

## **Public and Hybrid Cloud**

**NetApp Solutions** 

NetApp July 26, 2024

This PDF was generated from https://docs.netapp.com/us-en/netapp-solutions/ehc/ehc-overview.html on July 26, 2024. Always check docs.netapp.com for the latest.

# **Table of Contents**

Public and Hybrid Cloud	 	1
NetApp Hybrid Multicloud with VMware Solutions	 	1
VMware Sovereign Cloud	 	491
NetApp Hybrid Multicloud with Red Hat OpenShift Container workloads	 	493

# **Public and Hybrid Cloud**

## **NetApp Hybrid Multicloud with VMware Solutions**

## VMware for Public Cloud

## Overview of NetApp Hybrid Multicloud with VMware

Most IT organizations follow the hybrid cloud-first approach. These organizations are in a transformation phase and customers are evaluating their current IT landscape and then migrating their workloads to the cloud based on the assessment and discovery exercise.

The factors for customers migrating to the cloud can include elasticity and burst, data center exit, data center consolidation, end-of-life scenarios, mergers, acquisitions, and so on. The reason for this migration can vary based on each organization and their respective business priorities. When moving to the hybrid cloud, choosing the right storage in the cloud is very important in order to unleash the power of cloud deployment and elasticity.

## VMware Cloud options in Public Cloud

This section describes how each of the cloud providers support a VMware Software Defined Data Center (SDDC) and/or VMware Cloud Foundation (VCF) stack within their respective public cloud offerings.

## **Azure VMware Solution**



Azure VMware Solution is a hybrid cloud service that allows for fully functioning VMware SDDCs within the Microsoft Azure public cloud. Azure VMware Solution is a first-party solution fully managed and supported by Microsoft, verified by VMware leveraging Azure infrastructure. This means that when Azure VMware Solution is deployed, customer's get VMware's ESXi for compute virtualization, vSAN for hyper-converged storage, and NSX for networking and security, all while taking advantage of Microsoft Azure's global presence, class-leading data center facilities and proximity to the rich ecosystem of native Azure services and solutions.

## VMware Cloud on AWS



VMware Cloud on AWS brings VMware's enterprise-class SDDC software to the AWS Cloud with optimized access to native AWS services. Powered by VMware Cloud Foundation, VMware Cloud on AWS integrates VMware's compute, storage, and network virtualization products (VMware vSphere, VMware vSAN, and VMware NSX) along with VMware vCenter Server management, optimized to run on dedicated, elastic, bare-metal AWS infrastructure.

## Google Cloud VMware Engine



Google Cloud VMware Engine is an infrastructure-as-a-service (IaaS) offering built on Google Cloud's highly performant scalable infrastructure and VMware Cloud Foundation stack – VMware vSphere, vCenter, vSAN, and NSX-T. This service enables a fast path to the cloud, seamlessly migrating or extending existing VMware workloads from on-premises environments to Google Cloud Platform without the cost, effort ,or risk of rearchitecting applications or retooling operations. It is a service sold and supported by Google, working closely with VMware.



SDDC private cloud and NetApp Cloud Volumes colocation provides the best performance with minimal network latency.

#### Did you know?

Regardless of the cloud used, when a VMware SDDC is deployed, the initial cluster includes the following products:

- VMware ESXi hosts for compute virtualization with a vCenter Server appliance for management
- VMware vSAN hyper-converged storage incorporating the physical storage assets of each ESXi host
- · VMware NSX for virtual networking and security with an NSX Manager cluster for management

#### Storage configuration

For customers planning to host storage-intensive workloads and scale out on any cloud-hosted VMware solution, the default hyper-converged infrastructure dictates that the expansion should be on both the compute and storage resources.

By integrating with NetApp Cloud Volumes, such as Azure NetApp Files, Amazon FSx for NetApp ONTAP, Cloud Volumes ONTAP (available in all three major hyperscalers), and Cloud Volumes Service for Google Cloud, customers now have options to independently scale their storage separately, and only add compute nodes to the SDDC cluster as needed.

#### Notes:

- VMware does not recommend unbalanced cluster configurations, hence expanding storage means adding more hosts, which implies more TCO.
- Only one vSAN environment is possible. Therefore, all storage traffic will compete directly with production workloads.
- There is no option to provide multiple performance tiers to align application requirements, performance, and cost.
- It is very easy to reach the limits of storage capacity of vSAN built on top of the cluster hosts. Use NetApp Cloud Volumes to scale storage to either host active datasets or tier cooler data to persistent storage.

Azure NetApp Files, Amazon FSx for NetApp ONTAP, Cloud Volumes ONTAP (available in all three major hyperscalers), and Cloud Volumes Service for Google Cloud can be used in conjunction with guest VMs. This hybrid storage architecture consists of a vSAN datastore that holds the guest operating system and application binary data. The application data is attached to the VM through a guest-based iSCSI initiator or the NFS/SMB mounts that communicate directly with Amazon FSx for NetApp ONTAP, Cloud Volume ONTAP, Azure NetApp Files and Cloud Volumes Service for Google Cloud respectively. This configuration allows you to easily

overcome challenges with storage capacity as with vSAN, the available free space depends on the slack space and storage policies used.

Let's consider a three-node SDDC cluster on VMware Cloud on AWS:

- The total raw capacity for a three-node SDDC = 31.1TB (roughly 10TB for each node).
- The slack space to be maintained before additional hosts are added = 25% = (.25 x 31.1TB) = 7.7TB.
- The usable raw capacity after slack space deduction = 23.4TB
- The effective free space available depends on the storage policy applied.

For example:

- RAID 0 = effective free space = 23.4TB (usable raw capacity/1)
- RAID 1 = effective free space = 11.7TB (usable raw capacity/2)
- RAID 5 = effective free space = 17.5TB (usable raw capacity/1.33)

Thus, using NetApp Cloud Volumes as guest-connected storage would help in expanding the storage and optimizing the TCO while meeting the performance and data protection requirements.



In-guest storage was the only available option at the time this document was written. As supplemental NFS datastore support becomes available, additional documentation will be available here.

## **Points to Remember**

- In hybrid storage models, place tier 1 or high priority workloads on vSAN datastore to address any specific latency requirements because they are part of the host itself and within proximity. Use in-guest mechanisms for any workload VMs for which transactional latencies are acceptable.
- Use NetApp SnapMirror® technology to replicate the workload data from the on-premises ONTAP system to Cloud Volumes ONTAP or Amazon FSx for NetApp ONTAP to ease migration using block-level mechanisms. This does not apply to Azure NetApp Files and Cloud Volumes Services. For migrating data to Azure NetApp Files or Cloud Volumes Services, use NetApp XCP, BlueXP Copy and Sync, rysnc or robocopy depending on the file protocol used.
- Testing shows 2-4ms additional latency while accessing storage from the respective SDDCs. Factor this additional latency into the application requirements when mapping the storage.
- For mounting guest-connected storage during test failover and actual failover, make sure iSCSI initiators are reconfigured, DNS is updated for SMB shares, and NFS mount points are updated in fstab.
- Make sure that in-guest Microsoft Multipath I/O (MPIO), firewall, and disk timeout registry settings are configured properly inside the VM.



This applies to guest connected storage only.

### Benefits of NetApp cloud storage

NetApp cloud storage offers the following benefits:

- Improves compute-to-storage density by scaling storage independently of compute.
- Allows you to reduce the host count, thus reducing the overall TCO.
- Compute node failure does not impact storage performance.

- The volume reshaping and dynamic service-level capability of Azure NetApp Files allows you to optimize cost by sizing for steady-state workloads, and thus preventing over provisioning.
- The storage efficiencies, cloud tiering, and instance-type modification capabilities of Cloud Volumes ONTAP allow optimal ways of adding and scaling storage.
- Prevents over provisioning storage resources are added only when needed.
- Efficient Snapshot copies and clones allow you to rapidly create copies without any performance impact.
- · Helps address ransomware attacks by using quick recovery from Snapshot copies.
- Provides efficient incremental block transfer-based regional disaster recovery and integrated backup block level across regions provides better RPO and RTOs.

#### Assumptions

- SnapMirror technology or other relevant data migration mechanisms are enabled. There are many connectivity options, from on-premises to any hyperscaler cloud. Use the appropriate path and work with the relevant networking teams.
- In-guest storage was the only available option at the time this document was written. As supplemental NFS datastore support becomes available, additional documentation will be available here.



Engage NetApp solution architects and respective hyperscaler cloud architects for planning and sizing of storage and the required number of hosts. NetApp recommends identifying the storage performance requirements before using the Cloud Volumes ONTAP sizer to finalize the storage instance type or the appropriate service level with the right throughput.

#### **Detailed architecture**

From a high-level perspective, this architecture (shown in the figure below) covers how to achieve hybrid Multicloud connectivity and app portability across multiple cloud providers using NetApp Cloud Volumes ONTAP, Cloud Volumes Service for Google Cloud and Azure NetApp Files as an additional in-guest storage option.



## NetApp Solutions for VMware in Hyperscalers

Learn more about the capabilities that NetApp brings to the three (3) primary hyperscalers - from NetApp as a guest connected storage device or a supplemental NFS datastore to migrating workflows, extending/bursting to the cloud, backup/restore and disaster recovery.

Pick your cloud and let NetApp do the rest!





To see the capabilities for a specific hyperscaler, click on the appropriate tab for that hyperscaler.

Jump to the section for the desired content by selecting from the following options:

- VMware in the Hyperscalers Configuration
- NetApp Storage Options

#### NetApp / VMware Cloud Solutions

#### VMware in the Hyperscalers Configuration

As with on-premises, planning a cloud based virtualization environment is critical for a successful productionready environment for creating VMs and migration.

#### AWS / VMC

This section describes how to set up and manage VMware Cloud on AWS SDDC and use it in combination with the available options for connecting NetApp storage.



In-guest storage is the only supported method of connecting Cloud Volumes ONTAP to AWS VMC.

The setup process can be broken down into the following steps:

- Deploy and Configure VMware Cloud for AWS
- Connect VMware Cloud to FSx ONTAP

View the detailed configuration steps for VMC.

#### Azure / AVS

This section describes how to set up and manage Azure VMware Solution and use it in combination with the available options for connecting NetApp storage.



In-guest storage is the only supported method of connecting Cloud Volumes ONTAP to Azure VMware Solution.

The setup process can be broken down into the following steps:

- · Register the resource provider and create a private cloud
- · Connect to a new or existing ExpressRoute virtual network gateway
- · Validate the network connectivity and access the private cloud

View the detailed configuration steps for AVS.

#### **GCP / GCVE**

This section describes how to set up and manage GCVE and use it in combination with the available options for connecting NetApp storage.



In-guest storage is the only supported method of connecting Cloud Volumes ONTAP and Cloud Volumes Services to GCVE.

The setup process can be broken down into the following steps:

- · Deploy and Configure GCVE
- Enable Private Access to GCVE

View the detailed configuration steps for GCVE.

#### **NetApp Storage Options**

NetApp storage can be utilized in several ways - either as guest connected or as a supplemental NFS datastore - within each of the 3 major hyperscalers.

Please visit Supported NetApp Storage Options for more information.

#### AWS / VMC

AWS supports NetApp storage in the following configurations:

- · FSx ONTAP as guest connected storage
- Cloud Volumes ONTAP (CVO) as guest connected storage
- · FSx ONTAP as a supplemental NFS datastore

View the detailed guest connect storage options for VMC. View the detailed supplemental NFS datastore options for VMC.

#### Azure / AVS

Azure supports NetApp storage in the following configurations:

- Azure NetApp Files (ANF) as guest connected storage
- Cloud Volumes ONTAP (CVO) as guest connected storage
- Azure NetApp Files (ANF) as a supplemental NFS datastore

View the detailed guest connect storage options for AVS. View the detailed supplemental NFS datastore options for AVS.

#### GCP / GCVE

Google Cloud supports NetApp storage in the following configurations:

- Cloud Volumes ONTAP (CVO) as guest connected storage
- Cloud Volumes Service (CVS) as guest connected storage
- · Cloud Volumes Service (CVS) as a supplemental NFS datastore

View the detailed guest connect storage options for GCVE.

Read more about NetApp Cloud Volumes Service datastore support for Google Cloud VMware Engine (NetApp blog) or How to use NetApp CVS as datastores for Google Cloud VMware Engine (Google blog)

#### NetApp / VMware Cloud Solutions

With NetApp and VMware cloud solutions, many use cases are simple to deploy in your hyperscaler of choice. VMware defines the primary cloud workload use-cases as:

- Protect (includes both Disaster Recovery and Backup / Restore)
- Migrate
- Extend

AWS / VMC Browse the NetApp solutions for AWS / VMC	
Azure / AVS Browse the NetApp solutions for Azure / AVS	
GCP / GCVE Browse the NetApp solutions for Google Cloud Platform (GCP) / GCVE	

## Supported Configurations for NetApp Hybrid Multicloud with VMware

Understanding the combinations for NetApp storage support in the major hyperscalers.

	Guest Connected	Supplemental NFS Datastore
AWS	CVO FSx ONTAP Details	FSx ONTAP Details
Azur e	CVO ANF Details	ANF Details
GCP	CVO CVS Details	CVS Details

## Configuring the virtualization environment in the cloud provider

Details for how to configure the virtualization environment in each of the supported hyperscalers are covered here.

### AWS / VMC

This section describes how to set up and manage VMware Cloud on AWS SDDC and use it in combination with the available options for connecting NetApp storage.



In-guest storage is the only supported method of connecting Cloud Volumes ONTAP to AWS VMC.

The setup process can be broken down into the following steps:

- Deploy and Configure VMware Cloud for AWS
- Connect VMware Cloud to FSx ONTAP

View the detailed configuration steps for VMC.

#### Azure / AVS

This section describes how to set up and manage Azure VMware Solution and use it in combination with the available options for connecting NetApp storage.



In-guest storage is the only supported method of connecting Cloud Volumes ONTAP to Azure VMware Solution.

The setup process can be broken down into the following steps:

- · Register the resource provider and create a private cloud
- · Connect to a new or existing ExpressRoute virtual network gateway
- · Validate the network connectivity and access the private cloud

View the detailed configuration steps for AVS.

#### GCP / GCVE

This section describes how to set up and manage GCVE and use it in combination with the available options for connecting NetApp storage.



In-guest storage is the only supported method of connecting Cloud Volumes ONTAP and Cloud Volumes Services to GCVE.

The setup process can be broken down into the following steps:

- Deploy and Configure GCVE
- Enable Private Access to GCVE

View the detailed configuration steps for GCVE.

#### Deploy and configure the Virtualization Environment on AWS

As with on-premises, planning VMware Cloud on AWS is critical for a successful production-ready environment for creating VMs and migration.

This section describes how to set up and manage VMware Cloud on AWS SDDC and use it in combination

with the available options for connecting NetApp storage.



In-guest storage is currently the only supported method of connecting Cloud Volumes ONTAP (CVO) to AWS VMC.

The setup process can be broken down into the following steps:

VMware Cloud on AWS provides for a cloud native experience for VMware based workloads in the AWS ecosystem. Each VMware Software-Defined Data Center (SDDC) runs in an Amazon Virtual Private Cloud (VPC) and provides a full VMware stack (including vCenter Server), NSX-T software-defined networking, vSAN software-defined storage, and one or more ESXi hosts that provide compute and storage resources to your workloads.

This section describes how to set up and manage VMware Cloud on AWS and use it in combination with Amazon FSx for NetApp ONTAP and/or Cloud Volumes ONTAP on AWS with in-guest storage.



In-guest storage is currently the only supported method of connecting Cloud Volumes ONTAP (CVO) to AWS VMC.

The setup process can be broken down into three parts:

#### **Register for an AWS Account**

Register for an Amazon Web Services Account.

You need an AWS account to get started, assuming there isn't one created already. New or existing, you need administrative privileges in the account for many steps in this procedure. See this link for more information regarding AWS credentials.

#### **Register for a My VMware Account**

Register for a My VMware account.

For access to VMware's cloud portfolio (including VMware Cloud on AWS), you need a VMware customer account or a My VMware account. If you have not already done so, create a VMware account here.

#### Provision SDDC in VMware Cloud

After the VMware account is configured and proper sizing is performed, deploying a Software-Defined Data Center is the obvious next step for using the VMware Cloud on AWS service. To create an SDDC, pick an AWS region to host it, give the SDDC a name, and specify how many ESXi hosts you want the SDDC to contain. If you don't already have an AWS account, you can still create a starter configuration SDDC that contains a single ESXi host.

1. Log into the VMware Cloud Console using your existing or newly created VMware credentials.

	nware.com/cspigaraway/discovery hogoutaskipLogout	
Welcome to		
VMware Cloud Services		
Sign in with your VMware account		
Email address		
uumanegenel.com		
HEST		_
New to VMware Cloud?		
CREATE YOUR VMWARE ACCOUNT		
ENGLISH		
02021 Villaster Inc. Terms Property Children's Deback Dents		

2. Configure the AWS region, deployment, and host type and the SDDC name:

vmw VMware Cloud		Δ	0	Will Stown NetApp	
«	charges.				
8 Launchpad	v 1. SDDC Properties	Give your SDDC a name, choose a size, and specify the AWS region where it will be			
SDDCs		created.			
Subscriptions					
= Activity Log	AWS Region	US West (Oregon)			
🗄 Tools					
Developer Center	Deployment	Single Host O Multi-Host Stretched Cluster ()			
	Host Type	I I Local SSD) (1) Isen (Local SSD) (1)			
	SDDC Name	ntap-fsi-demo			
	Number of Hosts	1 D 1-host SDDCs expire in 60 days. LEARN MORE			
	Host Capacity	2 Sockets, 38 Cores, 512 GIB RAM, 10.37 TiB Storage			
	Total Capacity	2 Sockets, 36 Cores, 512 GiB RAM, 10.37 TiB Storage			
	SHOW ADVANCED CONF	IGURATION			
	NEXT				
	2 Connect to AWS	Specify the 4WS account that you want to connect your SDDC with			
	2. 55/1001 10 1010	chants and the second that has upped an and had and			
Lene .	A STATE OF A STATE OF				

3. Connect to the desired AWS account and execute the AWS Cloud Formation stack.

aws Services *	Q. Search for services, features, marketplace products, an	d docs (Option+5) 🖸 🎝 SSD-Administ	trator/WillStrive@netapp.com @ cloutheroes 🔻	Oregon 🔻 Su
ECloudFormation	Stacks ) Create stack			
Quick cr	ate stack			
Template				
Template URL https://vmwa /mo5ilohtrlie	sddc.s3.us-west-2.amazonaws.com/1eb9d184-a706-4489-abb8-6 1815b75mean%cc4bddd1iffoD7w.7v16fs36	592aad0a25d0		
Stack descript This template	in s created by VMware Cloud on AWS for SDDC deployment and main	iteriance. Please do not remove.		
Stack nam				
Stack name				
vmware-sdd Stock nome can	formation-a87f51c9-e5ac-4bb4-9d1e-9a3dabd197b7 clude letters (A-7 and a-ri, numbers (D-9), and dathes (-).			
Parameter Parameters are	fined in your template and allow you to input custom values when you create	or update a stack.		
Feedback English (US) 🔻		6 2006 - 2011, Amazin Web Solution, Inc. or its all	Mates All rights essences. Privacy Policy Tel	mis of Use Cookie
⊢ → C ŵ	O B https://us-west-2.conspie.aws amazon.com/ck	oudformation/home?region=us-west-2#/stacks/quickcreate/	7stackNamenymware-edd: 😋	<u> 2</u> 0
awsservices ▼	Q. Search for services, features, marketplace products, an	d dors (Option+5) 🖸 🗘 4 550-Administ	trator/WillStrive@inetapp.com@cloudheroes 🔻	Oregon 🔻 S
Stack nam				
Stack name			-	
Stack name can	formation-a87/51c9-e5ac-4bb4-9d1e-9a3dabd197b7 clude letters (A-Z and a-s), numbers (D-9), and dathes (-).			
Parameters are	fined in your template and allow you to input custom values when you create	or update a stock.		
	No parameter There are no parameters define	s d in your template		
Capabilities				
The follo This temp Check the	Ing resource(s) require capabilities: [AWS::IAM::Role] ste contains identity and Access Management (IAM) resources that r you want to create each of these resources and that they have the r	night provide entities access to make changes to your AWS a minimum required permissions. Learn more 🗹	account.	
Ditacke	wiedge that AWS CloudFormation might create IAM resources.			
Feedback English (US) 💌		e poor Just Amazon Web Service, by an un	Names All rights reserved. Privacy Policy	mis of Use - Doning
			and the second second	



Single-host configuration is used in this validation.

4. Select the desired AWS VPC to connect the VMC environment with.

			Well Stourp	
mw VMware Cloud		Δ 0	NetApp	
Launchpad     SoOCs     Subscriptions     Activity Log     Tools     Developer Center	Connect to AWS Aws Account ID 384ad01e-f5a7-3860-b1a7-3bf4d70db1db      O VPC and subnet Specify the VPC and the subnet to connect to your AWS account.      VPC ypc-0c6794aa6e67d2ddt (00.0.00/f6)      O     subnet VMsn (00.0.00/24, us west-2d, usw2-az4)      To leverage native AWS services on your SDDCs, deploy your AWS EC2 workbadts in the same availability zone to avoid cross AZ traffic charge.	Teer Annum Arcone Your AWS VPC Availability Zeer 1 (WTS School 1 BC2 Instance) AWS School 2 (AWS School 2 (AWS School 2 (AWS School 2) (AWS S		
	v 4. Configure Network Management Subnet (optional)			

5. Configure the VMC Management Subnet; this subnet contains VMC-managed services like vCenter, NSX, and so on. Do not choose an overlapping address space with any other networks that need connectivity to the SDDC environment. Finally, follow the recommendations for CIDR size notated below.

- → C 🏠	O B at B https://vmc.vmware.com/console/soldca/create/aws		\$	2 🛛	۲
vmw VMware Cloud			© 4	Will Stowe NetApp	
«	charges.		22		
) Launchpad	> SDDC Properties ntap-fsx-demo - 1)	Hosts - us-west-2			
3 Subscriptions	> 🧭 Connect to AWS Aws Account ID 3a	s4ad01e-f5a7-3860-b1a7-3bf4d70db1db			
) Tools	> O VPC and subnet VPC - vpc-0c6794	aa5e67d2dd1			
<ul> <li>Developer Center</li> </ul>	v 4. Configure Network Management Subn	et (optional)			
	<ul> <li>Specify a private subnet range (RFC 1918) to be used for</li> <li>Choose a range that will not overlap with other network</li> <li>Minimum CIDR sizes: /23 for up to 27 hosts, /20 for up to</li> <li>Reserved CIDRs: 10.0.0.0/75, 172.310.0/16.</li> </ul>	v vCenter Server, NSX Manager, and ESXi hosts. is or SDDC group members that connect to this SDDC. to 251 hosts, /16 for up to 4091 hosts.			
	Management Subnet ODB Risck Default 10.2.0.0.06				
	5. Review and Acknowledge Review and ack	nowledge cost before deployment			
DATH					

6. Review and acknowledge the SDDC configuration, and then click deploy the SDDC.



The deployment process typically takes approximately two hours to complete.

	eng a Soft 🛠 📑 VMvare Cloud - SDDCs 🛛 🗴 🔞 Bubnets I VPC Managemant Car 🛪 👎 3069 Disen Dr - Google Maps 🛪 🧕 AWS Ma	ionagement Consola 🛛 × 🕂 +		~
< → ଫ ଲ		Ŷ	\$ © (	<b>∂</b> =
vmw VMware Cloud		0 2	WE Stowe NetApp	
د ۵ Läunchpad	(SDDC)			
SDDCs	SDDC Groups		222.0-011000m	_
D Subscriptions	SODCs have been added and/or removed. Refresh the page to update the data.		Refresh.new	×
III Activity Log	reproving on AWS      DEPLOYING SDDC      Estimated time to completion: 86 Minutes			
	How easy was it for you to create your SDDC?			×
	Univ difficult 1 2 3 4 5 8 7 May anny			
	the SDDC is ready for use			
\fter completion				
After completion	,			
After completior	,			

<pre>stock stock stock make reade re</pre>	& Launchpart	Software-Defined Data Cente	88 I 🗐	CREATE BODC ACTIONS Y
BACK TO TOP GO TO GRID VIEW	■ 500Cs Subsciptions ■ Activity Log ← Tools •• Developer Carder	(SDDC) SDDC (SDDC) SDDC (Separative (SDDC) SDDC (Separative (SDDC)	Storage	
			BACK TO TOP SO TO GRID VIEW	

To connect VMware Cloud to FSx ONTAP, complete the following steps:

 With VMware Cloud deployment completed and connected to AWS VPC, you must deploy Amazon FSx for NetApp ONTAP into a new VPC rather than the original connected VPC (see the screenshot below). FSx (NFS and SMB floating IPs) is not accessible if it is deployed in the connected VPC. Keep in mind that ISCSI endpoints like Cloud Volumes ONTAP work just fine from the connected VPC.



2. Deploy an additional VPC in the same region, and then deploy Amazon FSx for NetApp ONTAP into the new VPC.

Configuration of an SDDC group in the VMware Cloud console enables the networking configuration options required to connect to the new VPC where FSx is deployed. In step 3, verify that "Configuring VMware Transit Connect for your group will incur charges per attachment and data transfers" is checked, and then choose Create Group. The process can take a few minutes to complete.

A landpadel Backgroupe A service and a service and a decidation for type agrice and a service and a decidation for type agrice and a service and a decidation for type agrice and a service and a decidation for type agrice and a service and a decidation for type agrice and a service and a decidation for type agrice and a service and a decidation for type agrice and a service and a decidation for type agrice and a service and a decidation for type agrice and a service and a decidation for type agrice and a service	<pre>i larcopad: i Boog, i Boog, i Boog, i Boog, i Boog, i Boog, i Boog, i Boog, i De wood of calify i I wood i I wood wood with the Station go the station i I wood i I wood i Wood i Wood i I wood i I wood i Wood i I wood i Wood i I wood i Wood i I woo</pre>	0	< Create SDDC	Group			
II BOOL       I Meet and Eleviderin       Central a none and deciperion for your group.         2 Addrop Constr       Addrop Constr       I Meet and Eleviderin       Meet and I Support         2 Addrop Constr       I Meet and Eleviderin       Meet and I       Meet and I         2 Addrop Constr       I Meet and Eleviderin       Meet and I       Meet and I         2 Addrop Constr       I Meet and Eleviderin       Meet and I       Meet and I         2 Addrop Constr       I Meet and Eleviderin       Meet and I       Meet and I         2 Addrop Constr       I Meet and Eleviderin       Meet and I       Meet and I         2 Addrop Constr       I Meet and Eleviderin       Meet and I       Meet and I         2 Meet and Meet and Eleviderin       Meet and I       Meet and I       Meet and I         2 Meet and Meet and Eleviderin       Meet and I       Meet and I       Meet and I         2 Meet and Meet and Eleviderin       Meet and I Meet and I       Meet and I       Meet and I         3 SOCI       Meet and Description       Meet and I Meet and I Meet and I       Meet and I       Meet and I         3 SOCI       Need and I Meet and I       Meet and I       Meet and I       Meet and I         3 Socicoton       Need and I Meet and I Mee	I work with the public work of the state and a sequence with the public of the publ	) Laurchpad	T A SHARE AND A				
<pre>1 Subcry Colog Subcry Colo</pre>	<pre>1 descriptions 2 restants togs 2 restants togs 2</pre>	500Cs	1. Name and Description	Create a name and description for your group			
<pre>k and k g g word S word S</pre>	<pre>k atory topy B with B wit</pre>	3 Subscriptions					
<pre>s out</pre>	<pre>* *** ********************************</pre>	E Activity Log	Name	1ddcgroup01			
And we want the first of the	<pre>set use use use use use use use use use use</pre>	: Developer Certer	Description	siddegroup01			
<pre>set vertex class to be set allow and the se</pre>	<pre>ver ver ver ver ver ver ver ver</pre>						
Image: Section (Section (S			NEXT				
Image: Control in the local state is the following before overing the SDDC Group:         Image: Control in the local state is a state in the following before overing the SDDC Group:         Image: Control in the local state is a state in the following before overing the SDDC Group:         Image: Control in the local state is a state in the following before overing the SDDC Group:         Image: Control in the local state is a state in the SDDC Group:         Image: Control in the local state is a state in the SDDC Group:         Image: Control in the local state is a state in the SDDC Group:         Image: Control in the local state is a state in the SDDC Group:         Image: Control in the local state is a state in the SDDC Group:         Image: Control in the local state is a state in the SDDC Group:         Image: Control in the local state is a state in the state in the local state is a state in the state in the local state is a state in the local state in the local state in the local state is a state in the local state in the local state is a state in the local state in the local state is a state in the local state is a state in the local state in the local state is a state in the local state in the local state is a	Interview		2. Membership	Members 1			
Image: control the log up and a sease of the ligboding parties control ging by 5000 CG to the grants       Light Model Control to your grants and the light Model Control to the grants         Image: Control to the log up and the light Model Control to the grants       Light Model Control to the log up and the log Control to the grants       Light Model Control to the log up and the log Control to the grants         Image: Model Control to the log up and the log Control to the grants       Light Model Control to the log up and the log Control to the grants       Light Model Control to the log up and the log Control to the grants         Image: Model Control to the log up and to control to the log up and the log Content to the log up and the log Control to the l	we where Code <ul> <li>C code in the life weight plane are used to the life back on the life</li></ul>		3. Acknowledgement				
Image: Start Star	<pre>ver ver ver ver ver ver ver ver ver ver</pre>		Passa confirm that you are	avera of the Informer before creation the \$200 G	nin.		
Image: Contract Address	<pre>set</pre>		Configuring VMware Trans	It Connect for your group will incur charges per attachm	ent and data transfers.		
Image: Second	Image: Section Control         Image: Section Control<		Create freedul rules to a	stabilish connectivity between the SODCs in the onup	La	sam More E	
Image: Control (Control (Contro) (Control (Control (Control (Control (Control (Control (Control (	Interviewed						
	Image:		CREATE GROUP				
Insert         Image: Sector Sect	Instrume         Image: Control         Image: Contro         Image: Control         Image: Control						
we we we we have a closed         Image: Closed	WMX/WX Clock         Image: Control image: Contro						
WWW/WY/N Cloud  C C Configuring Markes Ensure of the following before creating the group.  A definition  A definition  Beodesioner Center  C Configuring Markes Ensure of the following before creating the group.  A definition  A definition  Beodew and admonifiedge resultements before creating the group.  A definition  A de	www.         Image: Control of the second of the secon						
Image: VMW2P Cloud         Image:	American         Control         Contro         Control <thcontrol< th=""> <th< td=""><td></td><td></td><td></td><td></td><td></td><td></td></th<></thcontrol<>						
Instrume	we w						
With red Cloud       Image: Create SDDC Group         I canor hoped       Image: Create SDDC Group         I control       Image: Create SDDC Group<	With any cloud     O     O     Constant of partial     Constant of partia	Baata					
With You	Were closed   Cuancipued     Calculation        Calculation						
Create SDDC Group      Autor (Market Market Ma	Create SDDC Group      Autority Log      Sold      Create SDDC Group      Name and Description      Name						A Will Stowe
Create SDDC Group      Laurchper      Laurchpe	Create SDDC Group      Autoritype      Socce      Control      Co	www.VMware Cloud					<u> ()</u> () <u>назыка</u> ч
Laurchand SoCa SoCa SoCa SoCa SoCa SoCa SoCa SoCa	Lacerchael SOCA Select proofs Activity Log Tools Developer Center	mw VMware Cloud					Д 🕐 метАрр 🎽
S DOCG       If Name and Description         S Basis proces       2. Membership       Select SDDCs to be part of your group         S Tools       Image-basedwind       82946-022-924-9249-2424-2426-       US West (Oregan)         Developer Center       Image-basedwind       82946-022-924-9249-2424-2426-       US West (Oregan)         Image-basedwind       82946-022-924-9249-24240-       US West (Oregan)       114-0.04       10.45.0 0723         Image: Developer Center         Image: Developer Center       Image: Developer Center       Review and acknowledge requirements before creating the group       Image: Developer Center         Image: Developer Center       Image: Developer Center       Review and acknowledge requirements before creating the group       Image: Developer Center         Image: Developer Center       Image: Developer Center       Review and acknowledge requirements before creating the group       Image: Developer Center         Image: Developer Center       Review and acknowledge requirements before creating the group       Image: Developer Center       Image: Developer Center         Image: Developer Center       Review and acknowledge requirements before creating the group       Image: Developer Center       Image: Developer Center         Image: Developer Center       Re	SDOC4       I. Name and Description       I. Name and Description         Select pions       Activity Log       Select SDDCs to be part of your group         Tods       If there       1       Select SDDCs         Developer Center       If there       1       Select SDDCs       US West (Dregan)         Tods       If there       1       Select SDDCs       US West (Dregan)       114.0.14       10.45.0.023         Developer Center       If there       1       Select SDDCs       US West (Dregan)       114.0.14       10.45.0.023         Select SDDC Select Sele	me VMware Cloud	< Create SDDC	Group			Ц 🕘 негара ~
Activity Log  Activity Log  Tools  Developer Center  Activity Log  Center  Activity Log  Center  Activity Log  Ac	Alterity Log         Activity Log         Tools         Developer Center             If they-fixe-demo             Static typics             Static typics <td>Wware Cloud</td> <td>&lt; Create SDDC</td> <td>Group</td> <td></td> <td></td> <td>Д O неслер V</td>	Wware Cloud	< Create SDDC	Group			Д O неслер V
Control to be reader of the following before creating the group      Configuring VMwere Trankt Correctivity between the (2DCCs in the group)      Configuring VMwere Trankt Correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish correctivity between the (2DCCs in the group)      Create firewall indices to establish estab	Center (Ling)   Is Tools   Is Developer Center     Inter-fax-demo   StrateG22-924-4208-   US West (Cregan)     144.0.44   1045:0.023     Inter-fax-demo   StrateG22-924-4208-   US West (Cregan)     154.0.44   1045:0.023     Inter-fax-demo     StrateG22-924-4208-   US West (Cregan)     154.0.44   1045:0.023     Inter-fax-demo     StrateG22-924-4208-   US West (Cregan)     154.0.44   1045:0.023     Inter-fax-demo     StrateG2        104:0:1:1:1:1:1:1:1:1:1:1:1:1:1:1:1:1:1:	www.VMware Cloud	< Create SDDC	Group			ф O неларр ч
Image-basedeen 0       82%9/6/022/8/3/2/2/2/2/2/2/2/2/2/2/2/2/2/2/2/2/2	• Developer Center     • Developer Center     • Developer Center     • Integ-bac-demo     • Developer Center     • Developer Center <td>www.VMwane Cloud</td> <td>&lt; Create SDDC 1. Name and Description 2. Membership</td> <td>Group Name: sodcgroup01 Select SDDCs to be part of your group</td> <td></td> <td></td> <td>Д (Ø) негарр V</td>	www.VMwane Cloud	< Create SDDC 1. Name and Description 2. Membership	Group Name: sodcgroup01 Select SDDCs to be part of your group			Д (Ø) негарр V
	International and a standard of the following before creating the group: International and a standard of the following before creating the group: International and a standard of the following before creating the group: International and a standard of the following before creating the group: International and a standard of the following before creating the group: International and a standard of the following before creating the group: International and a standard of the following before creating the group: International and a standard of the following before creating the group: International and a standard of the following before creating the group: International and a standard of the following before creating the group: International and a standard of the following before creating the group: International and a standard of the following before creating the group: International and the group: In	www.vMware.Cloud & Launchpell 2 5000s 2 Subscriptions 2 Accelly Log 1 Tools	< Create SDDC 1. Name and Description 2. Membership Meme	Group Name: sodegroup01 Select SDDCs in the part of your group + Select	r Loodon	* Vester	C O NetApp *
Acknowledgement     Seview and acknowledge resultements before creating the group Passa confirm that you are aware of the following before creating the SDDC Group     Configuring VMware Transt Connect for your group will incur charges per attachment and data transfers.     Create frewel rules to establish connectivity between the SDDCs in the group.     Learn More [2]      CREATE GROUP	A Acknowledgement Review and acknowledge resultements before crienting the group Please confirm that you are aware of the following before creating the SDDC Group. Configuring Where Trankt Connect for your group will incur charges per attachment and data transfers. Create through	VMware Cloud     K     Laurchpet     Sooce     Sooce     Sooce     Acovity Lag     Tools     Developer Center	Create SDDC     Name and Description     Membership     Meme     me     me     me	Group Name: sodcgroup01 Select SDDCs to be part of your group	r Loodon US West (Chegon)	v Wester 194.0.14	Methods * Met
Acknowledgement      Beview and acknowledge resultancents before creating the group  Please confirm that you are aware of the following before creating this SOOC Group.      Configuring VMwere Trankt Connect for your group will incur charges per attachment and dela transfers.      Create frewall rules to establish connects/thy between the SODCs in the group.      Learn More [2]      CREATE GROUP		VMWare Cloud     K     Laurchpet     Sooce     Sooce     Sooce     Sooce     Control     Control     Tools     Developer Center	<ul> <li>Create SDDC</li> <li>Name and Description</li> <li>Membership</li> <li>Meme (constraint)</li>     &lt;</ul>	Group Name: sodcgroupOf Select SDDCs to be pert of your group	r Loodon US West (Chegan)	v Vestan 134.0.M	()     ()
	<ul> <li>Acknowledgement         Review and acknowledge read/rements before creating the group         Passe confirm that you are aware of the following before creating this SDDC Group.         Configuring VMwere Trankt Connect for your group will incur charges per attachment and data transfers.         Configuring VMwere Trankt connects/ty between the SDDCs in the group.         Learn More [?]         CRITATE GROUP         CRITATE CRITA</li></ul>	VMWare Cloud     K     Launchpad     SoOCa     SoOC	<ul> <li>Create SDDC</li> <li>Name and Description</li> <li>Membership</li> <li>Meme inter-fax-demo</li> <li>I</li> </ul>	Group Name: sodogroupOf Select SDDCs to be part of your group * Solect# S29a6622-92a1-920b- add3-9e4eb7a908c6	r Location US West (Criegan)	v Vester 1940 M	Management CDB     Management CDB     Model and CDB     Model
Please confirm that you are aware of the following before creating this SDDC Group.  Configuring VMwere Trankt Connect for your group will incur charges per attachment and data transfers.  C Create firewall makes to establish connects/by between the SDDCs in the group.  CREATE GROUP  CREATE GROUP	Please confirm that you are aware of the following before creating the SDDC Group.         Image: Configuring VMware Trankt Connect for your group will incur charges per attachment and data transfers.         Image: Create firewall make to establish connectivity between the SDDCs in the group.       Learn More [2]         CREATE GROUP	VMware Cloud     VMware Cloud     Caunchped     Soboce     Soboce     Subscriptions     Activity Log     Tools     Developer Center	<ul> <li>Create SDDC</li> <li>Name and Description</li> <li>Membership</li> <li>Neme</li> <li>Integ-fac-demo</li> <li>Integ-fac-demo</li> <li>Integration</li> </ul>	Group Name: sodcgroupOf Select SDDCs to be part of your group + Select 4 82966423-9241-4206- add3-9e4eb7a90826	r Location US West (Chegon)	v Vester 14 O.M	Managamand ODB     Managamand     Managamandand     Managamandand     Managamandand     Managamandand
Configuring VMwere Transit Connect for your group will incur charges per attachment and data transfers.  C Create Stewnill make to establish connectivity between the SODCs in the proup.  CREATE SHOUP	Configuring VMwere Trankt Connect for your group will incur charges per attachment and data transfers.  C Create firewall rules to establish connectivity between the RODCs in the group.  CREATE GROUP	VMware Cloud     Kaunchpel     Launchpel     Soboc     Subscriptions     Accenty Log     Tools     Developer Center	Create SDDC  Name and Description  C. Membership  Membership  Membership  Membership  Membership  A. Acknowledgement  A. Acknowledgement	Group Name: sodogroup01 Select SDDCs to be part of your group	r Loudon US West (Chegan)	T Vester 14.0.4	Management CDB     Management CDB     N45.0.023     Terrs pro page 100 - 1-10/10
Create firewal rules to establish correctivity between the SDDCs in the group.     Learn More [2]      CREATE GROUP	Croste Brewski nutics to establish connectivity between the SDDCs in the group.     Learn More [2]      CREATE GROUP	VMware Cloud     K     Laurotyset     Sooce     Subscriptions     Activity Lag     Tools     Developer Center	Create SDDC  Name and Description  Membership  Mem  Mem  Mem  Mem  Ker  Acknowledgement  Pesse confern that you are	Group Name: sodogroupOf Select SDDCs to be part of your group	r Looder US West (Oregan) creating the group	T Wester 114.0.14	Messgement CDB     Massgement CDB     Massgeme
CREATE GROUP	CREATE GROUP	VMware Cloud     K     Laurotypet     Sooce     Sooce     Acovity Lag     Tools     Developer Center	Create SDDC  Name and Description  C. Membership  Mem  Mem  Mem  Mem  Mem  Mem  Mem  Me	Group Name: sodegroup01 Select SDDCs to be part of your group	r Location US West (Criegan) critecting the group crip and data transfers.	v Vestan 134.0.14	С О негларо т Манараннан СПВ т0.45.0.0/23 Лета рег раде _100 — 1+10/10
CREATE GROUP	CREATE GROUP	WWWare Cloud     K     Laurochpait     Soboc     Soboc     Soboc     Activity Log     Soboc     Activity Log     Soboc     Developer Center	Create SDDC  Name and Description  Membership  Mem  Mem  Mem  Mem  Acknowledgement  Passa confer hav you are  Configuring Wiware Trans	Group Hame: sodicgroupOf Select SDDCs in the part of your group  Select SDDCs in the following before creating this SDDC Ge  Correct for your group will incur charges per attacher  stabilith correctivity between the SDDCs in the goop.	Location     US West (Chegon)  creating the group oup ent and data transfers.	* Venien 134.0.M	<ul> <li>Menagement CDB</li> <li>Menagement CDB</li> <li>10.45.0.0023</li> <li>Terms per page 100 = 1-1 of 10</li> </ul>
		WWare Cloud     K     Laurotypet     Sooce     Sooce     Sooce     Sooce     Activity Log     Tools     Developer Center	Create SDDC  Name and Description  Memilieration  Memilieration  Memilieration  Memilieration  Memilieration  Memilieration  Memilieration  Acknowledgement  Passa confer hav you are  Configuring Wiware Trans  Configuring Wiwa	Group Hame: sodicgroupOf Select SDDCs in the part of your group  Select SDDCs in the following before creating this SDDC Ge  Connect for your group will incur charges per attacher  stabilith connectivity between the SDDCs in the group.	Location     US West (Chegon)  creating the group oup ent and data transfers.	* Venier 134.0.M	(CDB
		WWare Cloud     K     Laurotypet     Sobc     Sobc     Sobc     Sobc     Activity Log     Sobc     Developer Center	Create SDDC  Name and Description  C. Membership  Mem  Inter-Sacients  KEK7  Acknowledgement  Passa confer har you are  Configuring Wiware Trans  C. Configuring	Group Name: sodicgroupOf Select SDDCs in the part of your group Select SDDCs in the part of your group Select SDDCs in the part of your group Select SDDCs in the solution Review and acknowledge resultements before means of the following before creating this SDDCs are a Connect for your group will incur charges per attached stabilith connectivity between the SDDCs in the group.	Location     US West (Criegon)  creating the group oup ent and data transfers.	* Version 134.0.M	A O Methods     Methods     Management COB     10.45 0 0023     Imma per page 100 - 1-1 of 10
		WWare Cloud     K     Launchpat     Sobc     Sobc	Create SDDC  Name and Description  Control of the strength  Control of	Group Name: sodicgroupOf Select SDDCs is the part of your group	Location     US West (Criegory     crienting the group     cop     ent and data transfers.     La	* Venier 134.0.M	Ф. О. нисларт • Манаранана ССПВ • 0.45.0.0023 Лята рат рада — 1-7 об 10
		VMWare Cloud     K     Laurotypet     SOOCe     Soloc     Activity Log     Tools     Developer Center	Create SDDC  Name and Description  Name Inter-Strip I	Group Name: sodicgroupOf Select SDDCs in the part of your group  Review and acknowledge resultancements before aware of the following before creating this SDDCs as a Connect for your group will incur charges per attached atablets connectivity between the SDDCs in the group.	Location     US West (Criegon)     creating the group     cop     ent and data transfers.     La	Version 134 0 M	<ul> <li>Манаратич СПВ</li> <li>10.45.0 0/23</li> <li>Тепа рег ради (100 - 1-10/10</li> </ul>
		Inner VMWare Cloud () Laurochpaet E 5000 Satescriptions E Activity Log Developer Center Developer Center	Create SDDC  Name and Description  Memilieration  Memilieration  Memilieration  Memilieration  Mexic  Acknowledgement  Acknowledgement  Configuring Wiware Trans  Configuring Wiware  Configuring Wiware Trans  Configuring Wiware  Confi	Group Name: sodicgroupOf Select SDDCs in the part of your group  Select SDDCs in the part of your group  Select SDDCs in the part of your group  Select SDDCs in the solid  Statistic Statistics  Review and acknowledge requirements before means of the following before creating this SDDCs as a Connect for your group will incur charges per attache stability connectivity between the SDDCs in the group.	Location     US West (Chegan)  creating the group oup ent and data transfers.	Version 1:24 0 M	Ф. О. наскара • Манаратичи ССВЯ • Манаратичи ССВЯ • 10.45.0.0/23 Пота рат ради — 1-10/10
		Inner VMWare Cloud () Laurotrpad E 5000 Sobscriptions E Activity Log D Tools Developer Center	Create SDDC  Name and Description  Meme  Meme M	Group Hame: sodicgroupOf Select SDDCs in the part of your group  Select SDDCs in the solid select SDDCs in the group  add8-beauting before creating this SDDCs in the group  add8-beauting before the SDDCs in the group.	Location     US West (Criegon)  creating the group oup ant and data transfers.	* Versten 1:4:0:M	Ф. О. меллен
		WWare Cloud     K     Laurotypet     SoDCe     SoDC	Create SDDC  A Name and Description  A Membership  Membership  Membership  Mext A Acknowledgement  Acknowledgement  Acknowledgement  Configuing Wiwere Trans  Configuing Wiwere Trans  Configuing Wiwere Trans  CREATE GROUP	Group Name: sodogroupOf Select SDDCs to be part of your group  * Select 300Cs to be part of your group  * Select 300Cs and advancements before acut-select and advancements before acut-select and advancements before acut-select and advancements before acut-select acutements before select acutements before acut-select acutements before acutem	r Looden US West (Criegon) crienting the group cup ent and data transfers.	* Versten 1/4 O.M	Management COM     Management COM     Modelson     M
		VMWare Cloud     K     Laurotyped     SOOCe     Solocciptions     Activity Log     Tools     Developer Center	Create SDDC	Group Name: sodogroupOf Select SDDCs to be part of your group  * Select 300Cs to be part of your group  * Select 300Cs and the part of your group  * Select 300Cs and the part of your group  * Select 300Cs and addressed and add	Location     US West (Criegon)  crienting the group oup ent and data transfers.	* Versien 1/4 Q.M	A O NetApp     Age     Automatic     Total Action
		VMWare Cloud     K     Laurotyped     SOOCe     Solocciptions     Activity Log     Tools     Developer Center	Create SDDC	Group Name: sodogroupOf Select SDDCs to be part of your group  * Select 300Cs to be part of your group  * Select 300Cs to be part of your group  * Select 300Cs to be part of your group  * Select 300Cs to be part of your group  * Select 300Cs to be part of your group  * Select 300Cs to be part of your group * Select 300Cs to be part of the following before creating the SDDCs to be part of your group will incur or group ar attacted attacted to your group will incur or group ar attacted attacted to your group will incur or groups ar attacted attacted to your group will incur or groups ar attacted attacted to your group will incur or groups ar attacted attacted to your group will incur or groups ar attacted attacted to your group will incur or groups ar attacted	Location     US West (Criegon)  crienting the group oup ent and data transfers.	* Version 1/4 Q.M	Макадетич ОПВ     10.45.0.0/23     Теть рег раде (100 - 1-1)//10
		WWare Cloud     K     Laurotypet     SODCe     SoDC	Create SDDC	Group Name: sodogroupOf Select SDDCs to be part of your group  * Select 300Cs to be part of your group  * Select 300Cs to be part of your group  * Select 300Cs to be part of your group  * Select 300Cs to be part of your group  * Select 300Cs to be part of your group  * Select 300Cs to be part of your group ** Select 300Cs	r Leader US West (Cregard creating the group cup ent and data transfers.	* Version 1:4 O.M	Ф О несклото с то на состати с состати то на состати с состати
		WWare Cloud     K     Laurotypet     SoDCe     SoDC	Create SDDC	Group Name: sodogroupOf Select SDDCs to be part of your group  * Select SDDCs to be part of your group * Select SDDCs to be part of your group * Select SDDCs to be part of your group * Select SDDCs to be part of your group * Select SDDCs to b	Location     US West (Criegon)  creating the group oup ent and data transfers.	* Version 1/4 Q.M	Макадатич ОПВ     10.45.0.0/23     Теть рег раде (100 - 2-1)//10
		WWare Cloud     K     Laurotyped     SODCe     Sobscriptions     Activity Log     Tools     Developer Center	Create SDDC	Group Name: sodogroupOf Select SDDCs to be part of your group  * Select SDDCs to be part of your group * Select SDDCs to be part of your group * Select SDDCs to be part of your group * Select SDDCs to be part of your group * Select SDDCs to b	r Leador US West (Cregan) crienting the group cup. ent and data transfers.	* Version 194 O.M	Ф О неллон ч * Макеренни (СПР 10.45.0.0/23 Тепа рет сере 100 – 1-10/10
		VMWare Cloud     K     Launctupae     SODCe     Sobscriptions     Activity Log     Tools     Developer Center	Create SDDC	Group Name: sodogroupOf Select SDDCs to be part of your group  * Select SDDCs to be part of your group * Select SDDCs to be part of your group * Select SDDCs to be part of your group * Select SDDCs to be part of your group * Select SDDCs to b	r Location US West (Cregari) creating the group out, and data transfers.	* Version 1/4 O.M	Management CDB     TO 45.0 00/23     Term pro page 100 1 1-10/10
		<pre>www.vMware.Cloud</pre>	Create SDDC	Group Name: sodegroupOf Select SDDCs to be part of your group  * Select SDDCs	r Loodon US West (Cregoru créenting the group cup ent and dela transfers.	* Version 14-0.H	Management CDB     TO 45.0 (0/23     Terrs pro page 100 1 - 110/10
		Internet VMWare Cloud       Image: Acceleration of the second	Create SDDC  Name and Description  Membership  Rem  Rem  Rem  Rem  Rem  Rem  Rem  Re	Group Name: sodegroupOf Select SDDCs to be part of your group  * Select SDDCs	r Loodon US West (Crearru créeting the group cup ant and data transfers.	* Version 14-0.H	Management CDB     TO 45.0 (0/23     Terrs are: page1001-10/10

