



Technical Report

SAP Applications on Microsoft Azure

Using NetApp Cloud Volumes ONTAP

Bernd Herth, Tobias Brandl, Nils Bauer, NetApp
Will Bratton, Geert van Teylingen, Microsoft
April 2019 | TR-4757

In partnership with



Abstract

This document provides best practices for leveraging NetApp® Cloud Volumes ONTAP® for SAP Applications and SAP HANA deployments. It includes different use cases from SAP shared file systems and specific performance considerations with SAP HANA on Cloud Volumes ONTAP.

TABLE OF CONTENTS

1	SAP on Microsoft Azure Overview	4
1.1	NetApp Values and Solutions on Microsoft Azure	4
1.2	Microsoft Certified Solutions for SAP	7
1.3	Storage Requirements	10
1.4	SAP Use Cases	11
2	SAP Shared File Systems	12
2.1	Shared File System for SAP NetWeaver Application Server	12
2.2	Legacy Shared Files Solutions for the Cloud	13
2.3	Providing Shared Files using Cloud Volumes ONTAP	14
2.4	Data Protection for SAP Shared Files	17
3	SAP anyDB	18
3.1	Storage Requirements for SAP anyDB Databases	19
3.2	Cloud Native Solutions	20
3.3	SAP anyDB on Cloud Volumes ONTAP	22
4	SAP HANA	24
4.1	Storage Requirements for SAP HANA	24
4.2	Certified Microsoft Azure Native Solutions	26
5	Example: Installation of SAP NetWeaver ABAP Stack Using SAP HANA on Cloud Volumes ONTAP	27
5.1	Installation Planning	28
5.2	Create and Deploy OnCommand Cloud Manager	29
5.3	Create and Configure Cloud Volumes ONTAP	36
5.4	Configure and Prepare the Required Storage	39
5.5	Prepare the Operating System and Mount the Volumes	44
5.6	SAP Installation	49
6	Data Protection and System Migration	49
6.1	Data Protection	49
6.2	System Migration	50
	Where to Find Additional Information	54
	Version History	55

LIST OF TABLES

Table 1) Azure SAP certified VMs.	7
Table 2) Microsoft Azure SAP certified HANA (large instances).	9
Table 3) NW5 NFS file systems.	28
Table 4) Volume and folder structure.	28
Table 5) SAP HANA basic volumes and file systems.	29
Table 6) Cloud Volumes ONTAP storage configuration.	29
Table 7) SAP HANA optional file system.	49

LIST OF FIGURES

Figure 1) Cloud Volumes ONTAP.	5
Figure 2) Microsoft Azure NetApp files.	6
Figure 3) NetApp ONTAP offerings.	7
Figure 4) SAP storage performance requirements and certification.	11
Figure 5) Shared files requirements.	12
Figure 6) Shared files for SAP.	13
Figure 7) Linux cluster NFS server.	14
Figure 8) Cloud Volumes ONTAP shared files for SAP.	15
Figure 9) Cloud Volumes ONTAP using Microsoft Azure resources.	16
Figure 10) Cloud Volumes ONTAP.	17
Figure 11) SnapCenter: Snapshot-based data protection.	18
Figure 12) Data protection.	20
Figure 13) Storage for an SAP database server.	21
Figure 14) Cloud Volumes ONTAP for SAP anyDB.	23
Figure 15) HANA single-host system.	24
Figure 16) SAP HANA multiple-host system.	25
Figure 17) TDI performance criteria.	26
Figure 18) SAP NetWeaver on SAP HANA using Cloud Volumes ONTAP.	27
Figure 19) SAP migration using SnapMirror technology.	51
Figure 20) SAP migration using Cloud Sync.	52
Figure 21) Migration automation using SAP LaMa.	54

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 60TB multiple-host configurations. In 2018, Microsoft introduced the Azure M-Series VMs with up to 4TB of memory; VMs with 12TB of memory are coming soon. 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.

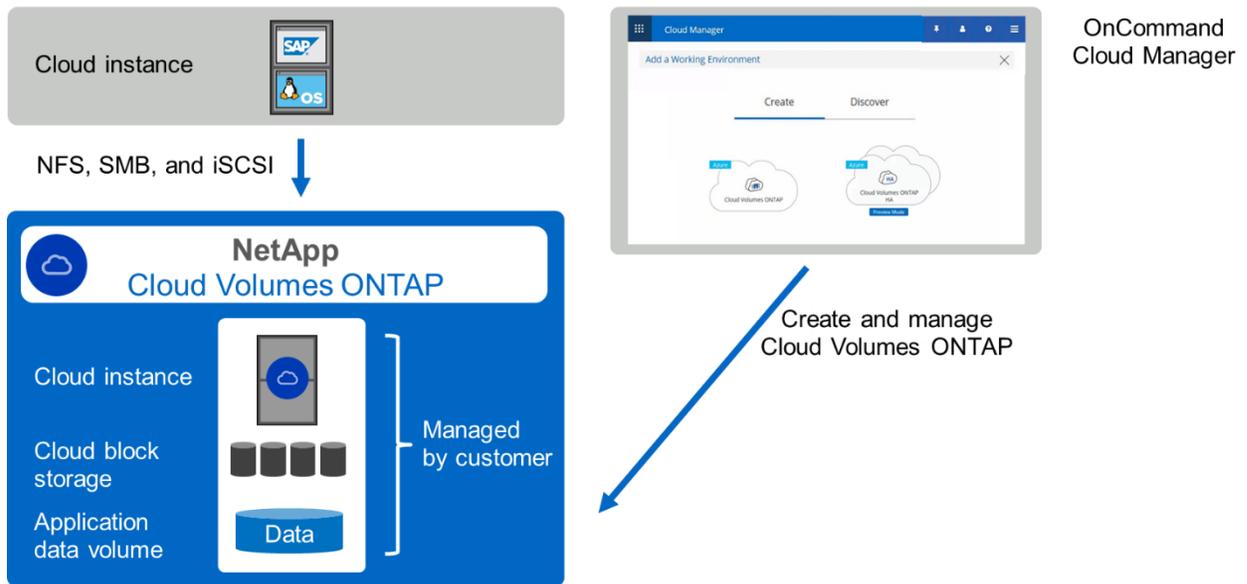
Many SAP customers want to move their SAP systems to the cloud don't want to relinquish the many NetApp benefits for their SAP projects and operations. Customers do not want to give up on the performance, reliability, and enterprise data management capabilities when they move these enterprise file-based 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 Cloud Volumes ONTAP only. For more information, see [TR-4746: SAP Applications on Microsoft Azure Using Azure NetApp Files](#).

Cloud Volumes ONTAP on Microsoft 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 hosted 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 install a small appliance NetApp OnCommand® Cloud Manager, as shown in Figure 1. OnCommand Cloud Manager 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 in their on-premises data centers.

Azure NetApp Files

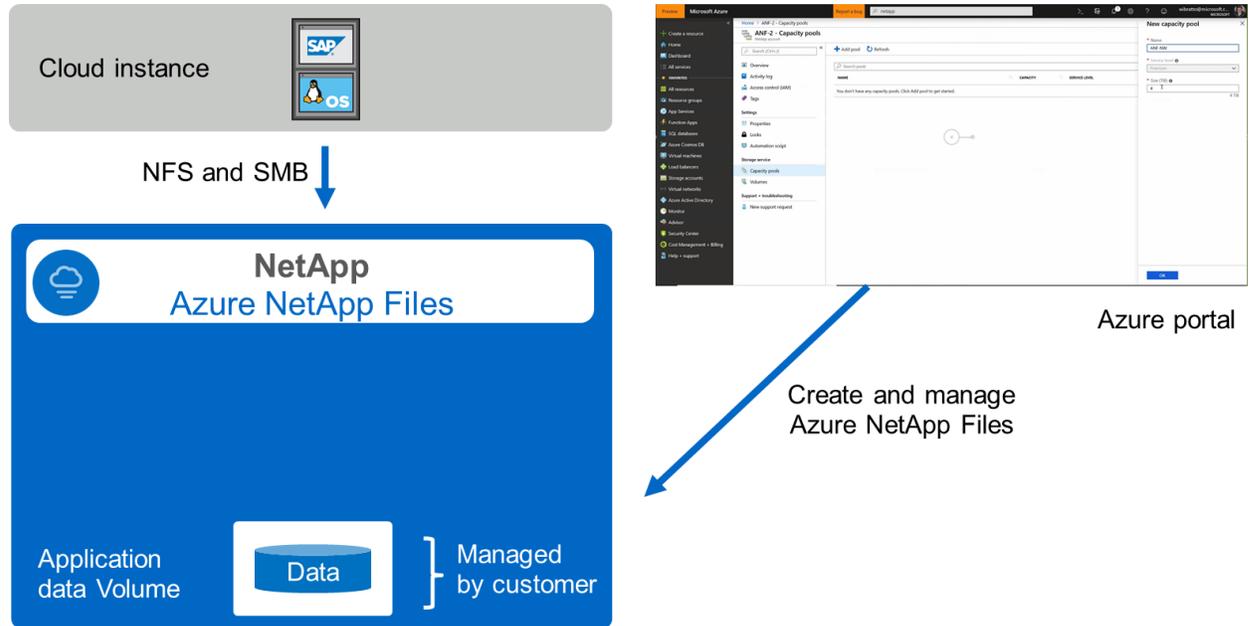
For the first time, Microsoft delivers an Azure native, first-party portal service for enterprise NFS or SMB file services based on NetApp ONTAP technology. This new development is driven by a strategic partnership between NetApp and Microsoft further extends the reach of NetApp world-class data services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

This Azure cloud-native data service delivers high performance, 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 object, as shown in Figure 3. 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 NetApp proven ONTAP storage hardware set 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.

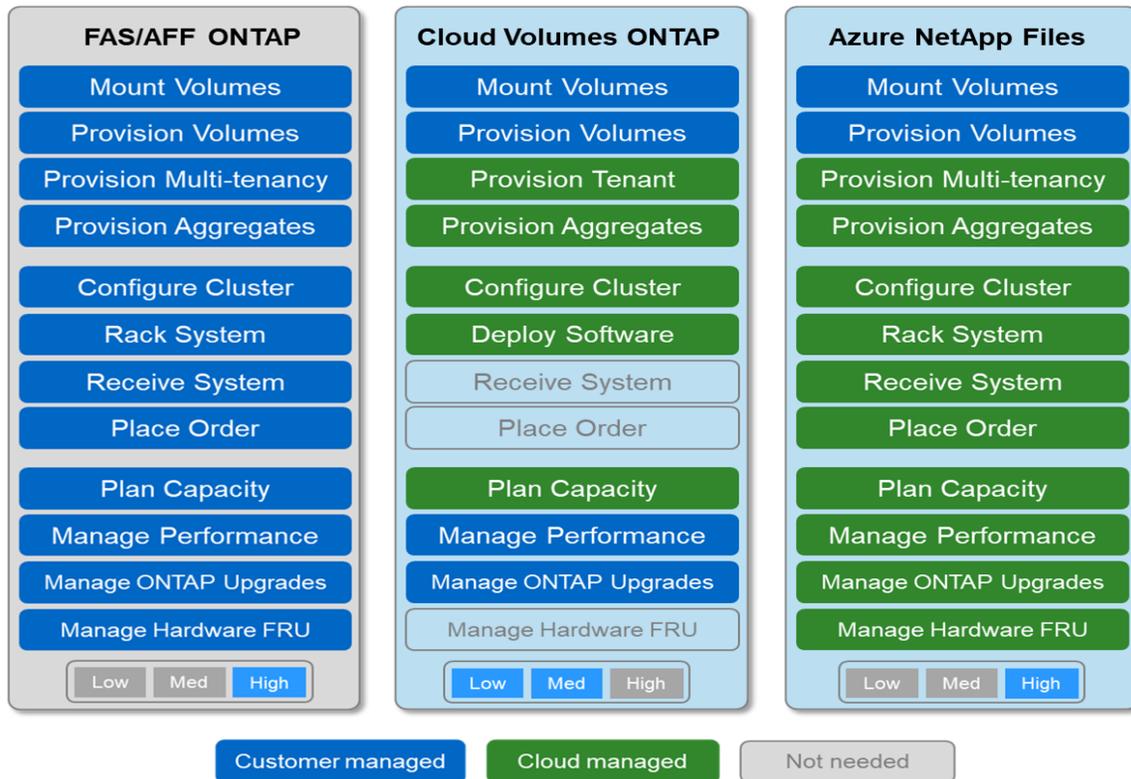
Figure 2) Microsoft Azure NetApp files.



Comparison

Figure 3 shows the different management levels for the various ONTAP-based systems, ranging from on-premises (FAS/AFF ONTAP) to infrastructure as a service (IaaS) (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.

Figure 3) NetApp ONTAP offerings.



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 (large instances) listed in Table 1 are certified for running SAP workloads.

Table 1) Azure SAP certified VMs.

VM Series	VM Type	VM Size	SAPS	Supported HANA Scenarios	SAP Certification
DS v2	DS11_v2	2 vCPU, 14 GiB	3,530	–	Certified (anyDB, App)
	DS12_v2	4 vCPU, 28 GiB	6,680	–	
	DS13_v2	8 vCPU, 56 GiB	12,300	–	
	DS14_v2	16 vCPU, 112 GiB	24,180	–	
	DS15_v2	20 vCPU, 140 GiB	30,430	–	

VM Series	VM Type	VM Size	SAPS	Supported HANA Scenarios	SAP Certification
GS	GS4	16 vCPU, 224 GiB	22,680	–	Certified (anyDB, App)
	GS5	32 vCPU, 448 GiB	41,670	OLAP, S4	GS5 certified for OLAP and controlled availability for S4
DS v3	D2s_v3	2 vCPU, 8 GiB	2,178	–	Certified (anyDB, App)
	D4s_v3	4 vCPU, 16 GiB	4,355	–	
	D8s_v3	8 vCPU, 32 GiB	8,710	–	
	D16s_v3	16 vCPU, 64 GiB	17,420	–	
	D32s_v3	32 vCPU, 128 GiB	34,840	–	
	D64s_v3	64 vCPU, 256 GiB	69,680	–	
ES v3	E2s_v3	2 vCPU, 16 GiB	2,178	–	Certified (anyDB, App)
	E4s_v3	4 vCPU, 32 GiB	4,355	–	
	E8s_v3	8 vCPU, 64 GiB	8,710	–	
	E16s_v3	16 vCPU, 128 GiB	17,420	OLTP, OLAP	HANA certification in roadmap App, anyDB certified
	E32s_v3	32 vCPU, 256 GiB	34,840	OLTP, OLAP	
	E64s_v3	64 vCPU, 432 GiB	70,050	OLTP/OLAP	
M	M32ts	32 vCPU, 192 GiB	33,670	OLTP	HANA, App, anyDB certified
	M32ls	32 vCPU, 256 GiB	33,300	OLTP	
	M64ls	64 vCPU, 512 GiB	66,600	OLTP	
	M64s	64 vCPU, 1,024 GiB	67,315	OLTP, OLAP	
	M64ms	64 vCPU, 1,792 GiB	68,930	OLTP	

VM Series	VM Type	VM Size	SAPS	Supported HANA Scenarios	SAP Certification
	M128s	128 vCPU, 2,048 GiB	134,630	OLTP, OLAP	
	M128ms	128 vCPU, 3,800 GiB	134,630	OLTP	

Table 2) Microsoft Azure SAP certified HANA (large instances).

HANA Large Instance	Threads	RAM	Storage	Workload
S96	96	768GB	3TB	OLAP, OLTP
S96m	96	1.5TB	6TB	OLTP
S192	192	2TB	8TB	OLAP, OLTP
S192m	192	4TB	16TB	OLTP
S192xm	192	6TB	16TB	OLTP TDiv5
S384	384	4TB	9,6TB	OLAP, OLTP
S384m	384	6TB	14,4TB	OLTP
S384xm	384	8TB	19,2TB	OLTP
S384xxm	384	12TB	28TB	OLTP TDiv5
S576m	576	12TB	28,8TB	OLTP
S576xm	576	18TB	41TB	OLTP TDiv5
S768m	768	16TB	38TB	OLTP
S768xm	768	24TB	56TB	OLTP TDiv5
S960m	960	20TB	48TB	OLTP

1.3 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

File shares require companies to comply with industry specific data security, data protection, and disaster recovery SLAs and regulations. NetApp 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 file shares are 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 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) 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 4.

Figure 4) SAP storage performance requirements and certification.

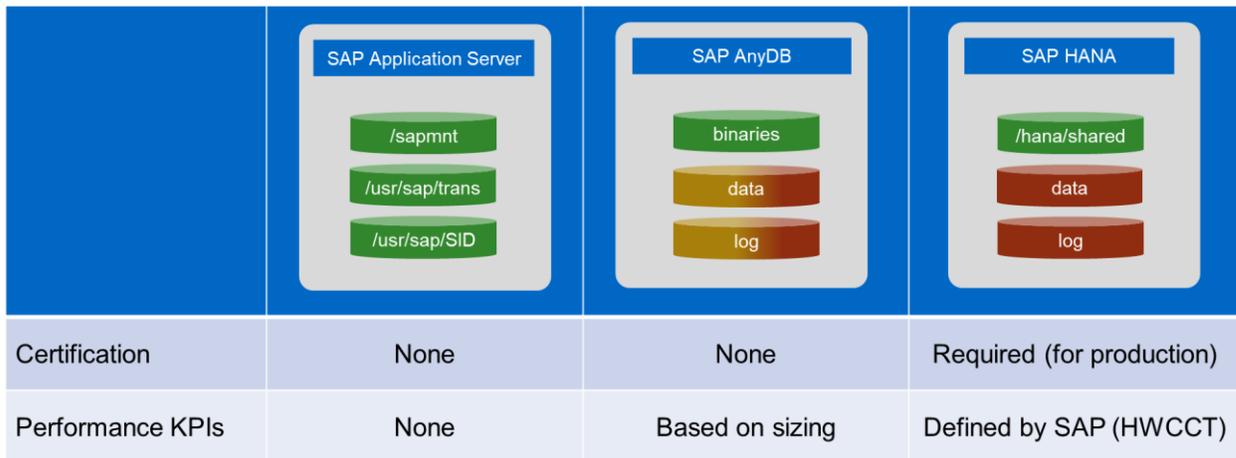


Figure 4 shows typical file systems that are used by a SAP application server, an SAP/AnyDB database server (such as Oracle, SQL, DB2, ASE, and MaxDB), and a SAP HANA server. The colors indicate the typical performance requirement for these types of file systems. Showing a low performance requirement 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 a SAP system in production, an important thing to remember 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. The SAP HANA certification includes the host and storage on which SAP HANA is running in production. Additionally, the certification includes a performance test 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](#).

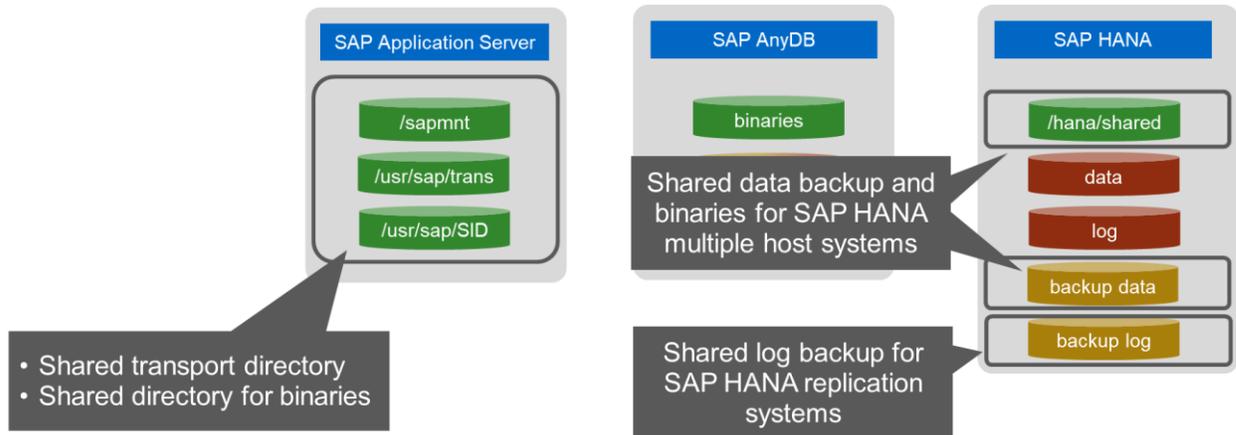
1.4 SAP Use Cases

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

- SAP shared files
- SAP databases (including SAP HANA)

As shown in Figure 5, shared files are needed for almost any SAP system landscape—starting with the typical candidates such as `/usr/sap/trans`, SAP HANA data volumes, and SAP HANA log volumes. 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.

Figure 5) Shared files requirements.



For all these file systems, it is evident that for production usage a highly available infrastructure should be provisioned.

Section 2, "SAP Shared File Systems," section 3, "SAP anyDB," and section 4, "SAP HANA" describe both use cases and how to position Cloud Volumes ONTAP in contrast to cloud native solutions.

2 SAP Shared File Systems

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, the Linux operating systems, or the SMB (in the case of a SAP system running on Windows operating system).

This section discusses the shared file requirements for an example setup using a SAP HANA multiple-host system with SAP HANA system replication and a SAP application server using clustered SAP ABAP SAP Central Services (ASCS) and a 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 6) Shared files for SAP.

