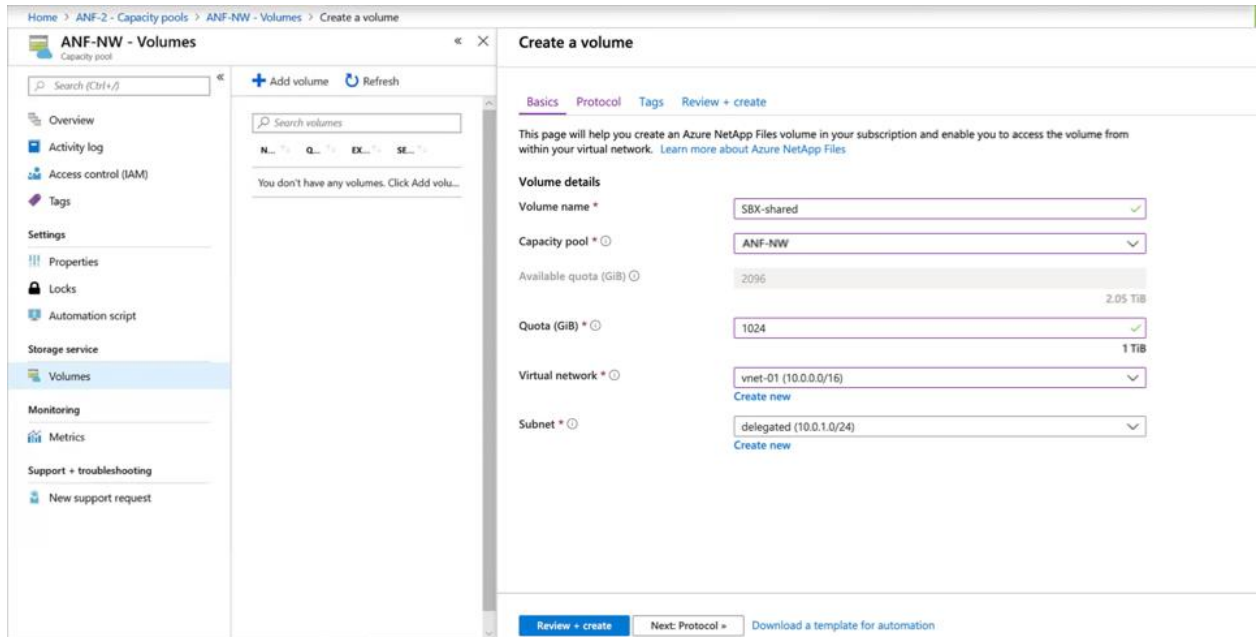
1. After both of the capacity pools have been created, select Capacity Pools > ANF-NW.

Note: This section provides options such as listing the volumes, monitoring and metrics, resizing the capacity pool, and so on.

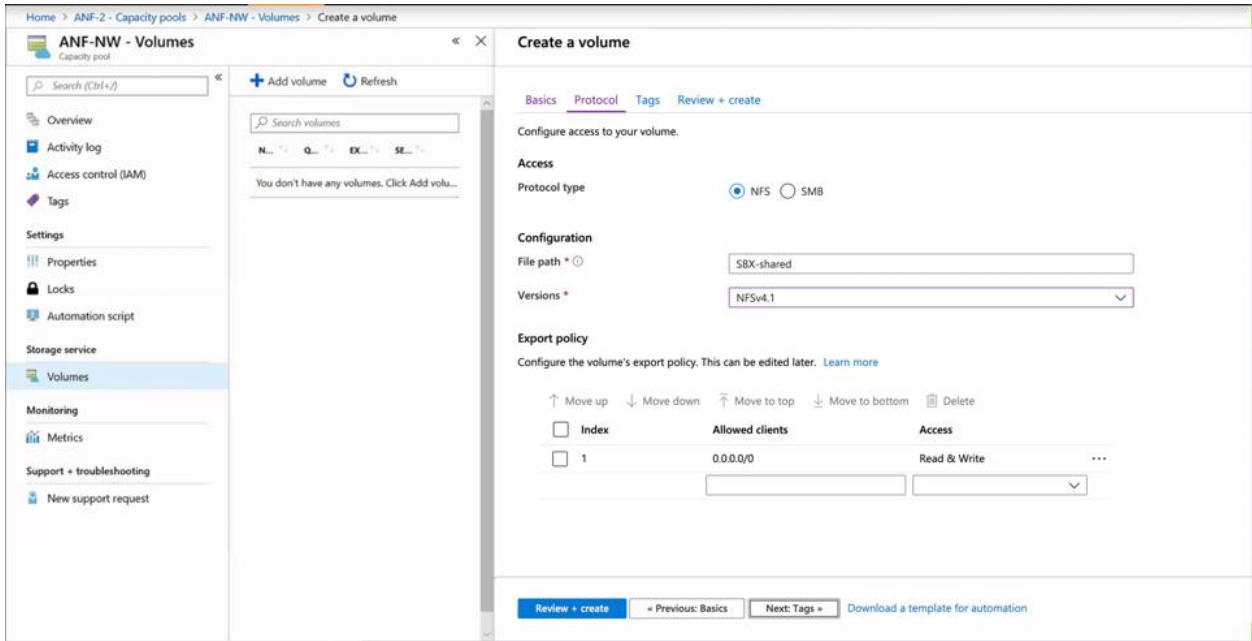
2. Select Volumes to show your current volumes, and then select Add Volume to create a new volume.