umur VMware Cloud		û ③ <sup>Will Stowe</sup> → 3
e.	< Create SDDC Group	
ő, Leunchped		
SDDC1	1. Name and Description Name: sddcgroup01	
El Subscriptions	2. Membership Members 1	
= Activity Log	3. Acknowledgement	
i⊞ Tools	Please confirm that you are aware of the following before creating this SDDC Group.	
Developer Center	Configuring VMware Transit Connect for your group will incur charges per attachment and data transfers.	
	Crosse frewst subs to establish connectivity between the SDDCs in the group.	
	CREATE GROUP	
C DARK		

3. Attach the newly created VPC to the just created SDDC group. Select the External VPC tab and follow the instructions for attaching an External VPC to the group. This process can take 10 to 15 minutes to complete.

WW MAKE Cloud     Will Store we	Image: Windows Cloud       Image: Windows Wind	WWWEICCOM       C       O       Will Stock of View o				
<ul> <li></li></ul>	(AL DDC Grass- SdCdcgroupO1          Samary Control Libring Direct Control External TOW Bouring Support         Samary Control Libring Direct Control External TOW Bouring Support         Samary Control Libring Direct Control External TOW Bouring Support         Samary Control Libring Direct Control External TOW Bouring Support         Samary Control Libring Direct Control External TOW Bouring Support         Samary Control Libring Direct Control External TOW Bouring Support         Samary Control Libring Direct Control External Tow Bouring Support         Samary Control Libring Direct Control External Tow Bouring Support         Samary Control Libring Direct Control External Tow Bouring Support         Samary Control Libring Direct Samary Dire		Wwware Cloud		C 0	Will Stowe NetApp
SDCS Summary Cleater Linking Direct Connect External IVPC External TGW Routing Support Seasorpoons Account Linking Direct Connect External TGW Routing Support Tools Developer Center	SDCS Summary vCentralizing Einst Connect Exemptitive Exemptitive Exemptitive Exemptitive Exemptified Routing Support Subscriptions Activity Log Proceedings Central	SOCC Serveron Active Linking Direct Colmic Baterial UNC Extend 1700 Bouing Balgott Serveron Active Linking Direct Colmic Baterial UNC Extend 1700 Bouing Balgott Toll Developed Colmic D Toll Active Toll Active D T	<li>K.</li>	e ALL SDDC Groups Sddcgroup01		ACTIONS ~
Suscriptions Activity Log Tools Developer Center	Subscriptors Accivity Log Tools Developed: Center	Subscriptor. Activity Log Tools Developer Center	SODCs	Summary vCenterLinking Elivect Connect External VPC External TGW Routing Support		
Accentry Log  Tools  Developer Center  AND ACCOUNT  ACCOUNT		Active Line Developed Control	Subscriptions			
Mode     Mode     Web Status	Crokk   • Doelsoner Center	Tork       wt0 desaurd D       immuno Stars Nare       issex       issex       itsex       i	E Activity Log	ADD ACCOUNT REMINE		
Developer Center	Developer Center	Developpin Cinter	2 Tools	AWS Scourt D = Resurce Share None = Share	<ul> <li>VPC Status</li> </ul>	
			- Developer Center	WMC-Group-vio/293000m-fbit7-4bid-b016-sept77bit/sited              •••••••••••••••••••••••••		

	Image: Addres/Auto-Automotive control and addres/a addres/addr	Ŷ	a ⊜ s :
vmw VMware Cloud		0 12	/II Stowe NetApp
ğ Launchpad	« c ALL SDCC Groups sddcgroup01		ACTIONS
S0004	Summary vCenterLinking Direct Connect External VPC External TOW Routing Support		
E Subscriptions	ADD ACCOUNT		
🗇 Tools	AVIS Account D = Resource Stars None = State	<ul> <li>WPC Distus</li> </ul>	
··· Developer Center	U w www.competitionality.competitionalit	Ξų.	

4. As part of the external VPC process, you are prompted through the AWS console to a new shared resource via the Resource Access Manager. The shared resource is the AWS Transit Gateway managed by VMware Transit Connect.



Services ¥	Q. Search for services, features, monketplac	re products, and docs [Option+5]	🖸 💠 SSD Administrate	e/Wil.Sloweginetapp.com @ cleachemes 💌	Oregen • Support
Resource Access × Manager	Resource Access Manager > Shared with me	Resource shares > Resource share 051a6	c5 Carle 4560 8531 e2935485660c	e_4560_853f_e2939	185650c)
* Shared by me	Details and information relating to this resource	share.	0816080 (05180165-081	6-4500-0551-62555	10500007
Resource shares Shared resources	Reject resource share Accept resource	rce share			
Shared with me	Summary				
Resource shares (1mr.tutor) Shared resources Principals	Name VMC-Group-dcx9300a-/5e7-6fa5-b016- ae6176a1e8a6	Owner 645453501102	Invitation date 2021/10/14	Status O Pending	
Permissions library Real	ARN am:awsramcus- west-2:464453001102:resource- share/051a6fc5-0a1e-4560-853/- e23394655b0;	Receiver 139763910815			
	<del>a</del>				

5. Create the Transit Gateway Attachment.

ount or across AWS accounts.	Decivitience (A.L.C. Bill, A.L.ASS MICHIE	the same
Details		
Name tag - optional Creates a tag with the key set to Name and the value set to the specified strin	g.	
my-transit-gateway-attachment		
Transit gateway ID Info		
tgw-001646b36ee07a2cb	•	
Attachment type info		
VPC	•	
VPC attachment		
Select and configure your VPC attachment.		
ONS support Info		
IPv6 support Info		
VPC ID		
Select the VPC to actach to the transit gateway		

6. Back on the VMC Console, Accept the VPC attachment. This process can take approximately 10 minutes to complete.

×	< AL 500C Brook	-
Laurenpad	sddcgroup01	0
500Cs	Summary vCenterElnking Direct.Connect External VPC External TGW Routing Support	
Subscriptions Activity Log		
Tools Developer Center	AWS Account 0 AWS Account 10 AWS Acc	×
	VPC ID     T     VHC on AWS Region     T     Tessak Galeway Absolvment ID     T     Boules     Service       Vpc:Odlc764bcc4956805     US West (Chrispin)     tow-attach-Go4883a9f82c67/064     Appl POVITER     PENDINC	

- 7. While in the External VPC tab, click the edit icon in the Routes column and add in the following required routes:
  - A route for the floating IP range for Amazon FSx for NetApp ONTAP floating IPs.
  - $\circ\,$  A route for the floating IP range for Cloud Volumes ONTAP (if applicable).
  - $\,\circ\,$  A route for the newly created external VPC address space.

×.	K ALL SDDC Groups					1	
Laurenpad	sddcgroup01					ACTIONS	1.90
500Cs	Summary Center Linking Di	rect Connect External VPC External	TGW Routing Support				
Subscriptions Activity Log	ADD ACCOUNT BENOUS						
Tools Developer Center	2405 Account 15 y	AWS Account ID : And Account ID : And Account ID : And Account ID : And Account ID Accou	0::-ស្រីក?-41a5 ឯវ៉ានិ-ae6176a1e8a6				>
		VPC/B T VPC-0dic/64/bcc49/be80/s	VHC on WVS Region T US West (Oregon)	Traval Grievay Attochroni ID 1gw-attach-Oc486356782c67864	<ul> <li>Routes</li> <li>196.19.255.0/24 <i>∂</i></li> </ul>	Statue AVALABLE	(T

8. Finally, allow bidirectional traffic firewall rules for access to FSx/CVO. Follow these detailed steps for compute gateway firewall rules for SDDC workload connectivity.

e Cout				
	sddcgroup01			
4				
(ana		Edit Routes Set of routes painting to this attachment @ 18819354.0294 @ 10216.0.0294 @ 10.2224		
		The profess can be addinized by somma, space or a now line	3 totis, 0 mailta ANCEL DONE	

9. After the firewall groups are configured for both the Management and Compute gateway, the vCenter can be accessed as follows:

											0.000	Neb	φp
Æ	The SODC will expire	m 54 days	LEARN	NORE								L	SCAL
	« ALL SDOCE								ii.	OPEN	WENT		7106
	lo ntan-fsx-de	emo l ve	K on AV	rs sooc 🙁 us west in	Credont				1	oren	e cerette	en je seu	TIGH
	Contrap Tox of												
	Summery Networking	& Security	Add	Oris Maintenanc	in Tre	sibleshooting Settings	Support						
	Overview	Gatev	vay F	irewall									
	Network	Manager	ment Ga	neway Comput	e Gatewa	у							
	Segments	and the second of the		ata ana relationna de la constante									
	VPN											- er vites	
	Tier-1 Gateways	+ ADD	BULE	(D-9399) *0	1990	(Element)			Filte	ic by N	ame, Pat	h and more	
	Transit Connect			Name	iD	Sources	Destinations	Services	Applied To	Att	lon		
	Security	4	Ó	elow internet Iro.	1019	🔛 vmc-sddc	Any	Any	All Uplinks		Allow		) 6
	Gateway Firewall					22 vmc-addc-2							
	Distributed Fireway	4	Ö	allow VMC to VPC	1017	tts emc-sdac	tonnecte.	Any	At uptres		Allow		) =
	Groups					wnc-sddc-2							
	Services	1		allow VPC to VNC	1016	11 Connecte.	22 www.eddc	Acty	All Uplines		Allow	·	) e
	Virtual Machines	1.0	č.			-	00 medada	Ami	and the second		We chi		
	Tools		ų.	arow to vincisity	TU ZZ	www.sddc-2	100	-375	Cospense		AJC/W	-	
	PFIX Port Mirroring	i i	Ö	all trom venetary?	1027	"" unctes2.v.	mc-sddc-2	Any	At Upmis		Allow		5.0
	Custom .	18	5	and the second second	and a	11 ( CONTRA	83 vmc sddc		2012/02/22		- Hereiter		1
	DNS	1	01	Default VTI Rule	1012	Ary.	Any	Any	VPN Tannel In.		ABOW		
	DHCP	4											101.00
	Global Configuration	1		Default Upleik Ru.		Any	Any	ANY	An uplinks		Drop	- 6	
	Direct Connect	CREFT	HER										
	Concepted VDC	10000000											

The next step is to verify that Amazon FSx ONTAP or Cloud Volumes ONTAP is configured depending on your requirements and that the volumes are provisioned to offload storage components from vSAN to optimize the deployment.

#### Deploy and configure the Virtualization Environment on Azure

As with on-premises, planning Azure VMware Solution is critical for a successful

production-ready environment for creating VMs and migration.

This section describes how to set up and manage Azure VMware Solution and use it in combination with the available options for connecting NetApp storage.

The setup process can be broken down into the following steps:

To use Azure VMware Solution, first register the resource provider within the identified subscription:

- 1. Sign in to the Azure portal.
- 2. On the Azure portal menu, select All Services.
- 3. In the All Services dialog box, enter the subscription and then select Subscriptions.
- 4. To view, select the subscription from the subscription list.
- 5. Select Resource Providers and enter Microsoft.AVS into the search.
- 6. If the resource provider is not registered, select Register.

Home > Subscriptions >					
Subscriptions «	Subscription		Resource providers		×
+ Add 📋 Manage Policies	P Search (Ctrl+/)		💎 Register 🏷 Unregister 🌔 Refresh		
View list of subscriptions for which you have	( Resource groups	^	₽ AVS		×
to manage Azure resources. To view subscriptions for which you have billing	Resources				
access, click here	Preview features		Provider	Status	
Don't see a subscription? Switch directories	Usage + quotas		Microsoft.AVS	Registering	
My role 🕕 Status 🛈	Policies				
8 selected V 3 selected V	Management certificates				
Apply	A My permissions				
Showing 1 of 1 subscriptions global Show only subscriptions selected in the	SE Resource providers				
subscriptions filter ①	i Deployments				
P Search	III Properties				
Subscription name 14	A Resource locks				
< Previous 1 V Next >	Support + troubleshooting	~			

	lan an
Provider	Status
Microsoft.OperationsManagement	Registered
Microsoft.Compute	Registered
Microsoft.ContainerService	Registered
Microsoft.ManagedIdentity	🥏 Registered
Microsoft.AVS	Registered
Microsoft.OperationalInsights	🧿 Registered
Microsoft.GuestConfiguration	O Registered

- 7. After the resource provider is registered, create an Azure VMware Solution private cloud by using the Azure portal.
- 8. Sign in to the Azure portal.
- 9. Select Create a New Resource.
- 10. In the Search the Marketplace text box, enter Azure VMware Solution and select it from the results.
- 11. On the Azure VMware Solution page, select Create.
- 12. From the Basics tab, enter the values in the fields and select Review + Create.

Notes:

- For a quick start, gather the required information during the planning phase.
- Select an existing resource group or create a new resource group for the private cloud. A resource group is a logical container in which the Azure resources are deployed and managed.
- Make sure the CIDR address is unique and does not overlap with other Azure Virtual Networks or onpremises networks. The CIDR represents the private cloud management network and is used for the cluster management services, such as vCenter Server and NSX-T Manager. NetApp recommends using a /22 address space. In this example, 10.21.0.0/22 is used.

Prerequisities Basics lags	Review and Create	
Project details		
Subscription * ②	SaaS Backup Production	×
Resource group * ④	(New) NimoAVSDemo	~
	Create new	
Private cloud details		
Resource name * 💿	nimoauspriv	Ŷ
Location * 💿	(US) East US 2	~~
Size of host * 💿	AV36 Trial	Ŷ
Number of hosts * ④	0	3
	Find	out how many hosts you nee
	There is no metering for the selected subscript data to display.	tion, region, and SKU. No cost
LIDK address block	d Real shares a second second at a second second beauty in the	
eromote en address for private clou	u vor vusser management, wake sure triese are uniqu itworks.	e and do not overlap with an
other Azure vnets or on-premise n		

The provisioning process takes approximately 4–5 hours. After the process is complete, verify that the deployment was successful by accessing the private cloud from the Azure portal. A status of Succeeded is displayed when the deployment is complete.

An Azure VMware Solution private cloud requires an Azure Virtual Network. Because Azure VMware Solution doesn't support on-premises vCenter, additional steps are required to integrate with an existing on-premises environment. Setting up an ExpressRoute circuit and a virtual network gateway is also required. While waiting for the cluster provisioning to complete, create a new virtual network or use an existing one to connect to Azure VMware Solution.

AVS Private cloud	***	
P Search (Ctrl+/)	x 🚺 Delete	
Overview	^ > Essentials	
Activity log	Resource group (change) NimoAVSDemo	Address block for private cloud 10.21.0.0/22
Access control (IAM)	Status Succeeded	Primary peering subnet 10.21.0.232/30
Diagnose and solve problems	Location East US 2	Secondary peering subnet 10.21.0.236/30
Settings	Subscription (change) SaaS Backup Production	Private Cloud Management network 10.21.0.0/26
A Locks	Subscription ID b58a041a-e464-4497-8be9-9048369ee8e1	vMotion network 10.21.1.128/25
Manage		Number of hosts
👷 Connectivity 🔍 Identity	Tags (change) Click here to add tags	
Clusters		

To create a new Azure Virtual Network (VNet), select the Azure VNet Connect tab. Alternatively, you can create one manually from the Azure portal by using the Create Virtual Network wizard:

- 1. Go to Azure VMware Solution private cloud and access Connectivity under the Manage option.
- 2. Select Azure VNet Connect.
- 3. To create a new VNet, select the Create New option.

This feature allows a VNet to be connected to the Azure VMware Solution private cloud. The VNet enables communication between workloads in this virtual network by automatically creating required components (for example, jump box, shared services such as Azure NetApp Files, and Cloud Volume ONTAP) to the private cloud created in Azure VMware Solution over ExpressRoute.

Note: The VNet address space should not overlap with the private cloud CIDR.

Search (Ctrl+/)	» (	🔄 Save 🕐 Refresh		
💁 Overview	^	Azure vNet connect Settings	ExpressRoute Public IP	
Activity log				
Access control (IAM)		This is an optional feature that allows private cloud. A vNet enables the cor	an Azure virtual network to be connected to nmunication between workloads in this virtua	your Azure VMware Solution I network (for example,
🔷 Tags		Jumpbox) to the private cloud create	d in Azure VMware Solution over ExpressRou lected. You can create a new vNet or use an	te. Only a vNet with a valid
Diagnose and solve problems		address space does not overlap with network.	your private cloud CIDR. Learn more about a	dding a subnet in a virtual
Settings		Virtual network		~
A Locks			Create new	
Manage		Address block for vnet	ж.	
👷 Connectivity		Address block for private cloud	10.21.0.0/22	D
udentity				
Clusters				

4. Provide or update the information for the new VNet and select OK.

This virtual network enab Azure VMware Solution o default address range an 172.16.0.0/16). Step 2: Ad (e.g. 172.16.1.0/24). Lean	les the communication between workloads in this virtual n wer an Express route. A default address range and a subne d subnet of this virtual network, follow these steps. Step 1: Id a subnet under "Subnets" with the name as "GatewaySu i more about virtual networks ⊡"	etwork (e.g. a Jumphost) to the private cloud creat t is selected for this virtual network. For changing Change the "Address Range" to desired range (e.e. bnet" and provide subnet's address range in CIDR	ted in the g. notation
Name *	nimoavspriv-vnet		2
Address space			
The virtual network's add	ress space specified as one or more address prefixes in CIE	DR notation (e.g. 10.0.0.0/16).	
Address range	Addresses	Overlap	
172,24.0.0/16	172.24.0.4 - 172.24.255.254 (65531 addresses)	None	Ē
	(0 Addresses)	None	
Subnets			
Subnets The subriet's address ran	ge in CIDR notation (e.g. 10.0.0.0/24). It must be contained	by the address space of the virtual network.	
Subnets The subnet's address ran Subnet name	ge in CIDR notation (e.g. 10.0.0.0/24). It must be contained Address range	by the address space of the virtual network. Addresses	
Subnets The subnet's address ran Subnet name GatewaySubnet	ge in CIDR notation (e.g. 10.0.0.0/24). It must be contained Address range 172.24.0.0/24	by the address space of the virtual network. Addresses 172.24.0.4 - 172.24.0.254 (251 addresses)	Û

The VNet with the provided address range and gateway subnet is created in the designated subscription and resource group.

If you create a VNet manually, create a virtual network gateway with the appropriate SKU and ExpressRoute as the gateway type. After the deployment is complete, connect the ExpressRoute connection to the virtual network gateway containing Azure VMware Solution private cloud using the authorization key. For more information, see Configure networking for your VMware private cloud in Azure.

 $(\mathbf{i})$ 

#### Validate the network connect and access to Azure VMware Solution private cloud

Azure VMware Solution does not allow you to manage a private cloud with on-premises VMware vCenter. Instead, jump host is required to connect to the Azure VMware Solution vCenter instance. Create a jump host in the designated resource group and sign in to the Azure VMware Solution vCenter. This jump host should be a Windows VM on the same virtual network that was created for connectivity and should provide access to both vCenter and the NSX Manager.

Create	e a vi	rtual mae	chine					
Basics	Disks	Networking	Management	Advanced	Tags	Review + create		
Create a image. Cr tab for fu	virtual ma omplete t Ill custom	achine that runs t he Basics tab the ization. Learn m	inux or Windows. In Review + create ore cl	Select an imag to provision a	e from A virtual m	zure marketplace or use your o lachine with default parameters	wn customized or review each	
Project d	ietails							
Select the manage i	e subscrip all your re	ition to manage esources.	deployed resource	rs and costs. U	ise resou	rce groups like folders to organ	ize and	
Subscript	tion * 💿		SaaS Bad	kup Productio	0		~	
B	esource (	aroup * 🛈	NimoAVSDemo				~	
		50.92	Create nev	N.				
Instance	details							
Virtual m	achine na	me * 💿	nimAV\$.94				~	
Region *	0		(US) East	US 2			~]	
Availabili	ty options	0	No infrastructure redundancy required			quired	~	
Image *	0		Wind	lows Server 20	12 R2 Da	tacenter - Gen2	~	
			See all ima	ages				
Azure Sp	ot instanc	e 🖸						
Size * 🖸	)		Standard	(_D2s_v3 - 2 vc	pus, 8 Gi	8 memory (\$130.67/month)	$\sim$	
			See all size	es			701	

After the virtual machine is provisioned, use the Connect option to access RDP.

Virtual machine	t
	▲ To improve security, enable just-in-time access on this VM. →
Overview	
Activity log	RDP SSH BASTION
Access control (IAM)	Connect with RDP
Tags	To connect to your virtual machine via RDP, select an IP address, optionally change the port number, and download the
Diagnose and solve problems	RDP file.
	IP address *
Settings	Public IP address (52:138:103:135)
2 Networking	Port number *
9 Connect	3389
Cite	Developed PDP File
Disks	Dowindau KDF File

Sign in to vCenter from this newly created jump host virtual machine by using the cloud admin user . To access the credentials, go to the Azure portal and navigate to Identity (under the Manage option within the private cloud). The URLs and user credentials for the private cloud vCenter and NSX-T Manager can be copied from here.

AVS Private cloud	5 Y Y SY		
Search (Ctrl+/)	« Login credentials		
Access control (IAM)	<ul> <li>vCenter credentials</li> </ul>		
🧳 Tags	Web client URL 💿	https://10.21.0.2/	Ð
Diagnose and solve problems	Admin username	cloudadmin@vsphere.local	D
Settings	Admin password ①		
A Locks	Certificate thumbprint ③	AE26B15A5CE38DC069D35F045F088CA6343475EC	Ð
Manage	NSX-T Manager credentials		
🤶 Connectivity	Web client URL ①	https://10.21.0.3/	D
🖳 Identity	Admin username	admin	ID
le Clusters	And Bennine C		10
Placement policies (preview)	Admin password 💿	P	
+ Add-ons	Certificate thumbprint ①	B2B722EA683958283EE159007246D5166D0509D3	Ð

In the Windows virtual machine, open a browser and navigate to the vCenter web client URL ("https://10.21.0.2/") and use the admin user name as **cloudadmin@vsphere.local** and paste the copied password. Similarly, NSX-T manager can also be accessed using the web client URL ("https://10.21.0.3/") and use the admin user name and paste the copied password to create new segments or modify the existing tier gateways.



vm vSphere Client	Menu V Q Search in all environments		C C	⊙ ✓ cloudadmi	nøvsherelocal ~ (
	😰 vc.beeb9fd29eab4cbea81e62.eas	tus2.avs.az	zure.com	ions -	
🗸 💋 vc.beeb9fd29eab4cbea8ie .	Sum. Mon. Confl., Permiss., Datace.,	Hosts & Cl.	V. Datast_	Netw_ Link	ed vCenter Server _ Extern
SDDC-Datacenter	Virtual Machines: 0			, cru	Free, 201-73 GHz
	Hosts: 3			Used: 10.0	2 OFE Capacity 247 75 Ore
				blarmry.	Free: 1.44.72
				Uturit: 246-	A1 GE Capacity: 1.10 TH
				Shrape	Feer 34.32 TB
				Used 7.61	18 Capacity, 41.92 TH
	Custom Attributes	~	Tags		^
	Attribute Value		Assigned Tag	Category	Description
		8			×.
Recent Tasks Alarms					
Task Name 🗢 Target	~ Status ~ Details ~ Initiator	~ Queue	ed For 🛛 👻 Start Tim	e 4 ~ Comple	tion Time 🗠 Server
Undepicy plug-in 💋 vc.beeb96	VMware vRops 129 VCompleted Client Plugin VSPHERE.0	OCALL 8 ms	08/12/20 AM	21, 11, 38, 11 08/12/2 AM	vc.beeb%fd2%cab

The Azure VMware Solution SDDC is now deployed and configured. Leverage ExpressRoute Global Reach to connect the on-premises environment to Azure VMware Solution private cloud. For more information, see Peer on-premises environments to Azure VMware Solution.

Deploy and configure the Virtualization Environment on Google Cloud Platform (GCP)

As with on-premises, planning Google Cloud VMware Engine (GCVE) is critical for a successful production-ready environment for creating VMs and migration.

This section describes how to set up and manage GCVE and use it in combination with the available options for connecting NetApp storage.

The setup process can be broken down into the following steps:

To configure a GCVE environment on GCP, login to the GCP console and access the VMware Engine portal.

Click on the "New Private Cloud" button and enter the desired configuration for the GCVE Private Cloud. On "Location", make sure to deploy the private cloud in the same Region/Zone where CVS/CVO is deployed, to ensure the best performance and lowest latency.

Pre-requisites:

- Setup VMware Engine Service Admin IAM role
- Enable VMWare Engine API access and node quota

Note: Private cloud creation can take between 30 minutes to 2 hours.

• Make sure that the CIDR range doesn't overlap with any of your on-premises or cloud subnets. The CIDR range must be /27 or higher.

Private Cloud name *	
NiMoGCVE	
Location *	
us-east4 > v-zone-a > VE Placer	nent Group 2 •
Node type *	
ve1-standard-72	
2x2.6 GHz, 36 Cores (72 HT), 768 ( 19.2 TB Raw, 3.2 TB Cache (All-Fla	SB RAM (h)
Node count *	
3	
(3to3)	
vSphere/vSAN subnets CIDR	range *
192.168.100.0	
IP Range: 192.168.100.0 - 192.168	103.255
HCX Deployment Network C	IDR range
192.168.104.0	
Once the Private Cloud is provisioned, configure private access to the Private Cloud for high-throughput and low-latency data-path connection.

This will ensure that the VPC network where Cloud Volumes ONTAP instances are running is able to communicate with the GCVE Private Cloud. To do so, follow the GCP documentation. For the Cloud Volume Service, establish a connection between VMware Engine and Cloud Volumes Service by performing a one-time peering between the tenant host projects. For detailed steps, follow this link.

Tenant P 👫 🗍 🌲	Service	÷	Region	*	Routing Mode	-	Peered Project ID 🌐	Peered VPC	$\frac{A}{T}$	VPC Peering Sta 🗘	<b>Region Status</b>
ke841388caa56b	VPC Network		europe-west3		Global		cv-performance-te	cloud-volumes-vpc		Active	<ul> <li>Connected</li> </ul>
jbd729510b3ebbf	NetApp CVS		europe-west3		Global		y2b6c17202afódc	netapp-tenant-vpc		Active	Connected

Sign in to vcenter using the CloudOwner@gve.local user. To access the credentials, go to the VMware Engine portal, Go to Resources, and select the appropriate private cloud. In the Basic info section, click the View link for either vCenter login info (vCenter Server, HCX Manager) or NSX-T login info (NSX Manager).



In a Windows virtual machine, open a browser and navigate to the vCenter web client URL ("https://10.0.16.6/") and use the admin user name as CloudOwner@gve.local and paste the copied password. Similarly, NSX-T manager can also be accessed using the web client URL ("https://10.0.16.11/") and use the admin user name and paste the copied password to create new segments or modify the existing tier gateways.

For connecting from an on-premises network to VMware Engine private cloud, leverage cloud VPN or Cloud Interconnect for appropriate connectivity and make sure the required ports are open. For detailed steps, follow this link.



## Deploy NetApp Cloud Volume Service supplemental datastore to GCVE

Refer Procedure to deploy supplemental NFS datastore with NetApp CVS to GCVE

# NetApp Storage options for Public Cloud Providers

Explore the options for NetApp as storage in the three major hyperscalers.

## AWS / VMC

AWS supports NetApp storage in the following configurations:

- FSx ONTAP as guest connected storage
- Cloud Volumes ONTAP (CVO) as guest connected storage
- FSx ONTAP as a supplemental NFS datastore

View the detailed guest connect storage options for VMC. View the detailed supplemental NFS datastore options for VMC.

# Azure / AVS

Azure supports NetApp storage in the following configurations:

- Azure NetApp Files (ANF) as guest connected storage
- Cloud Volumes ONTAP (CVO) as guest connected storage
- Azure NetApp Files (ANF) as a supplemental NFS datastore

View the detailed guest connect storage options for AVS. View the detailed supplemental NFS datastore options for AVS.

## GCP / GCVE

Google Cloud supports NetApp storage in the following configurations:

- Cloud Volumes ONTAP (CVO) as guest connected storage
- · Cloud Volumes Service (CVS) as guest connected storage
- · Cloud Volumes Service (CVS) as a supplemental NFS datastore

View the detailed guest connect storage options for GCVE.

Read more about NetApp Cloud Volumes Service datastore support for Google Cloud VMware Engine (NetApp blog) or How to use NetApp CVS as datastores for Google Cloud VMware Engine (Google blog)

## TR-4938: Mount Amazon FSx for ONTAP as a NFS datastore with VMware Cloud on AWS

This document outlines how to mount Amazon FSx for ONTAP as a NFS datastore with VMware Cloud on AWS.

Niyaz Mohamed, NetApp

# Introduction

Every successful organization is on a path of transformation and modernization. As part of this process, companies typically use their existing VMware investments to leverage cloud benefits and exploring how to migrate, burst, extend, and provide disaster recovery for processes as seamlessly as possible. Customers migrating to the cloud must evaluate the use cases for elasticity and burst, data-center exit, data-center consolidation, end-of-life scenarios, mergers, acquisitions, and so on.

Although VMware Cloud on AWS is the preferred option for the majority of the customers because it delivers unique hybrid capabilities to a customer, limited native storage options have restricted its usefulness for

organizations with storage-heavy workloads. Because storage is directly tied to hosts, the only way to scale storage is to add more hosts, which can increase costs by 35-40% or more for storage intensive workloads. These workloads need additional storage and segregated performance, not additional horsepower, but that means paying for additional hosts. This is where the recent integration of FSx for ONTAP comes in handy for storage and performance intensive workloads with VMware Cloud on AWS.

Let's consider the following scenario: a customer requires eight hosts for horsepower (vCPU/vMem), but they also have a substantial requirement for storage. Based on their assessment, they require 16 hosts to meet storage requirements. This increases the overall TCO because they must buy all that additional horsepower when all they really need is more storage. This is applicable for any use case, including migration, disaster recovery, bursting, dev/test, and so on.

This document walks you through the steps necessary to provision and attach FSx for ONTAP as a NFS datastore for VMware Cloud on AWS.



(i)

This solution is also available from VMware. Please visit the VMware Cloud Tech Zone for more information.

# **Connectivity options**

VMware Cloud on AWS supports both multi-AZ and single-AZ deployments of FSx for ONTAP.

This section describes the high-level connectivity architecture along with the steps needed to implement the solution to expand the storage in a SDDC cluster without the need for adding additional hosts.



The high-level deployment steps are as follows:

- 1. Create Amazon FSx for ONTAP in a new designated VPC.
- 2. Create an SDDC group.

- 3. Create VMware Transit Connect and a TGW attachment.
- 4. Configure routing (AWS VPC and SDDC) and security groups.
- 5. Attach an NFS volume as a datastore to the SDDC cluster.

Before you provision and attach FSx for ONTAP as a NFS datastore, you must first set up a VMware on Cloud SDDC environment or get an existing SDDC upgraded to v1.20 or above. For more information, see the Getting Started With VMware Cloud on AWS.



FSx for ONTAP is not currently supported with stretched clusters.

# Conclusion

This document covers the steps necessary to configure Amazon FSx for ONTAP with VMware cloud on AWS. Amazon FSx for ONTAP provides excellent options to deploy and manage application workloads along with file services while reducing the TCO by making data requirements seamless to the application layer. Whatever the use case, choose VMware Cloud on AWS along with Amazon FSx for ONTAP for rapid realization of cloud benefits, consistent infrastructure, and operations from on-premises to AWS, bidirectional portability of workloads, and enterprise-grade capacity and performance. It is the same familiar process and procedures used to connect storage. Remember, it is just the position of the data that changed along with new names; the tools and processes all remain the same, and Amazon FSx for ONTAP helps to optimize the overall deployment.

To learn more about this process, feel free to follow the detailed walkthrough video.

Amazon FSX for Ontap VMware Cloud

## NetApp Guest Connected Storage Options for AWS

AWS supports guest connected NetApp storage with the native FSx service (FSx ONTAP) or with Cloud Volumes ONTAP (CVO).