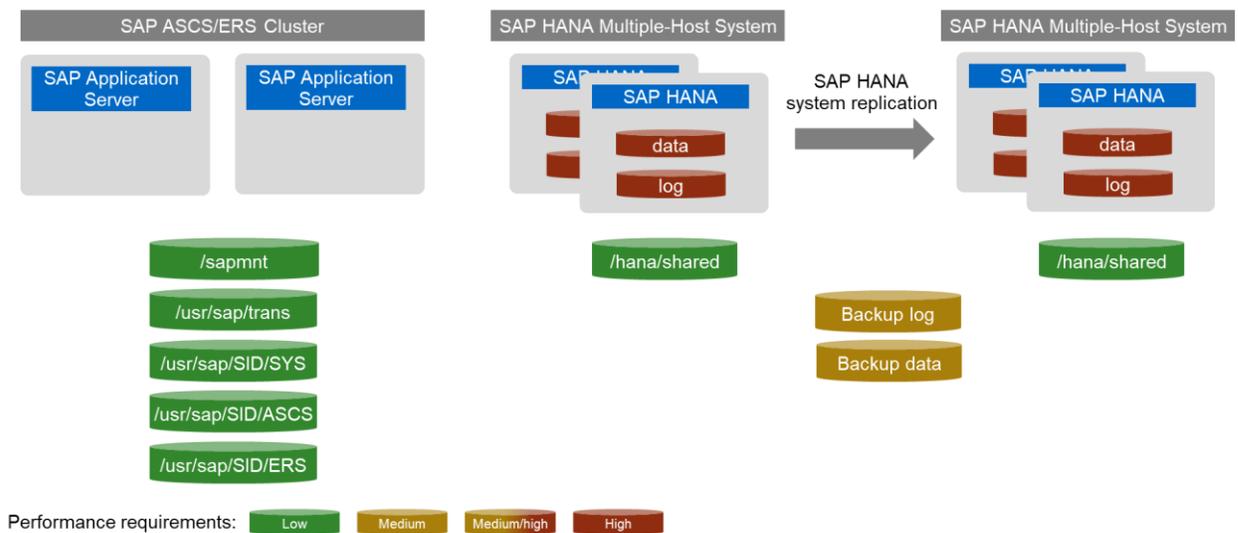


Figure 6 shows the following shared file systems that are required for this system landscape:

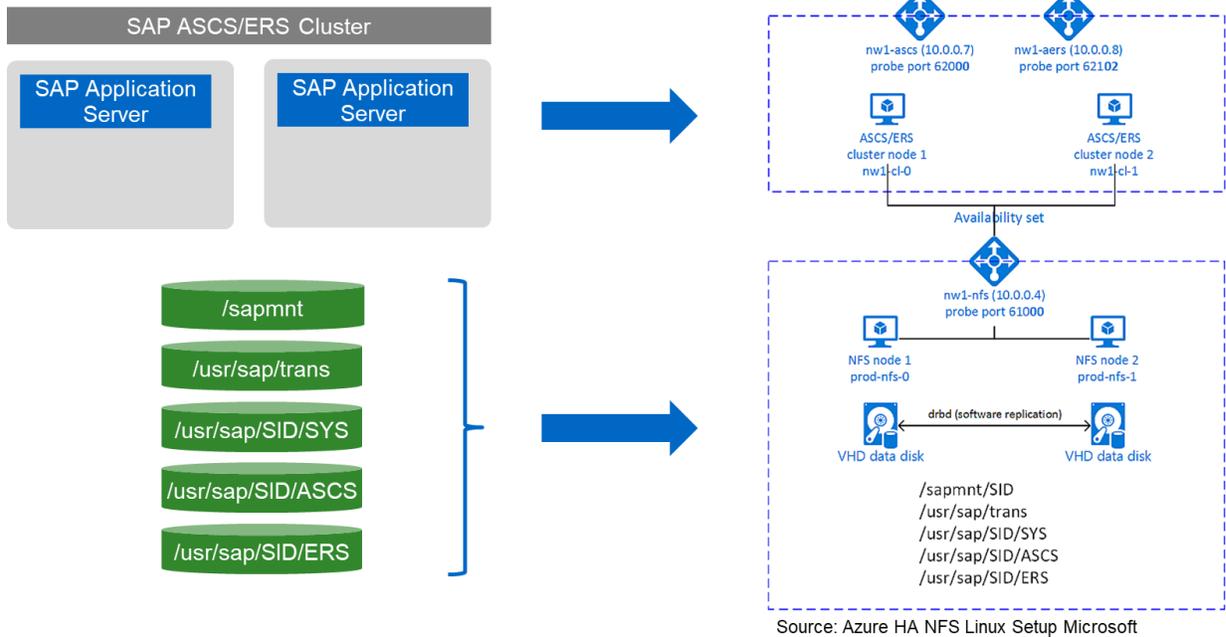
- **/sapmnt.** If you have more than one application instance, use the `/sapmnt` file system to store a common set of binaries and configuration files. The I/O pattern is reading the binaries and configuration files and writing view logs. In Figure 6, green indicates a lower performance requirement.
- **/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 be a high-available system as well so that the ERS locking table is preserved in case of a 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 a SAP HANA system replication setup, the location must be a shared location between both SAP HANA systems to allow for failover. Even in a SAP HANA single-host setup, this location should still be highly available. Also, for the log backup, a medium performance requirement is recommended.

For more information about the performance requirements for the SAP HANA database data and log volumes, see section 4, "SAP HANA."

2.2 Legacy Shared Files Solutions for the Cloud

A Linux cluster using a block replication device (as shown in Figure 7) is a commonly-used solution within the cloud to provide a highly available NFS service.

Figure 7) Linux cluster NFS server.



At a first glance, this solution looks appealing because the high degree of control that it appears to offer over the services provided by storage solutions, access protocols, file system features, and so on. 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. 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 Cloud Volumes ONTAP

This section explains how to use Cloud Volumes ONTAP to address the shared files requirements of the discussed example, as shown in Figure 8.

To provision the shared file systems listed in chapter 2.1, 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 backup logs and the optional file-based backup files

Figure 8) Cloud Volumes ONTAP shared files for SAP.

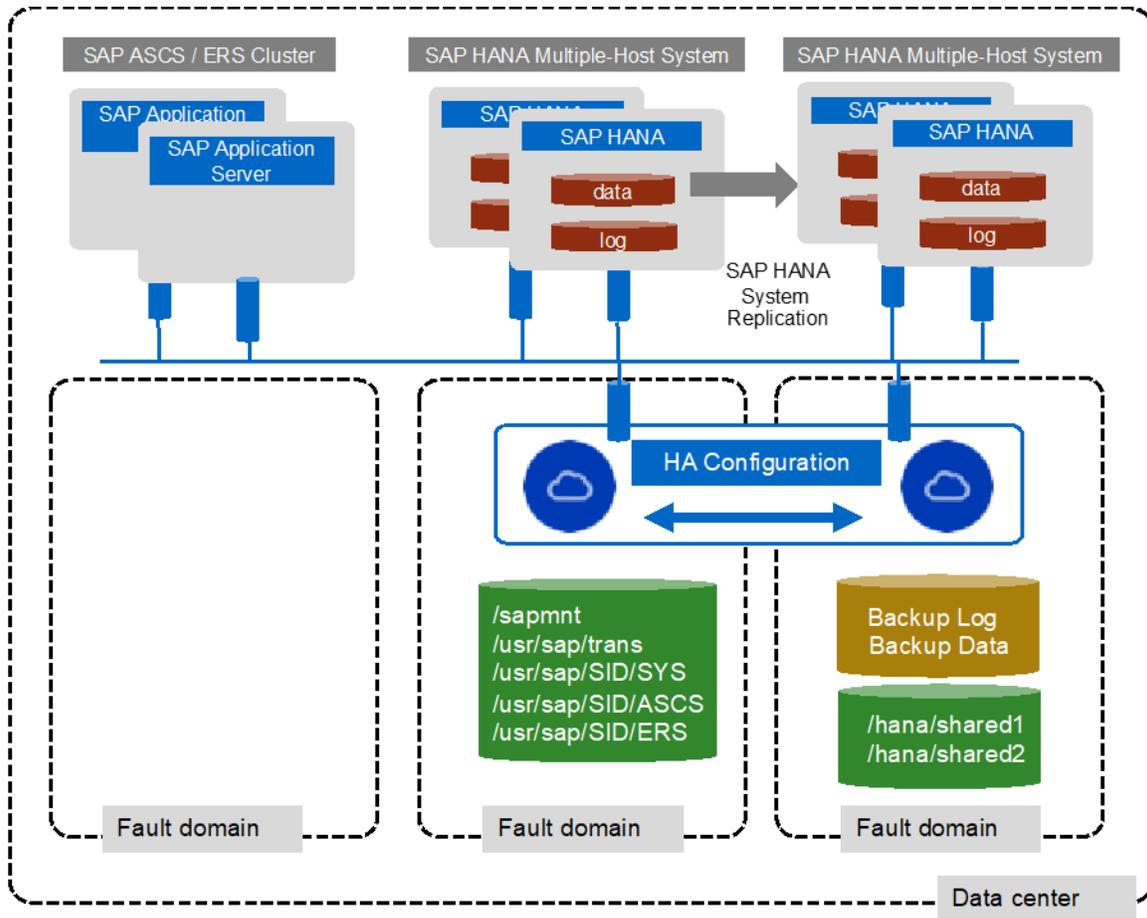
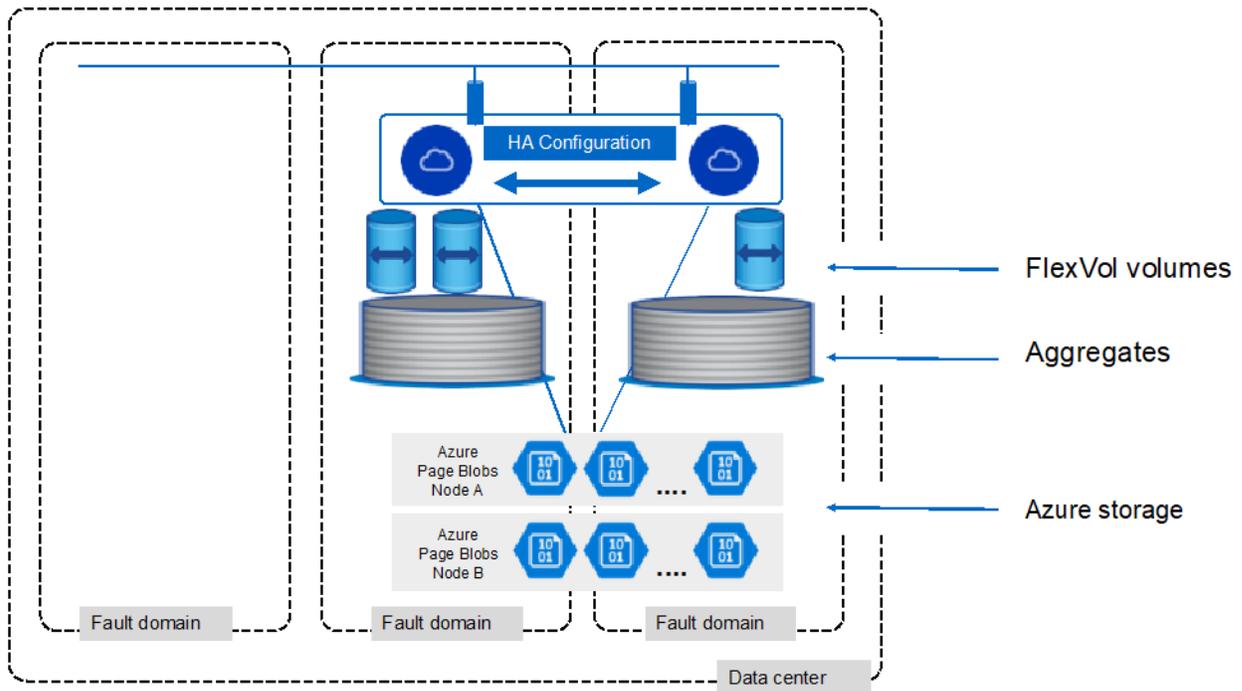


Figure 8 shows that the volumes are distributed over the two high available ONTAP instances. This is the logical view customers have, when using Cloud Volumes ONTAP.

High Availability

To understand how Cloud Volumes ONTAP enables highly available and redundant file systems, we have to look on the underlying infrastructure in which Cloud Volumes ONTAP is deployed.

Figure 9) Cloud Volumes ONTAP using Microsoft Azure resources.



Cloud Volumes ONTAP is using Microsoft Azure VMs and Azure storage as underlying persistence layer. Depending on the chosen Cloud Volumes ONTAP version, users can select from different sizes of Microsoft storage types. CVO single node uses LRS standard or premium disks (Ultra SSD will be supported in the future), CVO HA uses Azure page blobs, as shown in Figure 9. Each of the Cloud Volumes ONTAP instances has its own number of Microsoft Azure storage assigned to construct their [NetApp aggregates](#). NetApp aggregates form the storage container ONTAP is using to create [FlexVol volumes](#) to export NFS or SMB file systems.

Redundancy is achieved on the storage side using Microsoft Azure storage where each disk or page blob has three redundant copies in three distinct Microsoft Azure [fault domains](#). On the VMs, two ONTAP instances are built using an ONTAP-HA pair. In case one instance is failing, the second instance is taking ownership of all underlying resources such as storage, aggregates and FlexVols. The Cloud Volume ONTAP-HA configuration includes a Microsoft Azure Load Balancer, this ensures that clients can connect to their shared files without manual reconfiguration or downtime.

Capacity and Performance

The primary difference to a cloud native deployment, as shown in Figure 7, is that the volumes are attached to the client using NFS or SMB (file protocol) or iSCSI (block protocol) where the complexity of handling the physical storage and tuning for performance is completely removed from the client onto ONTAP.

Figure 10) Cloud Volumes ONTAP.

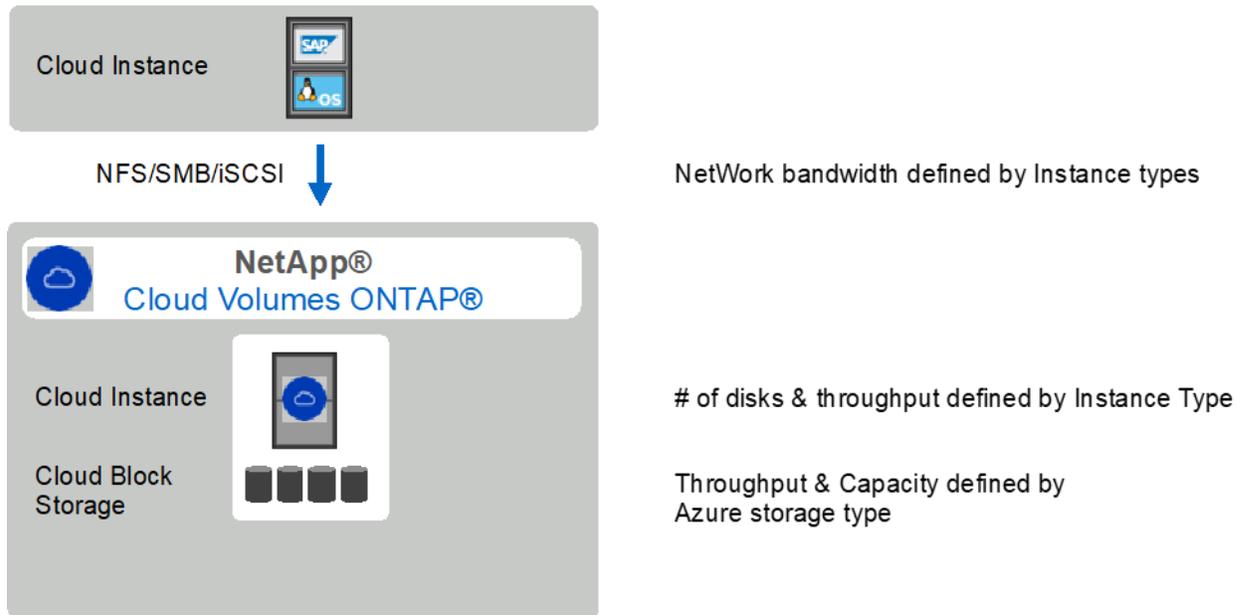


Figure 10 shows an architectural overview of a Cloud Volumes ONTAP single instance which is providing a data volume for the Cloud Instance.

The storage performance of a Cloud Volumes ONTAP volume can vary depending on the following factors:

- The underlying VM in which the Cloud Volume ONTAP is running. The VM type defines the number of disks and the maximum I/O the instance can provide.
For example, currently supported are VMs using instance types: DS3_v2 up to DS14_v2.
- The attached Cloud Block Storage and its performance and capacity restrictions.
For example, Cloud Volumes ONTAP supports, depending on the instance type, up to 63 Azure storage disks either HDD or SSD or page blobs with a maximum capacity of up to 252TB.
Note: For supported configurations, see [Cloud Volumes ONTAP Release Notes](#).
- The network bandwidth of the underlying VM for both the Cloud Instance and Cloud Volume ONTAP instance.
- The network connectivity within the Microsoft Azure region which connects the Cloud Instance and the Cloud Volumes ONTAP instance.

The physical capacity of the Cloud Volumes ONTAP instance depends on the number and size of the attached Azure storage.

A significant advantage of using Cloud Volumes ONTAP is the highly efficient data reduction, achieved through ONTAP's compression and deduplication capability. The benefit in storage efficiency (especially for shared files), facilitates the availability of a higher effective capacity, while being charged exclusively for the underlying physical attached storage capacity.

2.4 Data Protection for SAP Shared Files

In addition to performance, ease of management, and flexibility, SAP customers require enterprise grade data protection for their databases and shared file systems. Using Cloud Volumes ONTAP, customers can implement comprehensive data protection for their SAP shared file systems by using NetApp

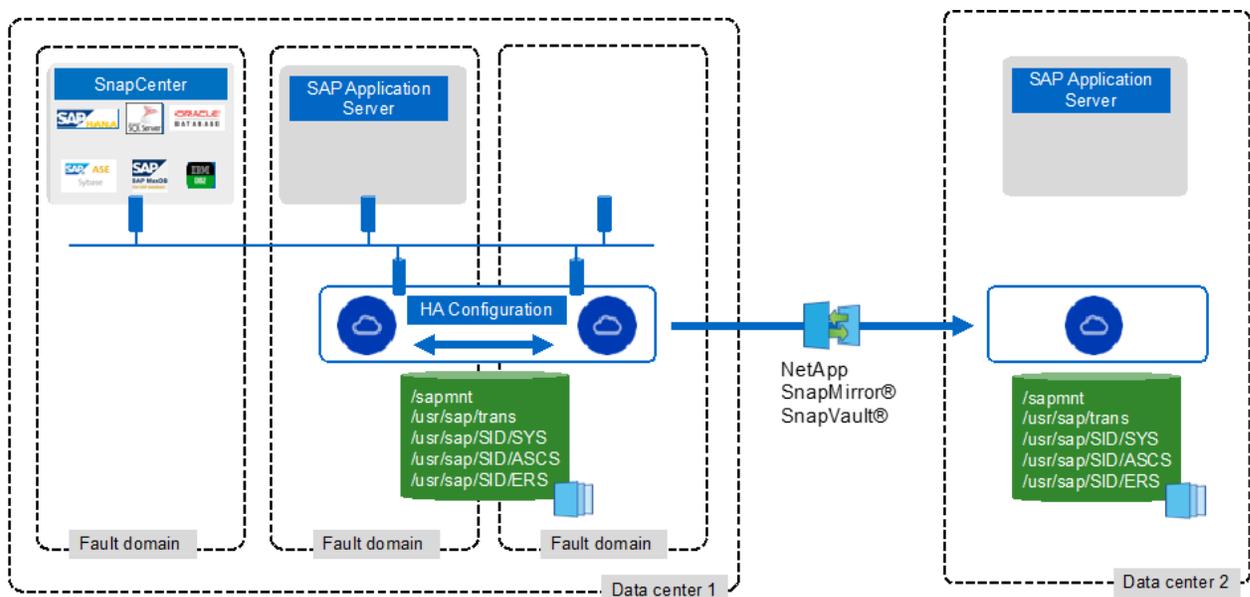
storage-based Snapshot technology in conjunction with OnCommand Cloud Manager, or other NetApp supported tools or products such as NetApp SnapCenter® software.

Snapshot Backups

NetApp Snapshot copies are extremely fast (running in seconds), space efficient, and have minimal impact on performance of the running SAP system regardless of the size of the file system.

Instead of full data copies, ONTAP's internal [WAFL](#) marks the blocks from the active file system to be part of the new Snapshot copy and ensures that whenever a block is changed, the new content will be written to an empty block, preserving the snapped data block and avoiding any additional I/O. In other words, Snapshot copies are pointers to data blocks, they allow both backup and restores to be extremely fast (minimizing recovery time objective) since only pointers are changed and no data is copied. See also the [NetApp Snapshot Technology](#) datasheet.

Figure 11) SnapCenter: Snapshot-based data protection.



The complete process of data protection using Snapshot technology can be automated using SnapCenter, as show in Figure 11. Snapshot copies can be scheduled hourly, daily weekly and monthly and for each schedule the number of Snapshot copies to be retained can be specified. To support a disaster recovery strategy, Snapshot copies can be replicated by using NetApp SnapMirror® technology into another Cloud Volumes ONTAP system in a different Azure region.

3 SAP anyDB

Many customers continue to run their SAP Business Suite and NetWeaver based applications on anyDB databases. The anyDB database encompasses all non-SAP HANA databases such as Oracle, MS SQL, ASE, DB2, or MaxDB.

SAP announced the end of support for the SAP Business Suite by 2025, which mandates that all SAP customers must make the transition to S/4HANA (SAP Business Suite 4) with its underlying SAP HANA database before this time.

To tackle this problem, customers are rethinking their data center strategy, which includes the option to move their SAP workloads to public clouds such as Microsoft Azure.

Using ONTAP technology simplifies data management and operations of large-scale, on-premises SAP systems. ONTAP also helps accelerate SAP deployments while reducing project risk by offering industry-leading solutions for storage cloning (to establish multiple copies of test environments) and using Snapshot technology for lightning fast backup and restore operations. With the availability of ONTAP technology in Microsoft Azure, customers can benefit from running their SAP anyDB workloads to simplify their journey to SAP HANA.