3. The Create a Volume wizard starts.



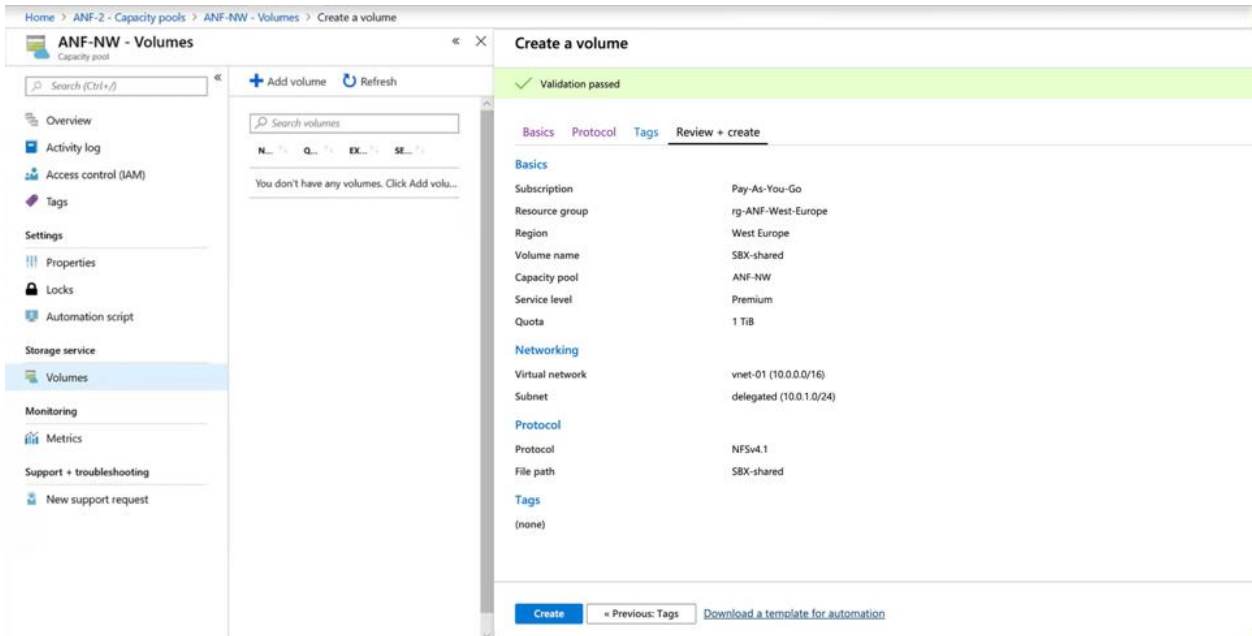
4. Specify the volume name, file path, and size (in gigabytes).
Specify the virtual network where your VMs belong. Azure NetApp Files requires you to specify a subnet within this virtual network where the storage IP addresses for the NFS exports are allocated. This is called a delegated subnet.
Note: For more information, see the [Delegate a Subnet to Azure NetApp Files](#) quick start guide.
5. To switch to the protocol selection, click Next and then Protocol.



6. Select the NFS protocol version NFSv4.1. You might also want to change the export policy. As a default, anyone will have read and write access.

Note: For SAP HANA data and log volume, only the protocol version NFSv4.1 is certified. You can use the protocol version NFSv3 for HANA shared and the SAP shared volumes; however, a common protocol was used in this setup.

7. Click Review + Create to start the entry validation.



8. After successful validation, click Create.

The volume list of the capacity pool shows the new volume.

Create Storage Volumes for the SAP HANA Database

To create a storage volume for the SAP HANA database, complete the following step:

1. Repeat steps 1–8 in the previous section for the following three volumes:
 - ANF-data-mnt00001
 - ANF-log-mnt00001
 - ANF-shared

For these volumes, use the ANF-HDB capacity pool instead of ANF-NW.

Note: Assign the delegated subnet only when you create the first volume in a capacity pool.

5.3 Prepare the Operating System and Mount the Volumes

To mount the required file systems to a host, complete these tasks:

1. Configure each client to support NFSv4.1.
2. Identify the mount options.
3. Mount the volume on a temporary mount point to create the relevant folders.
4. Create the required path structure on the OS level and set the permissions
5. Add the mount commands to `/etc/fstab` in order to make it persistent (this step is optional, but it is recommended).
6. Mount the volumes.

To install an SAP NetWeaver application server, SAP recommends that you use a specific SAP Linux edition. Linux editions, such as SLES for SAP, are preinstalled with an NFS client and many of the Linux settings required by SAP are preset.

Configure the Client OS to Support NFSv4.1 Mount

Clients that need to mount a NFSv4.1 volume must be mapped to the NFS domain that is configured for Azure NetApp Files. In the current version of Azure NetApp Files, the NFS domain is preconfigured to be `defaultv4iddomain.com`. Future versions might allow you to change this default value.