# **FSx ONTAP**

Amazon FSx for NetApp ONTAP is a fully managed service that provides highly reliable, scalable, highperforming, and feature-rich file storage built on NetApp's popular ONTAP file system. FSx for ONTAP combines the familiar features, performance, capabilities, and API operations of NetApp file systems with the agility, scalability, and simplicity of a fully managed AWS service.

FSx for ONTAP provides feature-rich, fast, and flexible shared file storage that's broadly accessible from Linux, Windows, and macOS compute instances running in AWS or on premises. FSx for ONTAP offers high-performance solid state drive (SSD) storage with submillisecond latencies. With FSx for ONTAP, you can achieve SSD levels of performance for your workload while paying for SSD storage for only a small fraction of your data.

Managing your data with FSx for ONTAP is easier because you can snapshot, clone, and replicate your files with the click of a button. In addition, FSx for ONTAP automatically tiers your data to lower-cost, elastic storage, lessening the need for you to provision or manage capacity.

FSx for ONTAP also provides highly available and durable storage with fully managed backups and support for cross-Region disaster recovery. To make it easier to protect and secure your data, FSx for ONTAP supports popular data security and antivirus applications.

## FSx ONTAP as guest connected storage

## Configure Amazon FSx for NetApp ONTAP with VMware Cloud on AWS

Amazon FSx for NetApp ONTAP files shares and LUNs can be mounted from VMs that are created within the VMware SDDC environment at VMware Cloud at AWS. The volumes can also be mounted on the Linux client and mapped on the Windows client using the NFS or SMB protocol, and LUNS can be accessed on Linux or Windows clients as block devices when mounted over iSCSI. Amazon FSx for the NetApp ONTAP file system can be set up quickly with the following steps.



Amazon FSx for NetApp ONTAP and VMware Cloud on AWS must be in the same availability zone to achieve better performance and avoid data transfer charges between availability zones.

To create and mount Amazon FSx for NetApp ONTAP file system, complete the following steps:

- 1. Open the Amazon FSx console and choose Create file system to start the file system creation wizard.
- 2. On the Select File System Type page, choose Amazon FSx for NetApp ONTAP, and then choose Next. The Create File System page appears.

P <b>tions</b> F5x for NetApp ONTAP	Amazon FSx for Windows File	Amazon FSx for Lustre	
F5x for NetApp ONTAP	Amazon FSx for Windows File	Amazon FSx for Lustre	11
FSX mazon FSx letApp ONTAP	FSX Amazon FSx for Windows File Server	FSX Amazon FSx for Lustre	
	Select file system type		
	トンステ mazon FSx letApp ONTAP	FSX mazon FSx letApp ONTAP FSX Amazon FSx for Windows File Server Select file system type	FSX mazon FSx letApp ONTAP Select file system type

1. In the Networking section, for Virtual Private Cloud (VPC), choose the appropriate VPC and preferred subnets along with the route table. In this case, vmcfsx2.vpc is selected from the dropdown.

reation method	
<ul> <li>Quick create</li> <li>Use recommended best-practice configurations.</li> <li>Most configuration options can be changed after the file system is created.</li> </ul>	<ul> <li>Standard create</li> <li>You set all of the configuration options, including specifying performance, networking, security, backups, and maintenance.</li> </ul>

1. For the creation method, choose Standard Create. You can also choose Quick Create, but this document uses the Standard create option.

File system name - optional	fo	
vmcfsxval2		
Maximum of 256 Unicode letters, wh	tespace, and numbers, plus + - = : /	
SSD storage capacity Info		
1024 0		
Minimum 1024 GB; Maximum 192 TE		
Provisioned SSD IOPS Amazon FSx provides 3 IOPS per GB needed.	f storage capacity. You can also provision additional SS	D IOPS as
Automatic (3 IOPS per GB o	SSD storage)	
O User-provisioned		
Throughput capacity Info The sustained speed at which the file burst to higher speeds for periods of	server hosting your file system can serve data. The file sime.	erver can also
E12 MP/s (Perommonded)		

1. In the Networking section, for Virtual Private Cloud (VPC), choose the appropriate VPC and preferred subnets along with the route table. In this case, vmcfsx2.vpc is selected from the dropdown.

Virtual Private Cloud (VPC) Info Specify the VPC from which your file system is accessible.	
vmcfsx2.vpc   vpc-0d1c764bcc495e805	v
VPC Security Groups Info Specify VPC Security Groups to associate with your file system's network interface.	
Choose VPC security group(s)	٧
5g-018896ea218164ccb (default) ×	
Preferred subnet Info Specify the preferred subnet for your file system.	
subnet02.sn   subnet-013675849a5b99b3c (us-west-2b)	
Standby subnet	
subnet01.sn   subnet-0ef956cebf539f970 (us-west-2a)	
VPC route tables Specify the VPC route tables associated with your file system.	
VPC's default route table	
Select one or more VPC route tables	
Endpoint IP address range Specify the IP address range in which the endpoints to access your file system will be created	
No preference	
Select an IP address range	



In the Networking section, for Virtual Private Cloud (VPC), choose the appropriate VPC and preferred subnets along with the route table. In this case, vmcfsx2.vpc is selected from the dropdown.

1. In the Security & Encryption section, for the Encryption Key, choose the AWS Key Management Service (AWS KMS) encryption key that protects the file system's data at rest. For the File System Administrative Password, enter a secure password for the fsxadmin user.

Incryption key Info WS Key Management Service (KMS) encryption key that protects	your file system data at re	st.
aws/fsx (default)		•
Description	Account	KMS key ID
Default master key that protects my FSx resources when no other key is defined	139763910815	72745367-7bb0-499c- acc0-4f2c0a80e7c5
ile system administrative password		
are system doministrative possible user which you can use t	a access the ONTAR CLLG	PEST ADI
assword for this file system's "fsxadmin" user, which you can use to Don't specify a password	o access the ONTAP CLI o	r REST API.
<ul> <li>assword for this file system's "fsxadmin" user, which you can use t</li> <li>Don't specify a password</li> <li>Specify a password</li> </ul>	to access the ONTAP CLI o	r REST API.
<ul> <li>assword for this file system's "fsxadmin" user, which you can use 1</li> <li>Don't specify a password</li> <li>Specify a password</li> <li>Password</li> </ul>	o access the ONTAP CLI o	r REST API.
<ul> <li>assword for this file system's "fsxadmin" user, which you can use 1</li> <li>Don't specify a password</li> <li>Specify a password</li> <li>Password</li> </ul>	o access the ONTAP CLI o	r REST API.
assword for this file system's "fsxadmin" user, which you can use t Don't specify a password Specify a password assword confirm password	o access the ONTAP CLI o	r REST API.

 In virtual machine and specify the password to use with vsadmin for administering ONTAP using REST APIs or the CLI. If no password is specified, a fsxadmin user can be used for administering the SVM. In the Active Directory section, make sure to join Active Directory to the SVM for provisioning SMB shares. In the Default Storage Virtual Machine Configuration section, provide a name for the storage in this validation, SMB shares are provisioned using a self-managed Active Directory domain.

Storage virtual machine name	
vmcfsxval2svm	
SVM administrative password Password for this SVM's "vsadmin" user, which you can use to access the ONTAP CLI or REST API.	
O Don't specify a password	
<ul> <li>Specify a password</li> </ul>	
Password	
******	
Confirm password	
******	
Active Directory Joining an Active Directory enables access from Windows and MacOS clients over the SMB protocol.	
<ul> <li>Do not join an Active Directory</li> </ul>	
Join an Active Directory	

 In the Default Volume Configuration section, specify the volume name and size. This is an NFS volume. For Storage Efficiency, choose Enabled to turn on the ONTAP storage efficiency features (compression, deduplication, and compaction) or Disabled to turn them off.

Default volume configuration	
Volume name	
val1	
Maximum of 203 alphanumeric characters, plus Junction path	
/vol1	
The location within your file system where your volume will be mounted.	
Volume size	
1024	0
Minimum 20 MiB; Maximum 104857600 MiB	
Storage enciency Select whether you would like to enable ONTAP storage efficiencies on your volume: deduplicat compression, and compaction.	ion,
Enabled (recommended)	
O Disabled	
Capacity pool tiering policy You can optionally enable automatic tiering of your data to lower-cost capacity pool storage.	
Auto	

1. Review the file system configuration shown on the Create File System page.

2. Click Create File System.

aws Services V	<b>Q</b> Search for services, features, marketplace	products, and docs	[Alt+S]	D A	nimo @ cloudheroes 🔻	Oregon 🔻 S	iupport
Amazon FSx ×	F5x > File systems						
File systems	File systems (3)			C Attach	Actions ¥	ireate file syste	em
Backups	Q Filter file systems					< 1 >	0
ONTAP							
Storage virtual machines Volumes	File system File sys name ♥	tem ID 🔺	File system 5 type ⊽	štatus 🗸 🗸	Deployment type ⊽	Storage type ⊽	6
Windows File Server	O fsxntapcifs	28399be9c1f9f	ONTAP (	O Available	Multi-AZ	SSD	
Data repository tasks	vmcfsxval2	acc5d0ac31017	ONTAP (	O Available	Multi-AZ	SSD	
FSx on Service Quotas 🖸	O fsxntapsql 🗗	b447ebd6082aa	ONTAP (	🔊 Available	Multi-AZ	SSD	
Network & security	Administration Storage vir	tual machines	Volume	s Backups	Tags		
SVM name ▼ fsxsmbtesting01 vmcfsxval2svm	SVM ID     ▼       svm-075dcfbe2cfa2ece9       svm-095db076341561212	Status ♥ ⊘ Created Created	Creatio 2021-14 +01:00 2021-14 +01:00	<b>n time</b> D-19 15:17:08 UT D-15 15:16:54 UT	Active rc FSXTE rc -	e Directory	⊽ AL
<ul> <li>vmcfsxval2svm</li> <li>x &gt; Storage virtual machi</li> <li>sxsmbtesting0</li> </ul>	svm-095db076341561212 nes > svm-075dcfbe2cfa2eces 1 (svm-075dcfbe	⊘ Created	2021-10 +01:00	0-15 15:16:54 U	TC _	Upda	ite
Summary							
SVM ID	Creation time	2		Active Direc	tory		
svm-075dcfbe2cfa2ece9	j 2021-10-19T	15:17:08+01:00	D	FSXTESTING	5.LOCAL		
SVM name	Lifecycle stat	е		Net BIOS na	ime		
fsxsmbtesting01 🗇	⊘ Created			FSXSMBTES	TING01		
	Subtype			Fully qualifi	ed domain name	6	
סוטט	DEFAULT			FSXTESTING	S.LOCAL		
4a50e659-30e7-11ec-ac4f- f3ad92a6a735	· · · · · · · · · · · · · · · · · · ·			Service acco	ount username		

File system ID

fs-040eacc5d0ac31017

administrator

Organizational unit distinguished name **CN=Computers** 

For more detailed information, see Getting started with Amazon FSx for NetApp ONTAP.

After the file system is created as above, create the volume with the required size and protocol.

- 1. Open the Amazon FSx console.
- 2. In the left navigation pane, choose File systems, and then choose the ONTAP file system that you want to create a volume for.
- 3. Select the Volumes tab.
- 4. Select the Create Volume tab.
- 5. The Create Volume dialog box appears.

For demo purposes, an NFS volume is created in this section that can be easily mounted on VMs running on VMware cloud on AWS. nfsdemovol01 is created as depicted below:

	>
File system	
fs-040eacc5d0ac31017   vmcfsxval2	
Storage virtual machine	
svm-095db076341561212   vmcfsxval2svm	•
Volume name	
nfsdemovol01	
Maximum of 203 alphanumeric characters, plus	
Junction path	
/nfsdemovol01	
The location within your file system where your volume will be mounted	L.
Volume size	
1024	(0)
Minimum 20 MiB; Maximum 104857600 MiB	
Storage efficiency Select whether you would like to enable ONTAP storage efficiencies on y compression, and compaction.	our volume: deduptication,
Storage efficiency Select whether you would like to enable ONTAP storage efficiencies on y compression, and compaction. C Enabled (recommended)	our volume: deduplication,
Storage efficiency Select whether you would like to enable ONTAP storage efficiencies on y compression, and compaction. Enabled (recommended) Disabled	our volume: deduplication,
Storage efficiency Select whether you would like to enable ONTAP storage efficiencies on y compression, and compaction. Cabled (recommended) Disabled Capacity pool tiering policy You can optionally enable automatic tiering of your data to lower-cost of	our volume: deduplication,

To mount the FSx ONTAP volume created in the previous step. from the Linux VMs within VMC on AWS SDDC, complete the following steps:

- 1. Connect to the designated Linux instance.
- 2. Open a terminal on the instance using Secure Shell (SSH) and log in with the appropriate credentials.
- 3. Make a directory for the volume's mount point with the following command:

```
$ sudo mkdir /fsx/nfsdemovol01
```

4. Mount the Amazon FSx for NetApp ONTAP NFS volume to the directory that is created in the previous step.

```
sudo mount -t nfs nfsvers=4.1,198.19.254.239:/nfsdemovol01
/fsx/nfsdemovol01
```

roat@ubuntu01:/fsx/nfsdemovol01# mount -t nfs 198.19.254.239:/nfsdemovol01 /fsx/nfsdemovol01

1. Once executed, run the df command to validate the mount.

🚱 vSphere - ubuntu01 - Summary ×	ubuntu01	× +	0
$\leftarrow \rightarrow c$	O A ≓ https://vcenter.s	ddc-52-37-127-104 vmwarevmc.com/ui/webconsole.html?vmld=vm-1003&vmName	=ubuntu018iseາ> ຊີງ
📵 Getting Started 🔋 EC2 Managem	ierit Con 🝓 New Tab		
ubuntu01	root@ubuntu01 Filesystem tmpfs /dev/mapper/ub tmpfs tmpfs /dov/sda2 tmpfs 170.16.0.22/m 190.19.254.23 root@ubuntu01 root@ubuntu01 root@ubuntu01	/fsk/nfsdemovol01# df IX-blocks Used Available Use% Mounted on 814396 1176 81220 1% /run untuvg-ubuntuiv 15412166 3666428 10943132 258 / 4071560 0 4071560 258 / 4071560 0 4071560 0% /run/lock 4096 0 4095 0% /sys/fs/cgroup 993220 254396 675512 28% /boot 814332 4 814388 1% /run/user/1000 814332 4 814388 1% /run/user/1000 814332 4 814388 1% /run/user/1000 5demovol01 9951472 4241735 571968 43% /fs/cgroup 995480 512 99548 1% /fs/cgroup	eyboard Layout View Fullscreen

Mount FSx ONTAP volume on Linux client

To manage and map file shares on an Amazon FSx file system, the Shared Folders GUI must be used.

- 1. Open the Start menu and run fsmgmt.msc using Run As Administrator. Doing this opens the Shared Folders GUI tool.
- 2. Click Action > All tasks and choose Connect to Another Computer.
- 3. For Another Computer, enter the DNS name for the storage virtual machine (SVM). For example, FSXSMBTESTING01.FSXTESTING.LOCAL is used in this example.



Tp find the SVM's DNS name on the Amazon FSx console, choose Storage Virtual Machines, choose SVM, and then scroll down to Endpoints to find the SMB DNS name. Click OK. The Amazon FSx file system appears in the list for the Shared Folders.

Management IP address

198.19.254.9

198.19.254.9

SMB IP address

198.19.254.9

iSCSI IP addresses

10.222.2.224, 10.222.1.94

NFS IP address

# Endpoints

Management DNS name

svm-075dcfbe2cfa2ece9.fs-040eacc5d0ac31017.fsx.us-

west-2.amazonaws.com

NFS DNS name

svm-075dcfbe2cfa2ece9.fs-040eacc5d0ac31017.fsx.us-

west-2.amazonaws.com

SMB DNS name

FSXSMBTESTING01.FSXTESTING.LOCAL

**iSCSI DNS** name

iscsi.svm-075dcfbe2cfa2ece9.fs-040eacc5d0ac31017.fsx.us-

west-2.amazonaws.com

1. In the Shared Folders tool, choose Shares in the left pane to see the active shares for the Amazon FSx file system.

	(E3) ( 20)					
mputer Management (FSXSM System Tools Task Scheduler Event Viewer Shared Folders Stares Stares Storage Device Manager Storage Windows Server Backup Storage Disk Management Services and Applications	VBTESTING01.FSXTESTING.LOCAL)	Share Name BacS B. pcS B. smbdemo B. testnimvol	Folder Path C:\ C:\smbdernovol01 C:\testnimvol	Type Windows Windows Windows	# Client Connections 0 1 1 0	Description
Now choose a ne	ew share and complete	the Create	a Shared Fol	der wiza	rd.	
eate A Shared Fol Name, Descriptio	der Wizard on, and Settings				×	
eate A Shared Fol Name, Description Specify how portion Type information all	der Wizard on, and Settings eople see and use this sha bout the share for users.	re over the r To modify ho	network. w people use t	the conten	x 22	
eate A Shared Fol Name, Description Specify how po Type information al offline, click Chang Share name:	der Wizard on, and Settings eople see and use this sha bout the share for users. e. nimtestsmb.01	re over the r To modify ho	network. w people use t	the conten	x 23	
eate A Shared Fol Name, Description Specify how po Type information al offline, click Chang Share name: Share path:	der Wizard on, and Settings eople see and use this sha bout the share for users. e. nimtestsmb01 \\FSXSMBTESTING01.F	To modify ho	network. w people use t LOCAL (nimtes)	the conten	x 22	
eate A Shared Fol Name, Description Specify how po Type information a offline, click Chang Share name: Share path: Description:	der Wizard on, and Settings eople see and use this sha bout the share for users. e. nimtestsmb01 \VFSXSMBTESTING01.F	To modify ho	network. w people use t LOCAL (nimtes)	tsmb01	x zz	
eate A Shared Fol Name, Description Specify how portion Type information al offline, click Chang Share name: Share path: Description: Offline setting:	der Wizard on, and Settings eople see and use this sha bout the share for users. e. nimtestsmb01 \VFSXSMBTESTING01.f Selected files and prog	To modify ho SXTESTING.	network. w people use t LOCAL (nimtes)	tsmb01	x while	
eate A Shared Fol Name, Descripti Specify how p Type information a offline, click Chang Share name: Share path: Description: Offline setting:	der Wizard on, and Settings eople see and use this sha bout the share for users. e. nimtestsmb01 \VFSXSMBTESTING01.F Selected files and prog	To modify ho SXTESTING.	network. w people use t LOCAL (nimtes)	tsmb01	t while	

< Back

Next >

Cancel

reate A Shared Folder Wizard	charing a characteristic de la	×
	Sharing was Successful	
	Status:	
22	You have successfully completed the Share a Folder Wizard.	~
	Summary:	×
	You have selected the following share settings on \ \FSXSMBTESTING01.FSXTESTING.LOCAL: Folder path: C:\nimtestsmb01 Share name: nimtestsmb01 Share path: \FSXSMBTESTING01.FSXTESTING.LOCAL \nimtestsmb01	~
	When I click Finish, run the wizard again to share and folder	ther
	To dose this wizard, dick Finish.	

To learn more about creating and managing SMB shares on an Amazon FSx file system, see Creating SMB Shares.

1. After connectivity is in place, the SMB share can be attached and used for application data. To accomplish this, Copy the share path and use the Map Network Drive option to mount the volume on the VM running on VMware Cloud on the AWS SDDC.

VMware Cloud Services - Log In $\times$	🚱 vSphere - vm	cdc01 - Summary ×	vmcdc01	× 📰 Sign	out	× +		C	- a
⊢ → C		os//vcenter.addc+	52-37-127-104.vmware	wmc.com/ui/webcontole.h	ntml/vmld=vm-10058cvr	nName	wmedet 67%	¢	⊚ 🛓
Getting Started 🔋 EC2 Managem	sent Con 🧕 New	Tab							C Other Bookm
edol/1							Enforce UI	Keyboard Leyest	lew Fullstream Send Ch1+AbH
Actos View Hulp		#12 <mark>9.</mark> **	Manage Unibdam	www.05(),118.7623438()(t))			- 0	×	~ 3
◆)名前(日月9)日前		4 The	PC + articlementP10108/02541	m - C		- 6	facet constants of the		
erepatie Management #SKMATELTAKOF PSHTESTAK § fytter Taob ③ Teak Scholder Feart Viewer W Bangd Pohlen	LOCAU Starefune gard gived gived gived gived gived gived gived	e Quel acres Destrop e Destrop e	Terre peus falder01 vimetes 1	Determotion Type 14/19/2011/02/04/4 Factories 15/22/2011/225.444 Factories	See				Actions Dismu More Actions anticipensed()
General     General     General     General     General     General     General     General	a tetrinu	S Decorante d'	Eadfarina C	1022/2011/2/2014 File failur					More Jahana
A Dence Manager Storage Wintexe Server Beitrage P Dia Management Server and Server		M DVD Drive (D.) 355, 31 Metzonik							

Connect a FSx for NetApp ONTAP LUN to a host using iSCSI

iSCSI traffic for FSx traverses the VMware Transit Connect/AWS Transit Gateway via the routes provided in the previous section. To configure a LUN in Amazon FSx for NetApp ONTAP, follow the documentation found here.

On Linux clients, make sure that the iSCSI daemon is running. After the LUNs are provisioned, refer to the detailed guidance on iSCSI configuration with Ubuntu (as an example) here.

In this paper, connecting the iSCSI LUN to a Windows host is depicted:

- 1. Access the NetApp ONTAP CLI using the management port of the FSx for the ONTAP file system.
- 2. Create the LUNs with the required size as indicated by the sizing output.

FsxId040eacc5d0ac31017::> lun create -vserver vmcfsxval2svm -volume
nimfsxscsivol -lun nimofsxlun01 -size 5gb -ostype windows -space
-reserve enabled

In this example, we created a LUN of size 5g (5368709120).

1. Create the necessary igroups to control which hosts have access to specific LUNs.

```
FsxId040eacc5d0ac31017::> igroup create -vserver vmcfsxval2svm -igroup
winIG -protocol iscsi -ostype windows -initiator ign.1991-
05.com.microsoft:vmcdc01.fsxtesting.local
FsxId040eacc5d0ac31017::> igroup show
Vserver
       Igroup Protocol OS Type Initiators
_____ ____
_____
vmcfsxval2svm
        ubuntu01 iscsi
                          linux iqn.2021-
10.com.ubuntu:01:initiator01
vmcfsxval2svm
        winIG
                   iscsi
                           windows iqn.1991-
05.com.microsoft:vmcdc01.fsxtesting.local
```

Two entries were displayed.

1. Map the LUNs to igroups using the following command:

FsxId040e /vol/nimf	acc5d0ac31017::> lun map -vserve sxscsivol/nimofsxlun01 -igroup w	r vmcfsx inIG	val2svm -	path
FsxId040e	acc5d0ac31017::> lun show			
Vserver Size	Path	State	Mapped	Туре
vmcfsxval	2svm			
5gb	/vol/blocktest01/lun01	online	mapped	linux
vmcfsxval	2svm			
5gb	/vol/nimfsxscsivol/nimofsxlun01	online	mapped	windows

Two entries were displayed.

1. Connect the newly provisioned LUN to a Windows VM:

To connect the new LUN tor a Windows host residing on VMware cloud on AWS SDDC, complete the following steps:

- a. RDP to the Windows VM hosted on the VMware Cloud on AWS SDDC.
- b. Navigate to Server Manager > Dashboard > Tools > iSCSI Initiator to open the iSCSI Initiator Properties dialog box.
- c. From the Discovery tab, click Discover Portal or Add Portal and then enter the IP address of the iSCSI target port.
- d. From the Targets tab, select the target discovered and then click Log On or Connect.
- e. Select Enable Multipath, and then select "Automatically Restore This Connection When the Computer Starts" or "Add This Connection to the List of Favorite Targets". Click Advanced.



The Windows host must have an iSCSI connection to each node in the cluster. The native DSM selects the best paths to use.

and a second i and the farmer of the second se		Connigoration		
Quick Connect	-	platu	5	
To discover and log on to a target using a basic connect DNS name of the target and then dick Quick Connect.	tion, type Qui	ck Connect		>
Target: 10.222.2.221	Ta pr	argets that are available for ovided are listed below. If r	connection at the IP address or DNS name that multiple targets are available, you need to conne	you ect
Discovered targets	to	each target individually.		
	0	onnections made here will be	e added to the list of Favorite Targets and an at	tempt
	to	restore them will be made e	every time this computer restarts.	
Name	Sta			
ign. 1992-08.com.netapp:sn. 264efe832dd911eca951c	ISP Cort D	iscovered targets		
		Name	Status	
		ign. 1992-08. com. netapp:sn	.f0c909af2dc611ecac4f Connected	
To connect using advanced options, select a target and	ithen P	ign. 1992-08. com. netapp:sn	.f0c909af2dc611ecac4f Connected	
To connect using advanced options, select a target and click Connect.	i then P	ign. 1992-08. com. netapp:sn rogress report	.f0c909af2dc611ecac4f Connected	
To connect using advanced options, select a target and dick Connect. To completely disconnect a target, select the target an then dick Disconnect.	i then P	ign. 1992-08. com. netapp:sn rogress report Login Succeeded.	.f0c909af2dc611ecac4f Connected	
To connect using advanced options, select a target and dick Connect. To completely disconnect a target, select the target an then dick Disconnect. For target properties, including configuration of session select the target and dick Properties.	i then P d s,	ign. 1992-08. com. netapp:sn rogress report Login Succeeded.	.f0c909af2dc611ecac4f Connected	

LUNs on the storage virtual machine (SVM) appear as disks to the Windows host. Any new disks that are added are not automatically discovered by the host. Trigger a manual rescan to discover the disks by completing the following steps:

- 1. Open the Windows Computer Management utility: Start > Administrative Tools > Computer Management.
- 2. Expand the Storage node in the navigation tree.
- 3. Click Disk Management.
- 4. Click Action > Rescan Disks.

	phone and a second s				(A		-				
Comparts Management (Lack) C. Stark Schadule C. Tak Schadule Device Management Device	Volume == 700 == Rien Volume (E == System Reserves = System Reserves	ut pve (b)	Lagent 1 Sergie Re Sergie Re Sergie Re Sergie Re	er FARSystem is NTFS is NTFS ic UDP is NTFS	n   Boha Heathy (Boot, Pege Tila, Casto Dumy, Prinney Parkbon Heathy (Prinney Parkbon) Heathy (Navel Parkbon) Heathy (System, Active, Prinney Parkbon)	Capacity 89.45 GB 5.95 GB 4.92 GB 540 MB	7129 GB 7129 GB 9.95 GB 9.96 9.96 9.96 9.98 9.98 9.98 9.98 9.98	55 Free 2015 100 % 2015 2115			
	Dea 1 Basic 8.95 GB Online	New Volume 9.00 (d) NTF5 Healthy (Frie	dia wy Partita								
	en Disk Basic 405.55 CB	499.9tt GB									2 MB Unafie
	Online										

When a new LUN is first accessed by the Windows host, it has no partition or file system. Initialize the LUN and, optionally, format the LUN with a file system by completing the following steps:

- 1. Start Windows Disk Management.
- 2. Right-click the LUN, and then select the required disk or partition type.
- 3. Follow the instructions in the wizard. In this example, drive F: is mounted.

$\leftarrow \rightarrow c$		08=	https://vcenter	sddc-52-37-127-104.vm	warevmc.com/ui/webconsole.	html/vmld=vm-1005i	IsymName+vmcdc0	80% 12		0	* =
🖢 Getting Started 🛛 🔒	EC2 Manage	ment Con 🤘	New Tab							Other B	lookmark
mede01							Enforce US K	eyboard Layout View	Fullscreen	Send Cel+	Až + Delet
Complie Hangered										1.1	0 . R.
file Artest Van Paly de els 1 de 1991 El 1991 no 1	200										
Computer Management Local	Tabana	[iarst]	Type   File Symmet   Salar		Coperity   Free Same   S. Free				*	free	_
Co Task Scheduler	- (C.) - New Universe (I.)	Single	9 2 1 + 1 The PC				- 0	×	0	d Management	
Shared follow	ATTLANTIC IN	UE, PVR (D) Simple Simple	4 W 1	ne PC			We heater				
d Devis Manape			v . Cashanna	- Folders (7)							
· If Window Investigation			Dotte /	2 Diparts	Dalita	Deuterte					
<ul> <li>By Service and Apphiations</li> </ul>			Douments /	Countrast.	Mair	The follow					
			al Form of		<b>S</b>	-					
			· · · ·	Tanan .							
			> @ Neuet	~ Devices and drives (4)							
	Rest. 104108	New Yolune (E) 3.08 GB N7PG		Load Date (C)	END Drive (D) SEL, YEMPING, DH 1/4, DVP	The lifete E					
		Hatty Poney faith		100000440(2.8)							
	er (bek Secie			· And all the of \$10.00							
	Coles	10.10 18		and done with think high	R(N)	a		0.00	-		
	-042	-		Vision dation	× .				-		
	AM IR Color	4.16 DE N791									
			Union Union administra					11.00	-		

# Cloud Volumes ONTAP (CVO)

Cloud volumes ONTAP, or CVO, is the industry-leading cloud data management solution built on NetApp's ONTAP storage software, available natively on Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP).

It is a software-defined version of ONTAP that consumes cloud-native storage, allowing you to have the same storage software in the cloud and on-premises, reducing the need to retrain you IT staff in all-new methods to

manage your data.

CVO gives customers the ability to seamlessly move data from the edge, to the data center, to the cloud and back, bringing your hybrid cloud together — all managed with a single-pane management console, NetApp Cloud Manager.

By design, CVO delivers extreme performance and advanced data management capabilities to satisfy even your most demanding applications in the cloud

# Cloud Volumes ONTAP (CVO) as guest connected storage

Cloud Volumes ONTAP shares and LUNs can be mounted from VMs that are created in the VMware Cloud on AWS SDDC environment. The volumes can also be mounted on native AWS VM Linux Windows clients, and LUNS can be accessed on Linux or Windows clients as block devices when mounted over iSCSI because Cloud Volumes ONTAP supports iSCSI, SMB, and NFS protocols. Cloud Volumes ONTAP volumes can be set up in a few simple steps.

To replicate volumes from an on-premises environment to the cloud for disaster recovery or migration purposes, establish network connectivity to AWS, either using a site-to-site VPN or DirectConnect. Replicating data from on-premises to Cloud Volumes ONTAP is outside the scope of this document. To replicate data between on-premises and Cloud Volumes ONTAP systems, see Setting up data replication between systems.



Use the Cloud Volumes ONTAP sizer to accurately size the Cloud Volumes ONTAP instances. Also, monitor on-premises performance to use as inputs in the Cloud Volumes ONTAP sizer.

1. Log into NetApp Cloud Central; the Fabric View screen is displayed. Locate the Cloud Volumes ONTAP tab and select Go to Cloud Manager. After you are logged in, the Canvas screen is displayed.



1. On the Cloud Manager home page, click Add a Working Environment and then select AWS as the cloud and the type of the system configuration.

 Cloud Ma	nager				Account Netapo PC		Workspace ~	Connector ~	٩	6)	0	8
Canwas	Replication	Backup & Restore	KBs Data Sense	File Cache	Compute	Sync	Ali Services (+8) 🗸					
Add Work	ing Environme	nt										×
		Micros	aft Azure Amazon W	(ish Services Choose Type)	coople Cloud Pletform		On-Premises					
		Cloud	CO Volumes ONTAP Engle Node	Cloud Volumes ON	NTAP HA	Amazon I	Sx for ONTAP					
												C

1. Provide the details of the environment to be created including the environment name and admin credentials. Click Continue.

Create a New Working Enviro	onment	Detai	ls and Credentials	
Previous Step	Instance Profile Credential Name	139763910815 Account ID	netapp.com-cloud-volumes Marketplace Subscription	Edit Credentials
	Details		Credentials	
	Working Environment N	lame (Cluster Name)	User Name	
	fsxcvotesting01		admin	
			Password	
	🕣 Add Tags	Optional Field Up to four tags	*******	
			Confirm Password	
			•••••	
<ol> <li>Select the add-o BlueXP backup</li> <li>Create a New Working Environment</li> </ol>	on services for Clo and recovery, and dronment	ud Volumes ONTA Cloud Insights. Cl Services	P deployment, including ick Continue.	BlueXP Classificatio
( Dat	a Sense & Compliance			-• -
Bac	kup to Cloud			-

1. On the HA Deployment Models page, choose the Multiple Availability Zones configuration.



Continue

1. On the Region & VPC page, enter the network information and then click Continue.

(III)

Monitoring

Previous Step	AWS Region			VPC		Security group	
	US West	Oregon	•	vpc-0d1c764bcc495e805 - 10.222.0.0/16	•	Use a generated security group	
	A Nor	do 1+		Noda 2:		Madiator	
	ECCERCI INCO	ae de		Node 2.		invediator:	
	Availability	y Zone		Availability Zone		Availability Zone	
	us-west-	-2a		us-west-2b		us-west-2c	÷
	Subnet			Subnet		Subnet	_
	10.222.1	.0/24	•	10.222.2.0/24	•	10.222.3.0/24	*
are a New Workin	g Environment			Connectivity 8, cc	HAuthon	tication	
Previous Step	g Environment			Connectivity & SS	H Authen	itication	
ate a New Workin Previous Step	g Environment	Nodes		Connectivity & SS	SH Authen	Itication	
Previous Step	g Environment	Nodes SSH Authentication M	Method	Connectivity & SS	SH Authen	Mediator	
Previous Step	g Environment	Nodes SSH Authentication M Password	Method	Connectivity & SS	SH Authen	Itication Mediator oup herated security group	
Previous Step	g Environment	Nodes SSH Authentication M Password	Method	Connectivity & SS	SH Authen Security Gr Use a ger Key Pair Na	Mediator oup herated security group	
Previous Step	g Environment	Nodes SSH Authentication M Password	Method	Connectivity & SS	Security Gr Use a ger Key Pair Na nimokey	Mediator oup herated security group	
Previous Step	g Environment	Nodes SSH Authentication M Password	Method	Connectivity & SS	Security Gr Use a ger Key Pair Na nimokey Internet Co Bublic IP	Mediator oup nerated security group me	
Previous Step	g Environment	Nodes SSH Authentication M Password	Method	Connectivity & SS	Security Gr Use a ger Key Pair Na nimokey Internet Co Public IP	Mediator oup herated security group ime innection Method address	
Previous Step	g Environment	Nodes SSH Authentication M Password	Method	Connectivity & SS	Security Gr Use a ger Key Pair Na nimokey Internet Co Public IP	Mediator oup herated security group hme innection Method address	-
Previous Step	g Environment	Nodes SSH Authentication M Password	Method	Connectivity & SS	SH Authen Security Gr Use a ger Key Pair Na nimokey Internet Co Public IP	Mediator oup herated security group ime innection Method address	-
Previous Step	g Environment	Nodes SSH Authentication M Password	Method	Connectivity & SS	SH Authen Security Gr Use a ger Key Pair Na Nimokey Internet Co Public IP	Itication Mediator oup herated security group ime innection Method address	
Previous Step Specify the	gEnvironment	Nodes SSH Authentication M Password	Method	Connectivity & SS	SH Authen Security Gr Use a ger Key Pair Na Nimokey Internet Co Public IP	Itication Mediator oup herated security group ime innection Method address	
Previous Step Specify the	gEnvironment	Nodes SSH Authentication M Password	Method	Connectivity & SS	Security Gr Use a ger Key Pair Na Nimokey Internet Co Public IP	Itication Mediator oup herated security group ime innection Method address	
Previous Step Specify the	gEnvironment	Nodes SSH Authentication M Password	Method	Connectivity & SS	SH Authen Security Gr Use a ger Key Pair Na nimokey Internet Co Public IP	Itication Mediator oup herated security group ime innection Method address	

Previous Step	nvironment		Floating IPs		
01156792352562579	Floating IP addresses a HA node	re required for cluster and SVM a s if failures occur. To access the d	access and for NFS and CIFS da lata from outside the VPC, you	ata access. These floating IPs a can set up an AWS transit g	can migrate betwee ateway.
	You mu	st specify IP addresses that are o	utside of the CIDR blocks for a	II VPCs in the selected AWS	region.
		Floating IP address for	cluster management		
		172.16.0.1			
		Floating IP address 1 fo	or NFS and CIFS data		
		172.16.0.2			
		Floating IP address 2 fo	or NFS and CIFS data		
		172.16.0.3			
		Floating IP address for	SVM management (Optional)		
		172.16.0.4			
			(HIMANAL)		
Continue.					
reate a New Working En	vironment	F	Route Tables		
	Name	Ado	Sitional information ()	Associate with Subpet	Tags
	EZ	Yes	rtb-00b2d30c3f68fdbdd	0 Subnets	1 Tags
	1 Route Tables   The main	n route table is the default for the	VPC		
			Continue		
1 On the Data B	-nervotion page, ch				
1. On the Data E	Encryption page, ch	oose AWS-managed	Continue encryption.		
1. On the Data E	Encryption page, ch	oose AWS-managed	Continue encryption.		
1. On the Data E	Encryption page, ch	oose AWS-managed	Continue		
1. On the Data E	Encryption page, ch	oose AWS-managed	Continue		
1. On the Data E	Encryption page, ch	oose AWS-managed	Continue encryption.		
1. On the Data E	Encryption page, ch	oose AWS-managed	Continue		
1. On the Data E	Encryption page, ch	oose AWS-managed	Continue encryption.		
1. On the Data E	Encryption page, ch	oose AWS-managed	Continue		

reate a new working crivin	onment	Data Encryptio	
Previous Step	AWS Ma	naged Encryption	
	AWS is responsi is handled by A	ible for data encryption and decryptior WS key management services.	n operations. Key management
	Default Master	Key: aws/ebs	ar Change Key
<ol> <li>Select the license Pay-As-You-Go o</li> </ol>	e option: Pay-As-You-Go or BY option is used.	Continue OL for using an existing lie	cense. In this example, the
reate a New Working Envi	ronmenic loud volumes on tai	P Charging Methous & r	NSS ACCOUNT
Cloud Volumes ONTAP C	Charging Methods	NetApp Support Site	Account (Optional)
Cloud Volumes ONTAP C Learn more about our charg	Charging Methods ging methods	NetApp Support Site Learn more about NetAp	Account <i>(Optional)</i> op Support Site (NSS) accounts
Cloud Volumes ONTAP C Learn more about our charg	Charging Methods gin <mark>g methods</mark> o by the hour	NetApp Support Site A Learn more about NetAp To register this Cloud Vo should add NetApp Supp	Account <i>(Optional)</i> op Support Site (NSS) accounts lumes ONTAP to support,you port Site Account.
Cloud Volumes ONTAP C Learn more about our charg	Charging Methods ging methods o by the hour m license	NetApp Support Site A Learn more about NetAp To register this Cloud Vo should add NetApp Supp Don't have a NetApp Sup finish deploying this syst Support Registration ont	Account <i>(Optional)</i> op Support Site (NSS) accounts lumes ONTAP to support,you port Site Account. oport Site account?Select go to tem.After its created,use the tion to create an NSS account
Cloud Volumes ONTAP C Learn more about our charg Pay-As-You-Ge TE Bring your ow 1. Select between s deployed on the V	Charging Methods ging methods to by the hour In license	NetApp Support Site / Learn more about NetAp To register this Cloud Vo should add NetApp Sup Don't have a NetApp Sup finish deploying this syst Support Registration ont ontinue s available based on the ty oud on AWS SDDC.	Account <i>(Optional)</i> op Support Site (NSS) accounts Jumes ONTAP to support, you port Site Account. opport Site account?Select go to tem.After its created, use the tion to create an NSS account.
Cloud Volumes ONTAP C Learn more about our charg Pay-As-You-Ge Cloud Volumes ONTAP C Pay-As-You-Ge Bring your ow 1. Select between s deployed on the V Create a New Working Environ	Charging Methods ging methods to by the hour In license everal preconfigured packages /Ms running on the VMware cla mment Preconfigured	NetApp Support Site A Learn more about NetAp To register this Cloud Vo should add NetApp Supp Don't have a NetApp Sup finish deploying this syst Support Registration ont ontinue a available based on the ty oud on AWS SDDC.	Account <i>(Optional)</i> op Support Site (NSS) accounts Jumes ONTAP to support,you port Site Account. opport Site account?Select go to tem.After its created,use the tion to create an NSS account
Cloud Volumes ONTAP C Learn more about our charg O Pay-As-You-Ge O Bring your ow 1. Select between s deployed on the V Create a New Working Enviror Select a pr	Charging Methods ging methods o by the hour m license everal preconfigured packages /Ms running on the VMware cle mment Preconfigured econfigured Cloud Volumes ONTAP system that best Preconfigured settings can be r	NetApp Support Site A Learn more about NetAp To register this Cloud Vo should add NetApp Supp Don't have a NetApp Sup finish deploying this syst Support Registration ont ontinue a available based on the ty oud on AWS SDDC. Packages	Account (Optional) op Support Site (NSS) accounts lumes ONTAP to support, you port Site Account. opport Site account?Select go to tem.After its created, use the tion to create an NSS account uppe of workload to be
Cloud Volumes ONTAP C Learn more about our charg O Pay-As-You-G O Bring your ow 1. Select between s deployed on the V Create a New Working Environ Select a pro-	Charging Methods ging methods o by the hour milcense everal preconfigured packages VMs running on the VMware clo mment Preconfigured econfigured Cloud Volumes ONTAP system that best Preconfigured settings can be r	NetApp Support Site A Learn more about NetAp To register this Cloud Vo should add NetApp Sup Don't have a NetApp Sup finish deploying this syst Support Registration ont ontinue available based on the ty oud on AWS SDDC. Packages	Account (Optional) op Support Site (NSS) accounts Jumes ONTAP to support,you port Site Account. oport Site account?Select go to tem.After its created,use the tion to create an NSS account ype of workload to be infiguration. Change Configuration

1. On the Review & Approve page, review and confirm the selections. To create the Cloud Volumes

eate a New Workin	ng Environm	ent		Review & Approve			
Previous Step	sting						Show API reques
ANT	us-west-2	на					
Concession of the local division of the loca							
This Cloud V	Volumes ONT	AP instance will be reg	istered with NetApp supp	ort under the NSS Account mchad.			
This Cloud V	Volumes ONT	AP instance will be reg oud Manazer will alloc	istered with NetApp supp ate the appropriate AWS	ort under the NSS Account mchad.	requirements. Mor	e information >	
This Cloud V	Volumes ONT/	AP instance will be reg oud Manager will alloc	istered with NetApp supp ate the appropriate AWS	ort under the NSS Account mchad. esources to comply with my above	requirements. Mor	e information >	
This Cloud V	Volumes ONT/ initiand that Clo rvlew	AP instance will be reg oud Manager will alloc Networking	istered with NetApp supp ate the appropriate AWS Storage	ort under the NSS Account mchad. esources to comply with my above	requirements. Mo	e information >	
This Cloud V	Volumes ONT/ rstand that Clo view tem:	AP instance will be reg oud Manager will alloc Networking Cloud Volume	istered with NetApp supp ate the appropriate AWS Storage Is ONTAP HA	ort under the NSS Account mchad. esources to comply with my above HA Deplo	requirements. Mor	e information > Multiple Availability Zones	
This Cloud 1 Tunder Over Storage Syst License Type	Volumes ONT/ Instand that Clo Inview tem: e:	AP instance will be reg oud Manager will alloc Networking Cloud Volume Cloud Volume	Istered with NetApp supp ate the appropriate AWS Storage Is ONTAP HA	ort under the NSS Account mchad. esources to comply with my above HA Deplo Encryptio	requirements. Mor yment Model: n:	e information > Multiple Availability Zones AWS Managed	

1. After Cloud Volumes ONTAP is provisioned, it is listed in the working environments on the Canvas page.



1. After the working environment is ready, make sure the CIFS server is configured with the appropriate DNS and Active Directory configuration parameters. This step is required before you can create the SMB volume.