This section describes the following topics:

- Storage requirements for SAP anyDB databases
- Cloud native solutions and its challenges
- Architecture and benefits of using Cloud Volumes ONTAP

3.1 Storage Requirements for SAP anyDB Databases

Certification

SAP customers must run their production SAP systems on certified infrastructures, irrespective whether SAP is running on-premise or in the public cloud. Hardware vendors and cloud providers test and benchmark their servers and VMs using SAP's Application Performance Standard (SAPS) – based on the Sales and Distribution workflow. Following the benchmark validation and other test cases, an infrastructure for example, VMs and storage can attain certification for running a SAP workload with a given number of SAPS.

Microsoft Azure's SAP certifications can be found on the following web site: [SAP certifications and configurations running on Microsoft Azure](#) and SAP's benchmark data is published at [SAP's Standard Application Benchmarks site](#). Table 1 lists all VMs certified at the time of writing this document. For up-to-date information, refer to [SAP Note 1928533- SAP Applications on Microsoft Azure: Supported Products and Microsoft Azure VM types](#)

While servers and VMs need to be certified, there is no certification for enterprise storage – however, a formal support declaration for anyDB running in production may be required e.g. Oracle on Microsoft Azure. This means, in the cloud, customer have the option to run their SAP workloads on NetApp storage to benefit from the respective storage features.

Performance

When performance is considered, enterprise storage must meet the requirements to sustain the SAP workload. A typical pattern for databases is to look at a 80/20 read/write I/O pattern, using a common block size such as 8k.

Customers already running SAP applications on-premise can take their I/O pattern and requirements, while customers new to SAP may start their sizing based on SAPS numbers. As a rule of thumb one can translate the number of SAPS to a required I/O with the following formula.

$$\text{I/O (80/20 r/w of 8k)} = 0.6 * \text{SAPS}$$

[TR-4691: Oracle Databases on ONTAP Cloud with Microsoft Azure](#) targets Oracle databases on Cloud Volumes ONTAP in Microsoft Azure in detail.

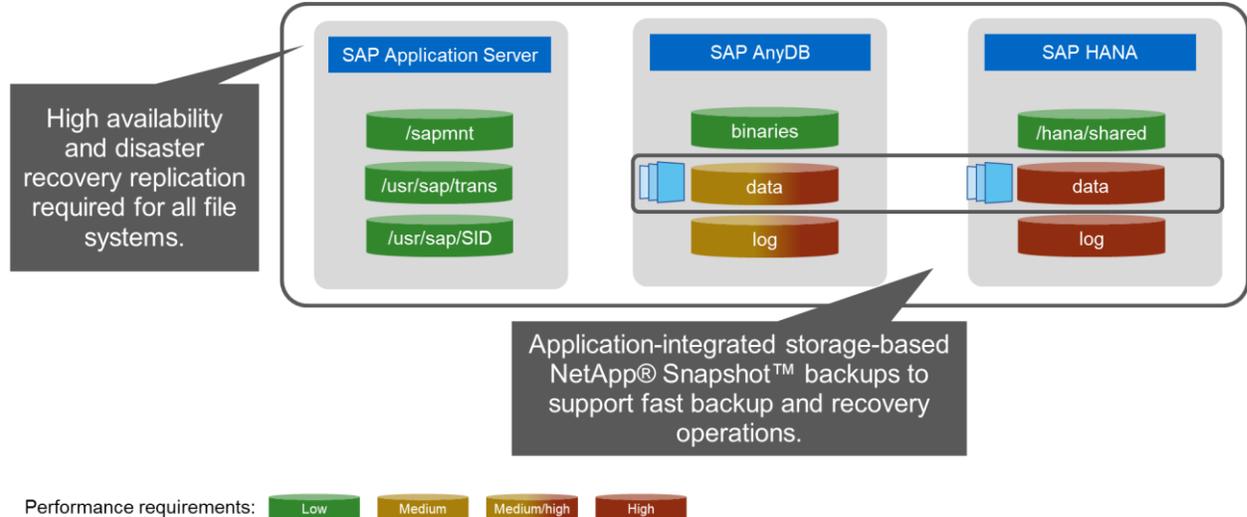
Connectivity

Both Cloud Volumes ONTAP as well as Azure NetApp Files offer enterprise NFS which is the first choice for database workloads in the Linux world. For customers running their SAP system on Windows, block devices are used. In the cloud this is possible using iSCSI and Cloud Volumes ONTAP.

Data Protection

An important point for databases is the need for storage-based data protection that ensures data consistency on the application layer.

Figure 12) Data protection.



This requires products such as NetApp SnapCenter® which ensures the database backup and storage-based Snapshot copies are synchronized. This application integration ensures that the storage Snapshot backup is usable for restore and recovery.

High Availability

For the majority of productive SAP systems, a highly available setup is mandatory. In order for this to be realized, SAP allows the NetWeaver System Central Services as well as the Replicated Enqueue Server to be implemented within a cluster setup to eradicate single points of failure. This is shown in Figure 7. Any shared file system used by an productive SAP system should be highly available.

3.2 Cloud Native Solutions

Microsoft Azure offers SAP customers automation capability for the deployment of Microsoft Azure infrastructure (IaaS) components, Linux clusters and the installation of an SAP software. The following selection of automation options describe automation with Terraform and Ansible and native Microsoft Azure Resource Manager templates:

- Terraform and Ansible:
 - [Automating SAP deployments in Microsoft Azure using Terraform and Ansible](#)
 - [GitHub : Azure/sap-hana](#)
- Native Microsoft Azure Resource Manager Templates:
 - [Accelerate your SAP on Azure HANA project with SUSE Microsoft Solution Templates](#)

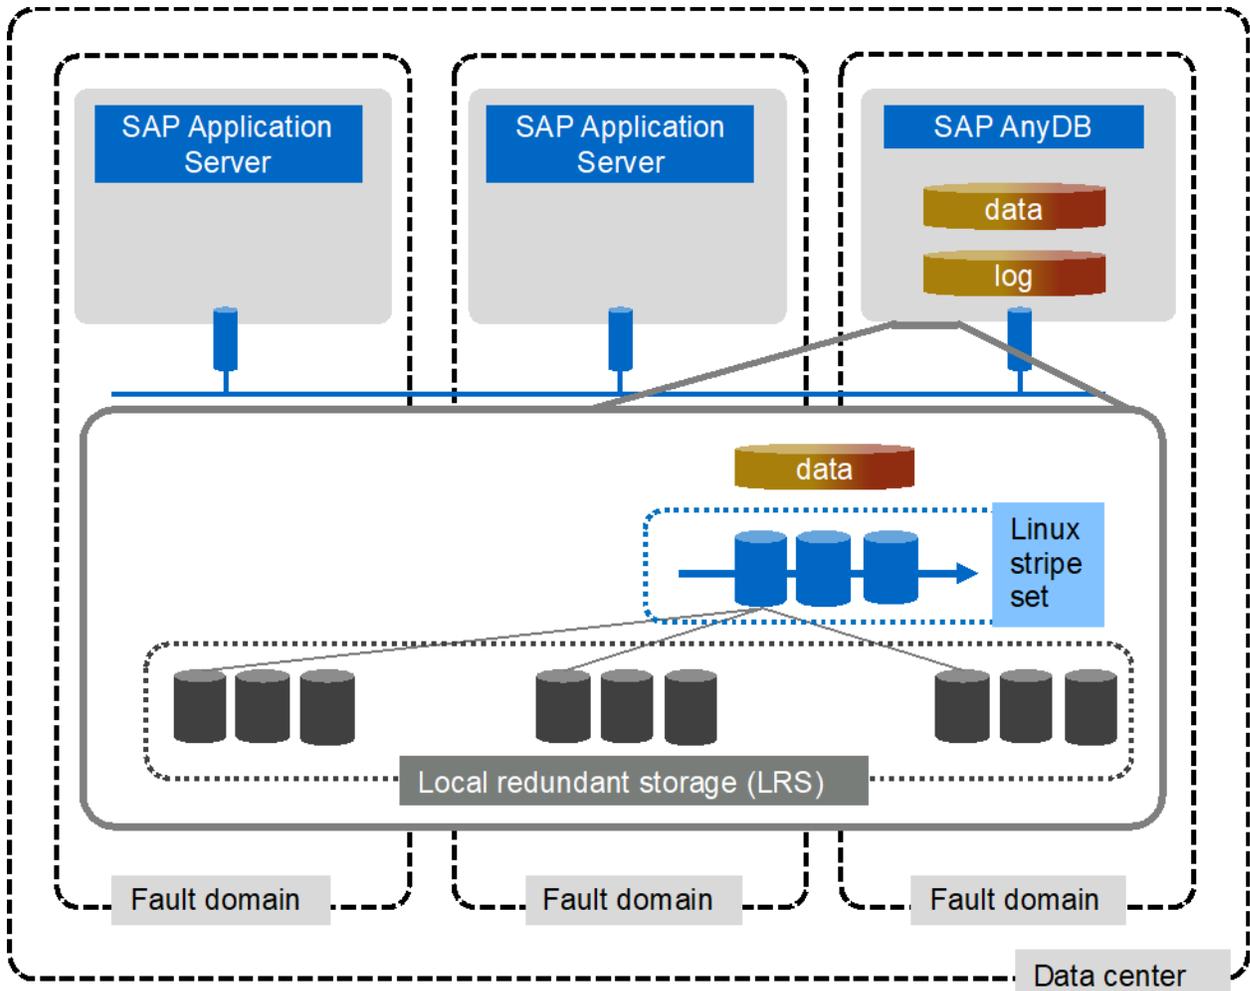
Customers should select the VM type depending on the SAP sizing requirements. Microsoft supports a broad range of SAP certified VMs, each of which provides a certain amount of CPU, memory, network, and storage I/O. The SAP value for each SAP certified VM is located in [SAP Note 1928533: SAP Applications on Microsoft Azure: Supported Products and Microsoft Azure VM types \(SAP login required\)](#).

For the storage performance requirements for the SAP database, different infrastructure components have an impact on the overall I/O performance:

- The VM type has limits regarding supported storage as well as the maximum storage IOPS it can provide.
- Each of storage disks, such as [Premium SSD](#) or [Ultra SSD](#), has storage KPIs with regards to IOPS, throughput, and size.

With these infrastructure constraints in mind, the customer must follow the Microsoft and SAP recommendations to select the right combinations to construct the infrastructure for the database server. Fortunately, the automation templates help build the necessary foundation, as shown in Figure 13.

Figure 13) Storage for an SAP database server.



Data Volume and Disk

Azure recommends that you use premium SSD for both production and nonproduction database disks. For more information, see [SAP Note 2367194: Using Azure Premium SSD Storage for SAP DBMS Instance \(SAP login required\)](#).

Use the following key figures to build the data disk:

- Required capacity
- Required performance, IOPS, and throughput

Assuming the VM type delivers sufficient storage throughput, the challenge is to find the right combination of premium disks to meet both requirements. For larger databases, a single premium disk is typically not sufficient.

Figure 13 shows three premium disks combined by using a Linux LVM stripe set to achieve the performance and capacity requirements. Each of the premium disk uses local redundant storage, three mirrored disks from three different fault domains to achieve the required redundancy. In this example, a total of nine disks are required to construct a single logical data disk for the SAP system.

Note: A customer pays for three Azure premium disks (not nine).

Log Volume Disk

The main requirements for the log volume include:

- High sequential write performance
- Low latency

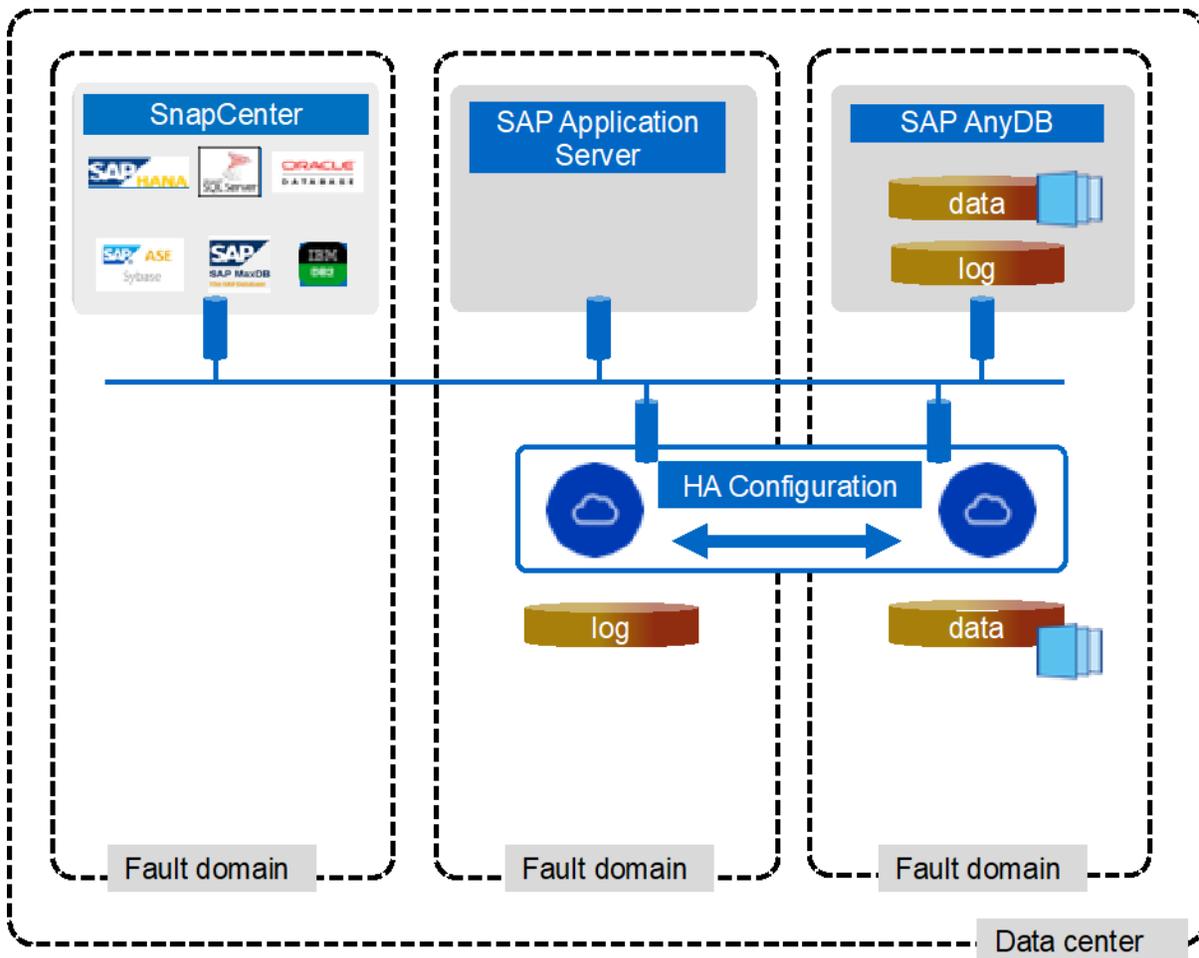
Premium SSD disks are generally sufficient for the log volume. However, in high-end cases where extreme low latency is required, [write acceleration](#) for the database LOG volume can be leveraged.

3.3 SAP anyDB on Cloud Volumes ONTAP

As described in section 2.3, "Providing Shared Files using Cloud Volumes ONTAP," is a fully featured ONTAP running in a cloud instance. It is simple to provision using NetApp OnCommand Cloud Manager, which provides management functions such as create, resize, backup and restore of volumes with a friendly UI.

Figure 14 shows a logical setup where Cloud Volumes ONTAP is configured in a HA configuration running on two cloud instances. With this configuration, transparent failover is provided if case one cloud instance becomes problematic.

Figure 14) Cloud Volumes ONTAP for SAP anyDB.



The DATA and LOG volumes can be created using OnCommand Cloud Manager, Cloud Volumes ONTAP CLI or other ONTAP workflow automation tools. As in any ONTAP environment, SAP customers can benefit from:

- Easy management and resize of a volume to adapt to a growing database
- NetApp volume-based encryption (NVE) to further protect the data. NVE is FIPS 140-2 compliant using AES 256-bit encryption on a volume granularity.
- Enhanced deduplication algorithm to reduce storage capacity
- Application integration via NetApp SnapCenter allowing the creation of application consistent, storage-based snapshot backups for all SAP supported databases.

Connectivity and Performance for Databases

As shown in Chapter 2.3 in section “Capacity and Performance” there are many factors influencing the overall performance required for a typical SAP anyDB scenario.

NetApp conducts database specific benchmarks which help to understand the specific database performance of Cloud Volumes ONTAP. The following technical report [TR-4691 Oracle Databases on ONTAP Cloud with Microsoft Azure](#) provides guidelines for using Cloud Volumes ONTAP with Oracle. Continuous improvements of Cloud Instance performance as well as ONTAP software are being made, as such technical reports will be updated periodically.

For details to configure SAP anyDB databases consult the database specific ONTAP technical reports for the required SAP database.

For NetApp customers using ONTAP for their SAP databases in their own datacenter, Cloud Volumes ONTAP is a perfect extension to Azure, allowing customers to easily migrate systems to the Azure using SnapMirror as well as enabling a hybrid operation. For more on the topic of data protection and system migration see chapter 6 Data Protection and System Migration.

4 SAP HANA

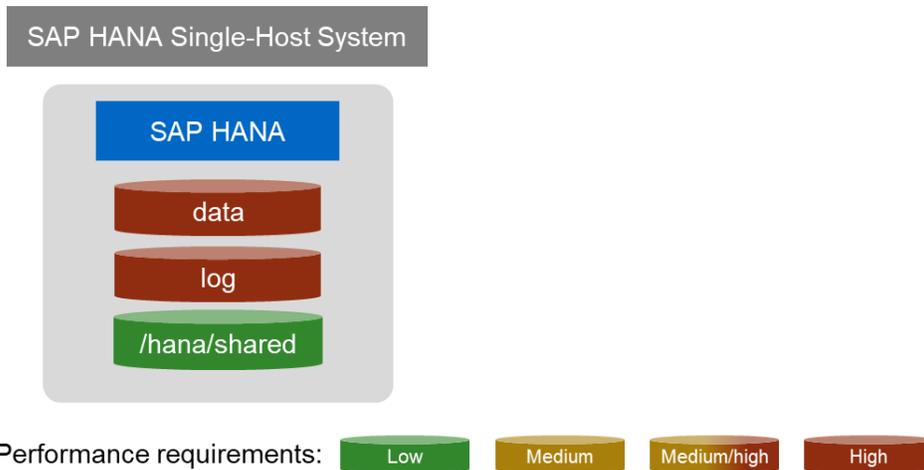
4.1 Storage Requirements for SAP HANA

SAP has published the [SAP HANA TDI-Storage Requirements](#) paper in which the SAP HANA storage requirements are detailed.

When looking at the storage requirements for SAP HANA the most basic configuration requires three different storage volumes as shown in Figure 15.

Single-Host SAP HANA

Figure 15) HANA single-host system.



- **/hana/shared.** This volume is used for the shared files such as binary, logs, configuration. For more information, see also section 2.1, "Shared File System for SAP NetWeaver Application Server." There are no specific storage requirements for the shared files filesystem for a single-host HANA system, it can be either local attached storage or NFS mounted storage. In case of a multiple-host HANA, the shared files system must be a shared NFS mounted file system.
- **DATA volume.** The HANA DATA volume needs to have the size of the memory allocated for the HANA database, e.g. the memory of the HANA VM in most of the setups. For the DATA as well as the LOG volume SAP defines performance criteria's and specific certification. For a Single-Host Systems it can be local storage or NFS mounted storage, for a multiple-host HANA system an external enterprise storage is recommended, although SAP HANA can be setup in "non-shared" mode, with local disks attached to each node of the multiple-host HANA system. In public cloud environments shared NFS volume is the only viable option for N+1 high-availability scenarios.
- **LOG volume.** The HANA LOG volume is required to persist the most recent redo logs. SAP recommend the size to be 50% of the HANA memory with a maximum size of 0.5TB. For a Single-Host System it can be local storage or NFS mounted storage, for a multiple-host HANA system an external enterprise storage is generally the way to go, although SAP HANA can be setup in "non-

shared” mode, with local disks attached to each node of the multiple-host HANA system. In public cloud environments shared NFS volume is the only viable option for N+1 high-availability scenarios.