1. To adapt the Linux OS, edit `/etc/idmapd.conf` and change the parameter `domain`.

```
[General]
Verbosity = 0
Pipefs-Directory = /var/lib/nfs/rpc_pipefs
Domain = defaultv4iddomain.com

[Mapping]
Nobody-User = nobody
Nobody-Group = nobody
```

Note: If the domain ID mapping is not correct, the permission for files on Azure NetApp volumes that are mounted on the OS might be displayed as 'nobody'.

2. When using NFS the Linux operating system, create a system variable that displays if the local ID mapping is disabled (which should be the default). To verify this value, run the following command as root user:

```
cat /sys/module/nfs/parameters/nfs4_disable_idmapping
Y
```

3. If this file does not contain Y, you can change it by running the following command:

```
echo "Y" > /sys/module/nfs/parameters/nfs4_disable_idmapping
```

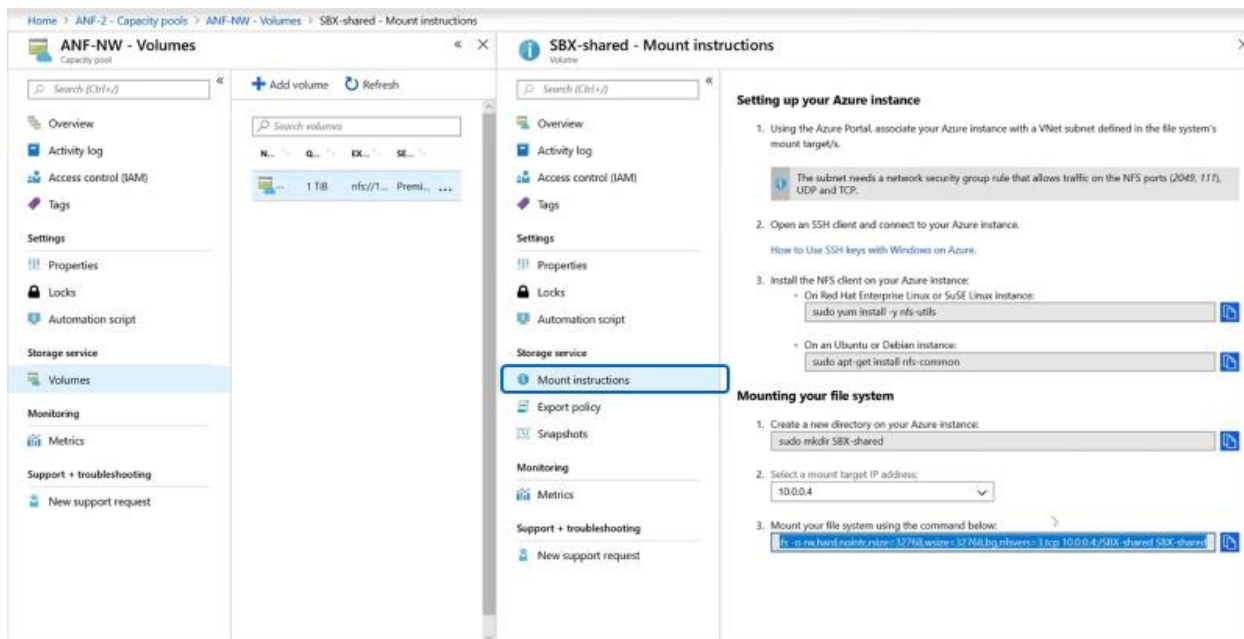
Note: You cannot create this file if it does not exist. Instead, temporarily mount an NFS volume to let the system create the file.

4. Reboot the system.

Prepare and Mount the Volume for the SAP NetWeaver Instance

The Azure portal for Azure NetApp Files provides valuable volume information such as instructions for customers to prepare the operating system and mount the volume to the host.

1. For mount instructions, select the volume. This opens the ANF Volumes view screen.
2. Select Mount Instructions to access the required instructions for preparing the operating system and mounting the volume.



3. Mount the volume over a temporary folder.

```
# sudo mount -t nfs -o  
rw,hard,intr,noatime,lock,rsize=65536,wsiz=65536,vers=4,minorversion=1,tcp,_netdev,sec=sys  
10.0.0.4:/SBX-shared /mnt
```

4. After the volume is successfully mounted, create the following folders and set the correct permissions.

```
# cd /mnt  
# mkdir usr_sap  
# mkdir trans  
# mkdir sapmnt  
# chmod 777 *  
# cd /  
# sudo umount /mnt
```

5. Create the mount points at the operating system level and set the permissions.

```
# mkdir -p /usr/sap/SBX  
# mkdir -p /usr/sap/trans  
# mkdir -p /sapmnt/SBX  
#  
# chmod -R 777 /sapmnt  
# chmod -R 777 /usr/sap
```

6. Add the entries into /etc/fstab via sudo /etc/fstab.

```
# cat /etc/fstab
...
10.0.0.4:/SBX-shared/usr_sap /usr/sap/SBX nfs
rw,hard,intr,noatime,lock,rsize=65536,wsiz=65536,vers=4,minorversion=1,tcp,_netdev,sec=sys 0 0
10.0.0.4:/SBX-shared/trans /usr/sap/trans nfs
rw,hard,intr,noatime,lock,rsize=65536,wsiz=65536,vers=4,minorversion=1,tcp,_netdev,sec=sys 0 0
10.0.0.4:/SBX-shared/sapmnt /sapmnt/SBX nfs
rw,hard,intr,noatime,lock,rsize=65536,wsiz=65536,vers=4,minorversion=1,tcp,_netdev,sec=sys 0 0
```