Volumes HA Status Cost Replications			0 0	C	٩	4-
Create a CIFS server		+ Advanced				
DNS Primary IP Address	Active Directory Doma	ain to Join				
192.168.1.3	fsktesting3ocal					
DNS Secondary IP Address (Optional)	Credentials authorized	d to join the domain				
Example: 127.0.0.1	Username	Password				

1. Select the CVO instance to create the volume and click the Create Volume option. Choose the appropriate size and cloud manager chooses the containing aggregate or use advanced allocation mechanism to place on a specific aggregate. For this demo, SMB is selected as the protocol.

Details & Protection		Protocol		
Volume Name:	Size (G8): 🛛 🛞	NFS	CIFS	iSCSI
smbdemovol01	100			
		Share name:	Permissions	52 <sup>°</sup>
Snapshot Policy:		smbdemovol01_sha	are Full Contr	• lo
default	•			
Default Policy		Users / Groups:		
		Everyone;		
		Valid users and groups i	separated by a semicolon	

1. After the volume is provisioned, it is availabe under the Volumes pane. Because a CIFS share is provisioned, you should give your users or groups permission to the files and folders and verify that those users can access the share and create a file.

INFO		CAPACITY	
Disk Type	GP2		1.67 MB
Tiering Policy	None	10 GB	EBS Used
Backup	OFF	Allocated	

- 1. After the volume is created, use the mount command to connect to the share from the VM running on the VMware Cloud in AWS SDDC hosts.
- 2. Copy the following path and use the Map Network Drive option to mount the volume on the VM running on the VMware Cloud in AWS SDDC.

(HA) fsxcvotesting01 (Multiple AZs)	AWS 🛛 🗄 AWS
Volumes HA Status Cost Replications	ى ك ( <del>0</del> )
Mount Volume smbdemovol01	
Access from inside the VPC using Floating IP	Access from outside the VPC using AWS Private IP
Auto failover between nodes	No auto-failover between nodes
The IP address automatically migrates between nodes if failures occur	The IP address does not migrate between nodes if failures occur
Go to your machine and enter this command	To avoid traffic between nodes, mount the volume by using the primary node's IP address:
\\172.16.0.2\smbdemovol01_share	\\10.222.1.100\smbdemovo101_share
	If the primary node opes offline, mount the volume by using the HA partner's IP address:



#### Connect the LUN to a host

To connect the Cloud Volumes ONTAP LUN to a host, complete the following steps:

1. On the Cloud Manager Canvas page, double-click the Cloud Volumes ONTAP working environment to create and manage volumes.

-----

2. Click Add Volume > New Volume, select iSCSI, and click Create Initiator Group. Click Continue.

	Details & Protection				Protocol		
	Volume Name:	Size	e (GB):		NFS	CIFS	iscsi
	nimofsxiscsicvo01	51	00				What about LUNs? ())
	Snapshot Policy:				Initiator Group 🛞		
	default		3 <b>-</b> 3		<ul> <li>Map Existing Ini</li> </ul>	tiator Groups 🔘 (	Create Initiator Group
	Default Policy				Operating System Ty	уре	
					Windows		•
					Select Initiator Grou	ps:	1 (of 3) Groups
					winiG   wi	ndows 5.com.microsoft-ymc	dc01 fsytestin
VMware Cloud - ntap-for-demo × → C Getting Started GEC2 Managem	vsphere - vmcdc01 - Summary ×       A ## https://vcentersidic-52-3       ent Con       New Tab	vmcdc01 37-127-104 vmwarevm	x .com/ui/webco	NetApp Cla	ad Manager X	+ e=vmcdc01 @0% 습	C - 0 C dther Ba
VMware Cloud - ntap-for-demo × Getting Started CC Managem cdco1 Set ver Manager •	Vsphere - vrnodc01 - Summary ×	vmcdc01 37-127-104 vmwarevm sent to be kep we	x	NetApp Cla	ad Manager X mild = vm=10055cvmNam	+ e=vmcdc01 80% 🟠	CO - C C d Other Bit View Fullicreen Send Other View Fullicreen Send Other Market View
VMware Cloud - ritap-for-demo X -  -  -  C  Getting Started  EC2 Managem CoCo1  Too Innge Server Manager	Vsphere - vrocdc01 - Summary ×	vmcdc01 37-127-104.vmwarevm settles.settles. Net settles.settles. Net settles.settle.pts	x Ccom/ul/webco	NetApp Cla nsole.html?v	ad Manager X	+ e=vmcdc01 80% 🛱 Enforce US Keyboard Layou 	Color Ba Color Ba View Fullaceen Send Oth A
VMware Cloud - ntap-fix-demo ×  Getting Started Getting Starte	Vsphere - vrnodc01 - Summary ×	vmcdc01 37-127-104.vmwarevm sont to chales we sont to chales we sont to chales we sont to chale to to to the sont to chale to to the sont	x Coom/Jul/Webbook	NetApp Cla nsole.html?v	ad Manager X nild i vm-1005.BivmNam	+ e=vmcdc01 80% 🟠 Enforce US Keyboard Layou	CO - C
VMware Cloud - ritap-for-demo × Getting Started C EC2 Managem rocol Setver Manager • Setver Manager • Setver Manager •	Vsphere - vrnodc01 - Summary ×	vmcdc01 37-127-104.vmwarevm arrest to calculate ware and the second to t	x Coom/Jul/Webcor c.com/Jul/Webcor Mazj 455 Mazj 455 Marine 27.211.44 10.111.44 1	NetApp Cla nsole.html?v	ad Manager X nild i vm-1005.BvmNam	+ e=vmcdc01 80% 🟠 Enforce US Keybuard Layou	CO - C
VMware Cloud - ntap-for-demo × → C Getting Started © EC2 Managem addot Server Manager • Veccoust Land Server Ad Server Ad Server Ad Server File and Strage Server File and Strage Server Weccoust Court server Manager • Weccoust Wecc	Vsphere - vrnodc01 - Summary ×   Vsphere - vrnodc01 - Summary ×   Vsphere - vrnodc01 - Summary ×  New Tab.  Dashbooard  Cont Cont  New Tab.  Dashbooard  Cont	vmcdc01 37-127-104.vmwarevm set international se	x Com/di/webco c.com/di/webco nazy 45 2155244 2155244 2155244 2155244 2155244 2155244 2155244 2155244 2155244 2155244 2155244 2155244 215525 215525 215525 215525 215525 215525 215525 215525 215525 215525 215525 215525 215525 215525 215525 215525 215525 215525 215525 2155555 2155555 2155555 2155555 2155555 2155555 2155555 2155555 2155555 2155555 21555555 21555555 2155555555	NetApp Cla nsole html?v	ad Manager X	+ e=vmcdc01 80% 🟠 Enforce US Keyboard Layou - C × + -	C - C C - C
VMware Cloud - map-fax-demo X	Vsphere - vrnodc01 - Summary ×	vmcdc01 37-127-104.vmwarevm area versionality www.example.com/station/ www.example.com/ www.example.com/station/ www.example.com/ www.e	KA23 429	NetApp Cla nsole html?v	ad Manager X	+ e=vmcdc01 80% 🟠 Enforce US Keybiard Layou	CO - C
VMware Cloud - ntap-for-demo ×  Getting Started Getting Starte	Vsphere - vrnodc01 - Summary ×	vmcdc01 37-127-104.vmwarevm some societies were soc	x Coom/Jul/Webco	NetApp Cla nsole html?v	ad Manager X nild i vm-1005.BvmNam	+ e=vmcdc01 80% 🟠 Enforce US Keybuard Layer	CO - C

1. After the volume is provisioned, select the volume, and then click Target IQN. To copy the iSCSI Qualified Name (IQN), click Copy. Set up an iSCSI connection from the host to the LUN.

To accomplish the same for the host residing on the VMware Cloud on AWS SDDC, complete the following steps:

a. RDP to the VM hosted on VMware cloud on AWS.

- b. Open the iSCSI Initiator Properties dialog box: Server Manager > Dashboard > Tools > iSCSI Initiator.
- c. From the Discovery tab, click Discover Portal or Add Portal and then enter the IP address of the iSCSI target port.
- d. From the Targets tab, select the target discovered and then click Log On or Connect.
- e. Select Enable Multipath, and then select Automatically Restore This Connection When the Computer Starts or Add This Connection to the List of Favorite Targets. Click Advanced.



The Windows host must have an iSCSI connection to each node in the cluster. The native DSM selects the best paths to use.

gets Dis	covery Favorite Targets Volume	s and Devices RADIUS Configuratio
uid: Conn	ect	
o discove NS name	and log on to a target using a basi of the target and then click Quick C	c connection, type the IP address or onnect.
arget:	172.24.2.9	Quick Connect.
iscovered	targets	
		Refresh
Name		Status
		1.000000
		1.00000
		199000
o connect	using advanced options, select a tu	arget and then Connect
To connect lick Conne To complet hen click C	using advanced options, select a ta ct. ely disconnect a target, select the t	arget and then Connect
is connect lick Conne is complet hen click C for target elect the	using advanced options, select a to ct. ely disconnect a target, select the t isconnect. properties, including configuration o arget and click Properties.	arget and then Connect arget and Deconnect of sessions, Properties

LUNs from the SVM appear as disks to the Windows host. Any new disks that are added are not automatically discovered by the host. Trigger a manual rescan to discover the disks by completing the following steps:

- 1. Open the Windows Computer Management utility: Start > Administrative Tools > Computer Management.
- 2. Expand the Storage node in the navigation tree.
- 3. Click Disk Management.
- 4. Click Action > Rescan Disks.

All down in the second second	Pictore and and	11.	1.	Con a	12		1000		Text
Compute Management Local  Compute Management Local  Compute Management Local  Compute Management  Compute	Values (10) 100 Values (1) 2055 States (1) 2055 State	Spread	A Type In Saulo In Saulo In Saulo In Saulo	File System NTPS NTPS NTPS NTPS	Statu Huathy Book, Fage File, Creab Dump, Huathy Primary Particul Huathy Original Particul Huathy Original Particul Huathy Original Action, Primary Bast	, Powery Partitions) Asses	1 Capacity 45 44 58 9 59 58 549 MB 549 MB	Free Special 71:32-68 8:45-68 8:45-68 9:MB 7115-648	ver (% Free 60 % 300 % 21 %
	Duk 0 Basic Store GB Online	pten Teserved IO ME NTPS Withy Cystem, J	ctive, Prin	nery Partition		IES 81-45 GB NTPS Healthy (Root, Pa	pr File, Cra	h Durry, Pos	Pennary Partices)
	- Diak 0 Basic Stoco Cit Online - Diak 1 Basic Stol Cit Online	pateon Reserved 40 MD NTTS eastby Cystem, A 100 Wolkene (B-) 90 QB NTPS authby (Promary 3	ctive, Print	nary Partition		453 35.46 GB MTES Healthy (Soot, Pa	pt File, Cre	h Durry, Por	Penney Faction)

When a new LUN is first accessed by the Windows host, it has no partition or file system. Initialize the LUN; and optionally, format the LUN with a file system by completing the following steps:

- 1. Start Windows Disk Management.
- 2. Right-click the LUN, and then select the required disk or partition type.
- 3. Follow the instructions in the wizard. In this example, drive F: is mounted.

⊢ → C	9	0 A #	https://worders	dd:-52-37-127-104 wm	warevers com/silvebconsole h	tenDumld.com-1005Ros	Name surrection 1876	4	69 🛃	
Getting Started	EC2 Manageme	et Con.	New lab						C) Other Bo	ookmark
modc01							Enforce US Keyboard	Lapout View Fullice	inn Sint Chil+A	R+Dekts
The Arten New Phile									- 8	1.10
** * * * *	0				40					
🔮 Computer Meragement Colum	m Ki	Laward Tar	Pel File Soliters   Status		[ Capacity [ Man Space ] 15 Man		- 0 ×		Did Meapment	-
E Burnet Tolean	at sendulute 21 210,00000,0000,00000	International Strength	the Lingdon sin	-			- 0		Mare Autors	•
A Incomerce	av lytter-baared	Serger #	+ + + + 9 - 54	PC		- 1	S Sand Park P			
- Di Mindens Saroar Bachur			v # Gottains	Process (2)	During	The Decements				
) (), Service and Applications			4 Institute of	0		13				
			E Folgers P	A Provinsile	Matt	Prive Prive				
			- Miller	Con Trans						
			· · · · Dit Dive (D) 100,10	- Devices and drives (4)						
	How New York	Values By		Land Date (C)	010 21-4 (D) 502,544905,214 (J, 217	New Yolune (E)				
	Online Alles	the Primary Partic		NUMBER	Table Allin	1313 Int. / 1113				
	-Data	_		<ul> <li>Interview</li> </ul>						
	animite etc.	14.18		<ul> <li>Network toohors (2) antideture(11) 198,1925</li> </ul>	48.00 embersolit, Are 31723433			2148 12400-00		
	-047			×	*					
	Arrie 4.0	CE MARK								
		in and in	Uners 10an-selected				1.00			

On the Linux clients, ensure the iSCSI daemon is running. After the LUNs are provisioned, refer to the detailed guidance on iSCSI configuration for your Linux distribution. For example, Ubuntu iSCSI configuration can be found here. To verify, run Isblk cmd from the shell.

To mount the Cloud Volumes ONTAP (DIY) file system from VMs within VMC on AWS SDDC, complete the following steps:

- 1. Connect to the designated Linux instance.
- 2. Open a terminal on the instance using secure shell (SSH) and log in with the appropriate credentials.
- 3. Make a directory for the volume's mount point with the following command.

\$ sudo mkdir /fsxcvotesting01/nfsdemovol01

4. Mount the Amazon FSx for NetApp ONTAP NFS volume to the directory that is created in the previous step.

```
sudo mount -t nfs nfsvers=4.1,172.16.0.2:/nfsdemovol01
/fsxcvotesting01/nfsdemovol01
root@ubuntu01:/fsx# mount -t nfs 172.16.0.2:/nfsdemovol01 /fsxcvotesting01/nfsdemovol01.
vbuntu01 - Summay x ubuntu01 x + 
 C O A =t https://voenter.sddc-52-37-127-104.vmwarevmc.com/ul/webconsole.html?vmlde.vm-1003&cm?Name=ubuntu01&cser
 C A =t https://voenter.sddc-52-37-127-104.vmwarevmc.com/ul/webconsole.html?vmlde.vm+1003&cm?Name=ubuntu01&cser
 C A =t https://voenter.sddc-52-37-127-104.vmwarevmc.com/ul/webconsole.html?vmlde.vm+1003&cm?Name=ubuntu01&cser
 C A =t https://sountsol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdemoviol/sdem
```

**Overview of ANF Datastore Solutions** 

Every successful organization is on a path of transformation and modernization. As part of this process, companies typically use their existing VMware investments while leveraging cloud benefits and exploring how to make migration, burst, extend, and disaster recovery processes as seamless as possible. Customers migrating to the cloud must evaluate the issues of elasticity and burst, data center exit, data center consolidation, end- of- life scenarios, mergers, acquisitions, and so on. The approach adopted by each organization can vary based on their respective business priorities. When choosing cloud-based operations, selecting a low- cost model with appropriate performance and minimal hindrance is a critical goal. Along with choosing the right platform, storage and workflow orchestration is particularly important to unleash the power of cloud deployment and elasticity.

# Use Cases

Although the Azure VMware solution delivers unique hybrid capabilities to a customer, limited native storage options have restricted its usefulness for organizations with storage-heavy workloads. Because storage is directly tied to hosts, the only way to scale storage is to add more hosts, which can increase costs by 35-40% or more for storage intensive workloads. These workloads need additional storage, not additional horsepower, but that means paying for additional hosts.

Let's consider the following scenario; a customer requires six hosts for horsepower (vCPU/vMem), but they also have a substantial requirement for storage. Based on their assessment, they require 12 hosts to meet storage requirements. This increases the overall TCO because they must buy all that additional horsepower when all they really need is more storage. This is applicable for any use case, including migration, disaster recovery, bursting, dev/test, and so on.

Another common use case for Azure VMware Solution is disaster recovery (DR). Most organizations do not have a fool- proof DR strategy, or they might struggle to justify running a ghost datacenter just for DR. Administrators might explore zero- footprint DR options with a pilot- light cluster or an on-demand cluster. They could then scale the storage without adding additional hosts, potentially an attractive option.

So, to summarize, the use cases can be classified in two ways:

- · Scaling storage capacity using ANF datastores
- Using ANF datastores as a disaster recovery target for a cost- optimized recovery workflow from onpremises or within Azure regions between the software-defined datacenters (SDDCs). This guide provides insight into using Azure NetApp Files to provide optimized storage for datastores (currently in public preview) along with best-in-class data protection and DR capabilities in an Azure VMware solution, which enables you to offload storage capacity from vSAN storage.



Contact NetApp or Microsoft solution architects in your region for additional information on using ANF datastores.

# VMware Cloud options in Azure

## **Azure VMware Solution**

The Azure VMware Solution (AVS) is a hybrid cloud service that provides fully functioning VMware SDDCs within a Microsoft Azure public cloud. AVS is a first-party solution fully managed and supported by Microsoft and verified by VMware that uses Azure infrastructure. Therefore, customers get VMware ESXi for compute virtualization, vSAN for hyper-converged storage, and NSX for networking and security, all while taking advantage of Microsoft Azure's global presence, class-leading data center facilities, and proximity to the rich ecosystem of native Azure services and solutions. A combination of Azure VMware Solution SDDC and Azure NetApp Files provides the best performance with minimal network latency.

Regardless of the cloud used, when a VMware SDDC is deployed, the initial cluster includes the following components:

- VMware ESXi hosts for compute virtualization with a vCenter server appliance for management.
- VMware vSAN hyper-converged storage incorporating the physical storage assets of each ESXi host.
- VMware NSX for virtual networking and security with an NSX Manager cluster for management.

# Conclusion

Whether you are targeting all-cloud or hybrid cloud, Azure NetApp files provide excellent options to deploy and
manage the application workloads along with file services while reducing the TCO by making the data requirements seamless to the application layer. Whatever the use case, choose Azure VMware Solution along with Azure NetApp Files for rapid realization of cloud benefits, consistent infrastructure, and operations across on-premises and multiple clouds, bi-directional portability of workloads, and enterprise-grade capacity and performance. It is the same familiar process and procedures used to connect the storage. Remember, it is just the position of the data that changed along with new names; the tools and processes all remain the same, and Azure NetApp Files helps in optimizing the overall deployment.

## Takeaways

The key points of this document include:

- You can now use Azure NetApp Files as a datastore on AVS SDDC.
- Boost the application response times and deliver higher availability to provide access workload data when and where it is needed.
- Simplify the overall complexity of the vSAN storage with simple and instant resizing capabilities.
- Guaranteed performance for mission-critical workloads using dynamic reshaping capabilities.
- If Azure VMware Solution Cloud is the destination, Azure NetApp Files is the right storage solution for optimized deployment.

## Where to find additional information

To learn more about the information described in this document, refer to the following website links:

• Azure VMware Solution documentation

https://docs.microsoft.com/en-us/azure/azure-vmware/

Azure NetApp Files documentation

https://docs.microsoft.com/en-us/azure/azure-netapp-files/

• Attach Azure NetApp Files datastores to Azure VMware Solution hosts (Preview)

https://docs.microsoft.com/en-us/azure/azure-vmware/attach-azure-netapp-files-to-azure-vmware-solution-hosts?tabs=azure-portal/

## NetApp Guest Connected Storage Options for Azure

Azure supports guest connected NetApp storage with the native Azure NetApp Files (ANF) service or with Cloud Volumes ONTAP (CVO).

## Azure NetApp Files (ANF)

Azure netApp Files brings enterprise-grade data management and storage to Azure so you can manage your workloads and applications with ease. Migrate your workloads to the cloud and run them without sacrificing performance.

Azure netApp Files removes obstacles, so you can move all of your file-based applications to the cloud. For the first time, you do not

have to re-architect your applications, and you get persistent storage for your applications without complexity.

Because the service is delivered through the Microsoft Azure Portal, users experience a fully managed service as part of their Microsoft enterprise Agreement. World-class support, managed by Microsoft, gives you complete peace of mind. This single solution enables you to quickly and easily add multiprotocol workloads. you can build and deploy both Windows and Linux file-based applications, even for legacy environments.

## Azure NetApp Files (ANF) as guest connected storage

## Configure Azure NetApp Files with Azure VMware Solution (AVS)

Azure NetApp Files shares can be mounted from VMs that are created in the Azure VMware Solution SDDC environment. The volumes can also be mounted on the Linux client and mapped on the Windows client because Azure NetApp Files supports SMB and NFS protocols. Azure NetApp Files volumes can be set up in five simple steps.

Azure NetApp Files and Azure VMware Solution must be in the same Azure region.

To create and mount Azure NetApp Files volumes, complete the following steps:

1. Log in to the Azure Portal and access Azure NetApp Files. Verify access to the Azure NetApp Files service and register the Azure NetApp Files Resource Provider by using the *az provider register --namespace Microsoft.NetApp –wait* command. After registration is complete, create a NetApp account.

For detailed steps, see Azure NetApp Files shares. This page will guide you through the step-by-step process.

	ources, services, and docs (G+/)		E.	Ð	P	?	Å
Home > Azure NetApp Files >							
Azure NetApp Files « HetApp (cloudcontrolproduction.com)	New NetApp account						
+ Create 🔞 Manage view 🗸 …	Name *						
Filter for any field	nimoAVSANFdemo	~					
Name 📬	Subscription						
	SaaS Backup Production	~					
	Resource group *						
	NimoAVSDemo	~					
	Create new						
NetApp accounts to display	Location *						
ire NetApp Files makes it easy to migrate and	East US 2	~					
complex, file-based applications with no code inge. With support for multiple protocols and grated data protection, storage management is simple, fast, and reliable.							
Create NetApp account							
Learn more of	Create Download a template for automation						

2. After the NetApp account is created, set up the capacity pools with the required service level and size.

For more information, see Set up a capacity pool.

Azure NetApp Files	nimoAVSANFde	mo	Capacity	pools –		Name *	1.5.5
+ Create 🔘 Manage view 🗸 …	P Search (Ctrl+/)	- 00	+ Add pool	<ul> <li>Refresh</li> </ul>		nimcappool	4
Filter for any field.	Azure NetApp Files	^	0 south resolu			Service level * 💿	
Name 1	Active Directory connections		Name	t. Canarity	†1. Service lev	Standard	~
in nimoAVSANFdemo ***	Storage service		You don't have	any canacity pools Click	Add pool to get started	Size (Ti8) * 🔘	
	E Capacity pools		Tou Gott Charte	any capacity prove circe	ndo pour to get surred	4	
	🗮 Volumes	Ъ.				Out have	4.1
	Data protection					O Manual	
	Snapshot policies					<ul> <li>Auto</li> </ul>	
	Storage service add-ons						
	RetApp add-ons						
	Automation						
< Page 1 V of 1 >	R Tasks (newiew)					Create Discard	

3. Configure the delegated subnet for Azure NetApp Files and specify this subnet while creating the volumes. For detailed steps to create delegated subnet, see Delegate a subnet to Azure NetApp Files.

		Add subnet	>
nimoavspriv-vn Virtual network	et   Subnets		
	e a	Name *	
Search (Ctrl+/)	« + Subnet -	anf.del	~
Overview	A Search subn	Subnet address range * ①	
Activity log		172.24.3.0/28	<i>~</i>
Access control (IAM)	Name ↑↓	172.24.3.0 - 172.24.3.15	(11 + 5 Azure reserved addresses)
Tags	GatewaySubne	Add IPv6 address space ①	
Diagnose and solve problems	VMSubnet	NAT gateway 💿	
	StorageSubnet	None	~
Settings	<	Network security group	
🤣 Address space		None	~
Ø Connected devices		Route table	
Subnets		None	~
ODoS protection			
🛖 Firewall			
Security		Save Cancel	

4. Add an SMB volume by using the Volumes blade under the Capacity Pools blade. Make sure the Active Directory connector is configured prior to creating the SMB volume.

Azure NetApp Files =	nimoAVSANFdemo NetApp account	Active Directory connections	Primary DNS* ③
+ Create 🙁 Manage view 👳 …	P Search (Ctrl+/) 4	🖉 Join 🕐 Refresh	172.24.1.5
Filter for any field	Activity log	DNS 11 AD DNS Domai 12 SMB Server	Secondary DNS
Name 1.	Access control (IAM)	No currently Johnard Arthus Philasterian	
nimoAVSANEdemo ····	<ul> <li>Taps</li> </ul>	<	AD DNS Domain Name * ③
	• 1031		nimodemo.com
	Settings		AD Site Name ①
	Quota		
	III Properties		SMB Server (Computer Account) Prefix * 🔘
	A Locks		nimsmb
	Azure NetApp Files		Organizational Unit Path 💿
	Active Directory connections		
	Storage service		( The second sec
< Page 1 V of 1 >	Capacity pools		Join

5. Click Review + Create to create the SMB volume.

If the application is SQL Server, then enable the SMB continuous availability.

NetApp account	mo	Volumes		. 6	Create a volume	**	9
,P Search (Ctrl+/)		+ Add volum	ne じ Refresh				
Azure NetApp Files	^			^	Basics Protocol Tags R	teview + create	
Active Directory connections		P Search vol	umes	- 1	This page will help you create an A	zure NetApp Files volume in your subscriptio	on and enable you to access the
Storage service		You don't have	ve any volumes. Click Add	vo	Volume details		
E Capacity pools		¢		>	Volume name *	nimvoltest1	9
🗟 Volumes	1			- 1	Capacity pool * 📀	nimcappool	<i>.</i>
Volumes Data protection	1				Capacity pool * 💿	nimcappool	~]
Volumes Data protection Snapshot policies	1				Capacity pool * 💿 Available quota (Gill) 🛈	nimcappool 4095	4.10
Volumes Data protection Snapshot policies Storage service add-ons					Capacity pool * 💿 Available quota (GiB) 💿 Quota (GiB) * 💿	nimcappool 4235 100	↓ 410 ✓

📃 nimoAVSAN	Fdemo	Volumes	339 339										
,P. Search (Ctrl+/)		+ Add volume	0	Refresh									
Quota	^	. <sup>O</sup> Search volum	65										
Properties		Name	$\uparrow_{\dot{\Phi}}$	Quota	÷ψ	Throughput	14	Protocol type	 Mount path		Service level	14	Capacity p
A Locks		🛒 nimsmbvo	12	100 GiB		1.6 Mi8/s		SMB	\\nimsmb-7c1c	nimode	Standard		nimcappo
· · · · · · · · · · · · · · · · · · ·		nimualtest	1	100 GiB		1.6 MB/s		NES-3	172 24 3.4 /nim	voltest1	Standard		nimcappor

To learn more about Azure NetApp Files volume performance by size or quota, see Performance considerations for Azure NetApp Files.

6. After the connectivity is in place, the volume can be mounted and used for application data.

To accomplish this, from the Azure portal, click the Volumes blade, and then select the volume to mount and access the mount instructions. Copy the path and use the Map Network Drive option to mount the volume on the VM running on Azure VMware Solution SDDC.

vstVM2						Enlo	ce US Keytoard Layout View Fullscree	n Se	nd Chi+	At+C
6. ·	northian :						<b>京</b> (大)			
(	🖯 🛈 - 🛛 Ser	ver Manager 🔸	Dashboard		)   🚩 Hange	linih Vin	2 mai			
i i i i i i i i i i i i i i i i i i i	Detboard Local Server (A) Server File and Storage Se	Current and a constraint of the constraint	viel	Units semantical Data semantical ACCENTY (ND) ANA ACCENTY (ND) ANA ACCENTY (ND) ANA ACCENTY (ND) ANA ACCENTY (ND) ANA	Servit normalismul Type Son The follow Set Deserves The Deserves The Deserves The Deserves					
		e Jama				1				

File Home Sh	nbvol2 are View			- 0	~
+ → - ↑ 🖳 🛚	nimsmb-7c1c.nimodemo.com\nimsmbvol2	~ 2	Search nimsmbv	012	۶
	Name	Date modified	Туре	Size	
A Quick access	nimfoo1	8/13/2021 10:21 AM	File folder		
Desktop	nimfoo2	8/13/2021 10:21 AM	File folder		
Downloads	nimfoo1	8/13/2021 10:21 AM	Text Document	0 K	В
Documents	* 📄 nimfoo2	8/13/2021 10:22 AM	Text Document	0 K	B
E Pictures	*				
This PC					
Network					

7. To mount NFS volumes on Linux VMs running on Azure VMware Solution SDDC, use this same process. Use volume reshaping or dynamic service level capability to meet the workload demands.

in our chan cho dell'en	at machtne	-> UI		n an	
lesystem	1K-DLOCKS	Used	Available	Usex	Mounted on
lev	8168112	6	8168112	6%	/dev
ipfs	1639548	1488	1638060	1%	/run
lev/sdaS	50824704	7982752	40310496	17%	1
ipfs	8197728	0	8197728	6%	/dev/shm
upfs .	5120	0	5120	0%	/run/lock
ipfs	8197728		8197728	8%	/sys/fs/cgroup
ev/loop0	56832	56832	0	100%	/snap/core18/2128
lev/loop2	66688	66688	6	100%	/snap/gtk-common-the
rs/1515					
ev/loop1	224256	224256	9	100%	/snap/gnome-3-34-18
72					
lev/loop3	52224	52224	.6	100%	/snap/snap-store/54
lev/loop4	33152	33152	0	100%	/snap/snapd/12704
lev/sda1	523248	4	523244	1%	/boot/efi
ofs	1639544	52	1639492	1%	/run/user/1000
lev/sr0	54738	54738	0	100%	/media/nimoadmin/VM
elools					

For more information, see Dynamically change the service level of a volume.

## Cloud Volumes ONTAP (CVO)

Cloud volumes ONTAP, or CVO, is the industry-leading cloud data management solution built on NetApp's ONTAP storage software, available natively on Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP).

It is a software-defined version of ONTAP that consumes cloud-native storage, allowing you to have the same storage software in the cloud and on-premises, reducing the need to retrain you IT staff in all-new methods to manage your data.

CVO gives customers the ability to seamlessly move data from the edge, to the data center, to the cloud and back, bringing your hybrid cloud together — all managed with a single-pane management console, NetApp Cloud Manager.

By design, CVO delivers extreme performance and advanced data management capabilities to satisfy even your most demanding applications in the cloud

# Cloud Volumes ONTAP (CVO) as guest connected storage

Cloud Volumes ONTAP shares and LUNs can be mounted from VMs that are created in the Azure VMware Solution SDDC environment. The volumes can also be mounted on the Linux client and on Windows client because Cloud Volumes ONTAP supports iSCSI, SMB, and NFS protocols. Cloud Volumes ONTAP volumes can be set up in a few simple steps.

To replicate volumes from an on-premises environment to the cloud for disaster recovery or migration purposes, establish network connectivity to Azure, either using a site-to-site VPN or ExpressRoute. Replicating data from on-premises to Cloud Volumes ONTAP is outside the scope of this document. To replicate data between on-premises and Cloud Volumes ONTAP systems, see Setting up data replication between systems.



Use Cloud Volumes ONTAP sizer to accurately size the Cloud Volumes ONTAP instances. Also monitor on-premises performance to use as inputs in the Cloud Volumes ONTAP sizer.

1. Log in to NetApp Cloud Central—the Fabric View screen is displayed. Locate the Cloud Volumes ONTAP tab and select Go to Cloud Manager. After you are logged in, the Canvas screen is displayed.



2. On the Cloud Manager home page, click Add a Working Environment and then select Microsoft Azure as the cloud and the type of the system configuration.

				and the second				
Canvas	Replication Back	ip & Restore	as Data Sensi	File Cache	Compute	Sync	All Services (	8) ~
Add New	Working Environment							×
		0						
		a	WS	<b></b>				
	MICTOSIT AZUTE	Armazon vi	eo services Go	ogle Cloud Platform	On-Premises			
	Choose Type							
		· · · · · · · · · · · · · · · · · · ·						
	6		0					
			<u> </u>		-			

3. When creating the first Cloud Volumes ONTAP working environment, Cloud Manager prompts you to deploy a Connector.



4. After the connector is created, update the Details and Credentials fields.

Managed Service Ide	SaaS Backup Prod	CMCVOSub	Edit Credentials
Credential Name	Azure Subscription	Marketplace Subscription	
Details		Credentials	
Working Environment Nam	e (Cluster Name)	User Name	
nimavsCVO		admin	
		Password	
		Continue	

5. Provide the details of the environment to be created including the environment name and admin credentials. Add resource group tags for the Azure environment as an optional parameter. After you are done, click Continue.

Details	Credentials
Working Environment Name (Cluster Name)	User Name
nimavsCVO	admin
	Password
Add Resource Group Tags     Optional Field	•••••
	Confirm Password
	*********

6. Select the add-on services for Cloud Volumes ONTAP deployment, including BlueXP Classification, BlueXP backup and recovery, and Cloud Insights. Select the services and then click Continue.

Data Sense & Compliance	<b>••</b> •
Backup to Cloud	• ~
(iii) Monitoring	<b>-</b>

7. Configure the Azure location and connectivity. Select the Azure Region, resource group, VNet, and subnet to be used.

Azure Region		Resource Group
East US 2		Create a new group     Use an existing group
Availability Zone	(Optional)	Resource Group Name
Select an Availability Zone	•	nimassCVO-rg
vNet nimoavspriv-vnet [ NimoAVSDemo	*	Security Group
Subnet		Generated security group     Ose existing security group
172.24.2.0/24	•	
		I have verified network connectivity between the Cloud Manager server and the selected there.

8. Select the license option: Pay-As-You-Go or BYOL for using existing license. In this example, Pay-As-



- Create a New Working Environment
   Preconfigured Cloud Volumes ONTAP system that best matches your needs, or create your own configuration. Preconfigured settings can be modified at a later time.
   Change Configuration

   Image: Configuration of the production workloads Up to 500GB of storage
   Image: Configuration of the production workloads
   Image: Configuration of the production workloads

   Image: Configuration of the production workloads
   Image: Configuration of the production workloads
   Image: Configuration of the production workloads
- 10. Accept the two agreements regarding activating support and allocation of Azure resources. To create the Cloud Volumes ONTAP instance, click Go.

		er en et allere	never a Approve
imaysCVO			
East US 2			
I understand that	in order to activate supp	oort, I must first register	Cloud Volumes ONTAP with NetApp. More information >
I understand that	Cloud Manager will alloc	ate the appropriate Azu	are resources to comply with my above requirements. More information >
Overview	Networking	Storage	

11. After Cloud Volumes ONTAP is provisioned, it is listed in the working environments on the Canvas page.

Canv	vas						Go to Tabular Vie
Add w	Vorking Environment	Sinds			on nima • On	IVSCVO	
		nimavsCVO Cloud Volumes ONTAP			DETAILS Cloud Volumes	ONTAP   A2	ure Single
					SERVICES	ration	
			G	•		Enter Work	Ing Environment

1. After the working environment is ready, make sure the CIFS server is configured with the appropriate DNS and Active Directory configuration parameters. This step is required before you can create the SMB volume.

imavsCVO nimavsCVO		Azur	e. fi	) Azure	Manage	ed Encry	ption
Volumes Replications		0	Ċ	C	٩	<b>1</b>	Ξ
Create a CIFS server	+ Advanced						
DNS Primary IP Address	Active Directory Domain to join						
172.24.1.5	nimodemo.com						
DNS Secondary IP Address (Optional)	Credentials authorized to join the domain						
	Contraction of the second seco						

 Creating the SMB volume is an easy process. Select the CVO instance to create the volume and click the Create Volume option. Choose the appropriate size and cloud manager chooses the containing aggregate or use advanced allocation mechanism to place on a specific aggregate. For this demo, SMB is selected as the protocol.

Details & Protection			Protocol		
Volume Name:	Size (GB);	0	NFS	CIFS	iSCSI
nimavssmbvol1	50				
			Share name:	Permissions:	
Snapshot Policy:			nimavssmbvol1_share	Full Control	
default		•			
③ Default Policy			Users / Groups:		
			Everyone;		

3. After the volume is provisioned, it will be availabe under the Volumes pane. Because a CIFS share is provisioned, give your users or groups permission to the files and folders and verify that those users can access the share and create a file. This step is not required if the volume is replicated from an on-premises environment because the file and folder permissions are all retained as part of SnapMirror replication.

ume 50 GB Alloca	ted 1.74 MB Total Used	(1.74 MB in Disk, 0 KB in Blo	b)
-			
nimavs	smbvol1		ONLINE
INFO		CAPACITY	
Disk Type	PREMIUM LRS		1 74 MP
Disk Type	PREMIUM_LRS	50 GP	1.74 MB Disk Used

- 4. After the volume is created, use the mount command to connect to the share from the VM running on the Azure VMware Solution SDDC hosts.
- 5. Copy the following path and use the Map Network Drive option to mount the volume on the VM running on Azure VMware Solution SDDC.

Volumes Replications

# Mount Volume nimavssmbvol1

Go to your machine and enter this command

👤   🖸 📒 🖛   nim	lavssm	bvol1_share						- 0	
File Home	Share	View							2
← → → ↑ 💆	1172.	24.2.8\nimivosml	bvol1_shate				 Search nimavssr	nbvol1_share	
		Name	0	Date modified	Туре	Size			
Desktop					This fol	der is empty.			
Downloads									
Documents	*								
Fictures	*								
💻 This PC									

#### Connect the LUN to a host

To connect the LUN to a host, complete the following steps:

- 1. On the Canvas page, double-click the Cloud Volumes ONTAP working environment to create and manage volumes.
- 2. Click Add Volume > New Volume and select iSCSI and click Create Initiator Group. Click Continue.

Details & Protection			Protocol		
Volume Name:	Size (GB):	Ð	NFS	CIFS	iSCSI
nimavsscsi1	500				What about LUNs? 🕕
Snapshot Policy:			Initiator Group 🍈		
default			Map Existing Init	liator Groups	Create Initiator Group
Default Policy			Initiator Group		
			avsvmlG		

3. After the volume is provisioned, select the volume, and then click Target IQN. To copy the iSCSI Qualified Name (IQN), click Copy. Set up an iSCSI connection from the host to the LUN.

To accomplish the same for the host residing on Azure VMware Solution SDDC:

- a. RDP to the VM hosted on Azure VMware Solution SDDC.
- b. Open the iSCSI Initiator Properties dialog box: Server Manager > Dashboard > Tools > iSCSI Initiator.
- c. From the Discovery tab, click Discover Portal or Add Portal and then enter the IP address of the iSCSI target port.
- d. From the Targets tab, select the target discovered and then click Log on or Connect.
- e. Select Enable multipath, and then select Automatically Restore This Connection When the Computer Starts or Add This Connection to the List of Favorite Targets. Click Advanced.

**Note:** The Windows host must have an iSCSI connection to each node in the cluster. The native DSM selects the best paths to use.