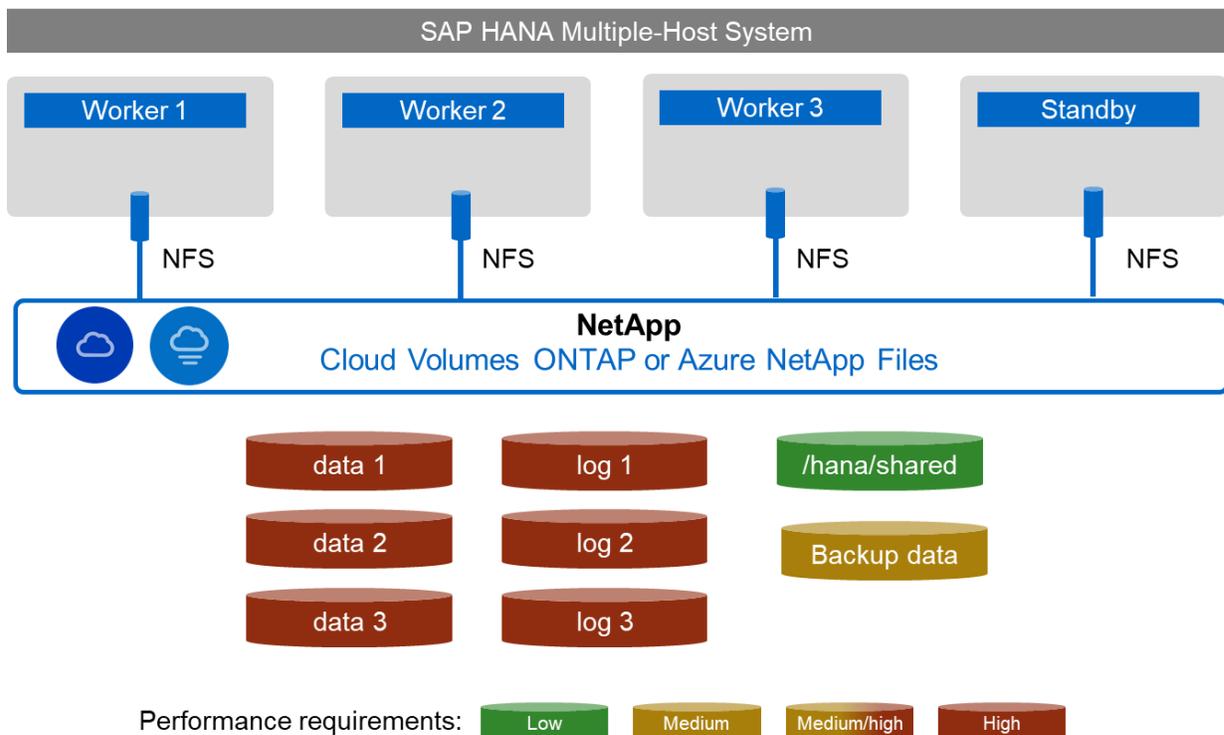
Multiple-Host SAP HANA

In cases where the memory size of a single server is not sufficient SAP HANA allows to combine the memory of multiple servers to run SAP HANA in a Multiple-Host configuration. In this setup, SAP HANA allows to configure a standby host that, in case of a host failure of one of the worker hosts, will take over the role of the failed server, as shown in Figure 16. This SAP HANA cluster mechanism requires that the `/hana/shared` file system must be available on all HANA hosts that require NFS. Also, the DATA and LOG volumes must be remounted on the standby host, which is also a perfect use case for NFS.

Additionally, SAP HANA multiple-host setups require for file-based backups, that a common backup volume is shared between all SAP HANA hosts:

- **Backup volume (data backup).** Depending on how many backup versions need to be kept a size of 2-3 times the HANA memory is recommended.
- **Backup volume (log backup).** SAP HANA automatically archives the redo logs from the log volume, to a shared filesystem where all hosts need to have access. 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 usage of the SAP HANA database and amount of changes that are created.

Figure 16) SAP HANA multiple-host system.



SAP HANA KPIs and Certification

In order to gain production support for SAP HANA, SAP requires that the underlying infrastructure provides sufficient performance. For on-prem infrastructures, SAP has created a Hardware Configuration Check Tool (HWCCT) and a set of performance metrics for the data and log volume, that customers can use to test if the setup fulfills the required storage performance.

Figure 17) TDI performance criteria.

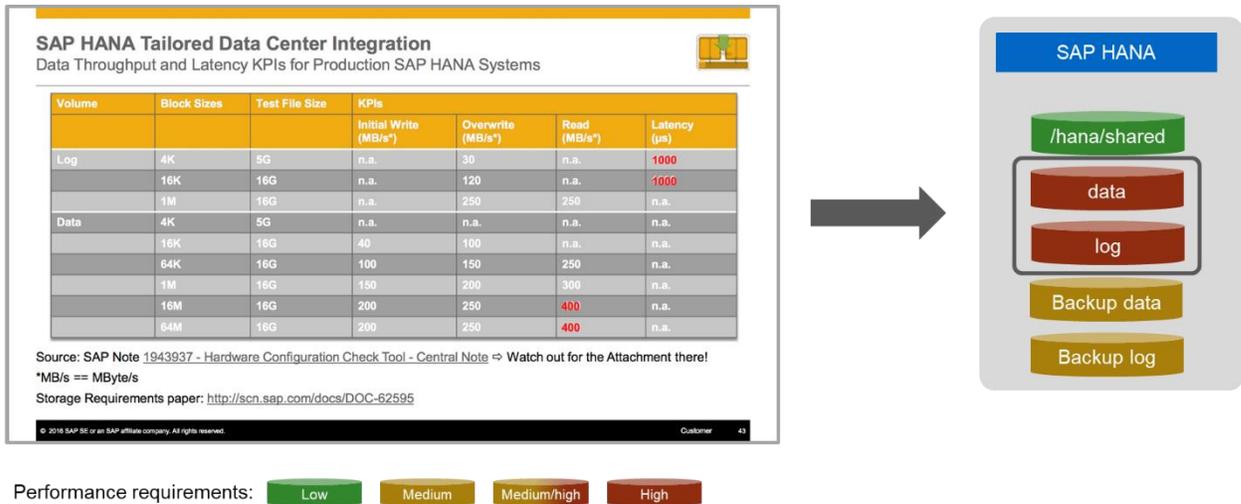


Figure 17 shows the published test criteria used for the SAP HANA Tailored Data Center Integration, that are valid for an on-premises setup where customers can build their SAP HANA infrastructure by using certified SAP HANA servers and certified SAP HANA storage.

These values are a performance guideline for cloud setups, but further 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 KPI's 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.

With Cloud Volumes ONTAP and the rich feature set of ONTAP as well as the available tool integration, we currently focus on non-production use cases of SAP HANA. Therefore, a certification of this setup is not required. Performance tests have shown that the performance of Cloud Volumes ONTAP is sufficient for most use cases and scenarios, including those that require a decent level of storage performance.

4.2 Certified Microsoft Azure Native Solutions

As part of Azure's offerings for SAP HANA, Microsoft has certified several Infrastructure as a Service components. An official list can be found at [Certified and Supported SAP HANA Hardware Directory](#).

The certified Microsoft Azure offerings for SAP HANA workloads on VMs range from 112GB (DS14v2) up to 4TB (M128ms). And, for HANA Large Instances, up to 20TB scale-up (S960m) and 60 TB scale-out (15 x S384).

Note: With [TDiv5](#) 24TB scale-up and 120TB scale-out is possible.

Note: IaaS platform certification include the certification of the combination of VM, network and storage.

The certified VMs for SAP HANA are using local storage which is optimized to meet the SAP HANA performance requirements. And, for M-SERIES VMs this is achieved by using [Write Accelerator](#) for the log volume and Linux LVM/mdadm stripe sets of premium SSD disks to provision the data volume and aggregate throughput.

Note: As at Q1CY19, Cloud Volumes ONTAP in combination with M-SERIES is not included in the official certified and supported SAP HANA hardware directory.

M-SERIES VMs can be used in combination with Cloud Volumes ONTAP to provision and run non-production SAP HANA workloads to leverage the benefits of ONTAP for optimizing SAP application

lifecycle workflows. The following examples demonstrate the use of Cloud Volumes ONTAP for SAP HANA for a non-production scenario e.g. development or test.

5 Example: Installation of SAP NetWeaver ABAP Stack Using SAP HANA on Cloud Volumes ONTAP

In this example, we run a complete system installation of an SAP NetWeaver 7.5 ABAP server with System ID (SID) NW5 together with an SAP HANA 2.0 database with SID CVO on two Azure VMs. A D12s_v2 VM is used for the SAP NetWeaver application server and a DS16s_v3 for the SAP HANA database.

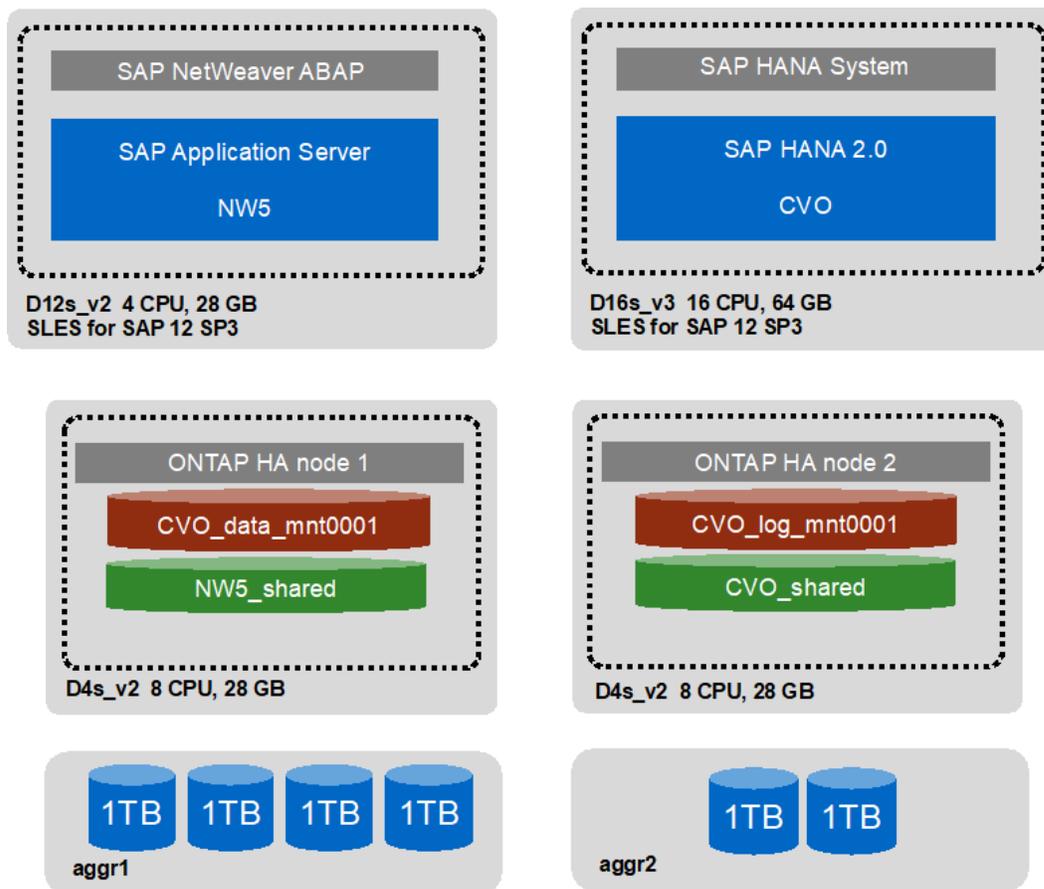
Note: In this example, the VM type is selected to meet the performance and memory requirements.

As the enterprise storage back-end, we use a Cloud Volumes ONTAP HA – two-node cluster installation where we deploy the required volumes for the SAP NetWeaver application server as well as SAP HANA database. The two ONTAP nodes are configured with 4x 1TB Azure premium disks for ONTAP node 1 and 2x 1TB premium disks for node 2. The underlying Azure premium storage is selected to provide sufficient capacity and I/O performance.

Note: Again, this configuration is only recommended for a nonproduction scenario. Cloud Volumes ONTAP is not certified SAP HANA in a production environment.

Figure 18 shows the installed components and storage setup for the solution.

Figure 18) SAP NetWeaver on SAP HANA using Cloud Volumes ONTAP.



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

Note: The specified Cloud Volumes ONTAP configuration can be expanded later to meet increasing capacity or performance requirements.

This example shows the following steps:

- Plan your installation in order to estimate the required capacity and performance.
- Create and deploy NetApp OnCommand Cloud Manager in order to deploy and manage the Cloud Volume ONTAP.
- Create and configure Cloud Volumes ONTAP by using the OnCommand Cloud Manager UI.
- Configure and prepare the underlying Azure storage for Cloud Volumes ONTAP in order to provide the required capacity, throughput and create the required volumes.
- Prepare the Azure VMs and NFS mounts in order to install SAP NetWeaver and SAP HANA.

5.1 Installation Planning

In order to plan the storage part of the installation, you must understand the capacity and performance requirements to determine the optimal Cloud Volumes ONTAP configuration.

File System Sizing and Volume Layout for the NetWeaver Instance

The following file systems need to be created. Table 3 NW5 NFS file systems show the mount path and the initial minimum size.

Table 3 NW5 NFS file systems.

Path	Size (GB)	Comment
/usr/sap/trans	100GB	Transport shared file system
/sapmnt/NW5	128GB	Shared executables for SID NW5
/usr/sap/NW5	128GB	Installation directory for SID NW5

It is best practice on many SAP NetWeaver installations on NetApp NFS to create a single volume with subfolders for file systems with similar protection and performance requirements. This configuration simplifies the storage management and data protection.

In this example, we created a single volume, `NW5_shared`, with a size of 1TB. Within this volume, we create three folders as mount points for the file systems.

Table 4) Volume and folder structure.

Volume	Size	Folder	Mount Point
NW5_shared	1TB	usr-sap	/usr/sap/NW5
–	–	trans	/usr/sap/trans
–	–	sapmnt	/sapmnt/NW5

File System Sizing and Storage Layout for the SAP HANA Instance

The basic file systems listed in Table 5 must be created for the SAP HANA instance. Table 5 shows the volume name, mount path, and the initial minimum size.

Table 5) SAP HANA basic volumes and file systems.

Volume	Mount Path	Required Capacity (GB)	Comment
CVO_data_mnt00001	/hana/data/CVO/mnt00001	1028GB	Data volume
CVO_log_mnt00001	/hana/log/CVO/mnt00001	512GB	Log volume
CVO_shared	/shared to /hana/shared /usr-sap to /usr/sap/CVO/home	512GB	For /hana/shared and /usr/sap/CVO/home

Note: HANA volume `CVO_shared` includes subfolders for all the HANA shared files and the home directory of the database administrative user.

Performance and Capacity Considerations

As shown in Figure 18, NetApp recommends that you distribute the volumes with higher performance and I/O requirements onto different ONTAP nodes. In this example, the HANA data volume is hosted on ONTAP HA node 1 while the HANA log volumes is hosted on node 2.

The capacity requirements for the HANA DATA volume is usually higher than for the log volume. Therefore, NetApp provisioned 4x 1TB Azure storage disks on node 1 and 2x 1TB Azure storage disks on node 2.

The number of Azure storage disks and the network connectivity (in case of page blobs) define the available capacity, the available I/O, and the throughput that the Azure storage is able to deliver. Customers can configure this by using OnCommand Cloud Manager in Advanced Storage Configuration mode.

Another factor that defines the overall performance of Cloud Volumes ONTAP is the chosen VM type, which is capped in terms of maximum number of disks that can be attached, disks IOPS, and throughput and the network bandwidth.

Table 6) Cloud Volumes ONTAP storage configuration.

Node	Instance Type	Aggregate	Azure Storage
ONTAP HA node 1	D4s_v2	aggr1	4 * 1TB
ONTAP HA node 2	D4s_v2	aggr2	2 * 1TB

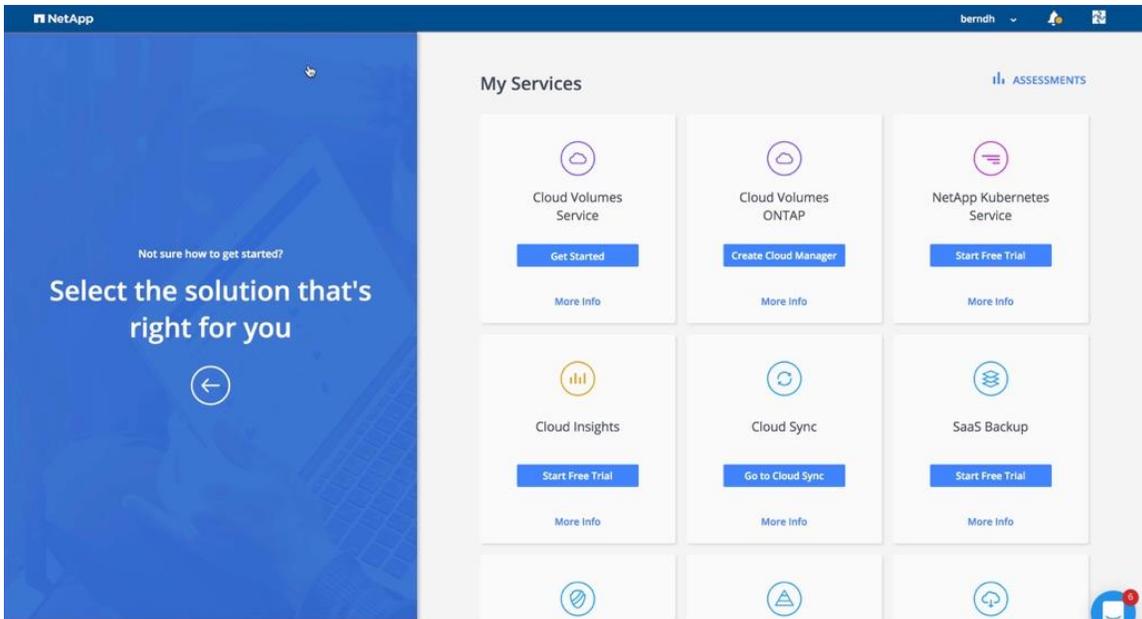
5.2 Create and Deploy OnCommand Cloud Manager

To create and deploy OnCommand Cloud Manager, complete the following steps:

1. As for all NetApp Cloud offerings [NetApp Cloud Central](#) is the central portal where customers can start. From here customers can log in into your user account or register as a new user.

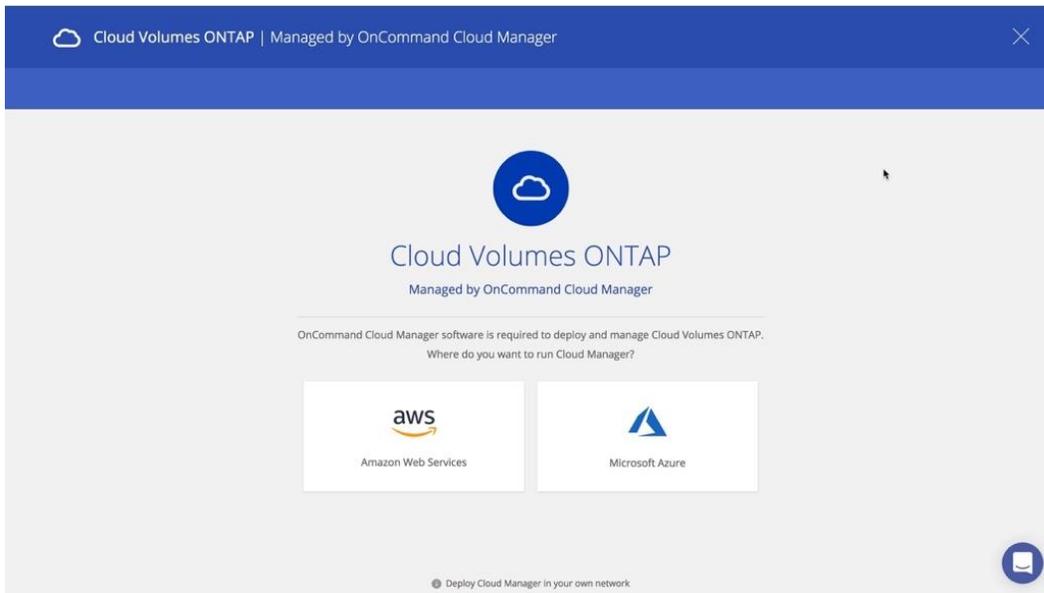


2. After you log in to Cloud Central, click My Services to directly navigate to your services.



All NetApp services are displayed. The Cloud Volumes ONTAP tile contains the Create Cloud Manager button, since Cloud Central recognizes that no OnCommand Cloud Manager instance has been created. As mentioned in the section titled, "Cloud Volumes ONTAP on Microsoft Azure," OnCommand Cloud Manager is required to deploy and manage Cloud Volumes ONTAP.

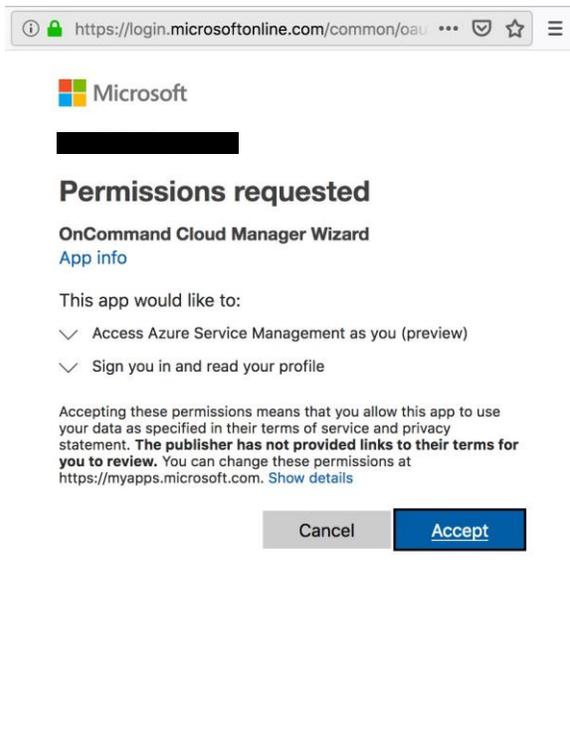
3. Click Create Cloud Manager to start the installation. The OnCommand Cloud Manager workflow starts.



4. Select the cloud where OnCommand Cloud Manager should be deployed.

Note: One OnCommand Cloud Manager can manage many different Cloud Volume ONTAP systems, even in different clouds or physical ONTAP clusters, as long as the network connectivity is available through VPN or ExpressRoute.

5. If you log in the first time with this account, Cloud Central must link your account to a Microsoft Azure user. Therefore, a popup window might ask to grant permission for this account to connect to your Azure account using Microsoft single sign on login.