7. And mount the volumes.

```
# sudo mount -a
```

8. Make sure that all the file systems have been mounted correctly and are visible at the host level.

```
# df -h
Filesystem                Size      Used Avail Use% Mounted on
.....
10.0.0.4:/SBX-shared/usr_sap    100T    320T    100T   1% /usr/sap/SBX
10.0.0.4:/SBX-shared/trans      100T    320T    100T   1% /usr/sap/trans
10.0.0.4:/SBX-shared/sapmnt      100T    320T    100T   1% /sapmnt/SBX
```

With these preparations, the volumes should be mounted and the VM is now ready for the SAP installation to be run.

Prepare the Operating System and Mount the Volumes for the SAP HANA Instance

To efficiently run the SAP HANA additional Linux kernel, adapt the following parameters:

1. Prepare the operating system with the specific SAP HANA settings, as described in [TR-4435: SAP HANA on NetApp AFF Systems with NFS Configuration Guide](#).
 - a. Adapt the kernel settings for the operating system. In this example, it's SUSE SLES 12.

Create a configuration file `91-NetApp-Hana.cfg` in `/etc/sysctl.d/`.

```
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.core.rmem_default = 16777216
net.core.wmem_default = 16777216
net.core.optmem_max = 16777216
net.ipv4.tcp_rmem = 65536 16777216 16777216
net.ipv4.tcp_wmem = 65536 16777216 16777216
net.core.netdev_max_backlog = 300000
net.ipv4.tcp_slow_start_after_idle=0
net.ipv4.tcp_no_metrics_save = 1
net.ipv4.tcp_moderate_rcvbuf = 1
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
```

- b. Prepare additional setting recommended by Microsoft for Azure. Create a configuration file `/etc/sysctl.d/ms-az.conf` with the following settings:

```
vi /etc/sysctl.d/ms-az.conf
# Add the following entries in the configuration file
ipv6.conf.all.disable_ipv6 = 1
net.ipv4.tcp_max_syn_backlog = 16348
net.ipv4.ip_local_port_range = 40000 65300
net.ipv4.conf.all.rp_filter = 0
sunrpc.tcp_slot_table_entries = 128
vm.swappiness=10
```

- c. Adjust the `sunrpc.tcp_max_slot_table_entries` value to 128 in `/etc/modprobe.d/sunrpc.conf`

```
options sunrpc tcp_max_slot_table_entries=128
```

2. Adjust the NFS domain settings as described in the previous section “Prepare and Mount the Volume for the SAP NetWeaver Instance”.
3. Create the required subdirectories in the ANF-shared volume.

Note: The storage IP addresses are created automatically, and each volume might have a unique IP address that is shown in the volumes overview and mount instructions.

```
# sudo mount 10.0.0.5:/ANF-shared /mnt
# cd /mnt
# mkdir shared
# mkdir usr_sap
# cd /
# sudo umount /mnt
```

4. Create the mount points and set the permissions for the directories.

```
# mkdir -p /hana/data/ANF/mnt00001
# mkdir -p /hana/log/ANF/mnt00001
# mkdir -p /hana/shared
# mkdir -p /usr/sap/ANF
#
# chmod -R 777 /hana/log/ANF
# chmod -R 777 /hana/data/ANF
# chmod -R 777 /hana/shared
# chmod -R 777 /usr/sap/ANF
```

5. Edit `/etc/fstab` to allow for automatic mounting of the volumes upon reboot. The mount options are described in [TR-4435: SAP HANA on NetApp AFF Systems with NFS Configuration Guide](#).

```
# cat /etc/fstab
...
10.0.0.5:/ANF-shared/usr_sap /usr/sap/ANF nfs
rw,hard,intr,noatime,lock,rsize=262144,wsiz=262144,vers=4,minorversion=1,tcp,_netdev,sec=sys 0 0
10.0.0.5:/ANF-shared/shared /hana/shared nfs
rw,hard,intr,noatime,lock,rsize=262144,wsiz=262144,vers=4,minorversion=1,tcp,_netdev,sec=sys 0 0
10.0.0.5:/ANF-data-mnt00001 /hana/data/ANF/mnt00001 nfs
rw,hard,intr,noatime,lock,rsize=262144,wsiz=262144,vers=4,minorversion=1,tcp,_netdev,sec=sys 0 0
10.0.0.5:/ANF-log-mnt00001 /hana/log/ANF/mnt00001 nfs
rw,hard,intr,noatime,lock,rsize=262144,wsiz=262144,vers=4,minorversion=1,tcp,_netdev,sec=sys 0 0
Mount the additional volumes with sudo mount -a.
```

Post-Installation Configuration

As described in the planning section (section 5.1), NetApp recommends that you create optional volumes for the file-based backup and automated log backup. After a successful installation, create and mount the volumes ANF-backup-data and ANF-backup-log to the default backup locations that are specified in Table 7 and start the first file-based backup to change the SAP HANA database to a standard logging mode. For example, you need to adapt these changes by using the backup console of SAP HANA Studio.