Quick C	Connect				
To disc DNS n	over and i arrie of the	og on to a target us target and then clo	ing a basic connection, k Quick Connect.	type the IP	address or
Target	2 17	2.24.2.9		Q	uick Connect
Discove	ered targe	tai			
	1000000				Refresh
Name				Status	
To con click Co	nect using	advanced options,	select a target and the	•	Correct
To con dick Ci To con then d	nect using prinect. ipletely dis lok Disconr	advanced options, connect a target, si lect.	select a target and ther elect the target and		Correct Decorrect
To con dick Ci To con then d For tar select	nect using sinect. spletely dis lok Disconr get proper the target	advanced options, connect a target, s lect. rises, including confl and dick Properties	select a target and ther slect the target and guration of sessions,		Connect Deconnect Properties
To con dick Ci To con then d For tar select For co the tar	nect using prinect, igk Disconr get proper the target infiguration get and th	advanced options, connect a target, si lect. rises, including confi and click Properties of devices associat en click Devices.	select a target and ther elect the target and guration of sessions, ed with a target, select		Connect Deconnect Properties Devices

LUNs on storage virtual machine (SVM) appear as disks to the Windows host. Any new disks that are added are not automatically discovered by the host. Trigger a manual rescan to discover the disks by completing the following steps:

- 1. Open the Windows Computer Management utility: Start > Administrative Tools > Computer Management.
- 2. Expand the Storage node in the navigation tree.
- 3. Click Disk Management.
- 4. Click Action > Rescan Disks.

curbates wasadement if pred	Volume	Layout	Type	File System	Status	Capacity	Free Space	% fras
System Tools	- (C)	Simple	Basic	NTFS	Healthy (Boot, Paga File, Crash Dump, Primary Partition	39.51 GB	24.99 GB	63 %
Task Scheduler	SSS_X64FREE_EN4-U	IS_DVP(D) Simple	BRDE	UDF.	Healthy (Primary Partition)	6.49 GB	0 MB	05
Event Viewer     Sanet Folders     Sanet Folders     Sec Utoers and Groups     Performance     Device Manager     Sanage     Sanage     Sanage     Sanage     Disk Management     Services and Applications	System Reserved	Sample	Basic	NTFS	Healthy System, Active, Primary Partition)	500 MB	169 MB	ыя.
	The Diak 0 Basic S 40.00 G8 St Ordine H	iystem Reserved 00 MB NTFS lealthy (System, Act	ive, Prim	ary Partition)	(C.) 39.51 GB NTFS Healthy (Boot, Page File, I	irash Dump,	Primary Parti	son)

When a new LUN is first accessed by the Windows host, it has no partition or file system. Initialize the LUN; and optionally, format the LUN with a file system by completing the following steps:

1. Start Windows Disk Management.

- 2. Right-click the LUN, and then select the required disk or partition type.
- 3. Follow the instructions in the wizard. In this example, drive E: is mounted

Computer Management (Local)		Design from the second		I a contra to the second	less 1				
System Tools System Tools Tool Scheduler Source Folders Source Folders Cocal Uses and Groups Cocal Uses and Groups Cocal Uses and Groups Cocal Uses and Groups Scorege Scorege Scorege Scorege Scorege Score Itackup Tools Management Score Score Itackup Tools Management	Velume (C2) = DBDial2 (F) = DBDial2 (F) = System Reserved System Reserved	Layout Type File System Simple Basic NTFS Simple Basic NTFS Simple Basic NTFS Simple Basic NTFS Simple Basic NTFS	Status Healthy (Boot, Page File, Crash Dump, Pilmary Partition) Healthy (Primary Partition) Healthy (Primary Partition) Healthy (System, Active, Primary Partition)	Capacity Free Space 1935 108 2355 08 499,87 499,73 68 5.97 68 5.93 68 5.49 68 0.MB 500 MB 169 MB	et % Free   61 % 100 % 100 % 100 % 34 %				
	Holak 1 Banc DBdisk (Ed 499.85 GB 499.87 GB N Online Healthy (Pri	TFS mary Partition()							
	The Disk 2 Basic 9.97 GB 9.97 GB NT Online Healthy (Pri	ij 5 mary Partition)							
1 🖓 📗 🗴 1 This PC								- 1	
le Computer V	Aeve								2
ie Computer N → ~ ↑ 💻 + Ti	New his PC					v ð	Search This PC		ں ر
Computer (	Alexe his PC Folders (6)		Documents	Dow	nloads	v [0]	Search This PC		-
Computer 1 Computer 1 Cuick access Desktop # Downloads # Documents # Pictures #	Aless Aless Aless Folders (6) Desktop Music		Documents	Dow Vide	nloads	v   0	Search This PC		,
Computer 1 Computer 1 Cuick access Desktop # Downloads # Documents # Pictures # nimoavsdemosn #	New his PC V Folders (6) Desktop Music V Devices and drive Local Disk (6)	es (4)	Documents Dictures DVD Drive (D.) CCC VALSEDEE ENAULE DUE	Dow Vide DBd	nloads os isk (E.)	v	Search This PC		2
Computer 1 Computer 1 Cuick access Desktop # Downloads # Documents # Pictures # nimoavsdemosn # This PC Network	New New Version PC Version Folders (6) Desktop Music Version Persion Devices and drive Local Disk (f) DBDisk2 (F) Version PC	es (4) C3 of 39.5 GB	Documents Pictures Pictures DVD Drive (D.) SSS_X64FREE_EN-US_DVP D bytes free of 6.49 GB	Vide	nloads os iak (E.) GB free of 499 GB	< 0	Search This PC		

Google Cloud VMware Engine Supplemental NFS Datastore with NetApp Cloud Volume Service

Customers can expand storage capacity on Google Cloud VMware Engine using NFS supplemental datastore with NetApp Cloud Volume Service.

## Overview

Authors: Suresh Thoppay, NetApp

Customers that requires additional storage capacity on their Google Cloud VMware Engine (GCVE) environment can utilize Netapp Cloud Volume Service to mount as supplemental NFS datastore. Storing data on NetApp Cloud Volume Service allows customers to replicate between regions to protect from diaster.



## Deployment steps to mount NFS datastore from NetApp CVS on GCVE

#### Provision CVS-Performance Volume

The NetApp Cloud Volume Service volume can be either provisioned by Using Google Cloud Console Using NetApp BlueXP portal or API

#### Mark that CVS volume as non-deletable

To avoid accidental deletion of volume while VM is running, ensure the volume is marked as nondeletable as shown in screenshot below. image::gcp\_ncvs\_ds02.png[NetApp CVS non-deletable option] For more info, please refer Creating NFS Volume documentation.

#### Ensure Private Connection on GCVE exists for NetApp CVS Tenant VPC.

To mount NFS Datastore, there should be a private connection exists between GCVE and NetApp CVS project.

For more info, please refer How to setup Private Service Access

#### Mount NFS datastore

For instructions on how to mount NFS datastore on GCVE, please refer How to create NFS datastore with NetApp CVS



As vSphere hosts are managed by Google, you don't have access to install NFS vSphere API for Array Integration (VAAI) vSphere Installation Bundle (VIB). If you need support for Virtual Volumes (vVol), please let us know.

If you like to use Jumbo Frames, please refer Maximum supported MTU sizes on GCP

## Savings with NetApp Cloud Volume Service

To learn more about your potential saving with NetApp Cloud Volume Service for your storage demands on GCVE, please check NetApp ROI Calculator

## **Reference Links**

- Google Blog How to use NetApp CVS as datastores for Google Cloud VMware Engine
- NetApp Blog A better way to migrate your storage-rich apps to Google Cloud

#### NetApp Storage Options for GCP

GCP supports guest connected NetApp storage with Cloud Volumes ONTAP (CVO) or Cloud Volumes Service (CVS).

## **Cloud Volumes ONTAP (CVO)**

Cloud volumes ONTAP, or CVO, is the industry-leading cloud data management solution built on NetApp's ONTAP storage software, available natively on Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP).

It is a software-defined version of ONTAP that consumes cloud-native storage, allowing you to have the same storage software in the cloud and on-premises, reducing the need to retrain you IT staff in all-new methods to manage your data.

CVO gives customers the ability to seamlessly move data from the edge, to the data center, to the cloud and back, bringing your hybrid cloud together — all managed with a single-pane management console, NetApp Cloud Manager.

By design, CVO delivers extreme performance and advanced data management capabilities to satisfy even your most demanding applications in the cloud

#### Cloud Volumes ONTAP (CVO) as guest connected storage

Cloud Volumes ONTAP shares and LUNs can be mounted from VMs that are created in the GCVE private cloud environment. The volumes can also be mounted on the Linux client and on Windows client and LUNS can be accessed on Linux or Windows clients as block devices when mounted over iSCSI because Cloud Volumes ONTAP supports iSCSI, SMB, and NFS protocols. Cloud Volumes ONTAP volumes can be set up in a few simple steps.

To replicate volumes from an on-premises environment to the cloud for disaster recovery or migration purposes, establish network connectivity to Google Cloud, either using a site-to-site VPN or Cloud Interconnect. Replicating data from on-premises to Cloud Volumes ONTAP is outside the scope of this document. To replicate data between on-premises and Cloud Volumes ONTAP systems, see xref:./ehc/Setting up data replication between systems.



Use Cloud Volumes ONTAP sizer to accurately size the Cloud Volumes ONTAP instances. Also monitor on-premises performance to use as inputs in the Cloud Volumes ONTAP sizer.

1. Log in to NetApp Cloud Central—the Fabric View screen is displayed. Locate the Cloud Volumes ONTAP tab and select Go to Cloud Manager. After you are logged in, the Canvas screen is displayed.

 Cloud Ma	nager			kccount ~ Ketapp_POC	Workspace cloud_heroes	<ul> <li>Connect</li> <li>fissawsc</li> </ul>	tor 💙	<b>↓ @ ⊘ ⊗</b>
Canvas	Replication	Backup & Restore	K8s	Data Sense	File Cache	Compute	Sync	All Services (+8) ~
Canv	as							Go to Canvas View
G Add W	orking Environment							

2. On the Cloud Manager Canvas tab, click Add a Working Environment and then select Google Cloud Platform as the cloud and the type of the system configuration. Then, click Next.

Ħ	Cloud Ma	nager				Account ~ Netapp_POC	Workspace cloud_hero	8	Connector 🛩	۵	۲	0	8
	Canvas	Replication	Backup & Restore	KBS	Data Sense	File Cache	Compute	Sync	All Services (+7) 🛩				
	Add Worl	ling Environme	nt										×
					aws	0	Ĩ	-					
			Microsoft Asure	Ana	an Web Services	Google Cloud Plat	form	On-Pretoises					
					Choos	e Type							
			0		G	ີ	1	ົ					
			Cloud Volumes ON	TAP	Cloud Volum	es ONTAP HA	Cloud Vol	umes Service					
			Single Node		Thigh Av	allability	High	wailability					
						1							-

3. Provide the details of the environment to be created including the environment name and admin credentials. After you are done, click Continue.

Previous Step	CV-Performance-Testing	HCLMainBillingAccountSubs		Calls Deplace
	Google Cloud Project	Marketplace Subscription		Edit Project
	Details		Credentials	
	Working Environment Name	(Cluster Name)	User Name	
	cvogcveva		admin	
		-	Password	
	Service Account			
	Notice: A Google Cloud se to use two features: backing	ervice account is required ing up data using Backup	Confirm Password	

4. Select or deselect the add-on services for Cloud Volumes ONTAP deployment, including Data Sense & Compliance or Backup to Cloud. Then, click Continue.

HINT: A verification pop-up message will be displayed when deactivating add-on services. Add-on services can be added/removed after CVO deployment, consider to deselect them if not needed from the beginning to avoid costs.

Previous Step			
	Data Sense & Compliance	-•	~
	Backup to Cloud		~
	WARNING:By turning off Backup to Cloud, future data recovery will not be possible in case of	f data corruption or loss	

5. Select a location, choose a firewall policy, and select the checkbox to confirm network connectivity to Google Cloud storage.

Previous Step	Location	Connectivity
	GCP Region	VPC
	europe-west3 •	cloud-volumes-vpc •
	GCP Zone	Subnet
	europe-west3-c •	10.0.6.0/24
		Firewall Policy
	I have verified connectivity between the target VPC and Google     Cloud storage	<ul> <li>Generated firewall policy</li> <li>Use existing firewall policy</li> </ul>
Select the lie reemium o Create a New W	cense option: Pay-As-You-Go or BYOL for us ption is used. Then, click on Continue. forking Environment Cloud Volumes ONTAP Cha	sing existing license. In this example, arging Methods & NSS Account
Select the lie reemium o Create a New W	cont cense option: Pay-As-You-Go or BYOL for us ption is used. Then, click on Continue. forking Environment Cloud Volumes ONTAP Cha	sing existing license. In this example, arging Methods & NSS Account
Select the liv reemium o Create a New W	Cont cense option: Pay-As-You-Go or BYOL for us ption is used. Then, click on Continue. forking Environment Cloud Volumes ONTAP Char bud Volumes ONTAP Charging Methods	sing existing license. In this example, arging Methods & NSS Account NetApp Support Site Account
Select the liv reemium o Create a New W T Previous Step Le	Cont cense option: Pay-As-You-Go or BYOL for us ption is used. Then, click on Continue. Norking Environment Cloud Volumes ONTAP Cha bud Volumes ONTAP Charging Methods	sing existing license. In this example, arging Methods & NSS Account NetApp Support Site Account Learn more about NetApp Support Site (NSS) account
Select the liv Freemium o Create a New W Previous Step Le	Cont cense option: Pay-As-You-Go or BYOL for us ption is used. Then, click on Continue. Yorking Environment Cloud Volumes ONTAP Char bud Volumes ONTAP Charging Methods arm more about our charging methods Pay-As-You-Go by the hour	sing existing license. In this example, arging Methods & NSS Account NetApp Support Site Account Learn more about NetApp Support Site (NSS) account NetApp Support Site Account
Select the live Freemium of Create a New W T Previous Step Le	Cont Contense option: Pay-As-You-Go or BYOL for us ption is used. Then, click on Continue. Norking Environment Cloud Volumes ONTAP Char bud Volumes ONTAP Charging Methods arm more about our charging methods O Pay-As-You-Go by the hour	tinue sing existing license. In this example, arging Methods & NSS Account NetApp Support Site Account Learn more about NetApp Support Site (NSS) account NetApp Support Site Account NetApp Support Site Account
Select the live Treemium of Create a New W T Previous Step Le	Cont Contense option: Pay-As-You-Go or BYOL for us ption is used. Then, click on Continue. Norking Environment Cloud Volumes ONTAP Char bud Volumes ONTAP Charging Methods arn more about our charging methods Pay-As-You-Go by the hour O Bring your own license	sing existing license. In this example, arging Methods & NSS Account NetApp Support Site Account Learn more about NetApp Support Site (NSS) account NetApp Support Site Account nethad - To add a new NetApp Support Site account, go to the Support - NSS Management tab.

7. Select between several preconfigured packages available based on the type of workload that will be deployed on the VMs running on VMware cloud on AWS SDDC.

HINT: Hoover your mouse over the tiles for details or customize CVO components and ONTAP version by clicking on Change Configuration.

Select a pra	configured Cloud Volumes ONTAP system that best Preconfigured settings can be n	matches your needs, or create your own c nodified at a later time.	onfiguration. Change Configuration
• <b>•</b> ••	50	\$0	<u>ę</u> ,
POC and small workloads Up to 500GB of storage	Database and application data production workloads	Cost effective DR Up to 500GB of storage	Highest performance production workloads

8. On the Review & Approve page, review and confirm the selections. To create the Cloud Volumes ONTAP instance, click Go.

Create a New Work	ing Environment	Review & Approve		
Previous Step     CVOgCVEVal	act3			Show API reques
This Cloud Volumes O	NTAP Instance will be registered with NetApp	support under the NSS Account mchad.	ar 5	
I understand that	Cloud Manager will allocate the appropriate C	SCP resources to comply with my above requirements. More in	formation >	
Overview	Networking Storage	F.F		
Overview	Networking Storage	Cloud Volumes ONTAP runs on:	n2-standard-4	
Overview Storage System: License Type:	Networking Storage Cloud Volumes ONTAP Cloud Volumes ONTAP Freemium	Cloud Volumes ONTAP runs on: Encryption:	n2-standard-4 Google Cloud Managed	

9. After Cloud Volumes ONTAP is provisioned, it is listed in the working environments on the Canvas page.



1. After the working environment is ready, make sure the CIFS server is configured with the appropriate DNS and Active Directory configuration parameters. This step is required before you can create the SMB volume.

HINT: Click on the Menu Icon (°), select Advanced to display more options and select CIFS setup.

Cvogcve01		GCP Managed Encryption
Volumes Replications		<u>©</u> ∪ c ⊙ ≁ Ξ
Create a CIFS server	+ Advanced	
DNS Primary IP Address	Active Directory Domain to join	
192.168.0.16	nimgeveval.com	
DN5 Secondary IP Address (Optional)	Credentials authorized to join the domain	
Example: 127.0.0.1	administrator	

2. Creating the SMB volume is an easy process. At Canvas, double-click the Cloud Volumes ONTAP working environment to create and manage volumes and click on the Create Volume option. Choose the appropriate size and cloud manager chooses the containing aggregate or use advanced allocation mechanism to place on a specific aggregate. For this demo, CIFS/SMB is selected as the protocol.

Ve	olume Name: cvogcvesmbvol01	Size (GB):	0	NFS	CIFS	iSCSI
	cvogcvesmbvol01	10				
				-	22	
				Share name:	Permissions:	
Sr	sapshot Policy:			cvogcvesmbvol01_share	Full Control	-
	default		•			
	Default Policy			Users / Groups:		
				Everyone;		
				Valid users and groups separate	ed by a semicolon	

3. After the volume is provisioned, it will be availabe under the Volumes pane. Because a CIFS share is provisioned, give your users or groups permission to the files and folders and verify that those users can access the share and create a file. This step is not required if the volume is replicated from an on-premises environment because the file and folder permissions are all retained as part of SnapMirror replication.

HINT: Click on the volume menu (°) to display its options.

INFO		CAPACITY	
Disk Type	PD-SSD		■ 1.84 MB
Tiering Policy	None	10 GB	Disk Used

4. After the volume is created, use the mount command to display the volume connection instructions, then connect to the share from the VMs on Google Cloud VMware Engine.

Volu	mes Re	eplications						
	Nount Volu	ıme cvog	cvesmbvo	101				
o to yo	ur machine an	d enter this c	command					
\\10.	0.6.251\cvog	cvesmbvol01	_share		Ţ	Сору		
py the	following pat	h and use th Cloud VMw	ne Map Netwo vare Engine	ork Drive	option to	mount th	e volum	e on the VI
0								
pecify th	e drive letter for th	e connection an	d the folder that y	ou want to co	nnect to:			
pecify th trive:	e drive letter for th	e connection an	d the folder that y	ou want to co	nnect to:			
pecify th Prive: older:	e drive letter for th Y: \\10.0.6.251\c Evample: \\ser	e connection an wogcvesmbvol0	d the folder that y	ou want to co	nnect to: rowse			
pecify th Drive: older:	e drive letter for th Y: \\10.0.6.251\c Example: \\ser Reconnect	e connection an wogcvesmbvol0 ver\share at sign-in	d the folder that y	ou want to co	nnect to: rowse			
pecify th Drive:	e drive letter for th Y: \\10.0.6.251\c Example: \\ser Reconnect Connect us	e connection an wogcvesmbvol0 ver\share at sign-in ing different crea	d the folder that y	ou want to co	nnect to: rowse			
pecify th )rive: older:	e drive letter for th Y: \\10.0.6.251\c Example: \\ser ☑ Reconnect ☑ Connect us <u>Connect to a V</u>	e connection an wogcvesmbvol0 ver\share at sign-in ing different cred Veb site that you	d the folder that y 1_share dentials i can use to store y	ou want to co	nnect to: rowse	55-		
pecify th Irive: older:	e drive letter for th Y: \\10.0.6.251\c Example: \\ser ☑ Reconnect ☑ Connect us <u>Connect to a \</u>	e connection an wogcvesmbvol0 ver\share at sign-in ing different cred Veb site that you	d the folder that y  I_share  dentials  can use to store y	rou want to co	nnect to: rowse	5.		

Once mapped, it can be easily accessed, and the NTFS permissions can be set accordingly. 💣 l 📝 📗 🕶 l Network - 0 X 📭 👳 🗌 🔄 😇 🕴 cvogcvesmbvol01\_share (\\10.0.6.251) (Y:) - 0 × Home Share View 0 4 ← → \* ↑ ★ > This PC > cvogcvesmbvol01\_share (\\10.0.6.251) (Y:) > ✓ Ŏ Search cvogcvesmbvol01\_sha... , P Net 1 Name Date modified Туре Size # Quick access foo1 11/9/2021 10:59 AM File folder Desktop 4 11/9/2021 10:59 AM File folder foo2 Downloads Documents ġ, Pictures Ŕ This PC

To connect the cloud volumes ONTAP LUN to a host, complete the following steps:

- 1. On the Canvas page, double-click the Cloud Volumes ONTAP working environment to create and manage volumes.
- 2. Click Add Volume > New Volume and select iSCSI and click Create Initiator Group. Click Continue.

	Details & Pr	rotection				Protocol				
	Volume Name:		Size	e (G8): 《		NFS	CIF	s	iscsi	
	cvogcvescsilur	101	1	0				What	about LUNs?	0
	Snapshot Policy:					Initiator Group				
	default					Map Existin	g Initiator Group	• Create	Initiator Grou	чþ
	Default Policy	1				Initiator Group				
						WinlG				
						Operating Syste	ет Туре			
						Windows			2	•
VMware Cloud - ntap-fr → C Getting Started ● EC	x-demo X 💋 VSphe	ere - vmcdc01 - Sur El Trittps://vcenti Mew Tab	nmary × vrncdc01 er.sddc-52-37-127-104	l vmwarevmc.cor	× n/u/webcon	NetApp Cloud Manage	r × + 10058cvmName=vms	:dc01 80% ☆	0	- 0 © 1
VMware Cloud - ntap-fr 	ar demo X 💋 VSphe	ere - vmcdc01 - Sur #1 https://vcent www.Tab	nmary X vncdc01 eraddc-52-37-127-104	witiwarevitic.com	× n/u/webcan	NetApp Cloud Manage	r × + 1005SovmName=vne Tede	nde O 1 80% 🖒	CO t View Fullscreen	- 5 Other Bo Send Orl+A
VMware Cloud - ntap-fr C Getting Statted cocol 	a-dema X 🕜 VSphe Q A a 2 Management Con ( anager • Dashbocan	ere - vincdc01 - Sur #3 https://vcent Wew Tab	nmary X vmcdc01 er sddc-52-37-127-104 er www.ender	Viriwareviric.com	× n n/u/webcon	NetApp Cloud Menage	r × + 10058xmName=vmr	rdc01 80% 🟠	C View Fullecterin	C Cther Bo
VMware Cloud - ntap-fr C Getting Started - EC cdc01 	ar-demo X 🕜 VSphe Q A 3 2 Management Con 1 arrager • Dashboar welcowe to server 4	ere + vmcdc01 - Sur #2 https://vcent www.Tab	nmary X vnccdc01 er scidc-52-37-127-104 e market in the intervention in the intervention of the intervention in the intervention of the intervention intervention of the intervention of the intervention intervention of the intervention of the intervention intervention of the intervention of the interventio	Comwarevmc.com	× n/u/webcon	NetApp Cloud Manage	r × + 10058vmName=vm Foto 	ror US Keyboard Layou	CO	- 0 Cother Bo Send Ox1+A Send Ox1+A
VMware Cloud - ntap-fr C Getting Statted	ar-dema X 🕜 VSphe Q A 3 2 Management Con 1 anager • Dashboan welcowe to server M	ere + vmcdc01 - Sur #2 https://vcent d	mmary X vmcdc01 er addc-52-87-127-104 er addc-52-87-104 er addc-52-87-127-104 er addc-52	Comwarevinc.com	× n n/u/webcon	NetApp Cloud Manage sole htmlPvmtd=vm-	r × + 10058xmName=vnx rote - & test - & test	roe US Keyboard Layou	C View Fulliscreen	- D L Coher Boo
VMware Cloud - ntap-fi Getting Statted   Cetting Statted  Cetting Statted  Cetting Statted  Cetting Statted  Cetting  Cetting Cetting  Cetting Cetting Cetting Cetting Cetting Cetting	andemo X 2 VSphe 2 Management Con 4 anager • Dashboar welcome to server with concerner	ere + vmcdc01 - Sur #2 https://vcent www.Tab	mmary X vncck01 eraddc-52-87-127-104 eraddc-52-87-127-124 eraddc-52-87-127-104 eraddc-52-87-1	Annwarevmc.com	x n n/u/webcon	NetApp Cloud Manage	r × + 10058cmName=unx €etc -∧ text - 0 best siteseed	ror US Keyboard Laysu	C Voer Fullecteri	- D Other Bo Send Chi+A
VMware Cloud - ntap-fi 	andemo X 😢 VSphe C A 1 2 Management Con 1 anager • Dashbaan vescouer to server M	ere + vmcdc01 - Sur ## https://vcent New Tab	Inmary X vmcdc01  rr sddc-52-37-127-104  r sdc-52-37-127-104	Announcements con which is to be the metal of the balance metal of the balance metal of the balance the metal of the	X National Action of the second secon	NetApp Cloud Manage sole html?vmld=vm-	r × + 10058xmName=vms Entr 	roo 15 Keyboard Layou	CO Verse Futureset	- C L
VMware Cloud - ntap-fr C Getting Started • Ec cclc01 Colore Color	ar-dema X 😢 VSphe C A 3 2 Management Con 1 arnager • Dashboran welcoms to sames ki coccuran	ere + vmcdc01 - Sur ## https://vcent www.Tab New Tab ************************************	nmary X vnccdc01 er sddc-52-37-127-104 er sdc-52-37-127-104 er sdc-52-37-127-10	Notice and the second s	x n/u/vebcon	NetApp Cloud Manage sole html?vmld=vm-	r × + 10058xmName=vm Entr - * har - *	ror IDS Keyboard Layou	C View Fullicreen	- D Uniter Boo
VMware Cloud - ntap-fr C Getting Started	a- demo X 😢 VSphe C A 3 2 Management Con 1 2 Management Con	ere + vmcdc01 - Sur ## https://vcent wr Tab New Tab ## or taba ## or taba	Inmary X vnccdc01  er sddc-52-37-127-104  er sddc-52-37-127-104  e of the first of the statement of the stat	Antimovarevinic.com origination and anti- mental press (1712-1842a) de antibiota de maistre anti- statistica anti-	x n/u/vebcon	NetApp Cloud Manage sole html?vmld=vm-	r × + 10058xmName=vmr Toto  - 	cor US Keyboard Layou	Volew Futhereen	- D there Boo
VMware Cloud - ntap-fr C Getting Statted	andemo X I VSphe C A 3 2 Management Con. 1 2 Management Con. 1 3 Manager • Dashboar welcoas to saves 4 welcoas to saves	ere + vmcdc01 - Sur ## https://vcent w New Tab w New Tab	Inmary X vmcclc01  ar addc-52-37-127-104  ar addc-52-37-127-104  br blance b	Anniwarevinic.com	X In A/us/veebcom	NetApp Cloud Manage sole html?vmid=vm-	r × + 10058xmName=vmx for 	ror US Keyboard Layou 	Voew Futhereers Voew Futhereers Voew	- C Uther Boo Send Cut+A Sand Ut+A Sand Ut+A
VMware Cloud - ntap-fr 	anager • Dashboan weicener Con. • anager • Dashboan weicener • Dashboan	ere + vincdoù 1 - Sur El https://vcent New Tab C Vincent C Vin	mmary X vmcclc01 eraddc-52-87-127-104 eraddc-52-87-127-124 eraddc-52-87-	Livernwarevenic.com	X In A/uv/webcom	NetApp Cloud Manage sole htmlPvmid=vm-	r × + 10058ovmName=vno (nd 	roe US Keyboard Laysu	C Vorus Futherreeft	- 0 Uther Boo
VMware Cloud - ntap-fi C Getting Started	andema X Visphe C A 3 2 Management Con 4 2 Management Con 4	ere + vmcdc01 - Sur ## https://vcent www.Tab	Inmary X vncck01 eraddc-52-37-127-104 eraddc-52-37-	Virmwarevinc.cor	X Tan n/uv/vebcom	NetApp Cloud Manage	r × + 10058xmName=vmr Ende 	ror IDS Keyboard Layou	CO Correction Contractions ()       Contractions ()       Contractions ()       Contractions ()         Contractions ()	- C Cher Bo C Other Bo Send Chi+M - C Sant Univ Nat
VMware Cloud - ntap-fi C Getting Started	an demo X I Sphe C A solution of the second	ere + vmcdc01 - Sur ## https://vcent w New Tab w New Tab	Inmary X vnccdc01  In soldc-52-37-127-104  In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In soldc-52-37-127-104 In sol	Announce of the second se	X Tar n/u/vebcon	NetApp Cloud Manage sole.html?vmld=vm-	e × + 10058xmName=vm Text	dadi ana di nor US Keyboard Layau 	C View Publicities	- G Uther Boo
Weare Cloud - ntap-fr         →       C         Getting Started       Image: Cloud - ntap-fr         Contract       Server: M         Accost       Server: M         Lad Server       Lad Server         Actors       Ones         In the end strange Server       Image: Cloud Server	an demo X Visphe Constructions 2 Management Con 2 Management Con 2 Management Con 2 Management Con 3 Anagement Con 3 Anagem	ere + vmcdc01 - Sur ## https://vcent w New Tab	Inmary X vnccdc01  Ir sddc-52-37-127-104  Ir sddc-52-37-127-104 Ir s	A vern ware vern c. cor or ung te to a ne- word to prove (1972-1982) ( add) (b) Der malifier Serie 2007-191 versionen 2017-1 versionen 2017-1	X Tee A/u/vebcon	NetApp Cloud Manage sole.html?vmld=vm-	r × + 10058xmName=vm Tete = 6 beer entered	core US Keyboard Layou	Vorw Futhereen	- G L

3. After the volume is provisioned, select the volume menu (°), and then click Target iQN. To copy the iSCSI Qualified Name (iQN), click Copy. Set up an iSCSI connection from the host to the LUN.

To accomplish the same for the host residing on Google Cloud VMware Engine:

- a. RDP to the VM hosted on Google Cloud VMware Engine.
- b. Open the iSCSI Initiator Properties dialog box: Server Manager > Dashboard > Tools > iSCSI Initiator.
- c. From the Discovery tab, click Discover Portal or Add Portal and then enter the IP address of the iSCSI

target port.

- d. From the Targets tab, select the target discovered and then click Log on or Connect.
- e. Select Enable multipath, and then select Automatically Restore This Connection When the Computer Starts or Add This Connection to the List of Favorite Targets. Click Advanced.



The Windows host must have an iSCSI connection to each node in the cluster. The native DSM selects the best paths to use.

←) → Ser	iSCSI Initiator	r Properties			-
<u> </u>	Targets Dig	covery Favorite Targets Volum	mes and Devices RA	COUS Configuration	
	Quick Conn	ect			
Dashboard	DNS name	r and log on to a target using a ba of the target and then click Quick	Connection, type	the P address or	
Local Server					
All Servers	Target:	10.0.6.253		Quick Connect	Į,
AD DS	Discovered	targets			
D DNS	10.1	1994 Ballion		Refresh	
	Name		Sta	tus	Ì
			Second and these		
	To connect click Conne To complet then click D	using advanced options, select a ct, ely disconnect a target, select the isconnect.	target and then	Carried Distanced	
	To connect dick Conne To complet then dick D For target select the 1	using advanced options, select a ct, illy disconnect a target, select the isconnect, properties, including configuration target and click Properties.	target and then o target and n of sessions,	Carrent Disconvect Properties	
	To connect dick Conne To complet then dick D For target select the 1 For configu the target	using advanced options, select a ct. Wy disconnect a target, select the isconnect. properties, including configuration arget and cick Properties. ration of devices associated with and then cick Devices.	target and then o target and nof sessions, a target, select	Connect Disconnect Properties Devices	

LUNs on storage virtual machine (SVM) appear as disks to the Windows host. Any new disks that are added are not automatically discovered by the host. Trigger a manual rescan to discover the disks by completing the following steps:

- 1. Open the Windows Computer Management utility: Start > Administrative Tools > Computer Management.
- 2. Expand the Storage node in the navigation tree.
- 3. Click Disk Management.
- 4. Click Action > Rescan Disks.



When a new LUN is first accessed by the Windows host, it has no partition or file system. Initialize the LUN; and optionally, format the LUN with a file system by completing the following steps:

- 5. Start Windows Disk Management.
- 6. Right-click the LUN, and then select the required disk or partition type.
- 7. Follow the instructions in the wizard. In this example, drive F: is mounted.



On the Linux clients, ensure the iSCSI daemon is running. Once the LUNs are provisioned, refer to the detailed guidance on iSCSI configuration with Ubuntu as an example here. To verify, run lsblk cmd from the shell.

niyaza	ntmubu01	:-\$	lsblk	ł., .,		
NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
loop0	7:0	0	55.4M	1	loop	/snap/core18/2128
loop1	7:1	Θ	219M	1	loop	/snap/gnome-3-34-1804/72
Loopz	7:2	0	65.1M	1	loop	/snap/gtk-common-themes/1515
Loop3	7:3	Θ	51M	1	loop	/snap/snap-store/547
loop4	7:4	0	32.3M	1	loop	/snap/snapd/12704
loop5	7:5	6	32.5M	1	loop	/snap/snapd/13640
loop6	7:6	θ	55.5M	1	loop	/snap/core18/2246
Loop7	7:7	0	4K	1	loop	/snap/bare/S
Loop8	7:8	0	65.2M	1	loop	/snap/gtk-common-themes/1519
sda	8:0	0	16G	0	disk	
-sda1	8:1	0	512M	. 0	part	/boot/efi
-sda2	8:2		1K	0	part	
-sda5	8:5	Θ	15.50	0	part	1
sdb	8:16	0	16	0	disk	

ntyaz@ntrubuð	1:~\$ df	-h				
Filesystem	Size	Used	Avail	Use%	Mounted on	
udev	1.9G	0	1.96	0%	/dev	
tmpfs	394M	1.5M	392M	1%	/run	
/dev/sda5	16G	7.6G	6.9G	53%	7	
tmpfs	2.0G	0	2.00	0%	/dev/shm	
tmpfs	5.0M	0	5.0M	0%	/run/lock	
tmpfs	2.0G	0	2.06	0%	/sys/fs/cgroup	
/dev/loop1	219M	219M	θ	100%	/snap/gnome-3-34-1804/72	
/dev/loop2	66M	66M	0	100%	/snap/gtk-common-themes/1515	
/dev/loop3	51M	51M	0	100%	/snap/snap-store/547	
/dev/loop0	56M	56M	0	100%	/snap/core18/2128	
/dev/loop4	33M	33M	.8	100%	/snap/snapd/12764	
/dev/sda1	511M	4.0K	511M	1%	/boot/efi	
tmpfs	394M	64K	394M	1%	/run/user/1000	
/dev/loop5	33M	33M	0	100%	/snap/snapd/13640	
/dev/loop6	56M	56M	6	100%	/snap/core18/2246	
/dev/loop7	128K	128K	0	100%	/snap/bare/5	
/dev/loop8	66M	66M	0	166%	/snap/gtk-common-themes/1519	
/dev/sdb	976M	2.6M	987M	1%	/mnt	

To mount the Cloud Volumes ONTAP (DIY) file system from VMs within Google Cloud VMware Engine, follow the below steps:

Provision the volume following the below steps

- 1. In the Volumes tab, click Create New Volume.
- 2. On the Create New Volume page, select a volume type:

NFO		CAPACITY	
Disk Type	PD-SSD		<b>6.08</b> GB
Tiering Policy	None	11.05 GB	Disk Used

3. In the Volumes tab, place your mouse cursor over the volume, select the menu icon (°), and then click Mount Command.

Volumes Replications



- 5. Connect to the designated Linux instance.
- 6. Open a terminal on the instance using secure shell (SSH) and log in with the appropriate credentials.
- 7. Make a directory for the volume's mount point with the following command.

\$ sudo mkdir /cvogcvetst

# root@nimubu01:~# sudo mkdir cvogcvetst

8. Mount the Cloud Volumes ONTAP NFS volume to the directory that is created in the previous step.

concentration of		1 MOUNT *T NTS	: 10 (	1.6.2	51.10	-VC	ac	ve				ICVETST
			. 10.0			- • •	ac	VC.			tor crog	leverat.
bu01						F	nforce	US Ke	shoar	Lawout	View Fullscrean	Servi Ohl+Alt+D
							10.00	0070	1000	Lay ou		
	Activities	s 🖸 Terminal 🕶	No	w16 12:42				~				
			roote	pnimubu01:		Q				۲		
		rootentnubu01:-# df								1		
	-	Filesysten	1K-blocks	Used	Available	Usex	Houn	ted o	<b>fi</b> )			
		udev	1978500		1978508	0%	/dev					
		topfs	402272	1432	400840	18	/run					
		/dev/soas	13929250	1032332	2011352	325	Idani	/ chm				
		teofs	5128		\$120	0%	Irun	/tock				
		tnofs	2011352		2011352	ON	1585	Itsie	group			
		/dev/loop8	128	128	0	100%	/sne	p/bar	15			
		/dev/loop1	56832	56832	. 0.	100%	/sna	p/cor	e18/2	128		
		/dev/loop2	56832	56832	0	100%	/sna	p/cor	e18/2	246		
		/dev/loop4	66688	66688	0	100%	/sea	p/gtk	conn	on		
	-0-	(dev () cont	53334	69994		1000	Inna			Cal.		
		547	Secc.	Percen		100.0	Visite.	ex server	h-arn	1997) 1997		
	Card of the local division of the local divi	/dev/loop5	66816	66816		100N	/sna	p/atk	conn	on e		
		thenes/1519								10 al		
		/dev/loop7	33280	33280	. 0	100%	/sna	p/sna	pd/13	640		
		/dev/loop8	224256	224256	0	100x	/sna	p/gno	ne-3-	34-		
		1864/72			Terrare and							
		/dev/sdal	323248		523244	12	/000	C/ert				
		/dev/sdb	\$15010816	42816812	446761228	O'E	Thom	e falv	az/cv	110		
				Constant of the	THOMAS .		100	100				
	1000	/dev/loop9	43264	43264	6	100N	/sna	p/sna	pd/13	831		
		10.0.6.251:/cvogcvenfsvol01	13199552	8577536	4622016	65%	1000	t/cvo	ocvet	st		

#### **Cloud Volumes Service (CVS)**

Cloud Volumes Services (CVS) is a complete portfolio of data services to deliver advanced cloud solutions. Cloud Volumes Services supports multiple file access protocols for major cloud providers (NFS and SMB support).

Other benefits and features include: data protection and restore with Snapshot; special features to replicate, sync and migrate data destinations on-prem or in the cloud; and consistent high performance at the level of a dedicated flash storage system.

#### Cloud Volumes Service (CVS) as guest connected storage

#### Configure Cloud Volumes Service with VMware Engine

Cloud Volumes Service shares can be mounted from VMs that are created in the VMware Engine environment. The volumes can also be mounted on the Linux client and mapped on the Windows client because Cloud Volumes Service supports SMB and NFS protocols. Cloud Volumes Service volumes can be set up in simple steps.

Cloud Volume Service and Google Cloud VMware Engine private cloud must be in the same region.

To purchase, enable and configure NetApp Cloud Volumes Service for Google Cloud from the Google Cloud Marketplace, follow this detailed guide.

To create and mount NFS volumes, complete the following steps:

1. Access Cloud Volumes from Partner Solutions within the Google cloud console.

			<ul> <li>CV-Performance-Ter</li> </ul>	ating • • • • • • • • • • • • • • • • • • •		× 5 0	÷ 3.0
A	Home	>	MMENDATIONS				CUSTOM
ŧ.	Pins appear here	×				1	
RT	NER SOLUTIONS		1	Compute Engine	I	<ul> <li>Google Cloud Platform status</li> <li>All services normal</li> </ul>	1
	Redis Enterprise Apache Kafka on Co				102%	Go to Cloud status dashboard	
	Databricks				40%	Billing	1
3	DataStax Astra				20%	Estimated charges For the billing period starting Nov 1, 2021	USD \$0.00
*	Elasticsearch Service		Backups	1215 1230 1245 1 PM	e	Take a tour of billing	
٦	Neo4j Aura Professi		Snapshots Active Directories	→ Bo to Compute Engine		→ View detailed charges	
5	Cloud Volumes	>	Volume Replication			III Monitoring	

2. In the Cloud Volumes Console, go to the Volumes page and click Create.

-			stores	t thereas										
2	Cloud Volumes	Volur	nes	CREA	TE DELE	TE								
2	Volumes	Quick n	eferer	nce for Cloud Volum										
2	Backups	Ŧ	Filter	Search for volume	es by name, ID, re	gion, etc.						0	m	
0	Snapshots			ID	Name	Region	Zone	Zone Redundancy	Life Cycle	Billi	ng Label	State	Detaile	6
Active Directories     Volume Replication	Active Directories Volume Replication		0	Dac8a83d- 03d8-c9db- 2aba- 190-7535445b	testnfsds01	europe- west3			available			Avai	lable fo	10
			0	330f35e2- b0c6-98b3- ec7a- 8dd4ea7ba00e	gcp-ve-ds4	europe- west3			available			Avail	lable fo	я
			0	7d0a6f0d- 3e0a-50c3- 5295- 5152040681fc	gcp-ve-ds3	europe- west3			available			Ayal	lable fo	ir.
			0	8cae6850- 0919-4eaf-	gcve-ds-2	europe- west3			available			Avail	lable fo	я

3. On the Create File System page, specify the volume name and billing labels as required for chargeback mechanisms.
| Cloud Volumes      | ← Create File System   |
|--------------------|--|
| 2 Volumes          |  |
| Backups            | Volume Name  |
| Snapshots          | nimCVNFSvol01  |
| Active Directories | A human readable name used for display purposes.   |
| Volume Replication | Billing Labels   |
|                    | Label your volumes for billing reports, queries.<br>Supported with CVS-Performance service type; can be set with CVS service type but not<br>available for billing at this time. |
|                    | + ADD LABEL  |

4. Select the appropriate service. For GCVE, choose CVS-Performance and desired service level for improved latency and higher performance based on the application workload requirements.