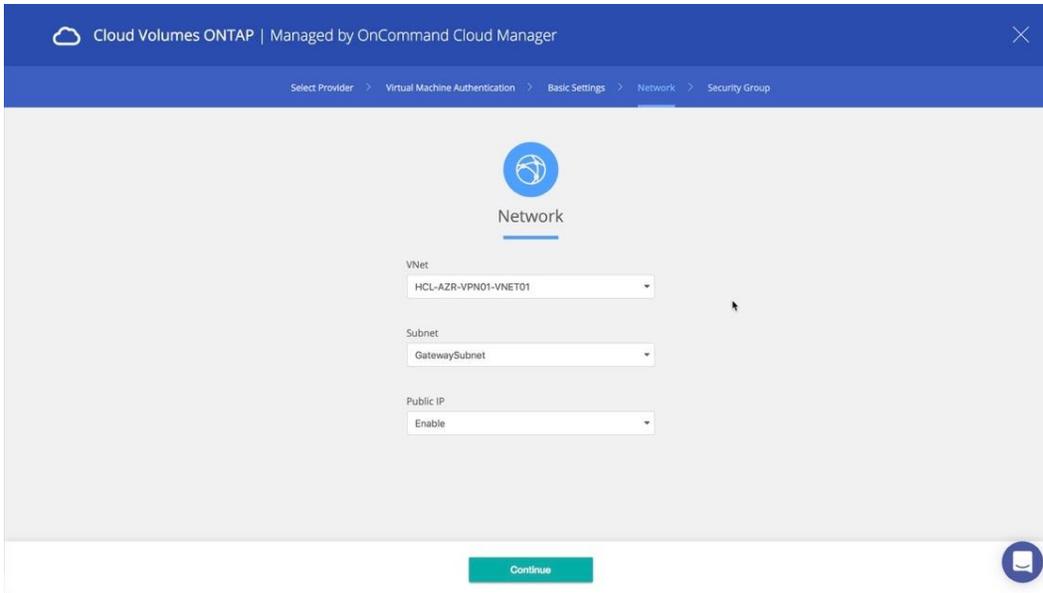
6. After you accept the request, select the Microsoft Azure tile to start the deployment in your Azure subscription. The Cloud Manager workflow continues.
7. Specify the name of you OnCommand Cloud Manager VMs, the login user and password you set to login to the OnCommand Cloud Manager Web UI.

The screenshot shows the 'Virtual Machine Authentication' step of the deployment wizard. The breadcrumb trail is 'Select Provider > Virtual Machine Authentication > Basic Settings > Network > Security Group'. The page title is 'Cloud Volumes ONTAP | Managed by OnCommand Cloud Manager'. A blue key icon is centered above the title. Below the title, there is a note: 'When prompted, accept the permissions request from Microsoft. You must also ensure that a role that meets Policy Requirements is attached to your Azure account.' The form contains the following fields: 'Cloud Manager VM Name' with the value 'bhCloudManager', 'User Name' with the value 'berndh', and 'Authentication Method' with 'Password' selected. There is an 'Enter Password' field with masked characters. A 'Continue' button is at the bottom center, and a help icon is at the bottom right.

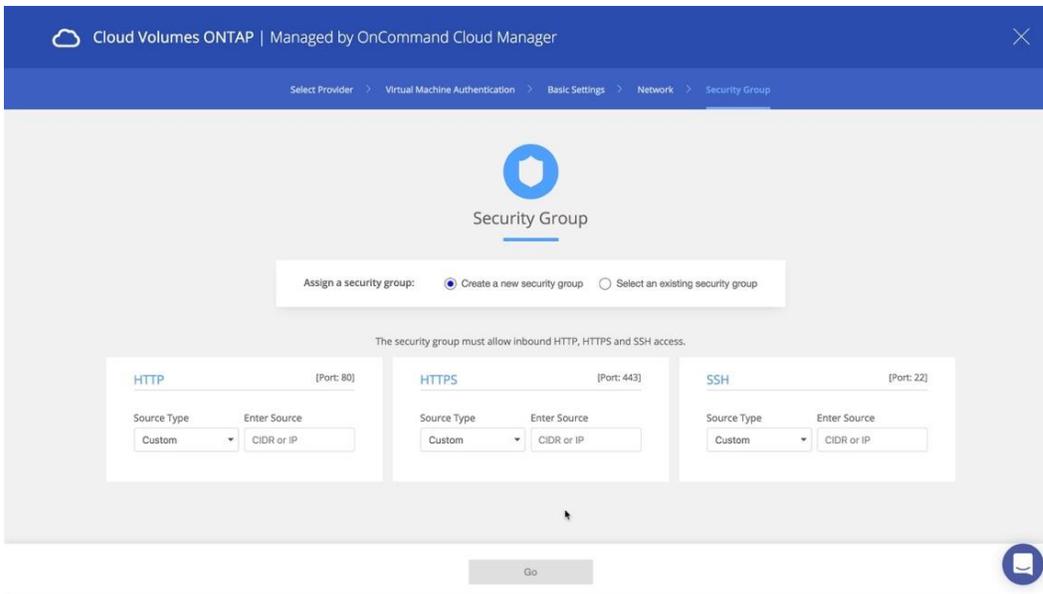
8. Select the subscription, region, and resource group you wish to use for the deployment. Click Continue.

The screenshot shows the 'Basic Settings' step of the deployment wizard. The breadcrumb trail is 'Select Provider > Virtual Machine Authentication > Basic Settings > Network > Security Group'. The page title is 'Cloud Volumes ONTAP | Managed by OnCommand Cloud Manager'. A blue gear icon is centered above the title. Below the title, there are three dropdown menus: 'Subscription' set to 'HCL Main Subscription', 'Region' set to 'East US 2', and 'Resource Group' set to 'bhcvorgus'. There are radio buttons for 'Create a new group' (unselected) and 'Use an existing group' (selected). A 'Continue' button is at the bottom center, and a help icon is at the bottom right.

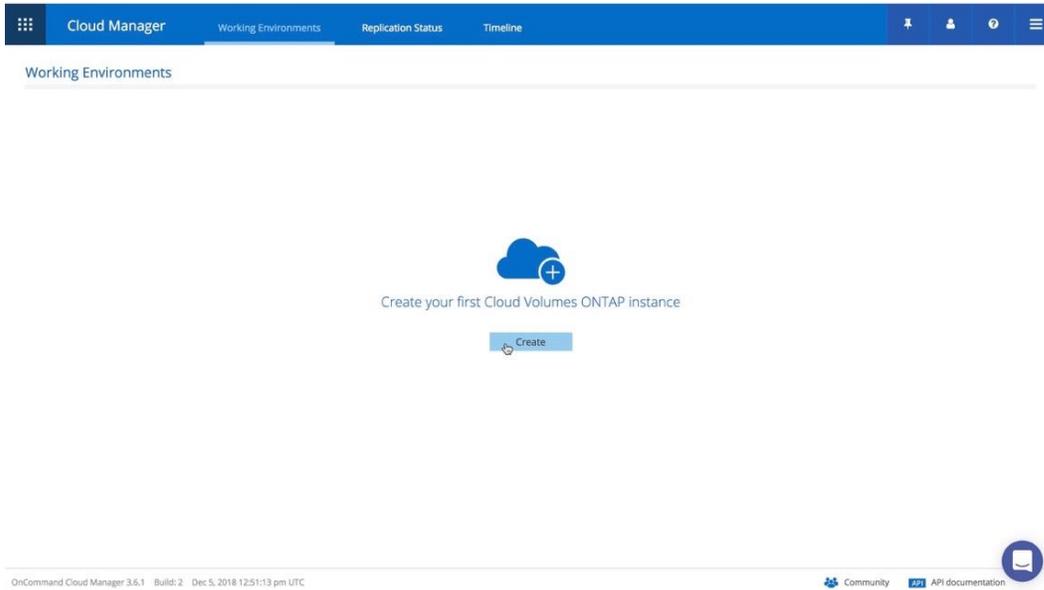
9. Specify the virtual network (vNet) you want to deploy OnCommand Cloud Manager as well as the subnet in which you need to assign a public IP address. Click Continue.



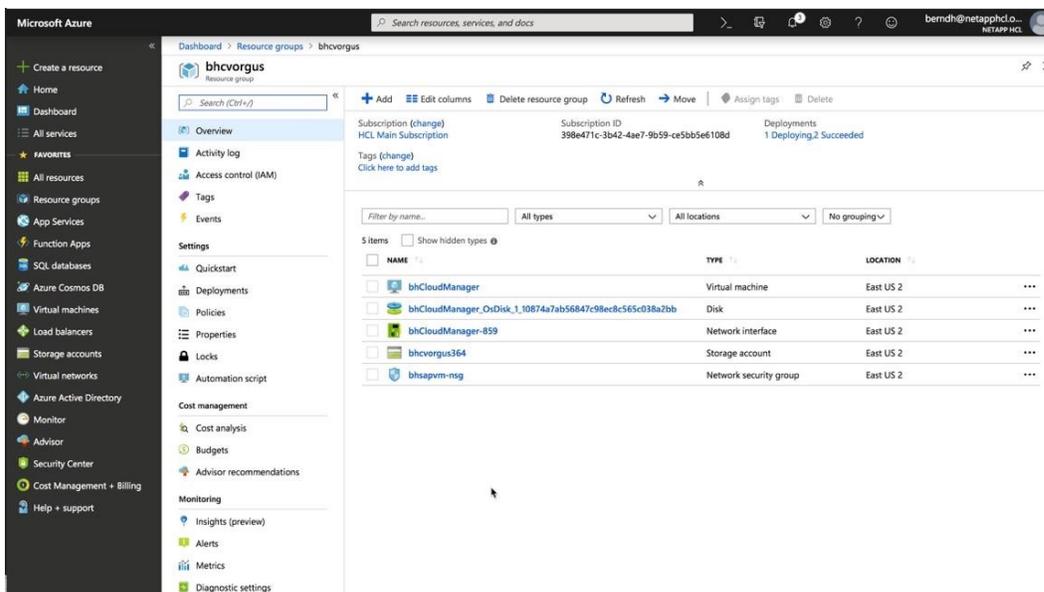
- Specify the network security group (access control list) to secure your OnCommand Cloud Manager instance.



- Select Go to start the deployment. The deployment takes approximately seven minutes to complete. After a successful deployment, the wizard will automatically open your OnCommand Cloud Manager web portal.



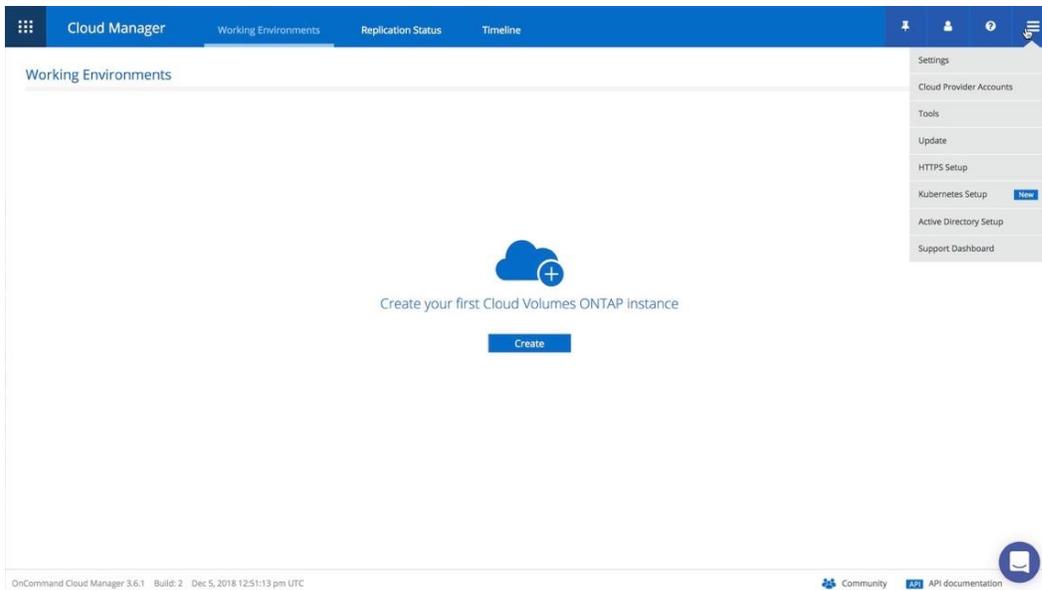
12. In the Azure portal, you can look for the objects that have been deployed by selecting the resource group where you deployed OnCommand Cloud Manager.



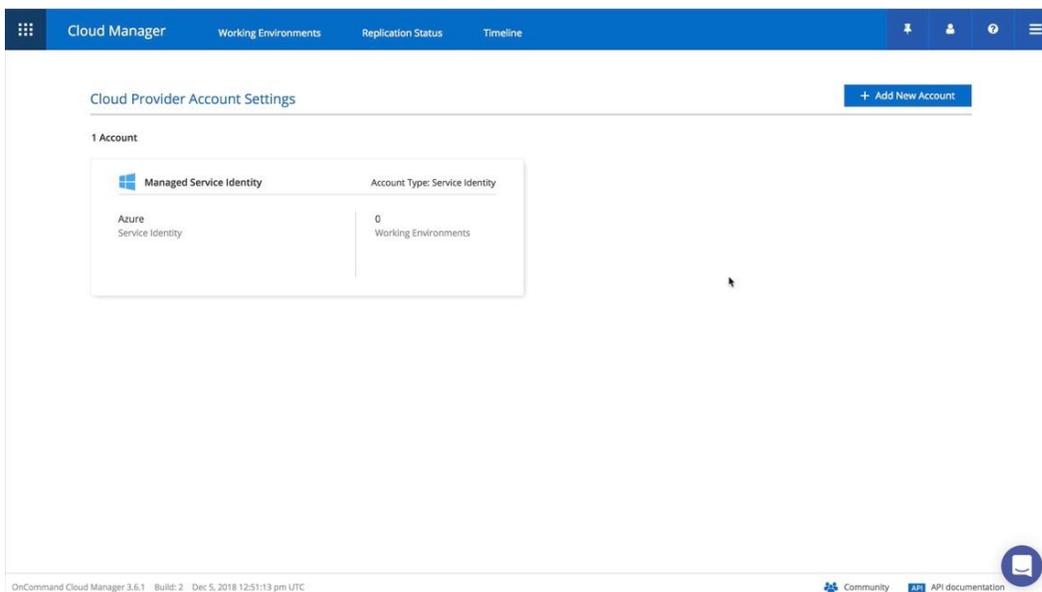
Add Microsoft Azure Subscription

As a final step you can configure OnCommand Cloud Manager - Cloud Providers to include all Azure subscription's you want to use with OnCommand Cloud Manager.

1. In OnCommand Cloud Manager, open the menu on the upper-right corner and select Cloud Provider Accounts.



2. Cloud Manager displays all of the configured cloud provider accounts. Select Add New Account to add a new account.

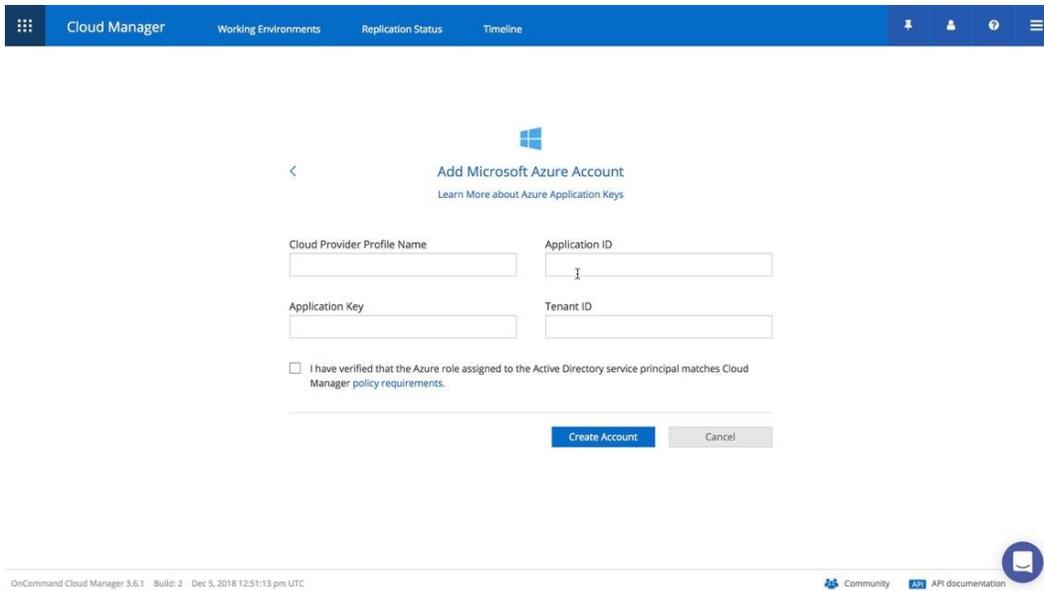


3. Enter the following Azure credentials:

- Application ID
- Application key
- Tenant ID

Specify the name of the new Cloud Provider profile. Click Create Account to save the settings.

Note: You might need to ask your company's Azure account administrator for this information.

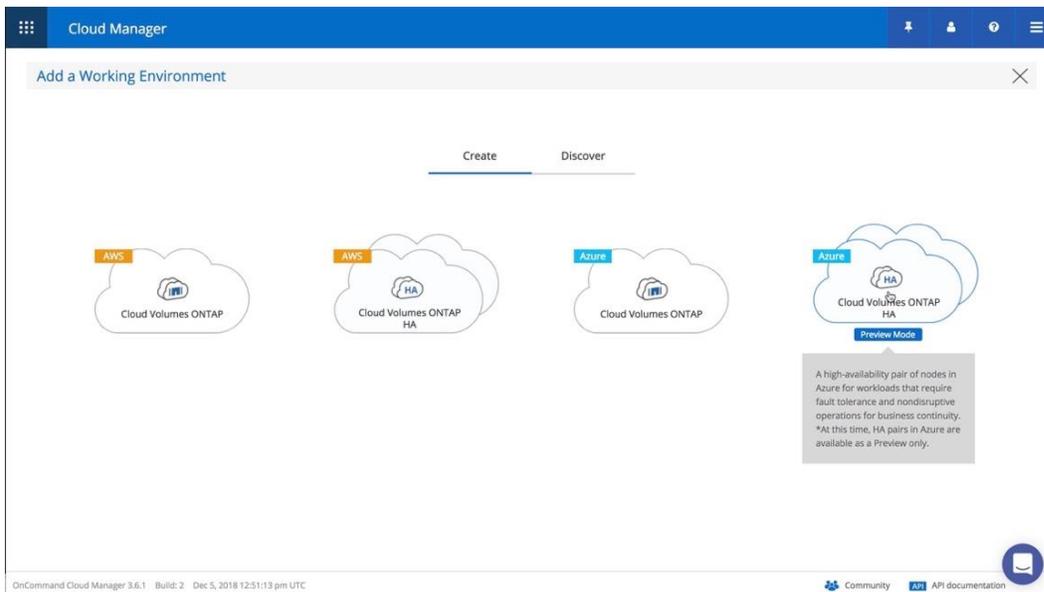


5.3 Create and Configure Cloud Volumes ONTAP

To create and configure Cloud Volume ONTAP, complete the following steps:

1. After you successfully create OnCommand Cloud Manager, deploy the first Cloud Volumes ONTAP instance by clicking Create in the OnCommand Cloud Manager Web UI.

This starts the Cloud Volumes ONTAP creation workflow.

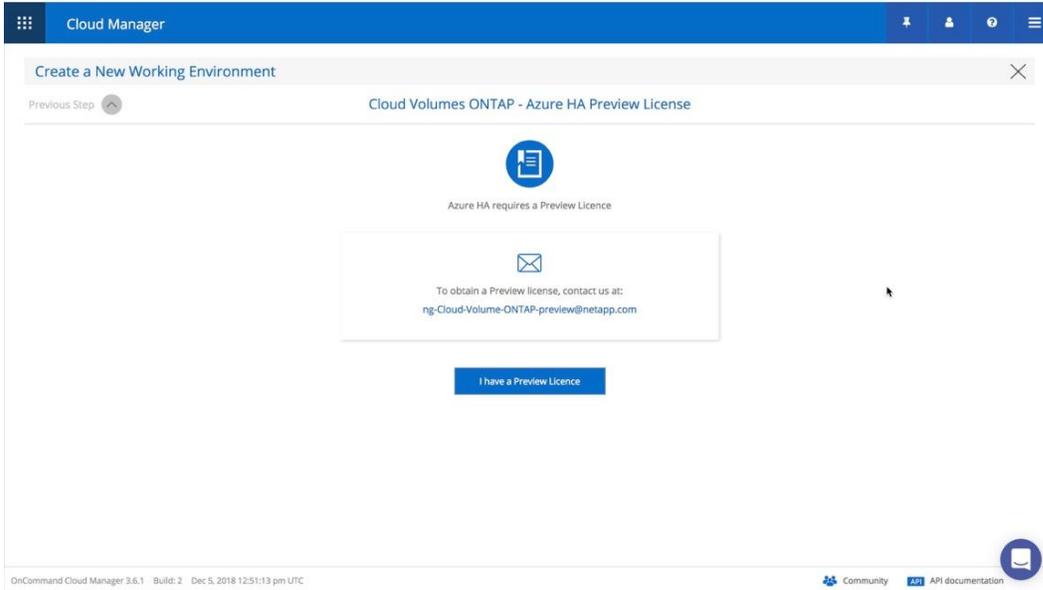


2. For Microsoft Azure there are two versions of Cloud Volumes ONTAP:
 - **Cloud Volumes ONTAP is a single node ONTAP.** This option is a perfect selection for noncritical systems. It's running only on one virtual instance without high availability.