To optimize SAP HANA for the underlying Azure NetApp Files storage, set the following SAP HANA parameters:

```
max_parallel_io_requests 128
async_read_submit on
async_write_submit_active on
async_write_submit_blocks all
```

For SAP HANA 1.0 versions up to SPS12, these parameters can be set during the installation of the SAP HANA database, as described in [SAP note #2267798 - Configuration of the SAP HANA database](#).

The parameters can also be configured after the SAP HANA database installation by using the hdbparam framework or global.ini. For more information, see SAP note [#2399079](#).

6 Example: Cost Optimized Deployments

As described in the previous chapters, Azure NetApp Files offers a perfect solution to deploy enterprise-grade storage for SAP shared files as well as volumes to support the performance requirements for SAP databases and log data. Customers following the best practice guides described in this document will benefit from the ease of deployment, superior performance, and simplified operations that can help to speed up your SAP projects and reduce risks through automated procedures.

Nevertheless, customers are asking how they can further optimize the use of Azure NetApp Files for their SAP projects in case they want to set up smaller SAP HANA development or test systems that might not require you to configure the storage for the demanding SAP HANA KPIs. Instead, the question that arises is how to optimize the costs but still keep the benefits of using Azure NetApp Files. In order to answer these questions and provide a guideline for customers, one has to understand the architectural concepts and the capacity and performance implications.

6.1 Know Your Limits

- Azure NetApp Files performance is defined by two factors: the service level to which the capacity pool is assigned and the size of the volume. Both factors can be used to calculate the maximum throughput in MBps for the volume.
- The volume size also defines the storage capacity of the volume, but this is only a soft limit. The actual consumption can grow beyond size of the volume, as Azure NetApp Files will not reject writes. Volumes will auto grow when consumption is more than the set size for a certain amount of time.
- New volumes can only be created when sufficient free (unassigned) capacity is available in the capacity pool. There is no overcommit possible. In case more volumes are needed or volumes need to be resized beyond the total size, the capacity pool first needs to be increased in size.
- The volume size limits are from 100GB up to 100TB. The single file limit is 16TB.
- To achieve the SAP HANA KPI limits for the data and log volume, customers need to allocate the appropriate sizes, typically from a capacity pool of Ultra service levels:
 - Data volume: 3.2TB Ultra
 - Log volume: 2TB Ultra

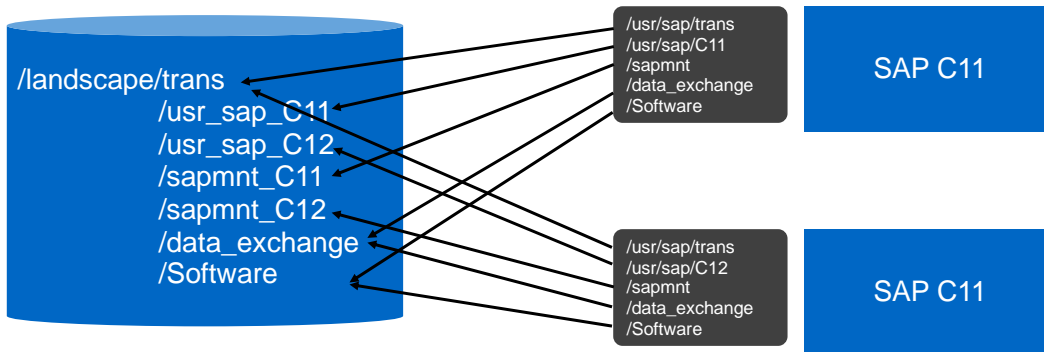
Note: This implies that for smaller SAP HANA systems, volumes might have to be overprovisioned.

Optimization for Shared Files Data

In the previous sections, some tuning options were recommended. Instead of using individual file systems (volumes) for `/usr/sap/trans`, `/usr/sap/SID` or `/hana/shared`, a common practice is to create only one volume for all shared files of a single system and create a folder structure within that volume that can be used to mount individual folders to the SAP defined mount points.

Instead of many volumes, each with a minimum of 100GB, only one larger volume needs to be created. This configuration not only simplifies maintenance and reduces storage overprovisioning, but it also provides more throughput for shared files with the increased overall size of the volume, as shown in Figure 28.

Figure 28) Cost optimization for shared files – single volume directories to client mountpoints mapping.