0	Cloud Volumes	← Create File System
	Volumes	Service Type
	Backups	Cloud Volumes Service is offered as two service types: CVS and CVS-Performance. Select the service type that matches your workload needs. Region availability 2 varies by
0	Snapshots	service type. Learn more
0	Active Directories	O CVS Offers volumes created with zonal high availability.
D	Volume Replication	<ul> <li>CVS-Performance</li> <li>Offers 3 performance levels and improved latency to address higher performance application requirements.</li> </ul>
		Volume Replication
		Secondary Select to create volume as a destination target for volume replication. Applicable only to CVS-performance volumes.

5. Specify the Google Cloud region for the volume and volume path (The volume path must be unique across all of cloud volumes in the project)



6. Select the level of performance for the volume.



7. Specify the size of the volume and the protocol type. In this testing, NFSv3 is used.

-		Volume Details
8	Volumes	Allocated Capacity *
Q	Backups	1024 GiB
_		Allocated size must be between 1 TiB (1024 GiB) and 100 TiB (102400 GiB)
0	Snapshots	Protocol Type *
0	Active Directories	NFSv3
Ø	Volume Replication	Make snapshot directory (.snapshot) visible
		Makes .snapshot directory visible to clients. For NFSv4.1 volumes (CVS-Performance only) the directory itself will not be listed but can be accessed to list contents, etc.
		Enable LDAP
		Enables user look up from AD LDAP server for your NFS volumes

8

HINT: If VPC peering has not been done, a pop-up button will be displayed to guide you through the peering commands. Open a Cloud Shell session and execute the appropriate commands to peer your VPC with Cloud Volumes Service producer. In case you decide to prepare VPC peering in beforehand, refer to these instructions.

9	Cloud Volumes	← Create File System
	Volumes	Network Details
	Backups	Provide the host project name when deploying in a shared VPC service project.
0	Snapshots	VPC Network Name *
0	Active Directories	Select the VPC Network from which the volume will be accessible. This cannot be changed later.
O	Volume Replication	Use Custom Address Range
		Reserved Address range netapp-addresses

9. Manage the Export policy rules by adding the appropriate rules and Select the checkbox for the corresponding NFS version.

Note: Access to NFS volumes won't be possible unless an export policy is added.

0	Cloud Volume	s 🔶	Create File S	System		
	Volumes	E	xport Policy			
	Backups	R	ules			
0	Snapshots		Item 1			~ ~ <b>i</b>
0	Active Directories		Allowed Clients 1 * 0.0.0.0/0			
	Volume Replication		Access			
			<ul> <li>Read &amp; Write</li> <li>Read Only</li> </ul>			
			Root Access			
			On			
			O off			
			Protocol Type (Sele	ect at least 1 of t	he below options	5)
			Must select for Proto NFSv4.1	ocol type NFSv3. C	Optional for Protoc	ol Type Both. Do not select for
			Allows Matchin	ng Clients for NF	SV3	
0. Click Save	to create the	volume.				
466ed9d9-	nimnfsdemods02	europe- Availabl west3	e for use CVS- Performance	Primary	Extreme	NFSv2: 10.53.0.4 /niminfademods02

Before preparing to mount the NFS volume, ensure the peering status of private connection is listed as Active. Once status is Active, use the mount command.

To mount an NFS volume, do the following:

- 1. In the Cloud Console, go to Cloud Volumes > Volumes.
- 2. Go to the Volumes page
- 3. Click the NFS volume for which you want to mount NFS exports.
- 4. Scroll to the right, under Show More, click Mount Instructions.

To perform the mounting process from within the guest OS of the VMware VM, follow the below steps:

- 1. Use SSH client and SSH to the virtual machine.
- 2. Install the nfs client on the instance.
  - a. On Red Hat Enterprise Linux or SuSE Linux instance:

sudo yum install -y nfs-utils

b. On an Ubuntu or Debian instance:

sudo apt-get install nfs-common

3. Create a new directory on the instance, such as "/nimCVSNFSol01":

```
sudo mkdir /nimCVSNFSol01
```

4. Mount the volume using the appropriate command. Example command from the lab is below:

```
sudo mount -t nfs -o rw,hard,rsize=65536,wsize=65536,vers=3,tcp
10.53.0.4:/nimCVSNFSol01 /nimCVSNFSol01
```

```
oot@vm1:-# sudo mkdlr nlmCVSNFSol01
oot@vm1:-# sudo mount -t nfs -o rw,hard,rsize=65536,wsize=65536,vers=3,tcp 10.53.0.4:/nimCVSNFSol01 /nimCVSNFSol01
```

root@vm1:~# df	an extension of the second	· · · · · · · · · · · · · · · · · · ·	and the second sec		1000 100 100	
Filesystem	1K-blocks	Used	Available	Use%	Mounted on	
udev	16409952	(10)0	16409952	0%	/dev	
tnpfs	3288328	1588	3286748	1%	/run	
/dev/sdb5	61145932	19231356	38778832	34%		
tripfs	16441628	6	16441628	8%	/dev/shm	
tnpfs	5120	6	\$120	6%	/run/lock	
tnpfs	16441628	a	16441628	0%	/sys/fs/cgroup	
/dev/loop0	128	128	0	100%	/snap/bare/5	
/dev/loop1	56832	56832	Ġ	100%	/snap/core18/2128	
/dev/loop2	66688	66688	.0	100%	/snap/gtk-comon-thenes/1515	
/dev/loop4	66816	66816	0	100%	/snap/gtk-connon-thenes/1519	
/dev/loop3	52224	52224	0	100%	/snap/snap-store/S47	
/dev/loop5	224256	224256		166%	/snap/gnone-3-34-1804/72	
/dev/sdb1	523248	100-100	523244	18	/boot/eft	
tnpfs	3288324	28	3288296	1%	/run/user/1008	
10.53.0.4:/gcve-ds-1	107374182400	1136086016	106238096384	2%	/base	
/dev/mapper/nfsprdvg1-prod01	419155968	55384972	363778996	14%	/datastore1	
/dev/loop8	33280	33280	0	100%	/snap/snapd/13270	
/dev/loop6	33280	33280	0	100%	/snap/snapd/13640	
/dev/loop7	56832	56832	6	100%	/snap/core18/2246	
10.53.0.4:/nimCVSNFSol01	107374182400	256	107374182144	1%	/ninCVSNFSol01	
root@vm1:-#						

For SMB volumes, make sure the Active Directory connections is configured prior to creating the SMB volume.

Active	Directory conr	ections	CREATE	T DELETE						
Create a V	Windows Active Dire	ectory connection to yo	our existing AD se	rver. This is a prerequisite	step before creating vol	umes with the SMB pro	tocol type. Learn	more (2		
포 Fil	ter Search for Act	ive Directory connection	ins by ID, useman	ne, DNS, netBIOS, region, e	tc.				0	ш
	Username	Domain	DNS Servers	NetBIOS Prefix	OU Path	AD Server Name	KDC IP	Region	Stat	tus
	administrator	nimgeveval.com	192,168.0.16	nimsmb	CN=Computers			europe-	- In	Use

Once the AD connection is in place, create the volume with the desired service level. The steps are like creating NFS volume except selecting the appropriate protocol.

- 1. In the Cloud Volumes Console, go to the Volumes page and click Create.
- 2. On the Create File System page, specify the volume name and billing labels as required for chargeback mechanisms.



## Volume Name

Name \* \_\_\_\_\_

nimCVSMBvol01

A human readable name used for display purposes.

### **Billing Label**

Label your volumes for billing reports, queries. Supported with CVS-Performance service type; can be set with CVS service type but not available for billing at this time.



3. Select the appropriate service. For GCVE, choose CVS-Performance and desired service level for improved latency and higher performance based on the workload requirements.

# **Create File System**

## Service Type

←

Cloud Volumes Service is offered as two service types: CVS and CVS-Performance. Select the service type that matches your workload needs. <u>Region availability</u> is varies by service type. Learn more is



Offers volumes created with zonal high availability.



Offers 3 performance levels and improved latency to address higher performance application requirements.

# **Volume Replication**

#### Secondary

Select to create volume as a destination target for volume replication. Applicable only to CVS-performance volumes.

4. Specify the Google Cloud region for the volume and volume path (The volume path must be unique across all of cloud volumes in the project)

# Create File System

## Region

Region availability varies by service type.

europe-west3	- 0
Volume will be provisioned in the region you select.	

Must be unique to the project.

nimCVSMBvol01

5. Select the level of performance for the volume.

C

Se	rvice Level	
Se	lect the performance level required for your workload.	
0	Standard Up to 16 MiB/s per TiB	
0	Premium Up to 64 MIB/s per TIB	
0	Extreme Up to 128 MiB/s per TiB	
s	inapshot 👻	
- T	he snapshot to create the volume from.	
Vo	lume Details	
Vo	lume Details	
Vo A 1	lume Details llocated Capacity * 024 GiB	
Vo A 1	lume Details llocated Capacity *	
Vo A 1	Ilume Details Ilocated Capacity * O24 GiB Ilocated size must be between 1 TiB (1024 GiB) and 100 TiB (102400 GiB)	
Vo A 1 A	Ilume Details Ilocated Capacity * O24 GiB Ilocated size must be between 1 TiB (1024 GiB) and 100 TiB (102400 GiB) Fotocol Type * MB T	
Vo A 1 A	Ilume Details Ilocated Capacity * O24 GiB Ilocated size must be between 1 TiB (1024 GiB) and 100 TiB (102400 GiB) Totocol Type * MB Totocol Type *	
Vo A 10 A	Ilume Details Ilocated Capacity * O24 GiB Ilocated size must be between 1 TiB (1024 GiB) and 100 TiB (102400 GiB) Totocol Type * MB Make snapshot directory (.snapshot) visible Makes .snapshot directory visible to clients. For NFSv4.1 volumes (CVS-Performance only), the directory itself will not be listed but can be accessed to list contents, etc.	
Vo A 11 A	Ilume Details Ilocated Capacity * O24 GiB Ilocated size must be between 1 TiB (1024 GiB) and 100 TiB (102400 GiB) Totocol Type * MB Makes snapshot directory (.snapshot) visible Makes .snapshot directory visible to clients. For NFSv4.1 volumes (CVS-Performance only), the directory itself will not be listed but can be accessed to list contents, etc. Enable SMB Encryption	
	Ilucated Capacity *	
	Iume Details         Ilocated Capacity *         024       GiB         Ilocated size must be between 1 TiB (1024 GiB) and 100 TiB (102400 GiB)         rotocol Type *         MB         Make snapshot directory (.snapshot) visible         Makes .snapshot directory visible to clients. For NFSv4.1 volumes (CVS-Performance only), the directory itself will not be listed but can be accessed to list contents, etc.         Enable SMB Encryption         Enable this option only if you require encryption of your SMB data traffic.         Enable CA share support for SQL Server. FSLogix	
	Ilure Details Illocated Capacity *	
	Ilune Details Ilocated Capacity * O24 GiB Ilocated Size must be between 1 TiB (1024 GiB) and 100 TiB (102400 GiB) Ilocated size must be between 1 TiB (1024 GiB) and 100 TiB (102400 GiB) Inotacol Type * MB Make snapshot directory (.snapshot) visible Makes .snapshot directory (.snapshot) visible Makes .snapshot directory visible to clients. For NFSv4.1 volumes (CVS-Performance only), the directory itself will not be listed but can be accessed to list contents, etc. Enable SMB Encryption Enable this option only if you require encryption of your SMB data traffic. Enable CA share support for SQL Server, FSLogix Enable this option only for SQL Server and FSLogix workloads that require continuous availability. Hide SMB Share	
	Ilume Details Illocated Capacity *	

	Shared V Provide th	PC configur ne host projec	ation It name w	hen deployi	ng in a shar	ed VPC ser	vice project.	
	PC Networ loud-volun	k Name * — nes-vpc						•
Se	elect the Vi ter.	PC Network f	rom whic	h the volume	e will be acc	essible. Th	is cannot be	changed
	Use Cust	om Address	Range					
R	eserved Ac etapp-add	ldress range resses						
~	SHOW SN	APSHOT PO	LICY					
S	AVE	CANCEL						
	< Save to	create the	volume.					
Click	6a4552ed-	nimCVSMBvol01	europe- west3	Available for use	CVS- Performance	Primary	Standard	SMB : \\nimamb-3830.nimgcveval.com\nimCVSMBvo

- 1. In the Cloud Console, go to Cloud Volumes > Volumes.
- 2. Go to the Volumes page
- 3. Click the SMB volume for which you want to map an SMB share.
- 4. Scroll to the right, under Show More, click Mount Instructions.

To perform the mounting process from within the Windows guest OS of the VMware VM, follow the below steps:

- 1. Click the Start button and then click on Computer.
- 2. Click Map Network Drive.
- 3. In the Drive list, click any available drive letter.
- 4. In the folder box, type:

\\nimsmb-3830.nimgcveval.com\nimCVSMBvol01

What n	etwork folder would you lik	te to map?				
Specify th	e drive letter for the connection and	i the folder that you want to	connect to:			
Drive:	Z:	~				
Folder:	\\10.53.0.4\nimcvsmbvpl01	~	Browse			
	Example: \\server\share					
	Beconnect at sign-in					
	Connect using different crea	lentials				
	Connect to a Web site that you	can use to store your docum	ents and pictures.			
To conne Click Fin I nimcvsmbv ie Share	ect every time you log on t iish. 1001 (\\10.53.0.4) (Z:) View	o your computer, sel	ect the Recon	nect at sig	n-in che	ck box. –
To conne Click Fin I nimcvsmbv ie Share	ect every time you log on t ish. vol01 (\\10.53.0.4) (Z:) View is PC > nimcvsmbvol01 (\\10.53.0.4) (Z	o your computer, sel	ect the Recon	nect at sig	n-in che	ck box. –
To conne Click Fin I nimevsmbv ie Share	ect every time you log on t ish. rol01 (\\10.53.0.4) (Z:) View is PC > nimcvsmbvol01 (\\10.53.0.4) (Z Name	o your computer, sele	ect the Recon	nect at sig	n-in che	ck box. –
To conne Click Fin I nimevsmbv e Share	ect every time you log on t ish. vol01 (\\10.53.0.4) (Z:) View is PC > nimcvsmbvol01 (\\10.53.0.4) (Z Name foo1	to your computer, sele	Type	nect at sig	n-in che	ck box. –
To conne Click Fin I nimevsmbv ie Share	ect every time you log on t ish. rol01 (\(10.53.0.4) (Z:) View is PC > nimcvsmbvol01 (\(10.53.0.4) (Z Name foo1 foo2	Date modified 11/1/2021 7:38 AM 11/1/2021 7:38 AM	Type File folder File folder	nect at sig	n-in che	ck box. –

Region Availability for Supplemental NFS datastores on AWS, Azure, and GCP

Learn more about the the Global Region support for supplemental NFS datastores on AWS, Azure and Google Cloud Platform (GCP).

## **AWS Region Availability**

£.

The availability of supplemental NFS datastores on AWS / VMC is defined by Amazon. First, you need to determine if both VMC and FSxN are available in a specified region. Next, you need to determine if the FSxN supplemental NFS datastore is supported in that region.

- Check the availability of VMC here.
- Amazon's pricing guide offers information on where FSxN (FSx ONTAP) is available. You can find that information here.
- Availability of the FSxN supplemental NFS datastore for VMC is coming soon.

While information is still being released, the following chart identifies the current support for VMC, FSxN and FSxN as a supplemental NFS datastore.

## Americas

AWS Region	VMC Availability	FSx ONTAP Availability	NFS Datastore Availability
US East (Northern Virginia)	Yes	Yes	Yes
US East (Ohio)	Yes	Yes	Yes
US West (Northern California)	Yes	No	No
US West (Oregon)	Yes	Yes	Yes
GovCloud (US West)	Yes	Yes	Yes
Canada (Central)	Yes	Yes	Yes
South America (Sao Paulo)	Yes	Yes	Yes

Last updated on: June 2, 2022.

## EMEA

AWS Region	VMC Availability	FSx ONTAP Availability	NFS Datastore Availability
Europe (Ireland)	Yes	Yes	Yes
Europe (London)	Yes	Yes	Yes
Europe (Frankfurt)	Yes	Yes	Yes
Europe (Paris)	Yes	Yes	Yes
Europe (Milan)	Yes	Yes	Yes
Europe (Stockholm)	Yes	Yes	Yes

Last updated on: June 2, 2022.

### Asia Pacific

AWS Region	VMC Availability	FSx ONTAP Availability	NFS Datastore Availability
Asia Pacific (Sydney)	Yes	Yes	Yes
Asia Pacific (Tokyo)	Yes	Yes	Yes
Asia Pacific (Osaka)	Yes	No	No
Asia Pacific (Singapore)	Yes	Yes	Yes
Asia Pacific (Seoul)	Yes	Yes	Yes
Asia Pacific (Mumbai)	Yes	Yes	Yes
Asia Pacific (Jakarta)	No	No	No
Asia Pacific (Hong Kong)	Yes	Yes	Yes

## **Azure Region Availability**

The availability of supplemental NFS datastores on Azure / AVS is defined by Microsoft. First, you need to determine if both AVS and ANF are available in a specific region. Next, you need to determine if the ANF supplemental NFS datastore is supported in that region.

- Check the availability of AVS and ANF here.
- Check the availability of the ANF supplemental NFS datastore here.

### **GCP Region Availability**

GCP region availability will be released when GCP enters public availability.

### Summary and Conclusion: Why NetApp Hybrid Multicloud with VMware

NetApp Cloud Volumes along with VMware solutions for the major hyperscalers provides great potential for organizations looking to leverage hybrid cloud. The rest of this section provides the use cases that show integrating NetApp Cloud Volumes enables true hybrid Multicloud capabilities.

### Use case #1: Optimizing storage

When performing a sizing exercise using RVtools output, it is always evident that the horsepower (vCPU/vMem) scale is parallel with storage. Many times, organizations find themselves in a situation where the storage space requires drives the size of the cluster well beyond what is needed for horsepower.

By integrating NetApp Cloud Volumes, organizations can realize a vSphere-based cloud solution with a simple migration approach, with no re-platforming, no IP changes, and no architectural changes. Additionally, this optimization enables you to scale the storage footprint while keeping the host count to least amount required in vSphere, but no change to the storage hierarchy, security, or files made available. This allows you to optimize the deployment and reduce the overall TCO by 35–45%. This integration also enables you to scale storage from warm storage to production-level performance in seconds.

### Use case #2: Cloud migration

Organizations are under pressure to migrate applications from on-premises data centers to the Public Cloud for multiple reasons: an upcoming lease expiration; a finance directive to move from capital expenditure (capex) spending to operational expenditures (opex) spending; or simply a top-down mandate to move everything to the cloud.

When speed is critical, only a streamlined migration approach is feasible because re-platforming and refactoring applications to adapt to the cloud's particular IaaS platform is slow and expensive, often taking months. By combining NetApp Cloud Volumes with the bandwidth-efficient SnapMirror replication for guest-connected storage (including RDMs in conjunction with application-consistent Snapshot copies and HCX, cloud specific migration (e.g. Azure Migrate), or third-party products for replicating VMs), this transition is even easier than relying on time-consuming I/O filters mechanisms.

#### Use case #3: Data center expansion

When a data center reaches capacity limits due to seasonal demand spikes or just steady organic growth,

moving to the cloud-hosted VMware along with NetApp Cloud Volumes is an easy solution. Leveraging NetApp Cloud Volumes allows storage creation, replication, and expansion very easily by providing high availability across availability zones and dynamic scaling capabilities. Leveraging NetApp Cloud Volumes helps in minimizing host cluster capacity by overcoming the need for stretch clusters.

### Use case #4: Disaster recovery to the cloud

In a traditional approach, if a disaster occurs, the VMs replicated to the cloud would require conversion to the cloud's own hypervisor platform before they could be restored – not a task to be handled during a crisis.

By using NetApp Cloud Volumes for guest-connected storage using SnapCenter and SnapMirror replication from on-premises along with public cloud virtualization solutions, a better approach for disaster recovery can be devised allowing VM replicas to be recovered on fully consistent VMware SDDC infrastructure along with cloud specific recovery tools (e.g. Azure Site Recovery) or equivalent third-party tools such as Veeam. This approach also enables you to perform disaster recovery drills and recovery from ransomware quickly. This also enables you to scale to full production for testing or during a disaster by adding hosts on-demand.

### Use case #5: Application modernization

After applications are in the public cloud, organizations will want to take advantage of the hundreds of powerful cloud services to modernize and extend them. With the use of NetApp Cloud Volumes, modernization is an easy process because the application data is not locked into vSAN and allows data mobility for a wide range of use cases, including Kubernetes.

### Conclusion

Whether you are targeting an all-cloud or hybrid cloud, NetApp Cloud Volumes provides excellent options to deploy and manage the application workloads along with file services and block protocols while reducing the TCO by making the data requirements seamless to the application layer.

Whatever the use case, choose your favorite cloud/hyperscaler together with NetApp Cloud Volumes for rapid realization of cloud benefits, consistent infrastructure, and operations across on-premises and multiple clouds, bidirectional portability of workloads, and enterprise-grade capacity and performance.

It is the same familiar process and procedures that are used to connect the storage. Remember, it is just the position of the data that changed with new names; the tools and processes all remain the same and NetApp Cloud Volumes helps in optimizing the overall deployment.

# VMware Hybrid Cloud Use Cases

## Use Cases for NetApp Hybrid Multicloud with VMware

An overview of the use cases of importance to IT organization when planning hybridcloud or cloud-first deployments.

## Popular Use Cases

Use cases include:

- Disaster recovery,
- Hosting workloads during data center maintenance, \* quick burst in which additional resources are required beyond what's provisioned in the local data center,
- VMware site expansion,

- Fast migration to the cloud,
- Dev/test, and
- Modernization of apps leveraging cloud supplemental technologies.

Throughout this documentation, cloud workload references will be detailed using the VMware use-cases. These use-cases are:

- Protect (includes both Disaster Recovery and Backup / Restore)
- Migrate
- Extend

## Inside the IT Journey

Most organizations are on a journey to transformation and modernization. As part of this process, companies are trying use their existing VMware investments while leveraging cloud benefits and exploring ways to make the migration process as seamless as possible. This approach would make their modernization efforts very easy because the data is already in the cloud.

The easiest answer to this scenario is VMware offerings in each hyperscaler. Like NetApp® Cloud Volumes, VMware provides a way to move or extend on-premises VMware environments to any cloud, allowing you to retain existing on-premises assets, skills, and tools while running workloads natively in the cloud. This reduces risk because there will be no service breaks or a need for IP changes and provides the IT team the ability to operate the way they do on-premises using existing skills and tools. This can lead to accelerated cloud migrations and a much smoother transition to a hybrid Multicloud architecture.

## Understanding the Importance of Supplemental NFS Storage Options

While VMware in any cloud delivers unique hybrid capabilities to every customer, limited supplemental NFS storage options have restricted its usefulness for organizations with storage-heavy workloads. Because storage is directly tied to hosts, the only way to scale storage is to add more hosts—and that can increase costs by 35–40 percent or more for storage intensive workloads. These workloads just need additional storage, not additional horsepower. But that means paying for additional hosts.

Let's consider this scenario:

A customer requires just five hosts for CPU and memory, but has a lot of storage needs, and needs 12 hosts to meet the storage requirement. This requirement ends up really tipping the financial scale by having to buy the additional horsepower, when they only need to increment the storage.

When you're planning cloud adoption and migrations, it's always important to evaluate the best approach and take the easiest path that reduces total investments. The most common and easiest approach for any application migration is rehosting (also known as lift and shift) where there is no virtual machine (VM) or data conversion. Using NetApp Cloud Volumes with VMware software-defined data center (SDDC), while complementing vSAN, provides an easy lift-and-shift option.

# NetApp Solutions for Amazon VMware Managed Cloud (VMC)

Learn more about the solutions that NetApp brings to AWS.

VMware defines the cloud workloads into one of three categories:

• Protect (including both Disaster Recovery and Backup / Restore)

- Migrate
- Extend

Browse the available solutions in the following sections.

## Protect

- Disaster Recovery with VMC on AWS (guest connected)
- Veeam Backup & Restore in VMC with FSx for ONTAP
- Disaster Recovery (DRO) with FSx for ONTAP and VMC
- Using Veeam Replication and FSx for ONTAP for Disaster recovery to VMware Cloud on AWS

## Migrate

Migrate Workloads to FSxN datastore using VMware HCX

# Extend

COMING SOON!!

## NetApp Solutions for Azure VMware Solution (AVS)

Learn more about the solutions that NetApp brings to Azure.

VMware defines the cloud workloads into one of three categories:

- Protect (including both Disaster Recovery and Backup / Restore)
- Migrate
- Extend

Browse the available solutions in the following sections.

### Protect

- Disaster Recovery with ANF and JetStream (supplemental NFS datastore)
- Disaster Recovery with ANF and CVO (guest connected storage)
- Disaster Recovery (DRO) with ANF and AVS
- Using Veeam Replication and Azure NetApp Files datastore for disaster recovery to Azure VMware Solution

## Migrate

• Migrate Workloads to Azure NetApp Files datastore using VMware HCX

## Extend

COMING SOON!!

## NetApp Solutions for Google Cloud VMware Engine (GCVE)

Learn more about the solutions that NetApp brings to GCP.

VMware defines the cloud workloads into one of three categories:

- Protect (including both Disaster Recovery and Backup / Restore)
- Migrate
- Extend

Browse the available solutions in the following sections.

#### Protect

- Application Disaster Recovery with SnapCenter, Cloud Volumes ONTAP and Veeam Replication
- Application Consistent Disaster Recovery with NetApp SnapCenter and Veeam Replication to NetApp CVS on GCVE

### Migrate

- Workload Migration using VMware HCX to NetApp Cloud Volume Service NFS datastore
- VM Replication using Veeam to NetApp Cloud Volume Service NFS datastore

### Extend

COMING SOON!!

# **NetApp Capabilities for AWS VMC**

Learn more about the capabilities that NetApp brings to the AWS VMware Cloud (VMC) - from NetApp as a guest connected storage device or a supplemental NFS datastore to migrating workflows, extending/bursting to the cloud, backup/restore and disaster recovery.

Jump to the section for the desired content by selecting from the following options:

- Configuring VMC in AWS
- NetApp Storage Options for VMC
- NetApp / VMware Cloud Solutions

## **Configuring VMC in AWS**

As with on-premises, planning a cloud based virtualization environment is critical for a successful productionready environment for creating VMs and migration.

This section describes how to set up and manage VMware Cloud on AWS SDDC and use it in combination with the available options for connecting NetApp storage.



In-guest storage is the only supported method of connecting Cloud Volumes ONTAP to AWS VMC.

The setup process can be broken down into the following steps:

• Deploy and Configure VMware Cloud for AWS

Connect VMware Cloud to FSx ONTAP

View the detailed configuration steps for VMC.

## NetApp Storage Options for VMC

NetApp storage can be utilized in several ways - either as guess connected or as a supplemental NFS datastore - within AWS VMC.

Please visit Supported NetApp Storage Options for more information.

AWS supports NetApp storage in the following configurations:

- FSx ONTAP as guest connected storage
- Cloud Volumes ONTAP (CVO) as guest connected storage
- FSx ONTAP as a supplemental NFS datastore

View the detailed guest connect storage options for VMC. View the detailed supplemental NFS datastore options for VMC.

## **Solution Use Cases**

With NetApp and VMware cloud solutions, many use cases are simple to deploy in your AWS VMC. Use cases are defined for each of the VMware defined cloud areas:

- Protect (includes both Disaster Recovery and Backup / Restore)
- Extend
- Migrate

Browse the NetApp solutions for AWS VMC

## Protecting Workloads on AWS / VMC

### TR-4931: Disaster Recovery with VMware Cloud on Amazon Web Services and Guest Connect

A proven disaster recovery (DR) environment and plan is critical for organizations to ensure that business-critical applications can be rapidly restored in the event of a major outage. This solution focuses on demonstrating DR use cases with a focus on VMware and NetApp technologies, both on-premises and with VMware Cloud on AWS.

Authors: Chris Reno, Josh Powell, and Suresh Thoppay - NetApp Solutions Engineering

## Overview

NetApp has a long history of integration with VMware as evidenced by the tens of thousands of customers that have chosen NetApp as their storage partner for their virtualized environment. This integration continues with guest-connected options in the cloud and recent integrations with NFS datastores as well. This solution focuses on the use case commonly referred to as guest-connected storage.

In guest-connected storage, the guest VMDK is deployed on a VMware-provisioned datastore, and application data is housed on iSCSI or NFS and mapped directly to the VM. Oracle and MS SQL applications are used to demonstrate a DR scenario, as shown in the following figure.



## Assumptions, pre-requisites and component overview

Before deploying this solution, review the overview of the components, the required pre-requisites to deploy the solution and assumptions made in documenting this solution.

### DR Solution Requirements, Pre-requisities and Planning

### Performing DR with SnapCenter

In this solution, SnapCenter provides application-consistent snapshots for SQL Server and Oracle application data. This configuration, together with SnapMirror technology, provides high-speed data replication between our on-premises AFF and FSx ONTAP cluster. Additionally, Veeam Backup & Replication provides backup and restore capabilities for our virtual machines.

In this section, we cover the configuration of SnapCenter, SnapMirror, and Veeam for both backup and restore.

The following sections cover configuration and the steps needed to complete a failover at the secondary site:

### Configure SnapMirror relationships and retention schedules

SnapCenter can update SnapMirror relationships within the primary storage system (primary > mirror) and to secondary storage systems (primary > vault) for the purpose of long-term archiving and retention. To do so, you must establish and initialize a data replication relationship between a destination volume and a source volume using SnapMirror.

The source and destination ONTAP systems must be in networks that are peered using Amazon VPC peering, a transit gateway, AWS Direct Connect, or an AWS VPN.

The following steps are required for setting up SnapMirror relationships between an on-premises ONTAP system and FSx ONTAP:



Refer to the FSx for ONTAP – ONTAP User Guide for more information on creating SnapMirror relationships with FSx.

For the source ONTAP system residing on-premises, you can retrieve the inter-cluster LIF information from System Manager or from the CLI.

1. In ONTAP System Manager, navigate to the Network Overview page and retrieve the IP addresses of Type: Intercluster that are configured to communicate with the AWS VPC where FSx is installed.

Buckets												
Qtrees												
Quotas		Network Interfaces	Portsets									
Storage VHs		4- 444								Q Search	number W Eiter (B Show Fil	-
Tiers		1.100								A DESIGN 2 D	HINDIG. 2 PAULT OF SOUTH PO	
NETWORK		Name	Status	Storage VM	IPspace	Address 0	Current Node	Current Port	Portset	Protocols	Туре	Thre
Overview		veeam_/repo	0	Backup	Default	10.61.181.179	E13A300_1	a0a-181		SMB/CIFS, NFS, S3	Data	0
Ethernet Ports		CM01	0		Default	10.61.181.180	E13A300_1	181-666			Cluster/Node Mgmt	0
FC Ports												1
EVENTS & JOBS	.**	HC_N3	0		Default	10.61.181.183	E13A300_1	261-606			Intercluster, Cluster/Node Mgmt	0
PROTECTION	*	HC_N2	٢		Default	10.61.181.184	E13A300_2	181-60e.			Intercluster, Cluster/Node Mgmt	0
With the second	1222	lif_ora_svm_614	0	ora_tvm	Default	10.61.181.185	E13A300_1	a0a-181		SMB/CIFS, NFS, FL	Data	0

2. To retrieve the Intercluster IP addresses for FSx, log into the CLI and run the following command:

FSx-Dest::> network interface show -role intercluster

	Logical	Status	Network	Current	Current	Is
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port	Home
FsxId0ae40e	08acc0dea67					
	inter 1	up/up	172.30.15.42/25	FsxId0ae40e08	acc0dea6'	7-01
					e0e	true
	inter 2	up/up	172.30.14.28/26	FsxId0ae40e08	acc0dea6	7-02
					e0e	true

To establish cluster peering between ONTAP clusters, a unique passphrase entered at the initiating ONTAP cluster must be confirmed in the other peer cluster.

1. Set up peering on the destination FSx cluster using the cluster peer create command. When prompted, enter a unique passphrase that is used later on the source cluster to finalize the creation process.

```
FSx-Dest::> cluster peer create -address-family ipv4 -peer-addrs
source_intercluster_1, source_intercluster_2
Enter the passphrase:
Confirm the passphrase:
```

2. At the source cluster, you can establish the cluster peer relationship using either ONTAP System Manager or the CLI. From ONTAP System Manager, navigate to Protection > Overview and select Peer Cluster.



- 3. In the Peer Cluster dialog box, fill out the required information:
  - a. Enter the passphrase that was used to establish the peer cluster relationship on the destination FSx cluster.

- b. Select Yes to establish an encrypted relationship.
- c. Enter the intercluster LIF IP address(es) of the destination FSx cluster.
- d. Click Initiate Cluster Peering to finalize the process.

Peer Cluster

Local	Rer
STORAGE VM PERMISSIONS	PASSPHRASE 🕐
All storage VMs (incl ×	•••••
Storage VMs created in the future also will be given permissions.	It cannot be determined from the passphrase wheth this relationship was encrypted. Is the relationship encrypted?
2	Yes No
	To generate passphrase, Launch Remote Cluste
	Intercluster Network Interfaces IP Addresses
	172.30.15.42
	172.30.14.28
	Cancel
	+ Add
4	

4. Verify the status of the cluster peer relationship from the FSx cluster with the following command:



#### Establish SVM peering relationship

The next step is to set up an SVM relationship between the destination and source storage virtual machines that contain the volumes that will be in SnapMirror relationships.

1. From the source FSx cluster, use the following command from the CLI to create the SVM peer relationship:

```
FSx-Dest::> vserver peer create -vserver DestSVM -peer-vserver
Backup -peer-cluster OnPremSourceSVM -applications snapmirror
```

- 2. From the source ONTAP cluster, accept the peering relationship with either ONTAP System Manager or the CLI.
- 3. From ONTAP System Manager, go to Protection > Overview and select Peer Storage VMs under Storage VM Peers.



- 4. In the Peer Storage VM's dialog box, fill out the required fields:
  - The source storage VM
  - The destination cluster
  - The destination storage VM



5. Click Peer Storage VMs to complete the SVM peering process.

SnapCenter manages retention schedules for backups that exist as snapshot copies on the primary storage system. This is established when creating a policy in SnapCenter. SnapCenter does not manage retention policies for backups that are retained on secondary storage systems. These policies are managed separately through a SnapMirror policy created on the secondary FSx cluster and associated with the destination volumes that are in a SnapMirror relationship with the source volume.

When creating a SnapCenter policy, you have the option to specify a secondary policy label that is added to the SnapMirror label of each snapshot generated when a SnapCenter backup is taken.



On the secondary storage, these labels are matched to policy rules associated with the destination volume for the purpose of enforcing retention of snapshots.

The following example shows a SnapMirror label that is present on all snapshots generated as part of a policy used for daily backups of our SQL Server database and log volumes.

Coloct cocondan		ention	antions	0
Select secondary	repli	cation	options	

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label	Custom Label - 1
	sql-daily
Error retry count	3 0

For more information on creating SnapCenter policies for a SQL Server database, see the SnapCenter documentation.

You must first create a SnapMirror policy with rules that dictate the number of snapshot copies to retain.

1. Create the SnapMirror Policy on the FSx cluster.

FSx-Dest::> snapmirror policy create -vserver DestSVM -policy
PolicyName -type mirror-vault -restart always

2. Add rules to the policy with SnapMirror labels that match the secondary policy labels specified in the SnapCenter policies.

```
FSx-Dest::> snapmirror policy add-rule -vserver DestSVM -policy
PolicyName -snapmirror-label SnapMirrorLabelName -keep
#ofSnapshotsToRetain
```

The following script provides an example of a rule that could be added to a policy:

```
FSx-Dest::> snapmirror policy add-rule -vserver sql_svm_dest -policy
Async SnapCenter SQL -snapmirror-label sql-ondemand -keep 15
```



Create additional rules for each SnapMirror label and the number of snapshots to be retained (retention period).

#### **Create destination volumes**

To create a destination volume on FSx that will be the recipient of snapshot copies from our source volumes, run the following command on FSx ONTAP:

```
FSx-Dest::> volume create -vserver DestSVM -volume DestVolName
-aggregate DestAggrName -size VolSize -type DP
```

Create the SnapMirror relationships between source and destination volumes

To create a SnapMirror relationship between a source and destination volume, run the following command on FSx ONTAP:

```
FSx-Dest::> snapmirror create -source-path
OnPremSourceSVM:OnPremSourceVol -destination-path DestSVM:DestVol -type
XDP -policy PolicyName
```

#### Initialize the SnapMirror relationships

Initialize the SnapMirror relationship. This process initiates a new snapshot generated from the source volume and copies it to the destination volume.

FSx-Dest::> snapmirror initialize -destination-path DestSVM:DestVol

Deploy and configure Windows SnapCenter server on-premises.

### Deploy Windows SnapCenter Server on premises

This solution uses NetApp SnapCenter to take application-consistent backups of SQL Server and Oracle databases. In conjunction with Veeam Backup & Replication for backing up virtual machine VMDKs, this provides a comprehensive disaster recovery solution for on-premises and cloud-based datacenters.

SnapCenter software is available from the NetApp support site and can be installed on Microsoft Windows systems that reside either in a domain or workgroup. A detailed planning guide and installation instructions can be found at the NetApp Documentation Center.

The SnapCenter software can be obtained at this link.

After it is installed, you can access the SnapCenter console from a web browser using *https://Virtual\_Cluster\_IP\_or\_FQDN:8146*.

After you log into the console, you must configure SnapCenter for backup SQL Server and Oracle databases.

To add storage controllers to SnapCenter, complete the following steps:

1. From the left menu, select Storage Systems and then click New to begin the process of adding your storage controllers to SnapCenter.

	NetApp SnapC	enter®	Þ		•	≅ 0	👤 scadmin Snap(	CenterAdmin 🛛 🖡 Sign Out
<		ONTA	P Storage					-
	Dashboard	Туре	ONTAP SVMs	• Search	n by Name		$\supset$	New Delars
	Resources	ONTA	AP Storage Connectio	ns				
-	Monitor		Name 41	IP	Cluster Name	User Nam	e Platform	Controller License
<b>a</b>	Reports		Backup	172.16.13.17	172.16.13.17		AFF	~
			<u>FS02</u>	172.16.13.17	172.16.13.17		AFF	×
•	Hosts		ora_svm	172.16.13.17	172.16.13.17		AFF	~
ł.	Storage Systems		ora svm dest		172.30.15.42		AFF	Not applicable
=	Settings		<u>sql_svm</u>	172.16.13.17	172.16.13.17		AFF	~
			sol svm_dest		172.30.15.42		AFF	Not applicable
A	Alerts		svm_HCApps		172.30.15.42		AFF	Not applicable

2. In the Add Storage System dialog box, add the management IP address for the local on-premises ONTAP cluster and the username and password. Then click Submit to begin discovery of the storage system.

Add Storage System	
Add Storage System	•
Storage System	10.61.181.180
Username	admin
Password	•••••
Event Management	System (EMS) & AutoSupport Settings
Send AutoSuppor	t notification to storage system
Log SnapCenter S	erver events to syslog
Dire Options : Pl	atform, Protocol, Preferred IP etc
Submit Cancel 3. Repeat this process to add the at the bottom of the Add Storag the FSx system as the secondar backup snapshots.	Reset Sx ONTAP system to SnapCenter. In this case, select More Options e System window and click the check box for Secondary to designate any storage system updated with SnapMirror copies or our primary

Platform	FAS	•	Secondary 🚯	
Protocol	HTTPS	•		
Port	443			
Timeout	60	seconds	0	
				0

### Add hosts to SnapCenter

The next step is adding host application servers to SnapCenter. The process is similar for both SQL Server and Oracle.

- 1. From the left menu, select Hosts and then click Add to begin the process of adding storage controllers to SnapCenter.
- 2. In the Add Hosts window, add the Host Type, Hostname, and the host system Credentials. Select the plug-in type. For SQL Server, select the Microsoft Windows and Microsoft SQL Server plug-in.

II Ne	etApp	• SnapCenter®					
>	Man	aged Hosts					
	Se	arch by Name		Add Host			
<b>I</b>		Name	臣	Host Type	Windows	•	
•		oraclesry_01.sddc.netapp.com		Host Name	sqlsrv-01.sddc.netapp.com		
		oraclesry 02.sddc.netapp.com		Credentials	sddc-jpowell	•	+
â		oraclesry_03.sddc.netapp.com					
A		oraclesrv_04.sddc.netapp.com		Select Plug-ins to In	stall SnapCenter Plug-ins Package 4.6 for Windo	ws	
ła –		oraclesry_05.sddc.netapp.com			Microsoft Windows		
		oraclesry_06.sddc.netapp.com			Microsoft SQL Server		
		oraclesry_07.sddc.netapp.com			Microsoft Exchange Server     SAP HANA		
		oraclesry_08.sddc.netapp.com		More Options : Po	ort, gMSA, Install Path, Custom Plug-Ins		
		oraclesry_09.sddc.netapp.com			94 (1997) - 1999 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 - 1997 -		
		oraclesry 10.sddc.netapp.com		Submit Cancel	]		

3. For Oracle, fill out the required fields in the Add Host dialog box and select the check box for the Oracle Database plug-in. Then click Submit to begin the discovery process and to add the host to SnapCenter.