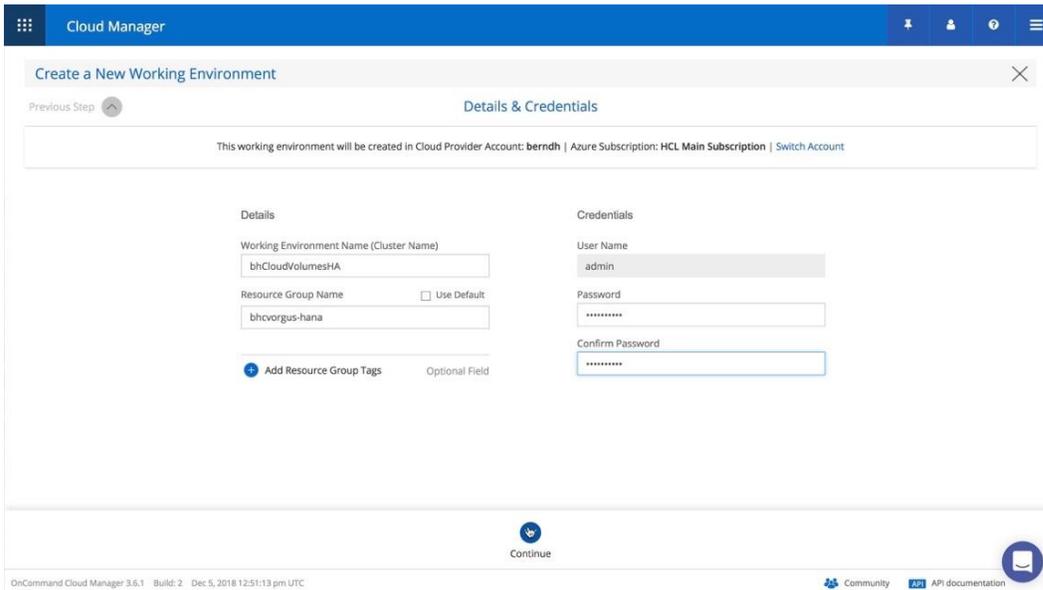
- **Cloud Volumes ONTAP HA.** Selecting this option will install a two node ONTAP cluster that offers high availability. In case one ONTAP instance fails, the second instance takes ownership of the storage and network configuration and allows an automated failover capability with zero downtime. The automated setup includes a Microsoft Azure Load Balancer to facilitate the cluster failover. As at Q1CY19, Cloud Volumes ONTAP HA is currently in preview and requires a specific preview license.

Note: In this example, Cloud Volumes ONTAP HA is selected.

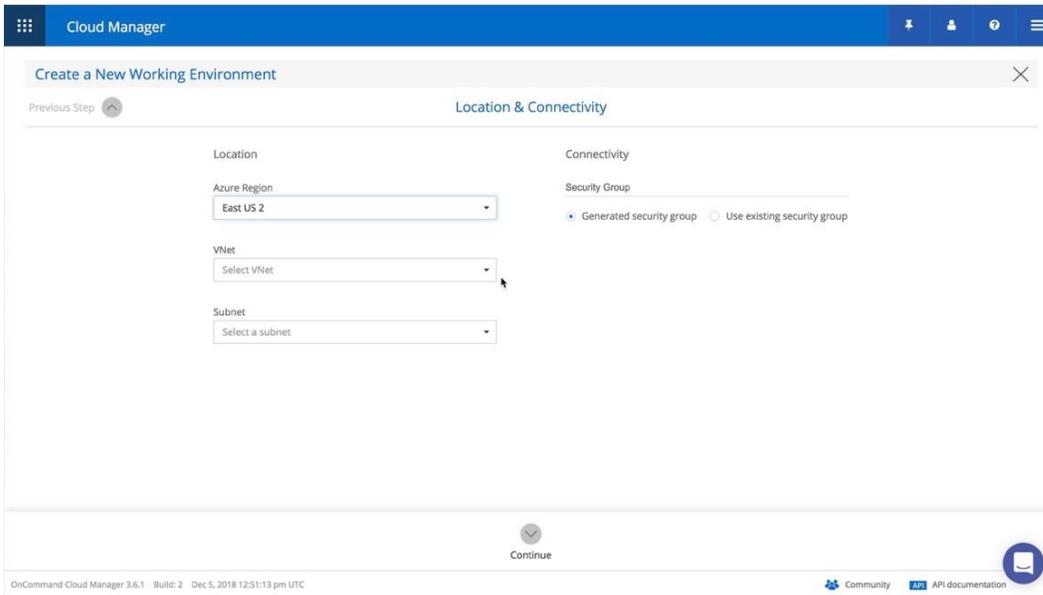
Click I Have a Preview License.



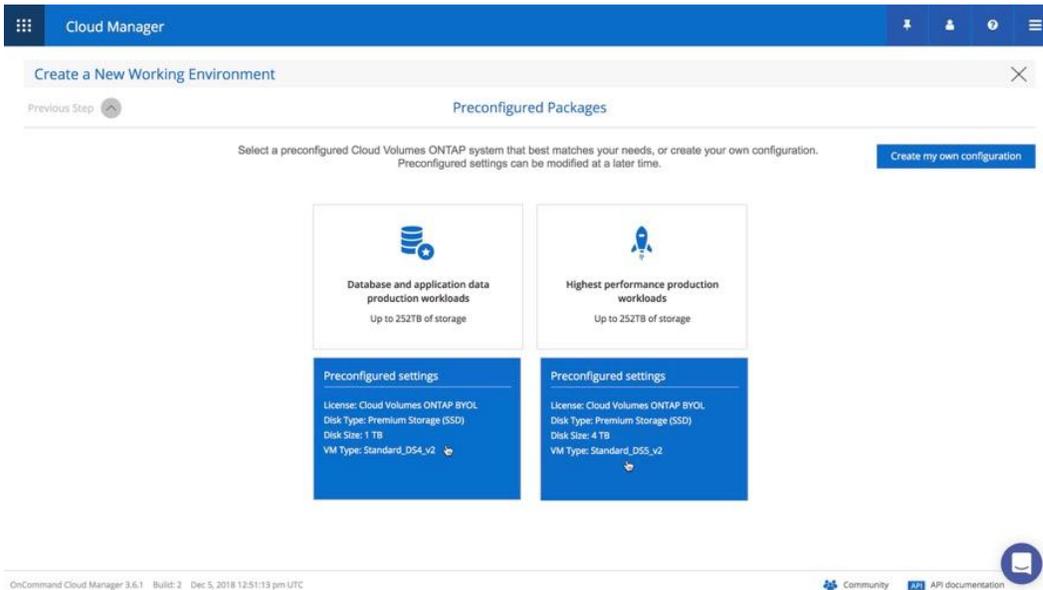
3. Specify the cluster name and the ONTAP administrator password. Select Continue.



4. Specify the Microsoft Azure region, VNet, and the subnet in which you want to deploy your Cloud Volumes ONTAP instances. You can specify an existing network security group or OnCommand Cloud Manager can create a new security group on your behalf. Select Continue.

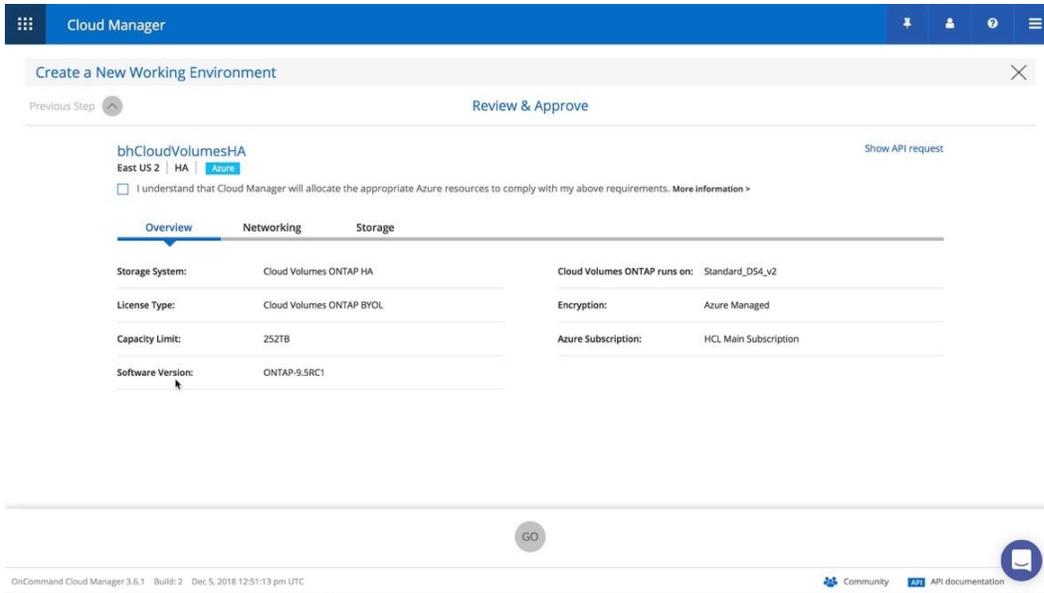


5. Enter the preview license code and then choose between two predefined packages for Cloud Volumes ONTAP HA. In this example, the package using 1TB drives is selected because the total capacity requirement is not that high, but we want to have more disks for better performance. Select the left preconfigured package “Database and application data production workloads”.

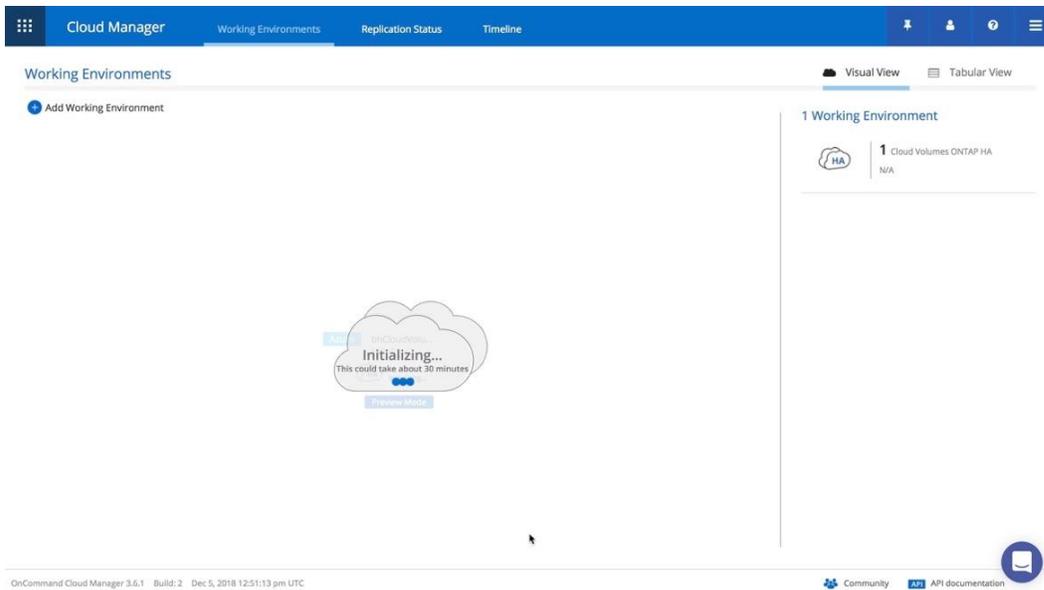


6. On the next two screens, select the NetApp support configuration, and if you want, create a first initial volume. We will skip this to continue the configuration after the two ONTAP instances have been created.

A summary page displays all the settings. Confirm that you understand that OnCommand Cloud Manager will create Microsoft Azure resources on your behalf. Click Go to start the creation.



7. OnCommand Cloud Manager starts the instance creation and the initial configuration of ONTAP. After approximately 30 minutes, Cloud Volumes ONTAP HA is up and running.



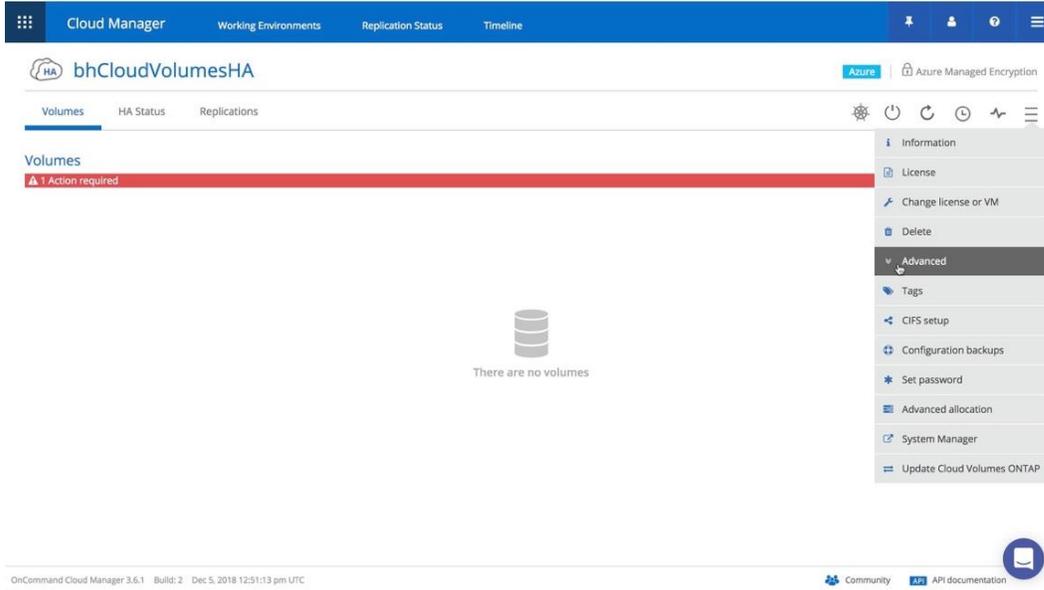
8. After the Cloud Volumes ONTAP cluster has been created, use Cloud Manager to familiarize yourself with navigation and the items Cloud Manager has provisioned.

5.4 Configure and Prepare the Required Storage

Before you can create volumes, configure Cloud Volumes ONTAP and add additional capacity. The initial package created one aggregate attached to the first Cloud Volumes ONTAP node containing only a single disk.

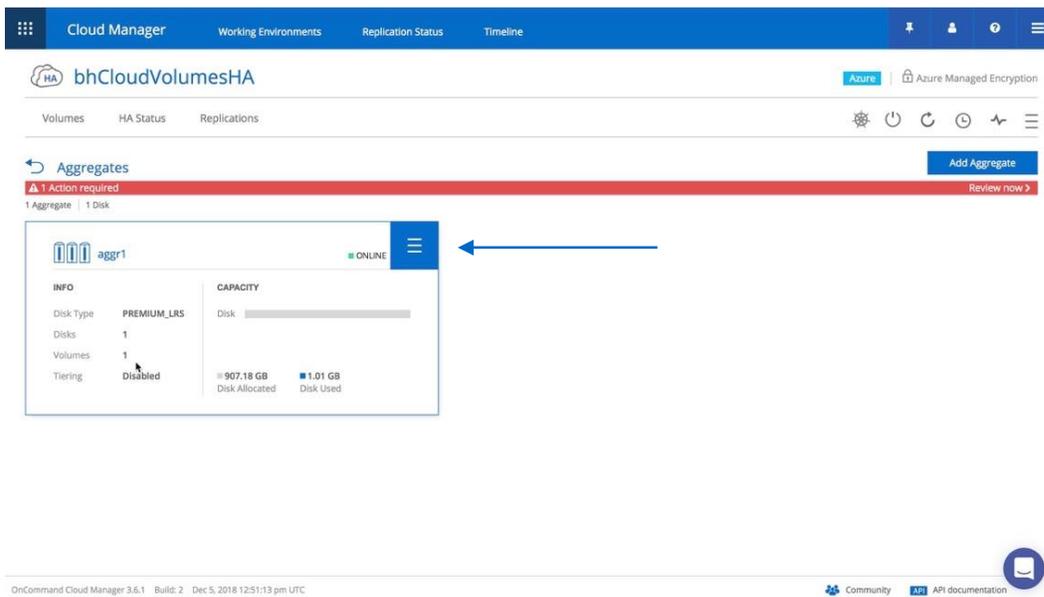
1. You can change the configuration by using the Advanced Storage Configuration menu. Select the double cloud icon in OnCommand Cloud Manager to navigate to Resource view and display the

available volumes. Select the menu, open the Advanced option, and then select Advanced Allocation to gain full control of the creation of aggregates and addition of disks.

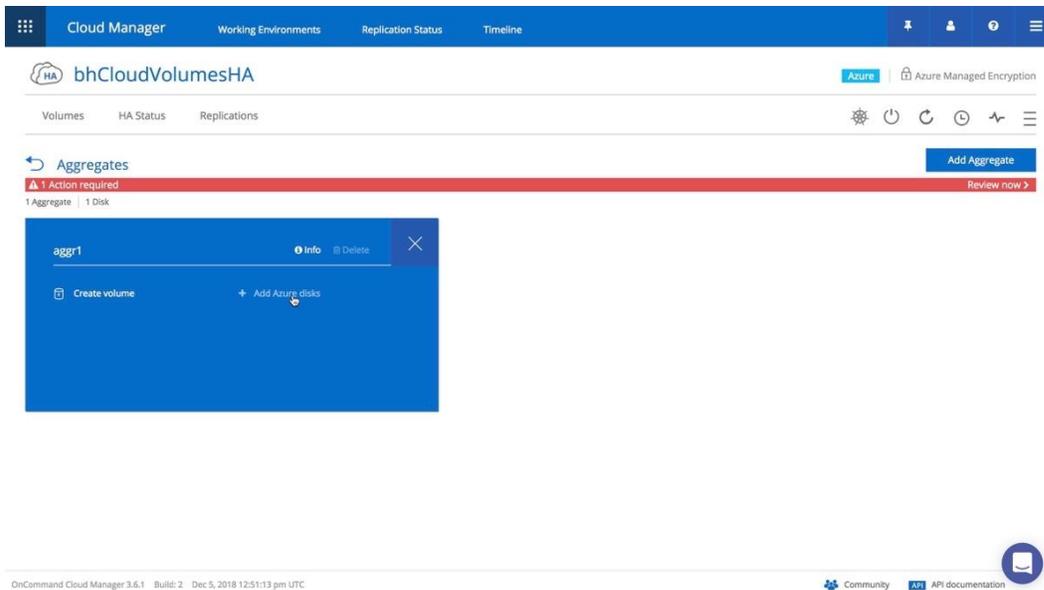


Note: In this example, the red warning bar 1 Action Required displays because only a four-week temporary license was installed. You can ignore the warning.

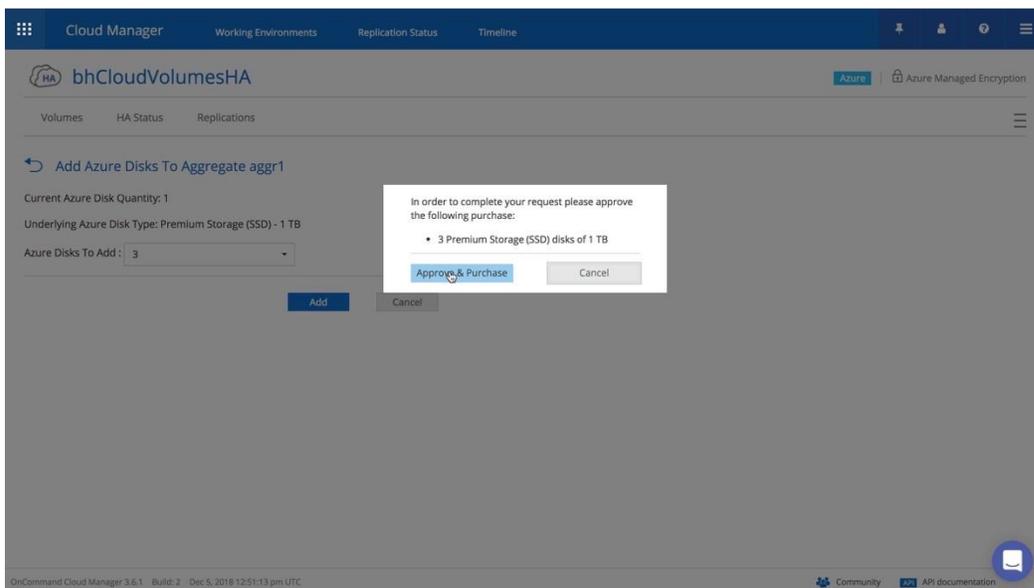
- The aggregate that was initially created, `aggr1`, is now displayed as a tile. Each aggregate has its own tile. Click the Tile Menu icon.



- The Aggregate menu is displayed. Click Add Azure Disks



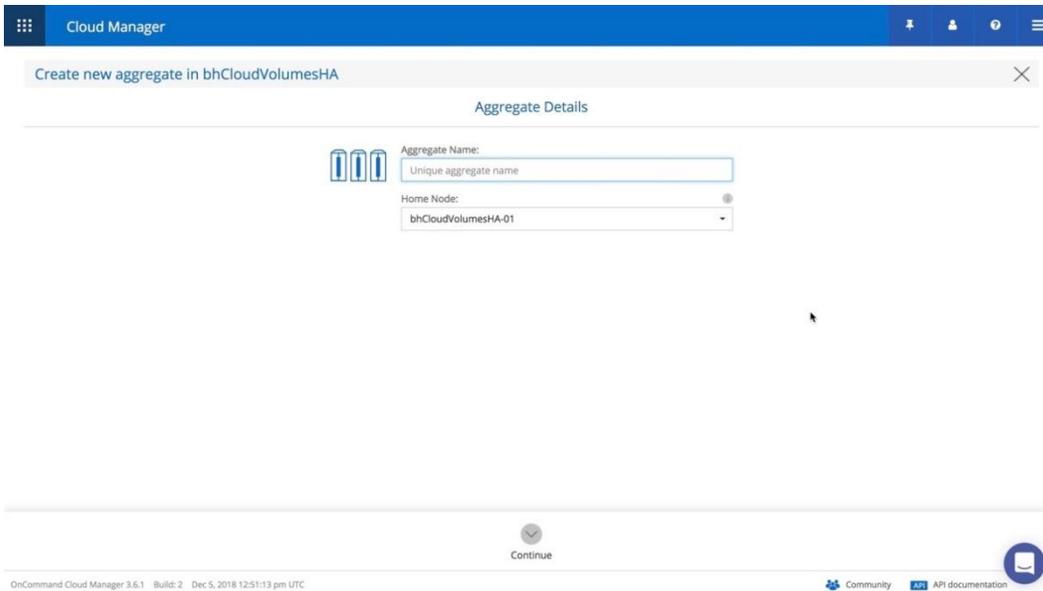
4. Select the number of disks you want to add and then approve.