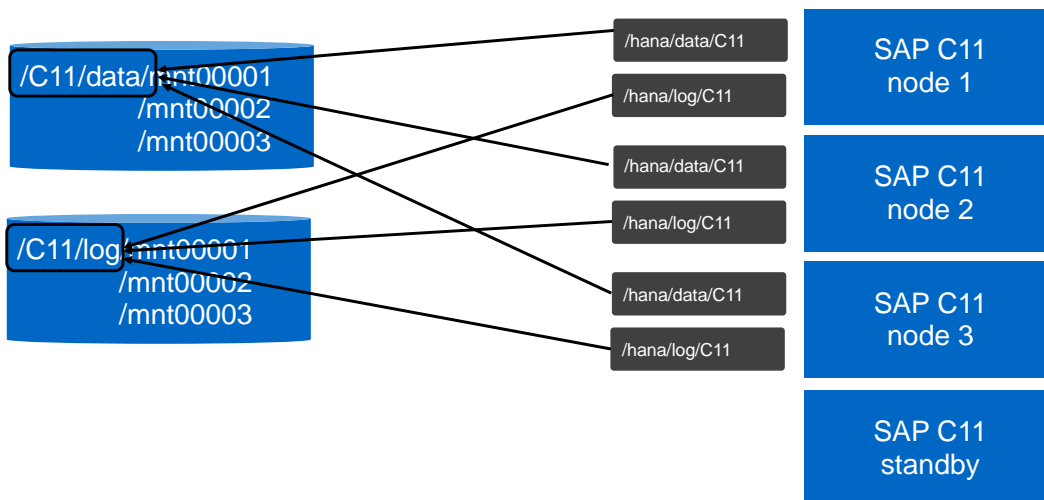
For complete system landscapes, consider grouping not only the shared files of a single system into one volume, but instead combining all shared files of a system landscape into a single volume. This will further optimize the cost performance ratio; however, it has an impact on the manageability. Using Snapshot copies for data protection always protects the shared files of all included systems, which can be an advantage; however, consider these factors:

- In case of a copy of a single system using Azure NetApp Files fast cloning technology, the cloned volume includes also all of these systems. This is not always a desired outcome as it might cause duplicate storage consumption.
- To avoid restoring one of the other systems unintentionally, a restore operation might need to be performed by using the `.snapshot` directory to manually copy the data from a Snapshot copy to the active file system. This process requires more copy transfer overhead and operator attention to avoid unintentional system restores.

6.2 Optimization for Database Files

To prevent losing the benefits of Snapshot-based data protection, NetApp recommends that you not combine data and log data into the same volume. Instead, consider combining data of different nodes of a scale-out SAP HANA database, as shown in Figure 29.

Figure 29) Cost optimization for scale-out – volume directories to client mountpoints mapping.



Structuring the volume layout shown in Figure 25 does not impact data protection because only data volumes from a single database instance are grouped together. However, it might have impact on performance, despite the cost benefit.

- Instead of deploying three volumes each with 3.2TB volume size of service level Ultra, the capacity when using a single volume can be reduced. This limits the overall throughput, but often for smaller test systems, the resulting throughput might be sufficient.
- Depending on the data and I/O distribution across the three data files (`mnt00001`, `mnt00002`, `mnt00003`), or in case of databases with multiple tenants where each tenant might store its data only on a single data file, this results in different I/O patterns for each of the individual tenants. In case not all tenants add load to the system at the same time, a combined larger data volume can still provide the same or better throughput for a tenant than smaller individual data volumes.

Note: It is important to understand that the performance impact depends on the usage pattern and configuration of the individual system, and is, therefore, not recommended for larger production systems.