Host Type	Linux	•	
Host Name	oraclesrv_11.sddc.netapp.com		
Credentials	root		
Select Plug-ins to In:	Stall SnapCenter Plug-ins Package 4.6 for Linux		e
Select Plug-ins to In:	SAP HANA		e
Select Plug-ins to In:	stall SnapCenter Plug-ins Package 4.6 for Linux         Oracle Database         SAP HANA         ort, Install Path, Custom Plug-Ins		

Policies establish the specific rules to be followed for a backup job. They include, but are not limited to, the backup schedule, replication type, and how SnapCenter handles backing up and truncating transaction logs.

You can access	policies ir	n the Settings	section of the	SnapCenter	web client.

n NetApp SnapCenter® 🔹 🚱 🗵										
<		Global Settings	Policies	Users and Access	Roles C	redential	Software			
	Dashboard									
<b>V</b>	Resources	Search by Name	2			New		lodity		
	Monitor	Name	44	Backup Type	Schedul	е Туре	Re	plication		
~2		SQL-Daily		Full and Log backup	Daily		Sn	apVault		
â	Reports	SQL-Hourly		Full and Log backup	Hourly		Sn	apVault		
A	Hosts	SQL-Hourly-Logs		Log backup	Hourly		Sn	apVault		
÷.	Storage Systems	SQL-OnDemand		Full and Log backup	On dema	and	Sn	apVault		
-		SQL-Weekly		Full and Log backup	Weekly		Sn	apVault		
-	Settings									
A	Alerts									

For complete information on creating policies for SQL Server backups, see the SnapCenter documentation.

For complete information on creating policies for Oracle backups, see the SnapCenter documentation.

## Notes:

- As you progress through the policy creation wizard, take special note of the Replication section. In this section you stipulate the types of secondary SnapMirror copies that you want taken during the backups process.
- The "Update SnapMirror after creating a local Snapshot copy" setting refers to updating a SnapMirror relationship when that relationship exists between two storage virtual machines residing on the same cluster.
- The "Update SnapVault after creating a local SnapShot copy" setting is used to update a SnapMirror relationship that exists between two separate cluster and between an on-premises ONTAP system and Cloud Volumes ONTAP or FSxN.

The following image shows the preceding options and how they look in the backup policy wizard.

Name	ACT 127 24 129	251 757				
Name	Select secondary replication options ()					
2 Backup Type	Update SnapMirror after creating a local Snapshot copy.					
	-					
3 Retention	Update SnapVault aft	er creating a	local Snaps	hot copy.		
3 Retention 4 Replication	Update SnapVault aft	er creating a	local Snaps	hot copy.	0	
3 Retention 4 Replication	Update SnapVault aft Secondary policy label Error retry count	er creating a Choose 3 0	local Snaps	hot copy.	0	

## Create SnapCenter Resource Groups

Resource Groups allow you to select the database resources you want to include in your backups and the policies followed for those resources.

- 1. Go to the Resources section in the left-hand menu.
- 2. At the top of the window, select the resource type to work with (In this case Microsoft SQL Server) and then click New Resource Group.

NetApp SnapCenter®							e scade	nin SnapCenterA	dmin 🛛 🗊 Sign Out
<		Micros	oft SQL Server						
	Dashboard	View	Resource Group	•	earch by name	e	Y	2	New Resource Group
0	Resources	19	Name	Resource Count	Tags		Policies	Last Backup	Overall Status
•	Monitor		SQLSRV-01	1			SQL-Daily SQL-Hourly	05/11/2022	Completed
ай	Reports						SQL- OnDemand	and	
A	Hosts		501 501 00				SQL-Weekiy	02/20/2022	Porto a
ł.	Storage Systems		SQLSKV-U2	1:			SQL-Daily 03/28/20 SQL-Hourly SQL-	03/28/2022	Falled
÷	Settings					O		OnDemand SQL-Weekly	
▲	Alerts		SQLSRV-03	1	1		SQL-Daily	05/11/2022	Completed

The SnapCenter documentation covers step-by-step details for creating Resource Groups for both SQL Server and Oracle databases.

For backing up SQL resources, follow this link.

For Backing up Oracle resources, follow this link.

## Deploy and configure Veeam Backup Server

Veeam Backup & Replication software is used in the solution to back up our application virtual machines and archive a copy of the backups to an Amazon S3 bucket using a Veeam scale-out backup repository (SOBR). Veeam is deployed on a Windows server in this solution. For specific guidance on deploying Veeam, see the Veeam help Center Technical documentation.

After you deploy and license the software, you can create a scale-out backup repository (SOBR) as target storage for backup jobs. You should also include an S3 bucket as a backup of VM data offsite for disaster recovery.

See the following prerequisites before getting started.

- 1. Create an SMB file share on your on-premises ONTAP system as the target storage for backups.
- 2. Create an Amazon S3 bucket to include in the SOBR. This is a repository for the offsite backups.
First, add the ONTAP storage cluster and associated SMB/NFS filesystem as storage infrastructure in Veeam.

1. Open the Veeam console and log in. Navigate to Storage Infrastructure and then select Add Storage.



- 2. In the Add Storage wizard, select NetApp as the storage vendor and then select Data ONTAP.
- 3. Enter the management IP address and check the NAS Filer box. Click Next.

	Data ONTAP storage by specifying DNS name or IP address.	
Name	Management server DNS name or IP address:	
Nume	10.61.181.180	
Credentials	Description	
NAS Filer	Created by SDDC\ipowell at 5/17/2022 10:34 AM.	
Analy		
Арріу		
Summary	Role: Block or file storage for VMware vSphere Block storage for Microsoft Windows servers NAS filer	
Add your credentials	< Previous Next > Finish Cances the ONTAP cluster.	el
New NetApp Data ONTAP Sto Credentials Specify account w	orage with storage administrator privileges.	
Name	Credentials:	
Name	Credentials:                HCIEUC\Admin (HCIEUC\Admin, last edited: 98 days ago)	
Name Credentials	Credentials:          Image accounts             Manage accounts	
Name Credentials NAS Filer	Credentials:  Credentials:  Add  Manage accounts  Protocol: HTTPS ~	
Name Credentials NAS Filer Apply	Credentials: Credentials: Credentials: Add Manage accounts Protocol: HTTPS ~ Port: 443	•••
Name Credentials NAS Filer Apply	Credentials: Credentials: Credentials: Add Manage accounts Protocol: HTTPS ~ Port: 443 •	
Name Credentials NAS Filer Apply Summary	Credentials: Credentials: Add Manage accounts Protocol: HTTPS ~ Port: 443	
Name Credentials NAS Filer Apply Summary	Credentials: HCIEUC\Admin (HCIEUC\Admin, last edited: 98 days ago) Add Manage accounts Protocol: HTTPS  Port: 443	
Name Credentials NAS Filer Apply Summary	Credentials: HCLEUC\Admin (HCLEUC\Admin, last edited: 98 days ago) Add Manage accounts Protocol: HTTPS ~ Port: 443 •	
Name Credentials NAS Filer Apply Summary	Credentials:	
Name Credentials NAS Filer Apply Summary	Credentials:	
Name Credentials NAS Filer Apply Summary	Credentials:	•••
Name Credentials NAS Filer Apply Summary	Credentials:    Manage accounts    Protocol:   HTTPS ✓   Port:   443	
Name Credentials NAS Filer Apply Summary	Credentials:      MclEUC\Admin (HClEUC\Admin, last edited: 98 days ago)      Manage accounts   Porte    HTTPS →    Port:    443	
Name Credentials NAS Filer Apply Summary	Credentials:    McLEUC\Admin (HCLEUC\Admin, last edited: 98 days ago)    Manage accounts   Protocol: HTTPS ✓ Port: 443	•••

New NetApp Data ONTAP St	orage	×
NAS Filer Specify how this	storage can be accessed by file backup jobs.	
Name	Protocol to use:	
Credentials		
NAS Filer	Create required export rules automatically	
	Volumes to scan:	
Apply	All volumes	Choose
Summary	Backup proxies to use:	
	Automatic selection	Choose
	< Previous Apply Finish	Cancel

6. Complete the Apply and Summary pages of the wizard and click Finish to begin the storage discovery process. After the scan completes, the ONTAP cluster is added along with the NAS filers as available resources.



7. Create a backup repository using the newly discovered NAS shares. From Backup Infrastructure, select Backup Repositories and click the Add Repository menu item.

	Backup Proxies
	Backup Repositories 2
	External Repositories
Þ	Scale-out Repositories
	WAN Accelerators
	Service Providers
4	🕂 SureBackup
	🕂 Application Groups
	📇 Virtual Labs
4	Canaged Servers
	VMware vSphere
	📑 Microsoft Windows

8. Follow all steps in the New Backup Repository Wizard to create the repository. For detailed information on creating Veeam Backup Repositories, see the Veeam documentation.

New Backup Repository

### Share

Type in UNC path to share (mapped drives are not supported), specify share access credentials and how backup jobs should write data to this share.

Name	Shared folder:
Share	Use \\server\folder format
Repository	This share requires access credentials:
Mount Server	R sddc\administrator (sddc\administrator, last edited: 85 days ago)
Review Apply Summary	Gateway server: Automatic selection The following server: veeam.sddc.netapp.com (Backup server) Use this option to improve performance and reliability of backup to a NAS located in a remote site.
	< Previous Next > Finish Cancel

 $\times$ 

The next step is to add the Amazon S3 storage as a backup repository.

1. Navigate to Backup Infrastructure > Backup Repositories. Click Add Repository.



2. In the Add Backup Repository wizard, select Object Storage and then Amazon S3. This starts the New Object Storage Repository wizard.

# Add Backup Repository

Select the type of backup repository you want to add.

_		
		~
_	_	
_	_	-1
_	-	-

Direct attached storage

Microsoft Windows or Linux server with internal or direct attached storage. This configuration enables data movers to run directly on the server, allowing for fastest performance.



Network attached storage

Network share on a file server or a NAS device. When backing up to a remote share, we recommend that you select a gateway server located in the same site with the share.



Þ

Deduplicating storage appliance

Dell EMC Data Domain, ExaGrid, HPE StoreOnce or Quantum DXi. If you are unable to meet the requirements of advanced integration via native appliance API, use the network attached storage option instead.



On-prem object storage system or a cloud object storage provider. Object storage can only be used as a Capacity Tier of scale-out backup repositories, backing up directly to object storage is not currently supported.

- 3. Provide a name for your object storage repository and click Next.
- 4. In the next section, provide your credentials. You need an AWS Access Key and Secret Key.

Account ACCOUNT AKIAX4H43ZT557HXQT2W (last edited: 107 days ago)	V Add
Manage cloud	loud accounts
Bucket AWS region:	
Global	
□ Use the following gateway server:	
Use the following gateway server:	
Use the following gateway server:	
Use the following gateway server: veeam.sddc.netapp.com (Backup server)	
Use the following gateway server: veeam.sddc.netapp.com (Backup server) Select a gateway server to proxy access to Amazon S3. If no gateway server	erver is specified, all scale-
Use the following gateway server: veeam.sddc.netapp.com (Backup server) Select a gateway server to proxy access to Amazon S3. If no gateway server backup reporting electron much bare direct Internet access	erver is specified, all scale-
Use the following gateway server: veeam.sddc.netapp.com (Backup server) Select a gateway server to proxy access to Amazon S3. If no gateway server backup repository extents must have direct Internet access.	erver is specified, all scale-
<ul> <li>Use the following gateway server:</li> <li>veeam.sddc.netapp.com (Backup server)</li> <li>Select a gateway server to proxy access to Amazon S3. If no gateway server backup repository extents must have direct Internet access.</li> </ul>	erver is specified, all scale-
Use the following gateway server: veeam.sddc.netapp.com (Backup server) Select a gateway server to proxy access to Amazon S3. If no gateway server backup repository extents must have direct Internet access.	erver is specified, all scale-
Use the following gateway server: veeam.sddc.netapp.com (Backup server) Select a gateway server to proxy access to Amazon S3. If no gateway server backup repository extents must have direct Internet access.	erver is specified, all scale-

×

Now that we have added our storage repositories to Veeam, we can create the SOBR to automatically tier backup copies to our offsite Amazon S3 object storage for disaster recovery.

1. From Backup Infrastructure, select Scale-out Repositories and then click the Add Scale-out Repository menu item.



- 2. In the New Scale-out Backup Repository provide a name for the SOBR and click Next.
- 3. For the Performance Tier, choose the backup repository that contains the SMB share residing on your local ONTAP cluster.

New Scale-out Backup Reposito	ry	×
Performance Tier Select backup repos	tories to use as the landing zone and for the short-term retention.	
Name	Extents:	Add
Performance Tier	Name	Add
Placement Policy		Kemove

- 4. For the Placement Policy, choose either Data Locality or Performance based your requirements. Select next.
- 5. For Capacity Tier we extend the SOBR with Amazon S3 object storage. For the purposes of disaster recovery, select Copy Backups to Object Storage as Soon as They are Created to ensure timely delivery of our secondary backups.

New Scale-out Backup Repo	isitory		
Capacity Tier Specify object st completely to re	orage to copy backups to for redundancy and DR purposes. Older backups can be moved to object duce long-term retention costs while preserving the ability to restore directly from offloaded back	t storage ups.	
Name	Extend scale-out backup repository capacity with object storage:		
Performance Tier	Amazon S3 Repo 🗸 🗸	Add	
Placement Policy	Define time windows when uploading to capacity tier is allowed	Window	
Capacity Tier	<ul> <li>Copy backups to object storage as soon as they are created Create additional copy of your backups for added redundancy by having all backups copied to the capacity tier as soon as they are created on the performance tier.</li> <li>Move backups to object storage as they age out of the operational restore window Reduce your long-term retention costs by moving older backups to object storage completely while preserving the ability to restore directly from offloaded backups.</li> </ul>		
Archive Tier			
Summary			
	Move backup files older than 14 📥 days (your operational restore window)	Override	
	Encrypt data uploaded to object storage     Password:     Manage passwords	Add	
	ivianage passwords		
	< Previous Next > Finish	Cancel	
-inally, select Apply	and Finish to finalize creation of the SOBR.		

### Create the scale-out backup repository jobs

The final step to configuring Veeam is to create backup jobs using the newly created SOBR as the backup destination. Creating backup jobs is a normal part of any storage administrator's repertoire and we do not cover the detailed steps here. For more complete information on creating backup jobs in Veeam, see the Veeam Help Center Technical Documentation.

### BlueXP backup and recovery tools and configuration

To conduct a failover of application VMs and database volumes to VMware Cloud Volume services running in AWS, you must install and configure a running instance of both SnapCenter Server and Veeam Backup and Replication Server. After the failover is complete, you must also configure these tools to resume normal backup operations until a failback to the on-premises datacenter is planned and executed.

### Deploy secondary Windows SnapCenter Server

SnapCenter Server is deployed in the VMware Cloud SDDC or installed on an EC2 instance residing in a VPC with network connectivity to the VMware Cloud environment.

SnapCenter software is available from the NetApp support site and can be installed on Microsoft Windows systems that reside either in a domain or workgroup. A detailed planning guide and installation instructions can be found at the NetApp documentation center.

You can find the SnapCenter software at this link.

### Configure secondary Windows SnapCenter Server

To perform a restore of application data mirrored to FSx ONTAP, you must first perform a full restore of the on-premises SnapCenter database. After this process is complete, communication with the VMs is reestablished and application backups can now resume using FSx ONTAP as the primary storage.

To achieve this, you must complete the following items on the SnapCenter Server:

- 1. Configure the computer name to be identical to the original on-premises SnapCenter Server.
- 2. Configure networking to communicate with VMware Cloud and the FSx ONTAP instance.
- 3. Complete the procedure to restore the SnapCenter database.
- 4. Confirm that SnapCenter is in Disaster Recovery mode to make sure that FSx is now the primary storage for backups.
- 5. Confirm that communication is reestablished with the restored virtual machines.

### Deploy secondary Veeam Backup & Replication server

You can install the Veeam Backup & Replication server on a Windows server in the VMware Cloud on AWS or on an EC2 instance. For detailed implementation guidance, see the Veeam Help Center Technical Documentation.

To perform a restore of virtual machines that have been backed up to Amazon S3 storage, you must install the Veeam Server on a Windows server and configure it to communicate with VMware Cloud, FSx ONTAP, and the S3 bucket that contains the original backup repository. It must also have a new backup repository configured on FSx ONTAP to conduct new backups of the VMs after they are restored.

To perform this process, the following items must be completed:

- 1. Configure networking to communicate with VMware Cloud, FSx ONTAP, and the S3 bucket containing the original backup repository.
- 2. Configure an SMB share on FSx ONTAP to be a new backup repository.
- 3. Mount the original S3 bucket that was used as part of the scale-out backup repository on premises.
- 4. After restoring the VM, establish new backup jobs to protect SQL and Oracle VMs.

For more information on restoring VMs using Veeam, see the section "Restore Application VMs with Veeam Full Restore".

### SnapCenter database backup for disaster recovery

SnapCenter allows for the backup and recovery of its underlying MySQL database and configuration data for the purpose of recovering the SnapCenter server in the case of a disaster. For our solution, we recovered the SnapCenter database and configuration on an AWS EC2 instance residing in our VPC. For more information on this step, see this link.

### SnapCenter backup prerequisites

The following prerequisites are required for SnapCenter backup:

- A volume and SMB share created on the on-premises ONTAP system to locate the backed-up database and configuration files.
- A SnapMirror relationship between the on-premises ONTAP system and FSx or CVO in the AWS account. This relationship is used for transporting the snapshot containing the backed-up SnapCenter database and configuration files.
- Windows Server installed in the cloud account, either on an EC2 instance or on a VM in the VMware Cloud SDDC.
- SnapCenter installed on the Windows EC2 instance or VM in VMware Cloud.

- Create a volume on the on-premises ONTAP system for hosting the backup db and config files.
- Set up a SnapMirror relationship between on-premises and FSx/CVO.
- Mount the SMB share.
- Retrieve the Swagger authorization token for performing API tasks.
- Start the db restore process.
- Use the xcopy utility to copy the db and config file local directory to the SMB share.
- On FSx, create a clone of the ONTAP volume (copied via SnapMirror from on-premises).
- Mount the SMB share from FSx to EC2/VMware Cloud.
- Copy the restore directory from the SMB share to a local directory.
- Run the SQL Server restore process from Swagger.

SnapCenter provides a web client interface for executing REST API commands. For information on accessing the REST APIs through Swagger, see the SnapCenter documentation at this link.

After you have navigated to the Swagger page, you must retrieve an authorization token to initiate the database restore process.

1. Access the SnapCenter Swagger API web page at *https://<SnapCenter Server IP>:8146/swagger/*.

SnapCenter A	API <sup>O</sup>	
[ Base URL: /api ]		
https://snapcenter.sddc.netapp.com;81	46/Content/swagger/SnapCenter.yaml	
Manage your SnapCenter Server To access the swagger documenta https://{SCV_hostname}.{SCV_host	using the SnapCenter API. ation of "SnapCenter Plug-in for VMware vSphere" API's, please use st_port}/api/swagger-ui.html	
xpand the Auth section	and click Try it Out.	
xpand the Auth section	and click Try it Out.	
Auth	and click Try it Out.	

3. In the UserOperationContext area, fill in the SnapCenter credentials and role and click Execute.

The		
IOKENNEVEREXPIRES boolean	Token never expires	
(query)	false v	
UserOperationContext * re object	uired User credentials	
(body)	Edit Value Model	
	<pre>{     "UserOperationContext": {         "UJer": {             "Name": "localhost\\scadmin",             "Passphrase": "NetApp321",             "Rolename": "SnapCenterAdmin"         }     } }</pre>	
		li.
	Cancel	
	Parameter content type application/json	

4. In the Response body below, you can see the token. Copy the token text for authentication when executing the backup process.

200	Response body
	"PlurinHama": null
	"Bostd": 0
	"RoleId": null,
	"JobIds": null
	$\mathbf{h}$
	"User": {
	"Token":
	$\label{eq:started} $$ $$ $$ $$ $$ $$ $$ $$ $$ $$ $$ $$ $$$
	$\label{eq:clfgrapg1} CLfgrapg1GmcagT08bgb5bMTx07EcdrAidzAXUDb3GyLCKtW0GdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQvo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQvo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQvo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQvo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQvo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQwo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQwo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQwo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQwo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQwo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQwo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQwo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQwo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQwo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQwo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQwo4v4RL0hhThhwPhVEWVGdwKzSeUwKj3uVupnk1E31skK6PRBv9RS8j0gHQwo4v4RL0hhThhwPhVEWVGdwKzSeUwKk6PRBv9RS8j0gHQwa4v4RL0hhThhwPhVEWVGdwKzSeUwKj4kgwkavkgkavkgkavkgkavkgkavkgkavkgkavkgka$
	9/23nFeJVF/p1Ev4vrV/zeZVTUHFHUN069XRe5cuW9nwyj4b015Y5FN3XDkjQ
	"Name": "SCAdmin",
	"TokenHashed": null,
	"Type": "",
	"TokenTime": "2022-03-22T14:21:57.3665661-07:00",
	"Id": "1",
	"FullName": "SCAdmin",
	"Host": null,
	"Author": null,
	"UserName": "",
	"Dowain": "", Downig
	"Passphrase": "",

Next go to the Disaster Recovery area on the Swagger page to begin the SnapCenter backup process.

1. Expand the Disaster Recovery area by clicking it.

# Disaster Recovery GET /4.6/disasterrecovery/server/backup POST /4.6/disasterrecovery/server/backup DELETE /4.6/disasterrecovery/server/backup Deletes the existing Snapcenter DR backup. POST /4.6/disasterrecovery/server/backup Deletes /4.6/disasterrecovery/server/restore Stats Stats POST /4.6/disasterrecovery/server/restore Stats Stats POST /4.6/disasterrecovery/server/restore Stats Stats

2. Expand the /4.6/disasterrecovery/server/backup section and click Try it Out.



3. In the SmDRBackupRequest section, add the correct local target path and select Execute to start the backup of the SnapCenter database and configuration.



The backup process does not allow backing up directly to an NFS or CIFS file share.

Name	Description
Token * required string	User authorization token
(header)	TUHFHUM069XRe5cuW9nwyj4b0I5Y5FN3XDkjQ==
SmDRBackupRequest * required	Parameters to take Backup
(body)	Edit Value Model
	<pre>{     "TargetPath": "C:\\SnapCenter_Backups\\" }</pre>
	Cancel Parameter content type application/json
	Execute

Log into SnapCenter to review log files when starting the database restore process. Under the Monitor section, you can view the details of the SnapCenter server disaster recovery backup.

Snap	Center Server disaster recovery backup	
~ *	SnapCenter Server disaster recovery backup	×
~	Precheck validation	
~	Disaster recovery backup of 'oraclesrv_04.sddc.netapp.com'	
~	Disaster recovery backup of SnapCenter Server 'SnapCenter.sddc.netapp.com'	
~	Disaster recovery backup of 'oraclesrv_02.sddc.netapp.com'	
~	Disaster recovery backup of 'oraclesrv_03.sddc.netapp.com'	
~	Disaster recovery backup of 'oraclesrv_05.sddc.netapp.com'	
~	Disaster recovery backup of 'oraclesrv_07.sddc.netapp.com'	
~	Disaster recovery backup of 'sqlsrv-02.sddc.netapp.com'	
4	Disaster recovery backup of 'sqlsrv-03.sddc.netapp.com'	
~	Disaster recovery backup of 'oraclesrv_10.sddc.netapp.com'	
1	Disaster recovery backup of 'sqlsrv-04.sddc.netapp.com'	
~	Disaster recovery backup of 'sqlsrv-01.sddc.netapp.com'	
~	Disaster recovery backup of 'sqlsrv-05.sddc.netapp.com'	
~	Disaster recovery backup of 'oraclesrv_09.sddc.netapp.com'	
~	Disaster recovery backup of 'sqlsrv-06.sddc.netapp.com'	
~	Disaster recovery backup of 'sqlsrv-07.sddc.netapp.com'	
<b>)</b> Tas :27:4	k Name: SnapCenter Server disaster recovery backup Start Time: 03/23/2022 10:27:11 AM End Time: 03/23/20 7 AM	22

Next you must move the backup from the local drive on the SnapCenter server to the CIFS share that is used to SnapMirror copy the data to the secondary location located on the FSx instance in AWS. Use xcopy with specific options that retain the permissions of the files.

Open a command prompt as Administrator. From the command prompt, enter the following commands:

```
xcopy <Source_Path> \\<Destination_Server_IP>\<Folder_Path> /O /X
/E /H /K
xcopy c:\SC_Backups\SnapCenter_DR \\10.61.181.185\snapcenter_dr /O
/X /E /H /K
```

### Failover

### Disaster occurs at primary site

For a disaster that occurs at the primary on-premises datacenter, our scenario includes failover to a secondary site residing on Amazon Web Services infrastructure using VMware Cloud on AWS. We assume that the virtual machines and our on-premises ONTAP cluster are no longer accessible. In addition, both the SnapCenter and Veeam virtual machines are no longer accessible and must be rebuilt at our secondary site.

This section address failover of our infrastructure to the cloud, and we cover the following topics:

- SnapCenter database restore. After a new SnapCenter server has been established, restore the MySQL database and configuration files and toggle the database into disaster recovery mode in order to allow the secondary FSx storage to become the primary storage device.
- Restore the application virtual machines using Veeam Backup & Replication. Connect the S3 storage that contains the VM backups, import the backups, and restore them to VMware Cloud on AWS.
- Restore the SQL Server application data using SnapCenter.
- Restore the Oracle application data using SnapCenter.

SnapCenter supports disaster recovery scenarios by allowing the backup and restore of its MySQL database and configuration files. This allows an administrator to maintain regular backups of the SnapCenter database at the on-premises datacenter and later restore that database to a secondary SnapCenter database.

To access the SnapCenter backup files on the remote SnapCenter server, complete the following steps:

- 1. Break the SnapMirror relationship from the FSx cluster, which makes the volume read/write.
- 2. Create a CIFS server (if necessary) and create a CIFS share pointing to the junction path of the cloned volume.
- 3. Use xcopy to copy the backup files to a local directory on the secondary SnapCenter system.
- 4. Install SnapCenter v4.6.
- 5. Ensure that SnapCenter server has the same FQDN as the original server. This is required for the db restore to be successful.

To start the restore process, complete the following steps:

- 1. Navigate to the Swagger API web page for the secondary SnapCenter server and follow the previous instructions to obtain an authorization token.
- 2. Navigate to the Disaster Recovery section of the Swagger page, select /4.6/disasterrecovery/server/restore, and click Try it Out.

POST	/4.6/disasterrecovery/server/restore Starts SnapCenter Server Restore.	
Starts SnapC	Center Server Restore.	
Parameters		Try it out

3. Paste in your authorization token and, in the SmDRResterRequest section, paste in the name of the backup and the local directory on the secondary SnapCenter server.

Description
Liser authorization token
KIYxOg==rMXzS7EPIGRzTXjfton6Q+JoNGpueQt
Parameters to take for Pastore
Parallelets to take for Restore
Edit Value Model
<pre>{     "BackupName": "SnapCenter.sddc.netapp.com_03-23-2022_12.38.00.6713",     "BackupPath": "C:\\SnapCenter\\" }</pre>

4. Select the Execute button to start the restore process.

5. From SnapCenter, navigate to the Monitor section to view the progress of the restore job.

	letApp Snap(	Center®		
<		Jobs	Schedules	Events Logs
	Dashboard	search	n by name	
0	Resources	Jobs - F	Filter	
•	Monitor	ID	Status	Name
<b>a</b>	Reports	20482	4	SnapCenter Server Disaster Recovery
		20481	~	SnapCenter Server disaster recovery backup
<b>A</b>	Hosts	20480	×	SnapCenter Server disaster recovery backup
80	Storage Systems	20475	~	Backup of Resource Group 'SQLSRV-09' with policy 'SQL-Hourly'
=	Settings	20474	~	Backup of Resource Group 'SQLSRV-05' with policy 'SQL-Hourly'
		20473	3	Backup of Resource Group 'OracleSrv_06' with policy 'Oracle-Hourly'
	Alerts	20472	×	SnapCenter Server disaster recovery backup

# Job Details

SnapCenter Server Disaster Recovery

- SnapCenter Server Disaster Recovery
- Prepare for restore job
- Precheck validation
- Saving original server state
- Schedule restore
- Repository restore
- Config restore
  - Reset MySQL password
- 6. To enable SQL Server restores from secondary storage, you must toggle the SnapCenter database into Disaster Recovery mode. This is performed as a separate operation and initiated on the Swagger API web page.
  - a. Navigate to the Disaster Recovery section and click /4.6/disasterrecovery/storage.
  - b. Paste in the user authorization token.
  - c. In the SmSetDisasterRecoverySettingsRequest section, change EnableDisasterRecover to true.
  - d. Click Execute to enable disaster recovery mode for SQL Server.

(header)       KIYxOg==rMXzS7EPIGRzTXjfton6Q+JoNGpueQt         SmSetDisasterRecoverySettingsRequest * required       Parameters to enable or disable the DR mode         object       Edit Value Model         (body)       { "EnableDisasterRecovery": true }	Token * <sup>required</sup>	User authorization token
SmSetDisasterRecoverySettingsRequest * required       Parameters to enable or disable the DR mode         object       Edit Value         (body)       Edit Value         * EnableDisasterRecovery*: true	(header)	KIYxOg==rMXzS7EPIGRzTXjfton6Q+JoNGpueQt
(body) Edit Value Model { "EnableDisasterRecovery": true }	SmSetDisasterRecoverySettingsRequest * required	Parameters to enable or disable the DR mode
<pre>{     "EnableDisasterRecovery": true }</pre>	(body)	Edit Value Model
		<pre>{     "EnableDisasterRecovery": true }</pre>

Restore application VMs with Veeam full restore

From the secondary Veeam server, import the backups from S3 storage and restore the SQL Server and Oracle VMs to your VMware Cloud cluster.

To import the backups from the S3 object that was part of the on-premises scale-out backup repository, complete the following steps:

1. Go to Backup Repositories and click Add Repository in the top menu to launch the Add Backup Repository wizard. On the first page of the wizard, select Object Storage as the backup repository type.

Add B Select the	ackup Repository type of backup repository you want to add.
0000	Direct attached storage Microsoft Windows or Linux server with internal or direct attached storage. This configuration enables data movers to run directly on the server, allowing for fastest performance.
	Network attached storage Network share on a file server or a NAS device. When backing up to a remote share, we recommend that you select a gateway server located in the same site with the share.
¥	Deduplicating storage appliance Dell EMC Data Domain, ExaGrid, HPE StoreOnce or Quantum DXi. If you are unable to meet the requirements of advanced integration via native appliance API, use the network attached storage option instead.
	Object storage On-prem object storage system or a cloud object storage provider. Object storage can only be used as a Capacity Tier of scale-out backup repositories, backing up directly to object storage is not currently supported.

2. Select Amazon S3 as the Object Storage type.





4. Select your pre-entered credentials from the drop-down list or add a new credential for accessing the cloud storage resource. Click Next to continue.

Name       Credentials:       Add         Account       Manage cloud accounts       Add         Bucket       AWS region:       Global         Summary       Global       Image cloud accounts         Use the following gateway server:       EC2AMAZ-3POTKQV (Backup server)       Select a gateway server to proxy access to Amazon S3. If no gateway server is specified, all scale-backup repository extents must have direct Internet access.	Specify AV	/S account to use for connecting to Amazon S3 storage bucket.
Account Accoun	Name	Credentials:
Account       Manage cloud accounts         Bucket       AWS region:         Global       Global	Account	👫 AKIAX4H43ZT53YJXPY2Y (last edited: 33 days ago) 🗸 Add
ucket       AWS region:         ummary       Global         Use the following gateway server:         EC2AMAZ-3POTKQV (Backup server)         Select a gateway server to proxy access to Amazon S3. If no gateway server is specified, all scale-backup repository extents must have direct Internet access.	ccount	Manage cloud accounts
Global         Immany	ucket	AWS region:
Use the following gateway server:  EC2AMAZ-3POTKQV (Backup server)  Select a gateway server to proxy access to Amazon S3. If no gateway server is specified, all scale-backup repository extents must have direct Internet access.		Global
EC2AMAZ-3POTKQV (Backup server) Select a gateway server to proxy access to Amazon S3. If no gateway server is specified, all scale- backup repository extents must have direct Internet access.		
Select a gateway server to proxy access to Amazon S3. If no gateway server is specified, all scale- backup repository extents must have direct Internet access.		Use the following gateway server:
		Use the following gateway server: EC2AMAZ-3POTKQV (Backup server)

5. On the Bucket page, enter the data center, bucket, folder, and any desired options. Click Apply.

Name	Data center:	
	US East (N. Virginia)	
Account	Bucket:	
Bucket	ehcveeamrepo	Browse
Summan	Folder:	
Junning	1000	Browse
	<ul> <li>Limit object storage consumption to: 10 ÷ TB</li> <li>This is a soft limit to help control your object storage spend. If the specified limalready running backup offload tasks will be allowed to complete, but no new</li> <li>Make recent backups immutable for: 30 ÷ days</li> <li>Protects backups from modification or deletion by ransomware, backers or modification or deletion by ransomware.</li> </ul>	nit is exceeded, tasks will be started

To import the backups from the S3 repository that was added in the previous section, complete the following steps.

1. From the S3 backup repository, select Import Backups to launch the Import Backups wizard.



2. After the database records for the import have been created, select Next and then Finish at the summary screen to start the import process.

	while we re preparing object storage repository.	
iport	Message	Duration
immary	Starting infrastructure item update process	0:00:16
	Creating database records for repository	0:00:04

3. After the import is complete, you can restore VMs into the VMware Cloud cluster.

Action type: Configuration Resynchronize Start time: 4/6/2022 3:01:30 PM nitiated by: EC2AMAZ-3POTKQV\vadmin End time: 4/6/2022 3:04:57 PM Log Message Starting backup repositories synchronization Enumerating repositories Found 1 repository Processing capacity tier extent of S3 Backup Repository 2 S3 Backup Repository: added 2 unencrypted Importing backup 2 out of 2	Duration
nitiated by: EC2AMAZ-3POTKQV\vadmin End time: 4/6/2022 3:04:57 PM Log Message Starting backup repositories synchronization Enumerating repositories Found 1 repository Processing capacity tier extent of S3 Backup Repository 2 S3 Backup Repository: added 2 unencrypted Importing backup 2 out of 2	Duration
Log         Message         Starting backup repositories synchronization         Enumerating repositories         Found 1 repository         Processing capacity tier extent of S3 Backup Repository 2         S3 Backup Repository: added 2 unencrypted         Importing backup 2 out of 2	Duration
Message Starting backup repositories synchronization Enumerating repositories Found 1 repository Processing capacity tier extent of S3 Backup Repository 2 S3 Backup Repository: added 2 unencrypted Importing backup 2 out of 2	Duration
<ul> <li>Starting backup repositories synchronization</li> <li>Enumerating repositories</li> <li>Found 1 repository</li> <li>Processing capacity tier extent of S3 Backup Repository 2</li> <li>S3 Backup Repository: added 2 unencrypted</li> <li>Importing backup 2 out of 2</li> </ul>	
<ul> <li>Enumerating repositories</li> <li>Found 1 repository</li> <li>Processing capacity tier extent of S3 Backup Repository 2</li> <li>S3 Backup Repository: added 2 unencrypted</li> <li>Importing backup 2 out of 2</li> </ul>	
<ul> <li>Found 1 repository</li> <li>Processing capacity tier extent of S3 Backup Repository 2</li> <li>S3 Backup Repository: added 2 unencrypted</li> <li>Importing backup 2 out of 2</li> </ul>	
<ul> <li>Processing capacity tier extent of S3 Backup Repository 2</li> <li>S3 Backup Repository: added 2 unencrypted</li> <li>Importing backup 2 out of 2</li> </ul>	
S3 Backup Repository: added 2 unencrypted Importing backup 2 out of 2	0:03:23
🛇 Importing backup 2 out of 2	0:03:20
	0:03:15
Sackup repositories synchronization completed successfully	

To restore SQL and Oracle virtual machines to the VMware Cloud on AWS workload domain/cluster, complete the following steps.

1. From the Veeam Home page, select the object storage containing the imported backups, select the VMs to restore, and then right click and select Restore Entire VM.

E Home Backup Instant Instant Disk Entire Virtual VM Guest Application Recovery Recovery WM Disks Files Files + Items + Restore	azon Microsoft Google C2 Azure laas CE Restore to Cloud Actions		
Home	Q. Type in an object name to search for	×	
▲ % Jobs ﷺ Backup ▲ Marka Backups	Job Name 1 SQL Servers	Creation Time 3/27/2022 1:00 AM 3/27/2022 1:00 AM	Re
Object Storage (Imported)     Last 24 Hours	SQLSRV-01	Instant recovery Instant disk recovery	
La success	SQLSRV-04         Image: Construction of the second se	Restore entire VM Restore virtual disks Restore VM files Restore guest files	
	SQLSRV-08	Restore to Amazon EC2 Restore to Microsoft Azure Restore to Google CE	
		Export backup Delete from disk	

2. On the first page of the Full VM Restore wizard, modify the VMs to backup if desired and select Next.

Virtual machines to restor	re:		_
Name	for instant lookup	Restore point	Add
SQLSRV-04	62.7 GB	less than a day ago (1:03 AM	
			Point
	Virtual machines to restor	Virtual machines to restore:          Image: Type in a VM name for instant lookup         Name       Size         Image: SQLSRV-04       62.7 GB	Virtual machines to restore:          Image: Size in a VM name for instant lookup         Name       Size Restore point         Image: SQLSRV-04       62.7 GB         Image: SQLSRV-04       62.7 GB         Image: SQLSRV-04       62.7 GB

3. On the Restore Mode page, select Restore to a New Location, or with Different Settings.

Full VM Restore	×				
Restore Mode Specify whethe	er selected VMs should be restored back to the original location, or to a new location or with different settings.				
Virtual Machines Restore Mode	Restore to the original location Quickly initiate the restore of selected VM to its original location, with the original name and settings. This option minimizes the chance of user input error.				
Host Resource Pool	Restore to a new location, or with different settings Customize the restored VM location, and change its settings. The wizard will automatically populate all controls with the original VM settings as the defaults.				
Datastore Folder	Staged restore Run the selected VM directly from backup files in the isolated DataLab to make changes to the guest OS or applications prior to placing the VM into production environment.				
Network	Pick proxy to use				
Summary					
	<ul> <li>Quick rollback (restore changed blocks only) Allows for quick VM recovery in case of guest OS software problem, or user error. Do not use this option when recovering from disaster caused by hardware or storage issue, or power loss.</li> </ul>				
	< Previous Next > Finish Cancel				

4. On the host page, select the Target ESXi host or cluster to restore the VM to.



5. On the Datastores page, select the target datastore location for both the configuration files and hard disk.

/irtual Machines	Files location:			
Restore Mode	File	Size	Datastore	Disk type
Host	Configuration files		WorkloadDatastore (VM	
	Hard disk 1 (SQLSR	100 GB	WorkloadDatastore (VM	Same as source
Resource Pool				
Datastore				
Folder				
Network				
Secure Restore				
Summary				

6. On the Network page, map the original networks on the VM to the networks in the new target location.

Network connections:		
Source	Target	
Management 181 (DSwitch)	Not connected	
Data - A - 3374 (DSwitch)	Not connected	
Data - B - 3375 (DSwitch)	Not connected	
	Network connections: Source SQLSRV-04 SQLSRV-04 Data - A - 3374 (DSwitch) Data - B - 3375 (DSwitch)	Network connections:       Target         Source       Target         Management 181 (DSwitch)       Not connected         Data - A - 3374 (DSwitch)       Not connected         Data - B - 3375 (DSwitch)       Not connected         Image: Connected       Image: Connected         Image: Connected



7. Select whether to scan the restored VM for malware, review the summary page, and click Finish to start the restore.

## **Restore SQL Server application data**

The following process provides instructions on how to recover a SQL Server in VMware Cloud Services in AWS in the event of a disaster that renders the on-premises site inoperable.

The following prerequisites are assumed to be complete in order to continue with the recovery steps:

- 1. The Windows Server VM has been restored to the VMware Cloud SDDC using Veeam Full Restore.
- 2. A secondary SnapCenter server has been established and SnapCenter database restore and configuration has been completed using the steps outlined in the section "SnapCenter backup and restore process summary."

After the restore of the VM is complete, you must configure networking and other items in preparation for rediscovering the host VM within SnapCenter.

- 1. Assign new IP addresses for Management and iSCSI or NFS.
- 2. Join the host to the Windows domain.
- 3. Add the hostnames to DNS or to the hosts file on the SnapCenter server.



If the SnapCenter plug-in was deployed using domain credentials different than the current domain, you must change the Log On account for the Plug-in for Windows Service on the SQL Server VM. After changing the Log On account, restart the SnapCenter SMCore, Plug-in for Windows, and Plug-in for SQL Server services.



To automatically rediscover the restored VMs in SnapCenter, the FQDN must be identical to the VM that was originally added to the SnapCenter on premises.

### Configure FSx storage for SQL Server restore

To accomplish the disaster recovery restore process for a SQL Server VM, you must break the existing SnapMirror relationship from the FSx cluster and grant access to the volume. To do so, complete the following steps.