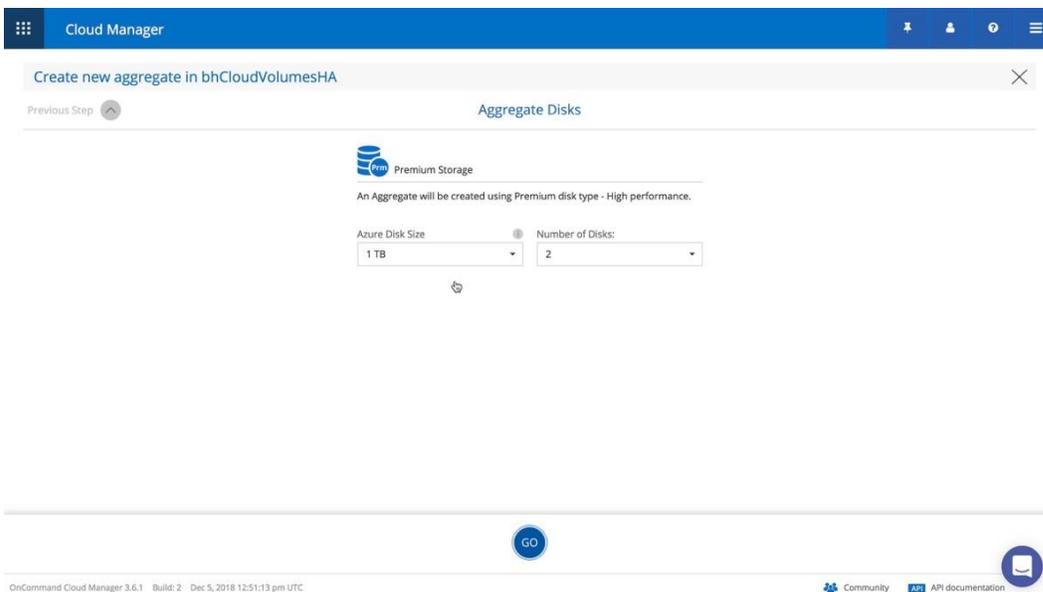
5. When the workflow returns to the aggregate view:
 - a. Wait until the number of disks have been added, and then click Add Aggregate to add a second aggregate.
 - b. Specify the new aggregate name `aggr2`.
 - c. Select the ONTAP node.

Note: In this example, the second cluster node is selected.

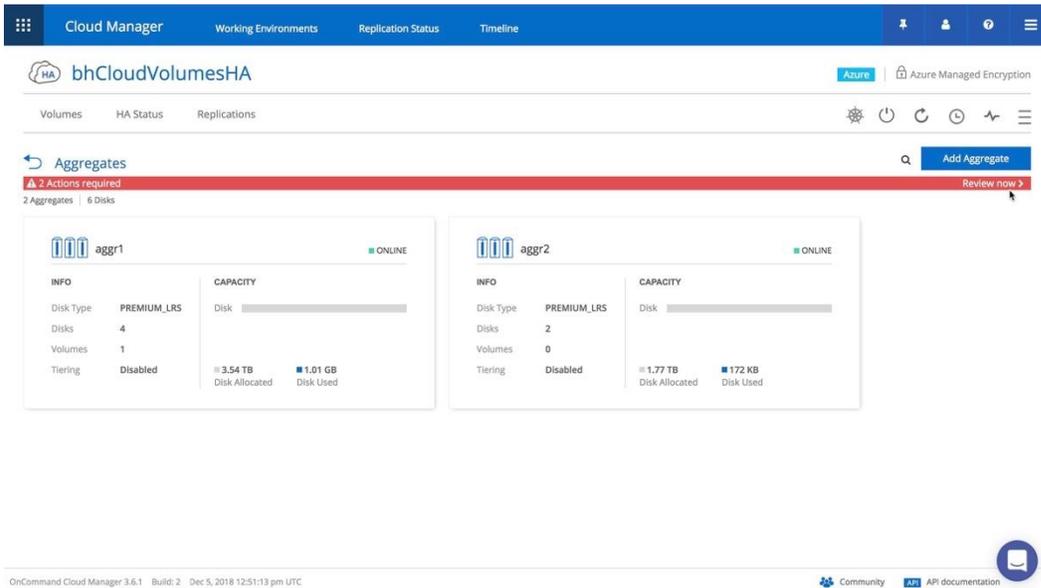
 - d. Click Continue.



6. Select the disk type and number of disks for the new aggregate. In this example, two 1TB premium disks were selected and the creation was started.



7. The workflow returns to the advanced allocation view, which now displays two aggregates.

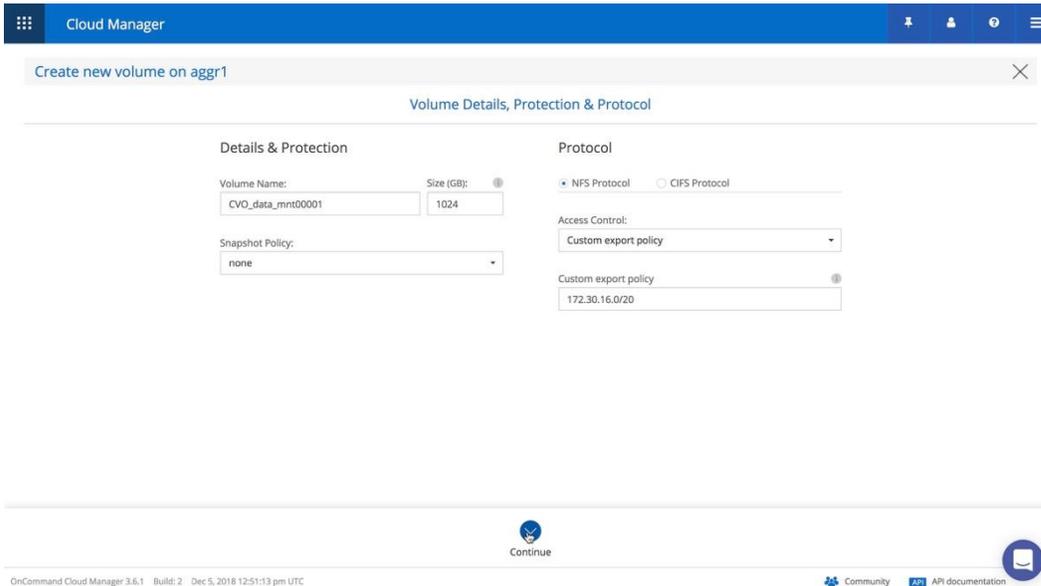


Create Volumes

As a final preparation, we will now create the volumes that are described in section 5.1, "Installation Planning." As stated, we will create the SAP HANA data volume on the 4TB aggr1 while we create the SAP HANA log volume on aggr2. The remaining volumes for the HANA shared file systems, as well as the NetWeaver shared files, are distributed on both Cloud Volume ONTAP instances. In addition, we will create a software share to store the SAP installation sources.

To create volumes, complete the following steps:

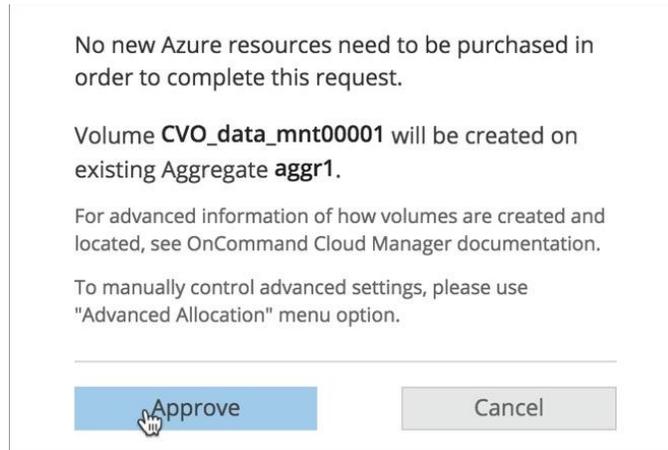
1. From Cloud Manager, select the aggr1 menu and select the Create Volume option.



2. Specify the volume name (see section 5.1, "Installation Planning"). For the HANA DATA and LOG volumes, the Snapshot policy should be None. For the shared file systems volumes, you can select a different Snapshot policy or select None. The export policy specifies the servers that are able to mount the volume by using an IP4 network CIDR; in this example, the whole subnet is selected.

Click Continue.

3. For the Cloud Manager volume confirmation, click Approve.



4. Repeat steps 1–3 for all remaining volumes. You should now have the following volumes available:
 - SAP HANA:
 - CVO_data_mnt00001
 - CVO_log_mnt00001
 - CVO_shared
 - SAP NetWeaver ABAP
 - NW5_shared
 - Software share
 - Software

5.5 Prepare the Operating System and Mount the Volumes

To mount the required file systems, complete the following tasks:

- Identify the mount options.
- Mount the volume on a temporary mount point to create the required folders.
- Create the required path structure on the OS level and set the permissions.
- Add the mount commands to `/etc/fstab`.
- Mount the volumes.

To install an SAP NetWeaver application server as well as the SAP HANA database, NetApp recommends using a specific SAP Linux edition. These Linux editions, such as SLES for SAP, are delivered with an installed NFS client and the relevant Linux tuning settings required by SAP. In this example, the following SAP editions were used (as shown in Figure 18):

- SLES for SAP 12 SP3 for the SAP NetWeaver instance NW5
- SLES for SAP 12 SP3 for the SAP HANA 2.0 SPS3 instance CVO

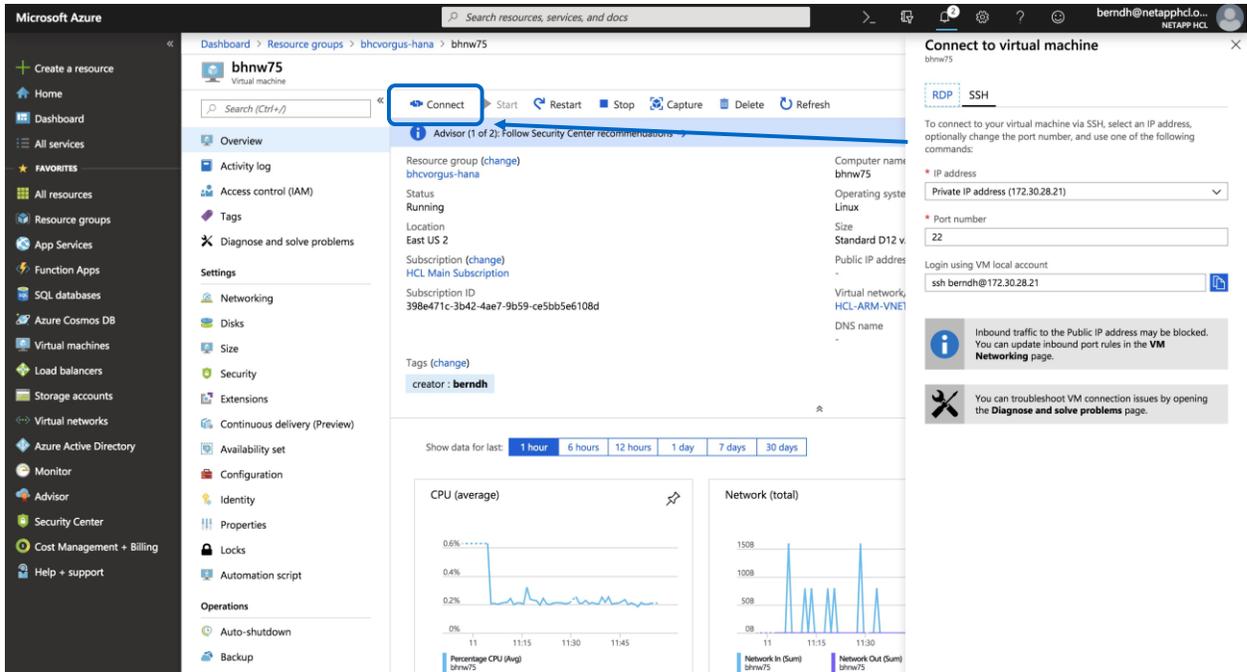
Note: SAP note 1928533 lists the supported operating systems for SAP Application on Microsoft Azure.

Prepare and Mount the Volume for the SAP NetWeaver Instance

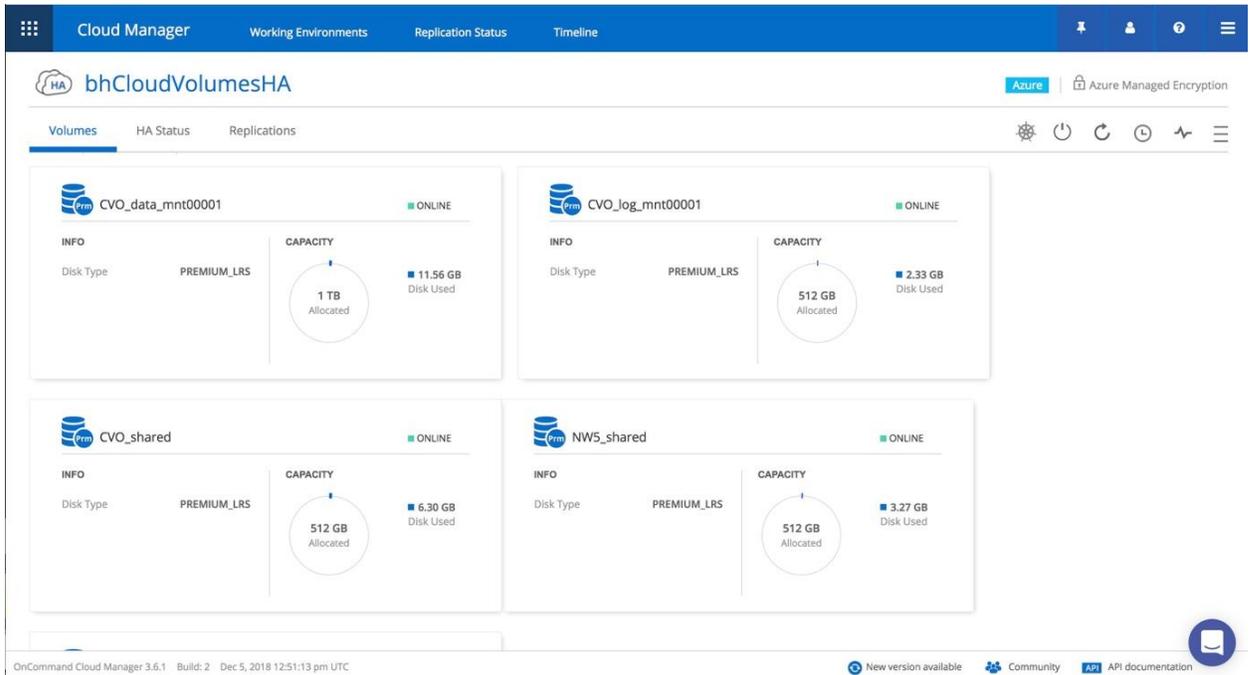
To prepare and mount the volume for the SAP NetWeaver instance, complete the following steps:

1. After the SAP NetWeaver instance is deployed, use the Azure portal to connect to the instance. Select the instance and click Connect to open the Connect to Virtual Machine property window on the right side. The `ssh` command to connect to the VM is displayed.

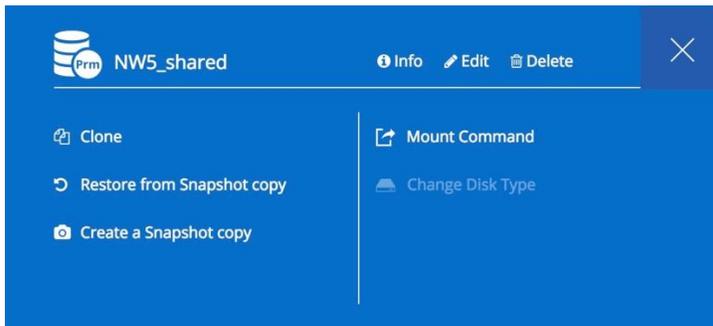
Note: The connection method varies depending on your configuration. In this example, we specified a user name and password to log in to the Linux operating system.



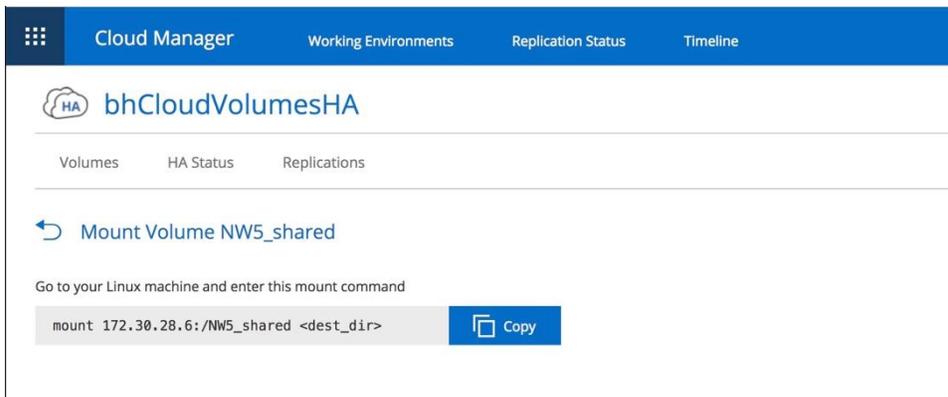
2. After you log in to the operating system, prepare the operating system (according to the SAP recommendations) to install SAP NetWeaver on Linux.
3. To find out the mount command for the volume, use the OnCommand Cloud Manager. Select Volume `NW5_shared` from the Volume view.



- Click the menu button that is displayed when you hover the mouse over the ONLINE icon. Select Mount Command.



- OnCommand Cloud Manager displays the mount command for this volume. You can cut and paste this command to use again in the terminal session.



6. According to NetApp guidelines, you must add additional mount options to mount the volume. As a first step, mount the volume to a temporary mount folder.

```
# sudo mount -t nfs -o rw,hard,intr,noatime,nolock,rsize=65536,wsiz=65536,nfsvers=3,tcp
10.0.0.4:/NW5_shared /mnt
```

7. After the volume is mounted, create the subfolders and set the rights.

```
# cd /mnt
# sudo mkdir usr-sap
# sudo mkdir trans
# sudo mkdir sapmnt
# sudo chmod 777 *
# cd /
# sudo umount /mnt
```

8. Set the mount points at the operating system level.

```
# sudo mkdir -p /usr/sap/NW5
# sudo mkdir -p /usr/sap/trans
# sudo mkdir -p /sapmnt/NW5
#
# sudo chmod -R 777 /sapmnt
# sudo chmod -R 777 /usr/sap
```

9. Add the entries into /etc/fstab by means of sudo /etc/fstab.

```
# sudo cat /etc/fstab
...
#
# NetWeaver 75 ABAP
#
172.30.28.6:/NW5_shared/trans /usr/sap/trans nfs
rw,vers=3,hard,timeo=600,rsiz=1048576,wsiz=1048576,intr,noatime,nolock 0 0
172.30.28.6:/NW5_shared/usr-sap /usr/sap/NW5 nfs
rw,vers=3,hard,timeo=600,rsiz=1048576,wsiz=1048576,intr,noatime,nolock 0 0
172.30.28.6:/NW5_shared/sapmnt /sapmnt/NW5 nfs
rw,vers=3,hard,timeo=600,rsiz=1048576,wsiz=1048576,intr,noatime,nolock 0 0
#
# Software share
#
....
```

10. Mount all of the required volumes.

```
# sudo mount -a
```

11. Verify whether all of the file systems are mounted.

```
berndh@bhnw75:~> sudo df
Filesystem                                1K-blocks    Used Available Use% Mounted on
devtmpfs                                  14393000      0 14393000  0% /dev
tmpfs                                      21599680      0 21599680  0% /dev/shm
tmpfs                                      14399784    18196 14381588  1% /run
tmpfs                                      14399784      0 14399784  0% /sys/fs/cgroup
/dev/sda2                                  29798460 2360292 25901452  9% /
/dev/sda1                                  999320      46552  900340  5% /boot
/dev/sdb1                                  206291640  60688 195728912  1% /mnt/resource
tmpfs                                      2879960      0  2879960  0% /run/user/1002
tmpfs                                      2879960      0  2879960  0% /run/user/1000
172.30.28.6:/NW5_shared/trans             510027392 3431040 506596352  1% /usr/sap/trans
172.30.28.6:/NW5_shared/usr-sap           510027392 3431040 506596352  1% /usr/sap/NW5
172.30.28.6:/NW5_shared/sapmnt            510027392 3431040 506596352  1% /sapmnt/NW5
172.30.28.6:/software                      1020054784 39444352 980610432  4% /mnt/software
```

The volumes are mounted and the VM is ready for the SAP installation.