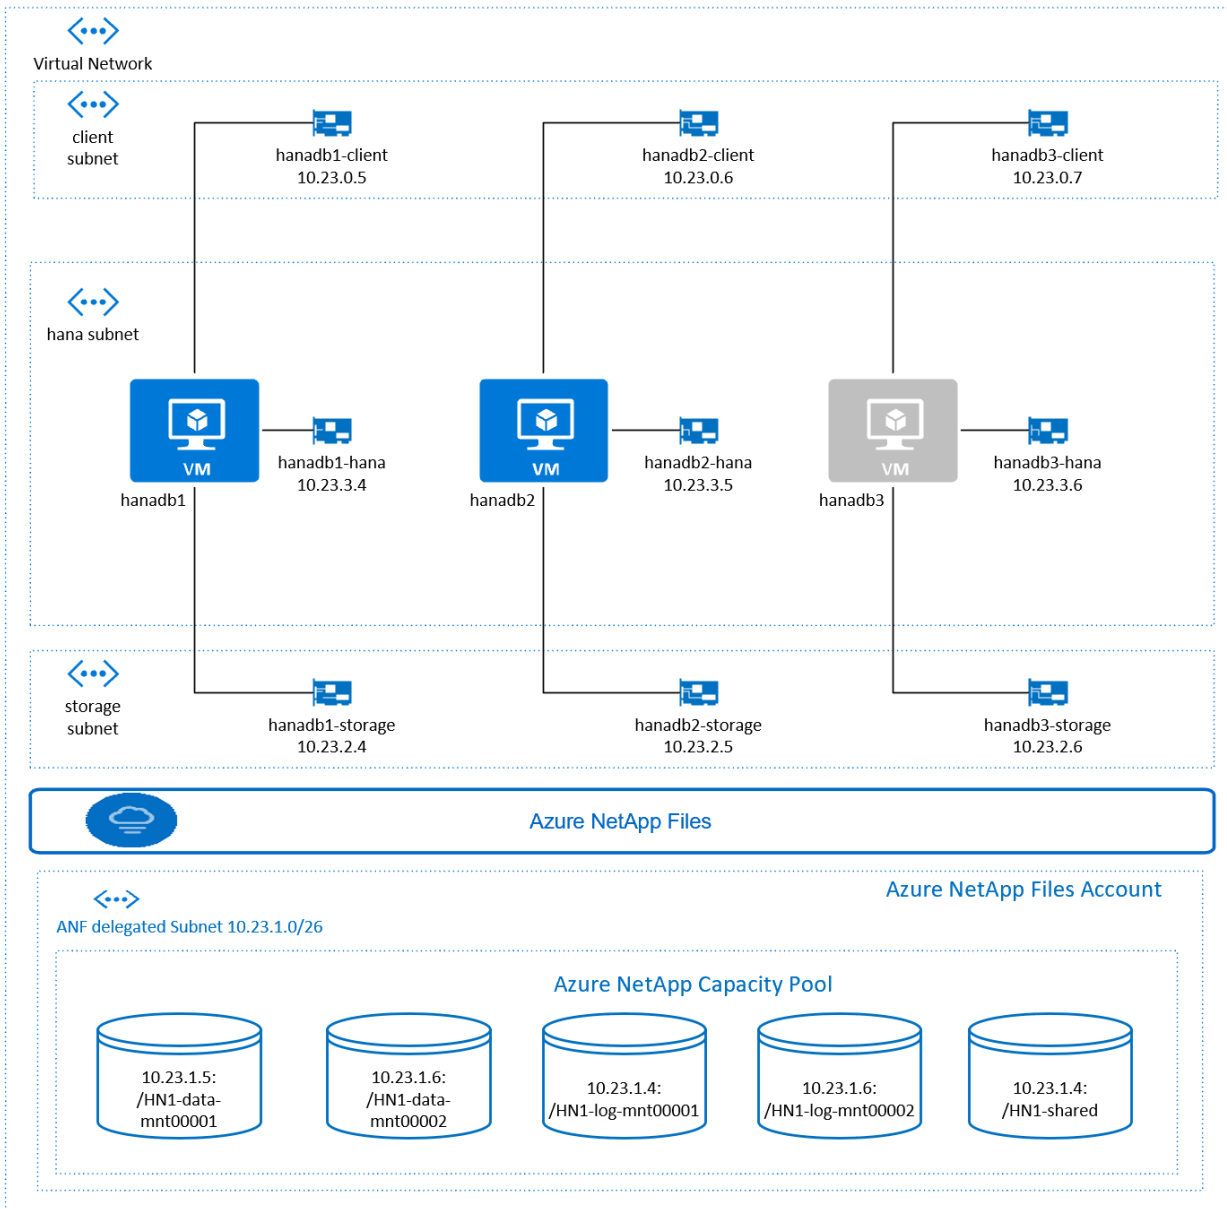
6.3 Optimization for Multiple SIDs

Combining data or log volumes from multiple independent SIDs into a single volume could be another option for reducing costs. However, this option has an impact on data protection and the ability to use Snapshot backups for restores or system copies, especially volume-level restores. This cost-reduction option should only be used for smaller development or test systems and by customers who are familiar with the concepts and architecture of Azure NetApp Files. This optimization option is not recommended for production environments.

7 Example: Installation of SAP HANA Scale-Out with Host Auto Failover

With Azure NetApp Files Azure, customers can deploy SAP HANA scale-out databases with standby nodes, which allows SAP HANA built-in cluster technology to automatically activate a standby node in the event of a failure of a database node. The ability to use SAP HANA with active standby node is a unique feature for cloud deployments where Azure NetApp Files NFSv4.1 locking mechanism enables these SAP HANA feature, as shown in Figure 30.

Figure 30) Typical SAP HANA scale-out with host auto failover setup.



Microsoft has provided detailed step-by-step guides on how to setup a SAP HANA scale-out deployment:

- [Deploy a SAP HANA scale-out system with standby node on Azure VMs by using Azure NetApp Files on SUSE Linux Enterprise Server](#)
- [Deploy a SAP HANA scale-out system with standby node on Azure VMs by using Azure NetApp Files on Red Hat Enterprise Linux Server](#)