Prepare the OS and Mount the Volumes for the SAP HANA Instance

To prepare the operating system and mount the volumes for the SAP HANA instance, complete the following steps:

1. Log in to the Azure instance provisioned for SAP HANA (see the previous section on how to identify the connection string in Azure portal). To efficiently run SAP HANA, you must set additional Linux kernel settings.
2. Prepare the operating system with specific settings for SAP HANA (as described in [TR-4435: SAP HANA on NetApp AFF Systems with NFS](#)):
 - a. Adapt the kernel settings for the operating system. In this example, SUSE SLES 12.
 - b. 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
```

- 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
```

3. Create the required subdirectories in the `CVO_shared` volume.

Note: For details about obtaining the IP and NFS export path for the volumes by using OnCommand Cloud Manager, see the previous section.

```
# sudo mount -t nfs -o rw,hard,intr,noatime,nolock,rsize=65536,wsiz=65536,nfsvers=3,tcp
10.0.0.5:/CVO_shared /mnt
# cd /mnt
# sudo mkdir shared
# sudo mkdir usr-sap
# sudo chmod 777 *
# cd /
# sudo umount /mnt
```

4. Create the mount points and set the permissions.

```
# sudo su -
# mkdir -p /hana/data/CVO/mnt00001
# mkdir -p /hana/log/CVO/mnt00001
# mkdir -p /hana/shared
# mkdir -p /usr/sap/CVO/home
#
# chmod -R 777 /hana/log/CVO
# chmod -R 777 /hana/data/CVO
# chmod -R 777 /hana/shared
# chmod -R 777 /usr/sap/CVO
```

5. Edit `/etc/fstab` to mount the volumes and run `# sudo mount -a`.

Note: The mount options are extracted from [TR-4435: SAP HANA on NetApp AFF Systems with NFS](#).

```
# cat /etc/fstab
...
```

```

10.0.0.5:/CVO_shared/usr-sap /usr/sap/CVO/home nfs
rw,hard,timeo=600,intr,noatime,nolock,rsize=65536,wsiz=65536,nfsvers=3,tcp 0 0
10.0.0.5:/CVO_shared/shared /hana/shared nfs rw,hard,
timeo=600,intr,noatime,nolock,rsize=65536,wsiz=65536,nfsvers=3,tcp 0 0
10.0.0.5:/CVO_data_mnt00001 /hana/data/CVO/mnt00001 nfs rw,hard,
timeo=600,intr,noatime,nolock,rsize=65536,wsiz=65536,nfsvers=3,tcp 0 0
10.0.0.5:/CVO_log_mnt00001 /hana/log/CVO/mnt00001 nfs rw,hard,
timeo=600,intr,noatime,nolock,rsize=65536,wsiz=65536,nfsvers=3,tcp 0 0
Mount the additional volumes with sudo mount -a.

```

Post-Installation Configuration (Optional)

As an optional step, create the volume `CVO_backup` to provide storage for file-based data backups and automated log backups. To simplify the configuration, mount the volume at the default backup location. This setting must be configured after you install SAP HANA. Also, `<nn>` must be replaced by the selected system number during the installation process. Table 7 lists the SAP HANA optional file system.

Table 7) SAP HANA optional file system.

Volume	Mount Path	Required Capacity (GB)	Comment
CVO_backup	/usr/sap/CVO/HDB<nn>/backup	4096GB	Default backup path

After a successful installation, you can create and mount the volume `CVO_backup` to the default backup location and start the first file-based backup to change the SAP HANA database to a standard logging mode.

Note: You can use a different mount point, but this requires an SAP HANA configuration change for the backup location.

5.6 SAP Installation

Use the software volume on Cloud Volumes ONTAP to download and extract the SAP required installation sources. You can proceed with the following tasks:

1. Install the SAP HANA database on the prepared database instance.
2. Install the SAP NetWeaver instance on the prepared application instance.
3. Optionally, install SAP HANA studio to manage the SAP HANA database and SAP GUI to login to the SAP NetWeaver instance.

Note: These steps are made in accordance with the SAP installation guides and procedures for the related version. For more information, see the relevant SAP installation guides.

6 Data Protection and System Migration

6.1 Data Protection

The main benefits of running SAP HANA on ONTAP are the advanced features and software integrations of ONTAP for optimizing SAP data protection tasks. Cloud Volumes ONTAP, from an ONTAP features and interfaces perspective, is identical to ONTAP running on a NetApp hardware storage system in an on-premises data center. Therefore, all tools and integrations work seamlessly with Cloud Volumes ONTAP in Microsoft Azure.

NetApp SnapCenter software offers enterprise-grade data protection capabilities for both SAP anyDB deployments as well as SAP HANA. SnapCenter software can be deployed on Microsoft Azure for Cloud Volumes ONTAP and integrated similarly as with on-premises. The following NetApp technical reports provide detailed descriptions on how to install, configure, and use SnapCenter for SAP HANA.

- [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#)
- [TR-4646: SAP HANA Disaster Recovery with Asynchronous Storage Replication](#)
- [TR-4667: Automating SAP System Copies Using the SnapCenter 4.0 SAP HANA Plug-In](#)
- [TR-4719: SAP HANA System Replication Backup and Recovery with SnapCenter](#)

6.2 System Migration

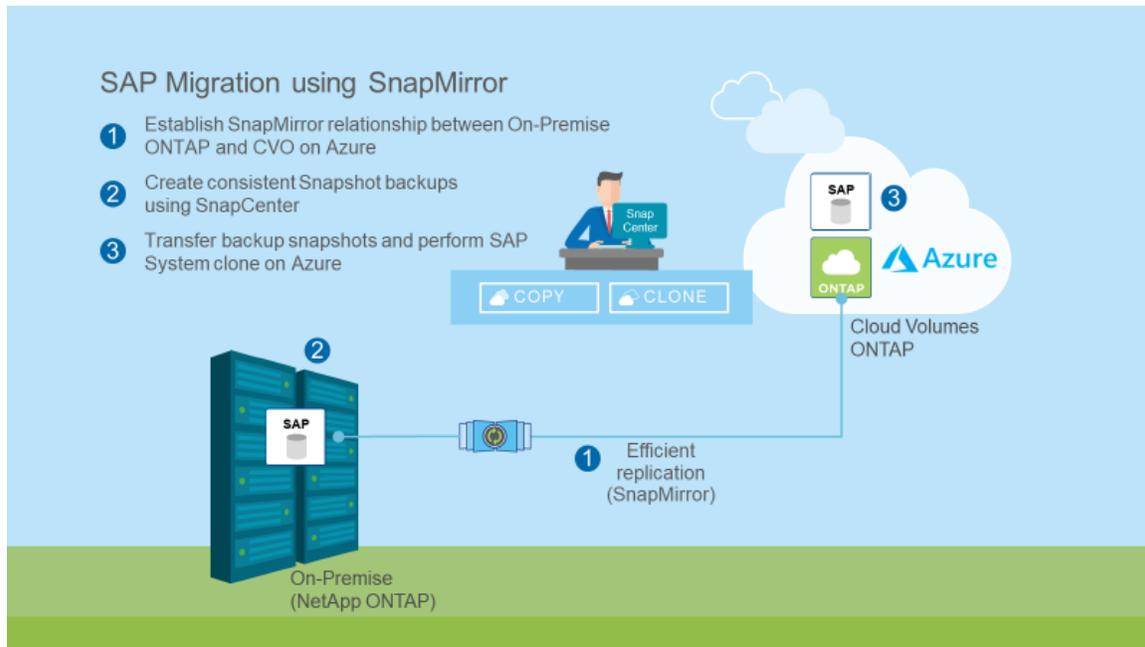
Up to this point, this document has described how to install and operate (from scratch) SAP systems on Microsoft Azure by using Cloud Volumes ONTAP as new systems. However, the majority of systems, especially test and development systems, are derived from existing systems of an on-premises SAP installation. Depending on the target use case and the specific SAP system type and database, there are various ways to migrate systems into Azure. This section describes three different approaches to migrate SAP systems leveraging different technologies of the NetApp Data Fabric to optimize the migration. In addition, it provides a short overview of how these three approaches can be further automated by using the SAP Landscape Management (SAP LaMa) tool.

For more information about NetApp integration into SAP LaMa, see [TR-4018: Integrating NetApp ONTAP Systems with SAP Landscape Management](#).

Migrating SAP Systems Using NetApp SnapMirror Technology

If your SAP landscape on-premises is already running on a NetApp ONTAP system, a convenient way to migrate SAP systems to Cloud Volumes ONTAP in Azure is to use SnapMirror replication. This is especially interesting if you plan to permanently operate your SAP landscape in a hybrid deployment with your production systems on-premises and some of your nonproduction systems on Azure with the need for regular system refreshes. In this example, you can use NetApp Snapshot backup functionality for your SAP systems integrated in NetApp SnapCenter, to create fast application-consistent backups of the on-premises systems. Those Snapshot copy backups can then be efficiently transferred to the Cloud Volumes ONTAP instances in Azure by using SnapMirror technology. After the backup Snapshot copies have been replicated, the backup images can be used to perform regular SAP system refreshes of the target systems in the Azure. Figure 19 shows a high-level overview of the workflow.

Figure 19) SAP migration using SnapMirror technology.



To configure SnapMirror between an on-premises system and a Cloud Volumes ONTAP instance on Azure, the required network connectivity for the cluster peering must be established. After this is complete, a SnapMirror replication can be initialized the same way as in a pure on-premises environment. A more convenient and easy way to configure the replication is to use NetApp OnCommand Cloud Manager that has been provisioned to install Cloud Volume ONTAP. Within OnCommand Cloud Manager, on-premises ONTAP systems can be discovered, registered, and managed. If the source ONTAP system and the Cloud Volumes ONTAP instance are both visible in OnCommand Cloud Manager, a SnapMirror replication can be initialized by a simple drag and drop of the source system on to the target system in the systems overview. A workflow then guides you through the additional required input steps and starts the initial baseline copy. For detailed information about SnapMirror configuration using OnCommand Cloud Manager, see [Cloud Manager and Cloud Volumes ONTAP Documentation](#) and navigate to Managing Data Across a Hybrid Cloud. After the initial baseline copy is finished, the ongoing data transfer using SnapMirror is bandwidth and time efficient, since only changed data on the storage block level needs to be replicated.

This approach for migrating SAP systems using SnapMirror can also be beneficial if you plan to perform a database migration as part of the move to Microsoft Azure. Use the method described above to repeatedly replicate your source systems to Azure, run the database test migrations within Azure to capture all required changes and perform necessary preparations to get ready for the final productive migration. This approach helps to reduce the on-premises infrastructure footprint—there is no need to build staging systems, while at the same time reduces risk and improves the overall quality of the migration.

Migrating SAP Systems from Non-NetApp Environments Using NetApp Cloud Sync

If your current on-premises SAP landscape is not running on ONTAP, using SnapMirror technology to migrate SAP systems is not possible. Even if you are currently running your landscape on ONTAP, if you're simply looking for a one-time copy of a SAP system or landscape to Azure, you won't gain the full benefits of SnapMirror because of the required initial baseline copy. However, there is another NetApp Data Fabric cloud service that can be used in these cases: NetApp Cloud Sync. Cloud Sync offers rapid and secure data synchronization. Whether you need to transfer files between on-premises NFS or CIFS

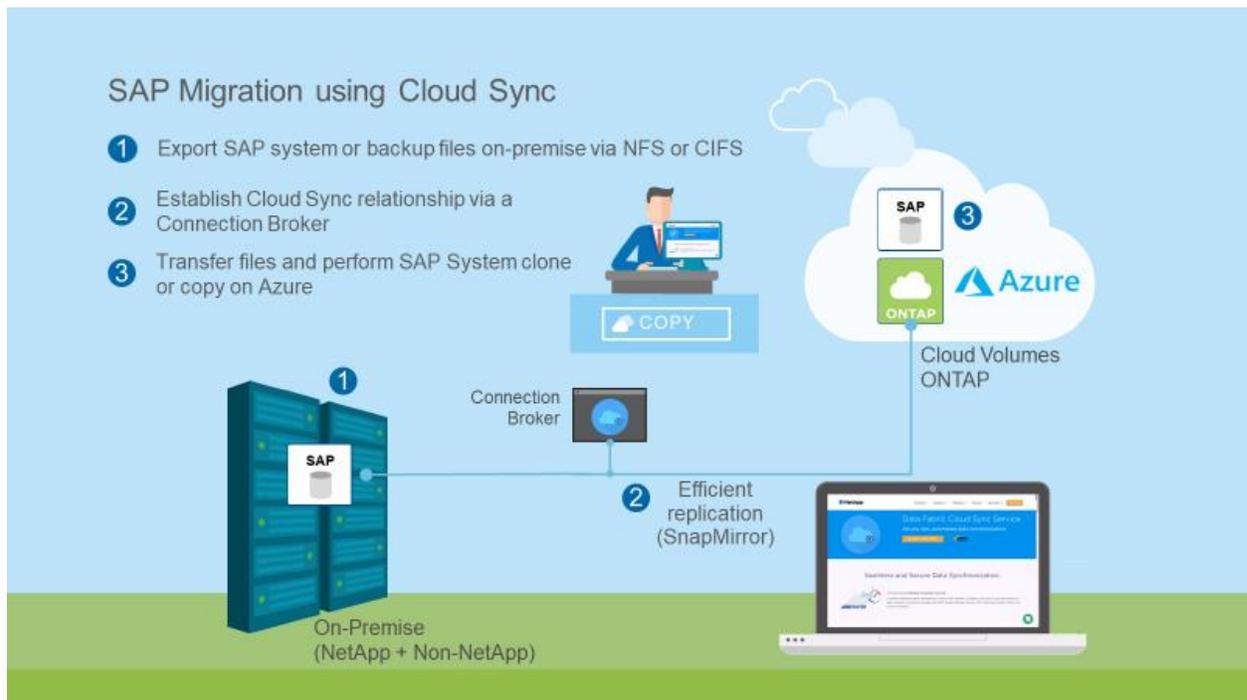
file shares, Microsoft Azure Blob, Amazon S3 object format, IBM Cloud Object Storage or NetApp StorageGRID® Webscale appliance, Cloud Sync moves the files where you need them quickly and securely ([Cloud Sync Synchronization](#)) and it can be managed by a simple cloud-like web interface or through well-defined REST API calls.

For SAP migration from on-premises to Cloud Volumes ONTAP in Azure, Cloud Sync can be efficiently used to copy all required SAP file systems and database files from any source NFS or CIFS share to the NFS volumes in Cloud Volumes ONTAP. Cloud Sync uses a Connection Broker instance to transfer data from source to target. Depending on the specific network layout, the connection broker can either be installed on-premises or in Azure. There are two approaches to move a system:

- For a nonproduction system, the required SAP and database file systems can be copied directly to the Cloud Volumes ONTAP volumes. If the source system is not stored on NFS volumes or CIFS shares, operating system features in Linux or Windows can be used to temporarily export the local file systems as NFS or CIFS shares.
Note: With this approach, the source system must be shut down during the transfer (offline copy). Therefore, this approach is only possible if the required downtime for the system is acceptable by the business.
- As an alternative, a regular file-based backup of the source system can be created and stored on any NFS or CIFS file share in the on-premise environment. Cloud Sync can then be used to efficiently transfer the backup files to Cloud Volumes ONTAP. After the files are available on Azure, they can be used to perform a homogenous SAP system copy using the available backup files.

The high-level process for such a migration is shown in Figure 20.

Figure 20) SAP migration using Cloud Sync.



Using file-based backup for migrating data to Microsoft Azure by using operating system tools such as rsync or robocopy. However, compared to those tools, Cloud Sync offers a more efficient and faster data transfer mechanism through massive parallelization and an advanced and enterprise-grade interface and REST API that allows easy integration into higher level tools and workflows.

SAP HANA Migration Using SAP HANA System Replication

SAP HANA is the most important database for SAP systems today; therefore, we use it as an example for a system migration using database specific tools and technologies. With SAP HANA, the HANA System Replication (HSR) feature can be used to migrate a SAP system from on-premises to Microsoft Azure. There are two main benefits of using HSR for the migration: failover times can be reduced significantly and HSR can be used independently of the compute and storage infrastructure in the on-premise datacenter.

As a first step, the replication target system must be installed on Azure. The description within this document should be used to prepare the Cloud Volumes ONTAP environment accordingly. The target HANA database can either be installed from scratch or based on a backup of the original on-premises database. In the latter case, one of the options mentioned in the previous sections (SnapMirror or Cloud Sync) could be used to transfer a Cloud Volumes ONTAP backup image to Microsoft Azure. When the target system is installed, it can be configured as a HSR target instance for the primary on-premises SAP HANA database. After the two systems are in sync, a SAP HANA failover operation can be triggered to hand over the operation to the target instance running on Cloud Volumes ONTAP in Azure. For more information about the installation and configuration of a SAP HANA System Replication scenario, see the [SAP HANA Administration Guide](#).

Migration Automation Using SAP Landscape Management

According to [SAP Landscape Management Solution Brief](#), SAP LaMa simplifies and standardizes the management of SAP operations with a powerful orchestration solution. You can automate repetitive, time-consuming administration tasks and tailor processes to your specific needs. In addition, you can centralize landscape operations and gain landscape-wide visibility through a single user interface. Some of the key features are:

- Centralized landscape management
- Automated system provisioning, such as copy and refresh
- Standard solution for SAP HANA and SAP S/4HANA

Since the earliest versions of SAP LaMa, powerful and deep integrations of NetApp functionality into SAP LaMa have been offered. All of the approaches for SAP migrations to Microsoft Azure can be integrated into and operated with SAP LaMa. In a similar fashion, Microsoft offers a SAP LaMa Adapter for Microsoft Azure that enables the automation of Microsoft Azure operations such as:

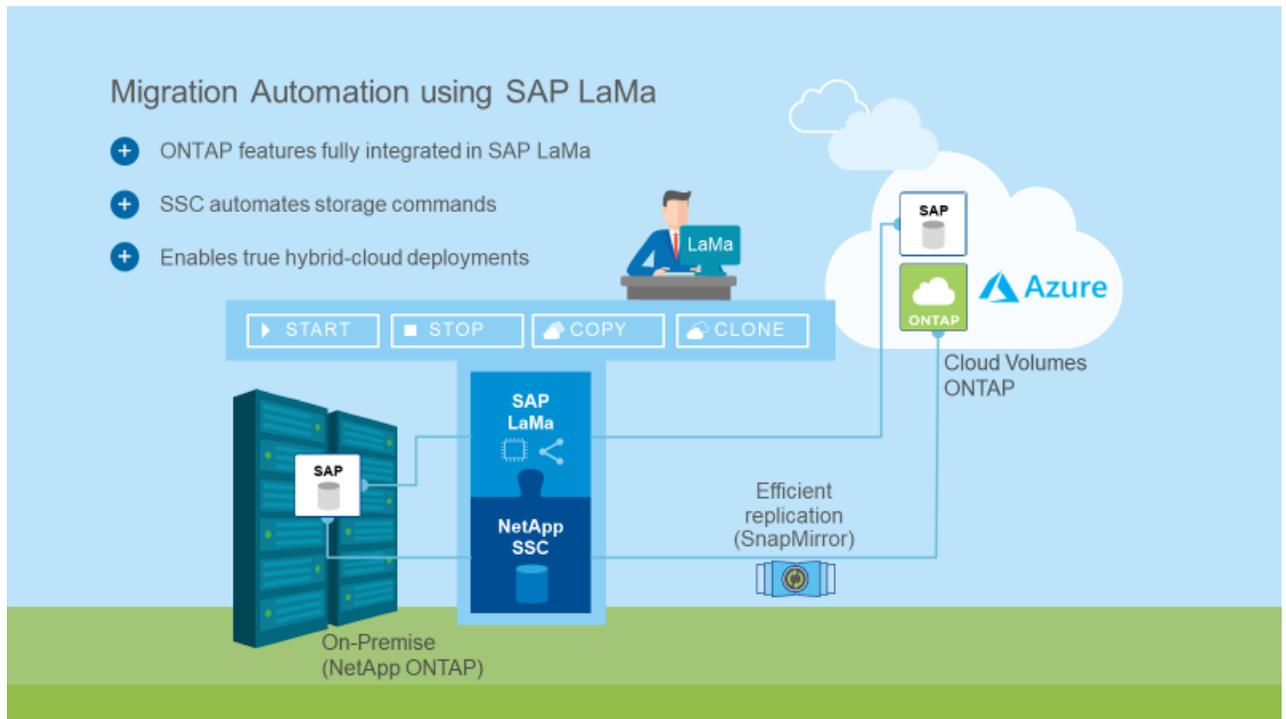
- Discover VMs
- Start/stop SAP systems and activate/deactivate VMs
- Microsoft Azure Template deployment
- System copy, clone, refresh and relocate

This section provides a brief overview of the different integrations.

The standard integration of NetApp ONTAP features, including Snapshot copies and SnapMirror replication, is provided by the NetApp Storage Services Connector. The NetApp Storage Services Connector plugs into the SAP LaMa enterprise storage interface and enables fast and efficient automated SAP system copy and refresh operations leveraging NetApp technology. The NetApp Storage Services Connector communicates with both the SAP LaMa system and one or more ONTAP environments. A Cloud Volumes ONTAP instance on Microsoft Azure can be seamlessly integrated in a hybrid cloud infrastructure managed by SAP LaMa and NetApp Storage Services Connector as, shown in Figure 21. After the Cloud Volumes ONTAP instance is registered in NetApp Storage Services Connector, the automated SAP system copies and refreshes can be performed from on-premises source systems to target systems running on Cloud Volumes ONTAP in Microsoft Azure.

For a detailed description of the NetApp Storage Services Connector installation and configuration with SAP LaMa, see [TR-4018: Integrating NetApp ONTAP Systems with SAP Landscape Management](#).

Figure 21) Migration automation using SAP LaMa.



Cloud Sync is not part of the NetApp Storage Services Connector plug-in; however, the migration process using Cloud Sync can be automated with SAP LaMa through the powerful custom extensibility interface. A description of this integration including a demo video can be found in the [NetApp blog post](#).

As previously mentioned, SAP LaMa is the standard solution from SAP to manage SAP HANA environments. This includes the configuration and operation of SAP HANA System Replication scenarios. New target system can be provisioned, and the system replication can be configured. Failover/takeover processes in an existing HSR configuration can be triggered by a single mouse click.

Where to Find Additional Information

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

- TR-4435: SAP HANA on NetApp AFF Systems with NFS
<https://www.netapp.com/us/media/tr-4435.pdf>
- NetApp Cloud Central
<https://cloud.netapp.com/home>
- Certified and Supported SAP HANA Hardware Directory
<https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/iaas.html#>
- SAP on Microsoft Azure: Get Started
<https://docs.microsoft.com/en-us/azure/virtual-machines/workloads/sap/get-started>
- SAP on Microsoft Azure Landing Page
<https://azure.microsoft.com/en-us/solutions/sap/>

Version History

Version	Date	Document Version History
Version 1.0	April 2019	Initial release.

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 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-4757-0419