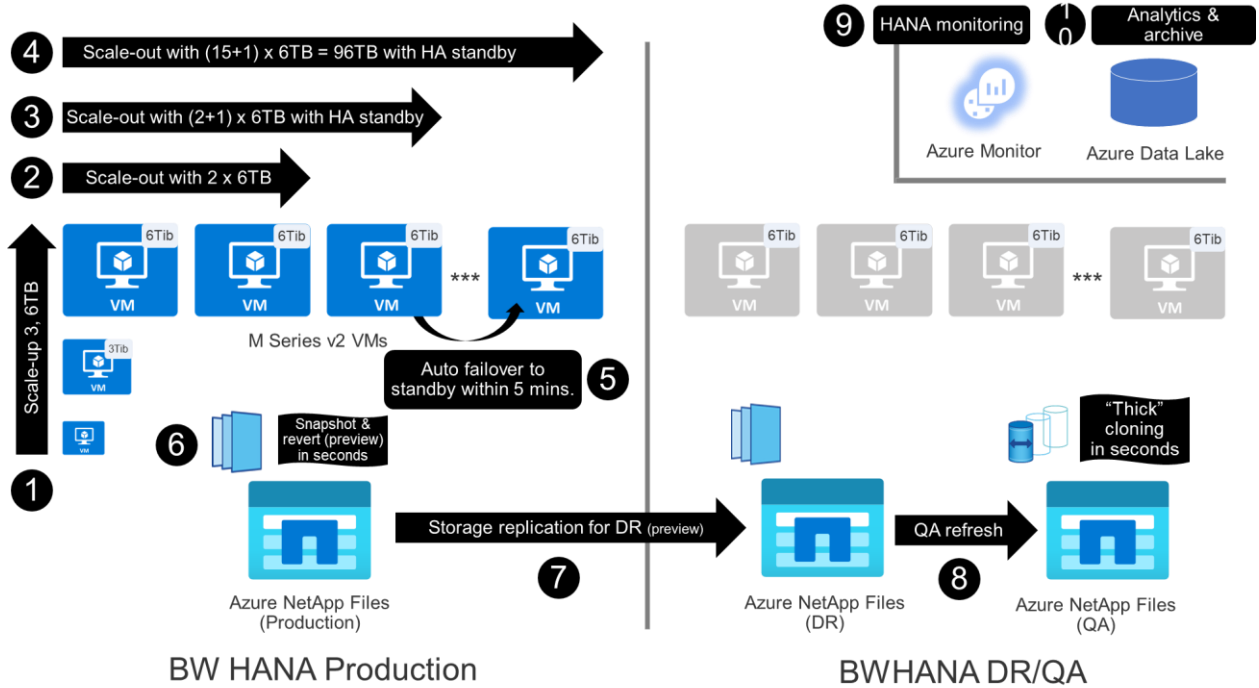
8 SAP BW HANA Landscape with Azure NetApp Files

This section describes the overall SAP BW/4 HANA deployment with various scale-up, scale-out and disaster recovery options in Azure. This deployment leverages all the technological options that have been described throughout this document. This will further help gain insight into the advanced possibilities

for HANA landscapes in Azure with the use of Azure NetApp Files. Most of this is equally applicable to non-BW environments as well (Figure 31).

8.1 Overview of SAP BW HANA Landscape with Azure NetApp Files

Figure 31) SAP BW HANA landscape with Azure NetApp files.



With Azure NetApp Files, you can obtain a series of possibilities and improvements for a SAP HANA landscape:

1. A typical BW HANA implementation can start with the smallest Mv2 VM with 3TB of RAM. Then the initially selected VM can be scaled-up to 6TB of RAM as the project proceeds and moves into production.
2. BW HANA can now also scale-out by adding one more 6TB VM using with NFSv4.1 on Azure NetApp Files as the underlying storage such that scale can double.
3. BW HANA scale-out with hot standby can be pursued by adding more VMs, fully certified and supported by SAP on Azure, also leveraging Azure NetApp Files' NFSv4.1.
4. Further highly available scale-out N+M HANA with up to 16 VMs and 96TB of total RAM is now fully supported.
Note: N is the number of active VMs; M is the number of standby VMs.
5. In the event that an outage where an active VM is automatically detected, automatic failover to standby VM completes within five minutes,.
6. An application-consistent Snapshot copy of a large HANA database or entire landscape can be taken in seconds, regardless of the database size. Point-in-time restores can be achieved in seconds as well.
7. Snapshot copies for HANA database files and log backup files at production region are replicated to the disaster recovery region without impacting production system performance, and without the need to run any VMs in the disaster recovery region.

Note: At the time of writing the Azure NetApp Files Cross Region, the replication feature was in private preview

8. Replicated volumes can be restored instantly from Snapshot copies to QA and nonproduction environments in seconds, without interrupting cross-region replication. This capability can also be used for disaster recovery testing without interrupting cross-region data protection.
9. Azure Log Analytics and Monitoring can monitor details at SAP HANA level as well as Azure infrastructure level.
10. Integration with Azure Data Lake further allows for innovations around reporting with PowerBI and IoT and Machine Learning services in Azure.

Where to Find Additional Information

To learn more about the information described in this document, refer to the following documents and/or websites:

- Certified and Supported SAP HANA Hardware Directory
<https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/iaas.html#categories=Microsoft%20Azure>
- NetApp Cloud Central
<https://cloud.netapp.com/home>
- SAP on Azure Getting Started
<https://docs.microsoft.com/en-us/azure/virtual-machines/workloads/sap/get-started>
- SAP on Azure Landing Page
<https://azure.microsoft.com/en-us/solutions/sap/>
- SAP HANA on NetApp AFF Systems with NFS Configuration Guide
<https://www.netapp.com/us/media/tr-4435.pdf>
- SAP HANA Storage Requirements
<https://www.sap.com/documents/2015/03/74cdb554-5a7c-0010-82c7-eda71af511fa.html>
- SAP Application solutions with Azure NetApp Files
<https://docs.microsoft.com/en-us/azure/azure-netapp-files/azure-netapp-files-solution-architectures#sap-application-solutions>
- SAP Applications on Microsoft Azure Using NetApp Cloud Volumes ONTAP
<https://www.netapp.com/us/media/tr-4757.pdf>

Version History

Version	Date	Document Version History
Version 1.0	February 2019	Initial release.
Version 2.0	March 2020	Revision update describing NFSv4.1, scale-out HANA, and SAP HANA data protection, cloning and on-prem to Azure migration

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 © 2019–2020 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.

TR-4746-0320