1. To break the existing SnapMirror relationship for the SQL Server database and log volumes, run the following command from the FSx CLI:

FSx-Dest::> snapmirror break -destination-path DestSVM:DestVolName

2. Grant access to the LUN by creating an initiator group containing the iSCSI IQN of the SQL Server Windows VM:

FSx-Dest::> igroup create -vserver DestSVM -igroup igroupName
-protocol iSCSI -ostype windows -initiator IQN

3. Finally, map the LUNs to the initiator group that you just created:

```
FSx-Dest::> lun mapping create -vserver DestSVM -path LUNPath igroup
igroupName
```

4. To find the path name, run the lun show command.
#### Set up the Windows VM for iSCSI access and discover the file systems

- 1. From the SQL Server VM, set up your iSCSI network adapter to communicate on the VMware Port Group that has been established with connectivity to the iSCSI target interfaces on your FSx instance.
- 2. Open the iSCSI Initiator Properties utility and clear out the old connectivity settings on the Discovery, Favorite Targets, and Targets tabs.
- 3. Locate the IP address(es) for accessing the iSCSI logical interface on the FSx instance/cluster. This can be found in the AWS console under Amazon FSx > ONTAP > Storage Virtual Machines.

Endpoints	
Management DNS name svm-045c077375d3d9799.fs-0ae40e08acc0dea67.fsx.us-east-1.amazonaws.com	Management IP address 198.19.254.53
NFS DNS name svm-045c077375d3d9799.fs-0ae40e08acc0dea67.fsx.us-east-1.amazonaws.com	NFS IP address 198.19.254.53
iSCSI DNS name iscsi.svm-045c077375d3d9799.fs-0ae40e08acc0dea67.fsx.us-east-1.amazonaws.com	iSCSI IP addresses 172.30.15.101, 172.30.14.49

4. From the Discovery tab, click Discover Portal and enter the IP addresses for your FSx iSCSI targets.

SCSI Init	iator Proper	ties				×
Targets	Discovery	Favorite Targets	Volumes and Devices	RADIUS	Configuration	
Targe The s	t portals system will lo	ok for Targets on fo	blowing portals:		Refresh	
Addr	ess	Port	Adapter	I	P address	
To ac	id <mark>a</mark> target p	ortal, dick Discover	Portal.	Disco	over Portal	
To re then	move a targ dick Remove	et portal, select the	address above and		Remove	

Discover larger Fortar	
Enter the IP address or DNS nam want to add.	e and port number of the portal you
To change the default settings of the Advanced button.	the discovery of the target portal, dick
To change the default settings of the Advanced button. IP address or DNS name:	Port: (Default is 3260.)

5. On the Target tab, click Connect, select Enable Multi-Path if appropriate for your configuration and then click OK to connect to the target.

argets	Discovery	Favorite Targets	Volumes and Devices	RADIUS	Configuration	
Ouick C	Connect					
To disc DNS na	over and log ame of the ta	on to a target usir arget and then click	ng a basic connection, t Quick Connect.	ype the IP	address or	
Target	:			Q	ick Connect	1
Discove	ered targets					
					Refresh	
Name	e			Status		1
ian. 19	992-08.com.	netapp:sn.5918b03	69ef411ecb007495	Inactive		
					1	
				_	1	
To con	nect using a	dvanced options, s	elect a target and then		1 Connect	
To con dick Co	nect using a onnect.	dvanced options, s	elect a target and then		Connect	
To con dick Co To th	nect using a onnect.	dvanced options, si arget	elect a target and then		1 Connect	
To con click Co To th Fo	nect using a onnect. onnect To Ta	dvanced options, si arget	elect a target and then		1 Connect	
To con click Co To th Fo Se 19	nect using a onnect. onnect To Ta rget name: 2-08.com.ne	dvanced options, so arget	elect a target and then	f45c:vs.6	1 Connect	
To con dick Co To Co Fo Ta Se 19 Fo	nect using a onnect. onnect To Ta rget name: 2-08.com.ne	dvanced options, si arget tapp:sn.5918b03f9	elect a target and then Def411ecb0074956fb75	f45c:vs.6	Connect	
To con click Co To th Fo Ta Se 19 Fo th	nect using a onnect. onnect To Ta rget name: 2-08.com.ne Add this con	dvanced options, si arget tapp:sn.5918b03f9 nection to the list o	elect a target and then Def411ecb0074956fb75 f Favorite Targets.	F45c:vs.6	Connect	
To con click Co To Co Fo Ta Se 19 Fo Ta	nect using a onnect. annect To Ta rget name: 2-08.com.ne Add this con This will mak connection e	dvanced options, se arget tapp:sn.5918b03f9 nection to the list o e the system auton	elect a target and then Def411ecb0074956fb75 f Favorite Targets. natically attempt to rest puter restarts.	f45c:vs.6 ore the	Connect	
To con dick Co To Co Fo Ta Se 19 Fo Ta	nect using a onnect. annect To Ta rget name: 2-08.com.ne Add this con This will mak connection e	dvanced options, se arget etapp:sn.5918b03f9 nection to the list o e the system auton every time this comp	elect a target and then Def411ecb0074956fb75 f Favorite Targets. natically attempt to rest puter restarts.	f45c:vs.6 ore the	Connect	
To con dick Co To Co Fo Ta Se 99 Fo th	nect using a onnect. onnect To Ta rget name: 2-08.com.ne Add this con This will mak connection e Enable multi-	dvanced options, so arget etapp:sn. 5918b03f9 nection to the list o e the system auton every time this comp path	elect a target and then Pef411ecb0074956fb75 f Favorite Targets. natically attempt to rest puter restarts.	f45c:vs.6 ore the	Connect	
To con dick Co To Co Fo Ta Se 99 Fo th	nect using a onnect. onnect To Ta rget name: 2-08.com.ne Add this con This will mak connection e Enable multi- dvanced	dvanced options, si arget etapp:sn. 5918b03f9 nection to the list o e the system auton every time this comp path 2	elect a target and then Def411ecb0074956fb75 f Favorite Targets. natically attempt to rest puter restarts.	f45c:vs.6 ore the	Connect	

6. Open the Computer Management utility and bring the disks online. Verify that they retain the same drive letters that they previously held.

Basic 579.98 GB Online	MSSQL_DATA (E:) 579.98 GB NTFS Healthy (Primary Partition)	
*O Disk 2 Basic		
99.98 GB Offline	Online	
	Properties	
0	Help	1. <del>6</del> .0

1. From the SQL Server VM, open Microsoft SQL Server Management Studio and select Attach to start the process of connecting to the database.



2. Click Add and navigate to the folder containing the SQL Server primary database file, select it, and click OK.

Locate Database Files - St	QLSRV-01				×
Database Data File location:	E:\MSSQL 20	9\MSSQL15.MSSQLSERVEF			2
C: C: SRECYCLE.BIN MSSQL 2019 MSSQL 15.MSS MSSQL MSSQL DATA System Volume Infor	QLSERVER	SQLHC01_01.mdf			
(£) <b>F</b> :					
t: File name: SQI	LHC01_01.mdf	Databa	ase Data Files('	*.mdf)	>

- 3. If the transaction logs are on a separate drive, choose the folder that contains the transaction log.
- 4. When finished, click OK to attach the database.



With the SnapCenter database restored to its previous state, it automatically rediscovers the SQL Server hosts. For this to work correctly, keep in mind the following prerequisites:

- SnapCenter must be placed in Disaster Recover mode. This can be accomplished through the Swagger API or in Global Settings under Disaster Recovery.
- The FQDN of the SQL Server must be identical to the instance that was running in the on-premises datacenter.
- The original SnapMirror relationship must be broken.
- The LUNs containing the database must be mounted to the SQL Server instance and the database attached.

To confirm that SnapCenter is in Disaster Recovery mode, navigate to Settings from within the SnapCenter web client. Go to the Global Settings tab and then click Disaster Recovery. Make sure that the Enable Disaster Recovery checkbox is enabled.

	letApp Snap(	Center®
<		Global Settings Policies Users and Access
	Dashboard	
9	Resources	Global Settings
•	Monitor	
<b>iii</b>	Reports	Hypervisor Settings 🚯
A	Hosts	Notification Server Settings 🚯
ł	Storage Systems	Configuration Settings ()
÷	Settings	Purge Jobs Settings
	Alerts	Domain Settings 🚺
		CA Certificate Settings 🚯
		Disaster Recovery
		Enable Disaster Recovery Apply

## **Restore Oracle application data**

The following process provides instructions on how to recover Oracle application data in VMware Cloud Services in AWS in the event of a disaster that renders the on-premises site inoperable.

Complete the following prerequisites to continue with the recovery steps:

- 1. The Oracle Linux server VM has been restored to the VMware Cloud SDDC using Veeam Full Restore.
- 2. A secondary SnapCenter server has been established and the SnapCenter database and configuration files have been restored using the steps outlined in this section "SnapCenter backup and restore process summary."

To make the secondary storage volumes hosted on the FSxN instance accessible to the Oracle servers, you must first break the existing SnapMirror relationship.

1. After logging into the FSx CLI, run the following command to view the volumes filtered by the correct name.

```
FSx-Dest::> volume show -volume VolumeName*
FsxId0ae40e08acc0dea67::> volume show -volume oraclesrv 03*
Vserver
         Volume
                      Aggregate
                                   State
                                              Type
                                                         Size Available Used%
ora svm dest
         oraclesrv_03_u01_dest
                      aggrl
                                   online
                                              DP
                                                        100GB
                                                                 93.12GB
                                                                            6%
ora svm dest
         oraclesrv_03_u02_dest
                      aggrl
                                   online
                                              DP
                                                        200GB
                                                                 34.98GB
                                                                            82%
ora svm dest
         oraclesrv 03 u03 dest
                      aggrl
                                              DP
                                                        150GB
                                                                 33.37GB
                                                                            778
                                   online
3 entries were displayed.
FsxId0ae40e08acc0dea67::>
```

2. Run the following command to break the existing SnapMirror relationships.

FSx-Dest::> snapmirror break -destination-path DestSVM:DestVolName

```
FsxId0ae40e08acc0dea67::> snapmirror break -destination-path ora_svm_dest:oraclesrv_03_u02_dest
Operation succeeded: snapmirror break for destination "ora_svm_dest:oraclesrv_03_u02_dest".
FsxId0ae40e08acc0dea67::> snapmirror break -destination-path ora_svm_dest:oraclesrv_03_u03_dest
```

Operation succeeded: snapmirror break for destination "ora svm dest:oraclesrv 03 u03 dest".

3. Update the junction-path in the Amazon FSx web client:

FSx > Volumes > fsvol-01167370e9b7aefa0 oraclesrv\_03\_u01\_dest (fsvol-01167370e9b7aefa0) Attach Actions 🔺 Update volume Summary Create backup Delete volume Volume ID Creation time SVM ID 2022-03-08T14:52:09-05:00 svm-02b2ad25c6b2e5bc2 fsvol-01167370e9b7aefa0 🗇 Lifecycle state Junction path Volume name ⊘ Created - 🗇 oraclesrv\_03\_u01\_dest Volume type Tiering policy name UUID ONTAP SNAPSHOT\_ONLY 3d7338ce-9f19-11ecb007-4956fb75f45c Size Tiering policy cooling period (days) 100.00 GB 🗇 2 File system ID fs-0ae40e08acc0dea67 Storage efficiency enabled Disabled Resource ARN arn:aws:fsx:useast-1:541696183547:volume/fs-0ae40e08acc0dea67/fsvol-01167370e9b7aefa0 🗇

4. Add the junction path name and click Update. Specify this junction path when mounting the NFS volume from the Oracle server.

# Update volume

# Junction path

# /oraclesrv\_03\_u01\_dest

The location within your file system where your volume will be mounted.

#### Volume size

102400

Minimum 20 MiB; Maximum 104857600 MiB

## Storage efficiency

Select whether you would like to enable ONTAP storage efficiencies on your volume: deduplication, compression, and compaction.

Snanshot Only	rage.
Snanshot Univ	-
Shapshot Only	

×

\$

In Cloud Manager, you can obtain the mount command with the correct NFS LIF IP address for mounting the NFS volumes that contain the Oracle database files and logs.

1. In Cloud Manager, access the list of volumes for your FSx cluster.

HCApps	Overview	Volumes			
	50 volume	S			
	Volum	e Name ‡	State	Storage VM 🔶	Disk Type
	oracle u02_d	srv_02_ est	<ul> <li>Online</li> </ul>	ora_svm_dest	SSD
	oracle u03_d	srv_02_ est	<ul> <li>Online</li> </ul>	ora_svm_dest	SSD
	oracle u01_d	srv_03_ est	<ul> <li>Online</li> </ul>	ora_svm_dest	SSD

2. From the action menu, select Mount Command to view and copy the mount command to be used on our Oracle Linux server.

Account ~	Information	or 🗸	
	Edit		
	Clone		
	Restore from Snapshot copy		
	Create a Snapshot copy		
Capacity Pool Us	Mount Command		
0 B	Change Tiering Policy		
0 B	Delete		
	Snapshot		
Go to your linux m	Mount Volume NFS oraclesrv_03_u01_dest achine and enter this mount comm	nand	
Mount Command			
mount 198.19.2	54.180:/oraclesrv_03_u01_dest <d< td=""><td>est_d</td><td>🗇 Сору</td></d<>	est_d	🗇 Сору
<u></u>			
Mount the NFS file system already exist on the Oracle From the Oracle Linux serv	to the Oracle Linux Server. The directorie Linux host. er, use the mount command to mount the	es for mounting NFS volumes	g the NFS share

FSx-Dest::> mount -t oracle server ip:/junction-path

Repeat this step for each volume associated with the Oracle databases.



To make the NFS mount persistent upon rebooting, edit the /etc/fstab file to include the mount commands.

5. Reboot the Oracle server. The Oracle databases should start up normally and be available for use.

#### Failback

Upon successful completion of the failover process outlined in this solution, SnapCenter and Veeam resume their backup functions running in AWS, and FSx for ONTAP is now designated as primary storage with no existing SnapMirror relationships with the original on-premises datacenter. After normal function has resumed on premises, you can use a process identical to the one outlined in this documentation to mirror data back to the on-premises ONTAP storage system.

As is also outlined in this documentation, you can configure SnapCenter to mirror the application data volumes from FSx for ONTAP to an ONTAP storage system residing on premises. Similarly, you can configure Veeam to replicate backup copies to Amazon S3 using a scale-out backup repository so that those backups are accessible to a Veeam backup server residing at the on-premises datacenter.

Failback is outside the scope of this documentation, but failback differs little from the detailed process outlined here.

#### Conclusion

The use case presented in this documentation focuses on proven disaster recovery technologies that highlight the integration between NetApp and VMware. NetApp ONTAP storage systems provide proven data-mirroring technologies that allow organizations to design disaster recovery solutions that span on-premises and ONTAP technologies residing with the leading cloud providers.

FSx for ONTAP on AWS is one such solution that allows for seamless integration with SnapCenter and SyncMirror for replicating application data to the cloud. Veeam Backup & Replication is another well-known technology that integrates well with NetApp ONTAP storage systems and can provide failover to vSphere-native storage.

This solution presented a disaster recovery solution using guest connect storage from an ONTAP system hosting SQL Server and Oracle application data. SnapCenter with SnapMirror provides an easy-to-manage solution for protecting application volumes on ONTAP systems and replicating them to FSx or CVO residing in the cloud. SnapCenter is a DR-enabled solution for failing over all application data to VMware Cloud on AWS.

### Where to find additional information

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

· Links to solution documentation

NetApp Hybrid Multicloud with VMware Solutions

## Veeam Backup & Restore in VMware Cloud, with Amazon FSx for ONTAP

Veeam Backup & Replication is an effective and reliable solution for protecting data in VMware Cloud. This solution demonstrates the proper setup and configuration for using Veeam Backup and Replication to backup and restore application VMs residing on FSx for ONTAP NFS datastores in VMware Cloud.

Author: Josh Powell - NetApp Solutions Engineering

## Overview

VMware Cloud (in AWS) supports the use of NFS datastores as supplemental storage, and FSx for NetApp ONTAP is a secure solution for customers who need to store large amounts of data for their cloud applications that can scale independent of the number of ESXi hosts in the SDDC cluster. This integrated AWS storage service offers highly efficient storage with all of the traditional NetApp ONTAP capabilities.

## **Use Cases**

This solution addresses the following use cases:

- Backup and restore of Windows and Linux virtual machines hosted in VMC using FSx for NetApp ONTAP as a backup repository.
- Backup and restore of Microsoft SQL Server application data using FSx for NetApp ONTAP as a backup repository.
- Backup and restore of Oracle application data using FSx for Netapp ONTAP as a backup repository.

## NFS Datastores Using Amazon FSx for ONTAP

All virtual machines in this solution reside on FSx for ONTAP supplemental NFS datastores. Using FSx for ONTAP as a supplemental NFS datastore has several benefits. For example, it allows you to:

- Create a scalable and highly available file system in the cloud without the need for complex setup and management.
- Integrate with your existing VMware environment, allowing you to use familiar tools and processes to manage your cloud resources.
- Benefit from the advanced data management features provided by ONTAP, such as snapshots and replication, to protect your data and ensure its availability.

This list provides the high level steps necessary to configure Veeam Backup & Replication, execute backup and restore jobs using FSx for ONTAP as a backup repository, and perform restores of SQL Server and Oracle VMs and databases:

- 1. Create the FSx for ONTAP file system to be used as iSCSI backup repository for Veeam Backup & Replication.
- 2. Deploy Veeam Proxy to distribute backup workloads and mount iSCSI backup repositories hosted on FSx for ONTAP.
- 3. Configure Veeam Backup Jobs to backup SQL Server, Oracle, Linux and Windows virtual machines.
- 4. Restore SQL Server virtual machines and individual databases.
- 5. Restore Oracle virtual machines and individual databases.

### Prerequisites

The purpose of this solution is to demonstrate data protection of virtual machines running in VMware Cloud and located on NFS Datastores hosted by FSx for NetApp ONTAP. This solution assumes the following components are configured and ready for use:

- 1. FSx for ONTAP filesystem with one or more NFS datastores connected to VMware Cloud.
- 2. Microsoft Windows Server VM with Veeam Backup & Replication software installed.
  - vCenter server has been discovered by the Veeam Backup & Replication server using their IP address or fully qualified domain name.
- 3. Microsoft Windows Server VM to be installed with Veeam Backup Proxy components during the solution deployment.
- 4. Microsoft SQL Server VMs with VMDKs and application data residing on FSx for ONTAP NFS datastores. For this solution we had two SQL databases on two separate VMDKs.
  - Note: As a best practice database and transaction log files are placed on separate drives as this will improve performance and reliability. This is in part due to the fact that transaction logs are written sequentially, whereas database files are written randomly.
- 5. Oracle Database VMs with VMDKs and application data residing on FSx for ONTAP NFS datastores.
- 6. Linux and Windows file server VMs with VMDKs residing on FSx for ONTAP NFS datastores.
- 7. Veeam requires specific TCP ports for communication between servers and components in the backup environment. On Veeam backup infrastructure components, the required firewall rules are automatically created.

For a full listing of the network port requirements refer to the Ports section of the Veeam Backup and Replication User Guide for VMware vSphere.

### **High Level Architecture**

The testing / validation of this solution was performed in a lab that may or may not match the final deployment environment. For more information, please refer to the following sections.



## Hardware / Software Components

The purpose of this solution is to demonstrate data protection of virtual machines running in VMware Cloud and located on NFS Datastores hosted by FSx for NetApp ONTAP. This solution assumes the following components are already configured and ready for use:

- · Microsoft Windows VM's located on an FSx for ONTAP NFS Datastore
- · Linux (CentOS) VM's located on an FSx for ONTAP NFS Datastore
- Microsoft SQL Server VM's located on an FSx for ONTAP NFS Datastore
  - Two databases hosted on separate VMDK's
- Oracle VM's located on an FSx for ONTAP NFS Datastore

### **Solution Deployment**

In this solution we provide detailed instructions for deploying and validating a solution utilizing Veeam Backup and Replication software to perform backup and recovery of SQL Server, Oracle, and Windows and Linux file server virtual machines in a VMware Cloud SDDC on AWS. The Virtual Machines in this solution reside on a supplemental NFS datastore hosted by FSx for ONTAP. In addition, a separate FSx for ONTAP file system is used to host iSCSI volumes that will be used for Veeam backup repositories.

We will go over FSx for ONTAP file system creation, mounting iSCSI volumes to be used as backup repositories, creating and running backup jobs, and performing VM and database restores.

For detailed information on FSx for NetApp ONTAP refer to the FSx for ONTAP User Guide.

For detailed information on Veeam Backup and Replication refer to the Veeam Help Center Technical Documentation site.

For considerations and limitations when using Veeam Backup and Replication with VMware Cloud on AWS, refer to VMware Cloud on AWS and VMware Cloud on Dell EMC Support. Considerations and Limitations.

## **Deploy Veeam Proxy server**

A Veeam proxy server is a component of the Veeam Backup & Replication software that acts as an intermediary between the source and the backup or replication target. The proxy server helps to optimize and accelerate data transfer during backup jobs by processing data locally and can use different Transport Modes to access data using VMware vStorage APIs for Data Protection or through direct storage access.

When choosing a Veeam proxy server design it is important to consider the number of concurrent tasks and the transport mode or type of storage access desired.

For sizing the number of proxy servers, and for their system requirements, refer to the Veeam VMware vSphere Best Practice Guide.

The Veeam Data Mover is a component of the Veeam Proxy Server and utilizes a Transport Mode as a method for obtaining VM data from the source and transferring it to the target. The transport mode is specified during the configuration of the backup job. It is possible to increase the efficiency backups from NFS datastores by using direct storage access.

For more information on Transport Modes refer to the Veeam Backup and Replication User Guide for VMware vSphere.

In the following step we cover deployment of the Veeam Proxy Server on a Windows VM in the VMware Cloud SDDC.

In this step the Veeam Proxy is deployed to an existing Windows VM. This allows backup jobs to be distributed between the primary Veeam Backup Server and the Veeam Proxy.

- 1. On the Veeam Backup and Replication server, open the administration console and select **Backup Infrastructure** in the lower left menu.
- 2. Right click on Backup Proxies and click on Add VMware backup proxy... to open the wizard.



3. In the Add VMware Proxy wizard click the Add New... button to add a new proxy server.

erver	Choose server:	-
	VeeamSrv (Backup server)	Add New
raffic Rules	Proxy description:	
ummary		
	Transport mode:	
	Automatic selection	Choose
	Connected datastores:	
	Automatic detection (recommended)	Choose
	Max concurrent tasks:	

• Fill out the DNS name or IP address

4.

- $\circ$  Select an account to use for Credentials on the new system or add new credentials
- Review the components to be installed and then click on Apply to begin the deployment

Name	Message	Duration
Cradantials	Starting infrastructure item update process	0:00:03
STEGICITURIS	Collecting hardware info	
leview	Detecting operating system	
	🖉 Detecting OS version	
Apply	🖉 Creating temporary folder	
	Package VeeamTransport.msi has been uploaded	0:00:05
oummary	Package VeeamGuestAgent_x86.msi has been uploaded	
	🙁 Package VeeamGuestAgent_x64.msi has been uploaded	
	📀 Package VeeamLogBackupService_x86.msi has been uploaded	0:00:01
	Package VeeamLogBackupService_x64.msi has been uploaded	
	Installing package Transport	0:00:19

5. Back in the **New VMware Proxy** wizard, choose a Transport Mode. In our case we chose **Automatic Selection**.

Server	Transport Mode	×
Managed Se	Backup proxy transport mode:	lux servers added to the
Server Traffic Rules	Automatic selection Data retrieval mode is selected automatically by analyzing backup proxy configuration and reachable VMFS and NFS datastores. Transport modes allowing for direct storage access will be used whenever possible.	2/2022 9 🗸 🛛 Add New
Apply Summary	<ul> <li>Direct storage access         Data is retrieved directly from shared storage, without impacting production hosts. For block storage, backup proxy server must be connected into SAN fabric via hardware or software HBA, and have VMFS volumes mounted.     </li> <li>Virtual appliance         Data is retrieved directly from storage through hypervisor I/O stack by hot adding backed up virtual disks to a backup proxy VM. Datastores containing protected VMs must be connected to a host running backup proxy VM.     </li> <li>Network         Data is retrieved from storage through hypervisor network stack using NBD protocol over host management interface. This mode has no special setup requirements. Recommended for 10 Gb Ethernet or faster.     </li> <li>Options         Failover to network mode if primary mode fails, or is unavailable         Enable host to proxy traffic encryption in Network mode (NBDSSL)     </li> </ul>	Choose

6. Select the Connected datastores that you want the VMware Proxy to have direct access to.

### New VMware Proxy

#### Server

Choose a server for VMware backup proxy. You can choose between any Microsoft Windows or Linux servers added to the Managed Servers which are not assigned a VMware backup proxy role already.

Server	Choose server:	
	veeamproxy.demozone.com (Created by VEEAMSRV\Administrator at 12/22/2022 9 🗸	Add New
Traffic Rules	Proxy description:	
Apply	Created by VEEAMSRV\Administrator at 12/22/2022 9:11 PM.	
c		
summary		
summary	Transport mode:	
summary	Transport mode: Direct storage access	Choose
Summary	Transport mode: Direct storage access Connected datastores:	Choose

199

X

Select objects:			62	er
<ul> <li>✓ I Hosts and Di</li> <li>✓ I Hosts and Di</li> <li>✓ I Vcenter.s</li> <li>DS01</li> <li>DS02</li> </ul>	sks ddc-52-3 <mark>4-17</mark> -99.vmwar	evmc.com		)2:
Type in an obje	ct name to search for		Q	
	- DL	OK	Cancel	

7. Configure and apply any specific network traffic rules such as encryption or throttling that are desired. When complete click on the **Apply** button to complete the deployment.

Server	Network traffic rules Throttling is global,	s control encryption an with set bandwidth spl	d throttling of it equally acro	network traffic based o ss all backup proxies fa	on the destination. Iling into the rule.
Traffic Rules	The following netwo	ork traffic rules ap <mark>ply</mark> to	this proxy:		
Apply	Name	Encryption Enabled	Throttling Disabled	Time period	Wiense
Summary					
	Manage network tra	ffic rules			

## Configure storage and Backup Repositories

The primary Veeam Backup server and Veeam Proxy server have access to a backup repository in the form of direct connected storage. In this section we cover creating an FSx for ONTAP file system, mounting iSCSI LUNs to the Veeam servers and creating Backup Repositories.

Create an FSx for ONTAP file system that will be used to host the iSCSI volumes for the Veeam Backup Repositories.

1. In the AWS console, Go to FSx and then **Create file system** 



2. Select Amazon FSx for NetApp ONTAP and then Next to continue.

Amizon PSX for Nebapp On the	Amazon FSx for OpenZFS	Amazon FSx for Windows File Server	<ul> <li>Amazon FSx for Lustre</li> </ul>
FSXa	FSX-	FSx□	FSXa
Amazon FSx for NetApp ONTAP	Amazon FSx for OpenZFS	Amazon FSx for Windows File Server	Amazon FSx for Lustre
Broadly accessible from Linux Windows ar	nd macOS compute instances and containers (nu	nning on AWS or on-premises) via industry-standard NF	S, SMB, and iSCSI protocols.
troubly accessing from come, from any, in	그는 것은 가방에서 이번 것 같아. 이는 것	The second se	
Provides ONTAP's popular data manageme Deliver: hundred: of theurands of IOPS with	ent capabilities like Snapshots, SnapMirror (for d	ata replication), FlexClone (for data cloning), and data co	ompression / deduplication.
Provides ONTAP's popular data manageme Delivers hundreds of thousands of IOPS wi Offers highly-available and highly-durable	ent capabilities like Snapshots, SnapMirror (for d th consistent sub-millisecond latencies, and up t multi-AZ SSD storage with support for cross-res	ata replication), FlexClone (for data cloning), and data co to 3 GB/s of throughput. aion replication and built-in, fully managed backups.	ompression / deduplication.

3. Fill in the file system name, deployment type, SSD storage capacity and the VPC in which the FSx for ONTAP cluster will reside. This must be a VPC configured to communicate with the virtual machine network in VMware Cloud. Click on **Next**.

Creation method	
Quick create Use recommended best-practice configurations. Most configuration options can be changed after the file system is created.	Standard create You set all of the configuration options, including specifying performance, networking, security, backups, and maintenance.
Quick configuration	
File system name - optional Info BackupFSxN	
Maximum of 256 Unicode letters, whitespace, and numbers, plus +	-=:/
Deployment type Info O Multi-AZ Single-AZ	
SSD storage capacity Info 4096 GiB Minimum 1024 GiB; Maximum 192 TiB	
Virtual Private Cloud (VPC) Info	
Demo-FsxforONTAP-VPC   vpc-05596abe79cb653b7	<b> ▲</b>
Storage efficiency Select whether you would like to enable ONTAP's storage efficiency Enabled (recommended)	y features: deduplication, compression, and compaction
) Disabled	

4. Review the deployment steps and click on **Create File System** to begin the file system creation process.

Create and configure the iSCSI LUNs on FSx for ONTAP and mount to the Veeam backup and proxy servers. These LUNs will later be used to create Veeam backup repositories.



Creating an iSCSI LUN on FSx for ONTAP is a multi-step process. The first step of creating the volumes can be accomplished in the Amazon FSx Console or with the NetApp ONTAP CLI.



For more information on using FSx for ONTAP, see the FSx for ONTAP User Guide.

1. From the NetApp ONTAP CLI create the initial volumes using the following command:

```
FSx-Backup::> volume create -vserver svm_name -volume vol_name
-aggregate aggregate_name -size vol_size -type RW
```

2. Create LUNs using the volumes created in the previous step:

```
FSx-Backup::> lun create -vserver svm_name -path
/vol/vol_name/lun_name -size size -ostype windows -space-allocation
enabled
```

3. Grant access to the LUNs by creating an initiator group containing the iSCSI IQN of the Veeam backup and proxy servers:

FSx-Backup::> igroup create -vserver svm\_name -igroup igroup\_name -protocol iSCSI -ostype windows -initiator IQN



To complete the preceding step you will need to first retrieve the IQN from the iSCSI initiator properties on the Windows servers.

4. Finally, map the LUNs to the initiator group that you just created:

```
FSx-Backup::> lun mapping create -vserver svm_name -path
/vol/vol_name/lun_name igroup igroup_name
```

5. To mount the iSCSI LUNs, log into the Veeam Backup & Replication Server and open iSCSI Initiator Properties. Go to the **Discover** tab and enter the iSCSI target IP address.

scorer rarger oran	× Infigur	ation
nter the IP address or DNS name and port number of the por ant to add.	tal you resh	
change the default settings of the discovery of the target p e Advanced button.	oortal, dick Idress	
address or DNS name: Port: (Default is 3	260.)	
0.49.0.154 3260		
Advanced OK	Cancel	h;
then dick Remove.	Remove	
SNS servers		
The system is registered on the following iSNS servers:	Refresh	
Name		
To add an iSNS server, dick Add Server.	Add Server	

6. On the **Targets** tab, highlight the inactive LUN and click on **Connect**. Check the **Enable multi-path** box and click on **OK** to connect to the LUN.

rgets (	Discovery	Favorite Targets	Volumes and Devices	RADIUS	Configuration	
Quick Co	nnect					
To disco DNS nam	iver and log ne of the ta	) on to a target usir arget and then click	ng a basic connection, t Quick Connect.	ype the IP	address or	
Target:				Qu	uid: Connect	
Discover	ed targets			-		-1
					Refresh	
Name	92-08.com.r	netapp:sn.d9aad3c	d818011edbfcd87a	Status Inactive		
Name iqn. 199	92-08.com.t	netapp:sn.d9aad3c	d818011edbfcd87a	Status Inactive		
Name iqn:199 To conne click Con	92-08.com.r ect using ad inect.	netapp:sn.d9aad3c	elect a target and then	Status Inactive	Connect	
Name iqn:199 To conne click Con To comp then clid	ect using ad nnect. letely disco	netapp:sn.d9aad3c dvanced options, so onnect a target, sel ct.	elect a target and then	Status Inactive	Connect Disconnect	
Name iqn. 199 To conne dick Con To comp then did For targ select th	ect using ad nect. letely disco k Disconnect et propertie ne target ar	netapp:sn.d9aad3d dvanced options, si onnect a target, sel ct. es, including config nd dick Properties.	d8 180 1 1 edbfcd87a elect a target and then ect the target and uration of sessions,	Status Inactive	Connect Disconnect Properties	

7. In the Disk Management utility initialize the new LUN and create a volume with the desired name and drive letter. Check the **Enable multi-path** box and click on **OK** to connect to the LUN.

Computer Management (Local V	dume	Lavout Type	File System	Statur		
<ul> <li>Computer Management (Local Version of the second sec</li></ul>	New Simple Volum Format Partition To store data Choose wheth Do not Format File : Alloc Volu File : Alloc Volu Disk 1 asic 899.98 GB	Layout Type e Wizard on this partition, yo her you want to fon format this volume this volume with th system: sation unit size: me label: Perform a quick form shable file and folde g.98 GB	File System ou must format nat this volume e following set NTFS Default Backup_ nat er compression	I Status It first. It firs	ant to use.	hary Partitio

8. Repeat these steps to mount the iSCSI volumes on the Veeam Proxy server.

In the Veeam Backup and Replication console, create backup repositories for the Veeam Backup and Veeam Proxy servers. These repositories will be used as backup targets for the virtual machines backups.

1. In the Veeam Backup and Replication console click on **Backup Infrastructure** in the lower left and then select **Add Repository** 

	Repository Tools	
<b>≣</b> • Home	Backup Repository	
Add Edit Repository Reposit Manage Reposito	Rescan rory ry Tools	
Backup Infrastruct	ture	
Backup Pro	xies	
Backup Rep	positories	
Scale-out R	epositories	
Co WAN Accel	erators	
Service Prov	viders	
SureBackup	) 	
Managed S	ervers	
A Home		
Inventory		
Backup Infras	tructure	
Storage Infras	tructure	
Tape Infrastru	cture	
Files		
		Ľ⊛ <b>≥</b>

2. In the New Backup Repository wizard, enter a name for the repository and then select the server from the drop-down list and click on the **Populate** button to choose the NTFS volume that will be used.

Name	Repository server:			
Conver	veeamproxy.demozone.com (Crea	ted by VEEAMSRV\Administrator at 12	/22/2022 9 🗸	Add New
Server	Path	Capacity	Free	Populate
Repository	@ C:\	89.4 GB	74 GB	
Mount Server	⊂ E\	1.9 TB	1.9 TB	
Keview				
Apply				
Summany				
Janninary				

- 3. On the next page choose a Mount server that will be used to mount backups to when performing advanced restores. By default this is the same server that has the repository storage connected.
- 4. Review your selections and click on **Apply** to start the backup repository creation.

	3	on server veeamproxy.demozone.com:
	Component name	Status
erver	Transport	already exists
lenository	vPower NFS	will be installed
	Mount Server	will be installed
Aount Server		
	Search the repository for existing backups	and import them automatically
	Search the repository for existing backups	and import them automatically the catalog

## Configure Veeam backup jobs

Backup jobs should be created utilizing the the Backup Repositories in the previous section. Creating backup jobs is a normal part of any storage administrator's repertoire and we do not cover all of the steps here. For more complete information on creating backup jobs in Veeam, see the Veeam Help Center Technical Documentation.

In this solution separate backup jobs were created for:

- Microsoft Windows SQL Servers
- Oracle database servers
- · Windows file servers
- Linux file servers

- 1. Enable application-aware processing to create consistent backups and perform transaction log processing.
- 2. After enabling application-aware processing add the correct credentials with admin privileges to the application as this may be different than the guest OS credentials.

Specify	Oracle a	ccount w	ith SYS	DBA	privileges: 🕕	100	
🔧 Us	e guest O	S credent	ials			~	Add
1.0					Manage acco	unts	
Archive	ed logs:						
O Do	not delet	e archive	d logs				
Del	ete logs o	lder than	: 24		hours		
() Del	ete logs o	ver:	10	A	GB		
	kun loas	ever."	15		minutes		
Dot Doc	nin lan hi	ereij.	1.2		initiates		
(®)	Until the	correspo	ndina	imac	ie-level backup is d	eleted	
0	Keep onl	y last	5	day	s of log backups		
Loc	shipping	servers:					
Au	tomatic s	election				11	Choose
1							

3. To manage the retention policy for the backup check the **Keep certain full backups longer for archival purposes** and click the **Configure...** button to configure the policy.

Con	figure GFS		×	
	Keep weekly full backups for: 15 🚔 weeks			10:3 ~
	If multiple full backups exist, use the one from:	Sunday	~	backup
	Keep monthly full backups for: 12 📮 months			
	Use weekly full backup from the following week of a month:	First	~	Configure
	Keep yearly full backups for: 1 🔅 years			
	Use monthly full backup from the following month:	January	$\sim$	
				Ve recommend to ma d off-site.
¢.,	ove As Default OK	Cance	l.	

## **Restore Application VMs with Veeam full restore**

Performing a full restore with Veeam is the first step in performing an application restore. We validated that full restores of our VMs powered on and all services were running normally.

Restoring servers is a normal part of any storage administrator's repertoire and we do not cover all of the steps here. For more complete information on performing full restores in Veeam, see the Veeam Help Center Technical Documentation.

### **Restore SQL Server databases**

Veeam Backup & Replication provides several options for restoring SQL Server databases. For this validation we used the Veeam Explorer for SQL Server with Instant Recovery to execute restores of our SQL Server databases. SQL Server Instant Recovery is a feature that allows you to quickly restore SQL Server databases without having to wait for a full database restore. This rapid recovery process minimizes downtime and ensures business continuity. Here's how it works:

- Veeam Explorer mounts the backup containing the SQL Server database to be restored.
- The software **publishes the database** directly from the mounted files, making it accessible as a temporary database on the target SQL Server instance.
- While the temporary database is in use, Veeam Explorer **redirects user queries** to this database, ensuring that users can continue to access and work with the data.
- In the background, Veeam **performs a full database restore**, transferring data from the temporary database to the original database location.
- Once the full database restore is complete, Veeam Explorer **switches user queries back to the original** database and removes the temporary database.

1. In the Veeam Backup and Replication console, navigate to the list of SQL Server backups, right click on a server and select **Restore application items** and then **Microsoft SQL Server databases...** 



2. In the Microsoft SQL Server Database Restore Wizard select a restore point from the list and click on **Next**.

estore Point eason ummary	VM name: sql_srv_wkld_1 VM size: 43.9 GB O Restore from the latest available backup Restore from this restore point:	Original ho	st: vcenter.sddc-44-235-223-88.vm.
	Created	Туре	Backup
	🕒 less than a day ago (9:44 PM Tuesday	Increment	SQL Server Backups

3. Enter a **Restore reason** if desired and then, on the Summary page, click on the **Browse** button to launch Veeam Explorer for Microsoft SQL Server.

Microsoft	SQL	Server	Database	Restore
-----------	-----	--------	----------	---------

estore Point	Summary:
eason ummary	VM name: sql_srv_wkld_1 Restore point: Current: sql_srv_wkld_1 less than a day ago (9:07 PM Tuesday 1/10/2023)

4. In Veeam Explorer expand the list of database instances, right click and select **Instant recovery** and then the specific restore point to recover to.

i ∎• Home	Database		sql_srv_wkld_	1 as of less than a day ago (9:07 PM Tuesday 1/10/2023) - Veeam Explorer for Microsoft SQL Serv
Instant Recovery • tant Recovery	Publish Database * Publish Publish	Restore Schema + tore	Export Files * Schema *	
Databases			Database Info	
SQLSRV-(	)1 It instance		Name: Backup created:	DATA_01 1/10/2023 9:07 PM
0/ 0/	Instant recovery	instant recovery	of the state of Tuesday 1/10 to an server	0/2023, 9:07 PM to SQLSRV-01
	Restore database +		Available Restore	Period
	Export backup       •         Export files       •         Export schema       •		Database Files Primary database file E:\MSSQL 2019\MSSQL 1	15.MSSQLSERVER\MSSQL\DATA\DATA_01.mdf
			Secondary database and E\MSSQL 2019\MSSQL E\MSSQL 2019\MSSQL E\MSSQL 2019\MSSQL E\MSSQL 2019\MSSQL	d log files 15.MSSQLSERVER\MSSQL\LOGS\DATA_log.ldf 15.MSSQLSERVER\MSSQL\DATA\DATA_02.ndf 15.MSSQLSERVER\MSSQL\DATA\DATA_03.ndf 15.MSSQLSERVER\MSSQL\DATA\DATA_04.ndf

5. In the Instant Recovery Wizard specify the switchover type. This can either be automatically with minimal downtime, manually, or at a specified time. Then click the **Recover** button to begin the restore process.
| Specify switchover type:         |   |                        |
|----------------------------------|---|------------------------|
| Auto                             |   |                        |
| Switchover will be per<br>ready. | formed automatically with minimal possible downtim    | e once the database is |
| O Manual                         |   |                        |
| Switchover can be per            | formed manually at any point in time after the databa | ise is ready.          |
| Scheduled at:                    | 1/10/2023 10:16 PM                                    |                        |
|                                  |   |                        |
|                                  |   |                        |
|                                  |   |                        |
|                                  |   |                        |
|                                  | Back  | Cancel                 |
|                                  |   |                        |

Databases	Instant Recovery Inf	0	
Instant Recovery (1)	Status	Starting (restored)	
DATA 01	SQL Serven	SQLSRV-01	
<ul> <li>SQLSRV-01</li> </ul>	Target names	DATA 01	
- E Default Instance	Target point in time:	1/10/2023 9:07 PM	
DATA 01	Restore point:	sql_sry_widd_1	
DASA_02	Switchover mode	Auto	
	Database Files		
	Ratur	Persistent	
	Primary database file ENMSSQL 2019/MSSQL 15.	MSSGLSERVER/MSSGL/0ATA/0ATA_01.mdf	
	EVMSSQL 2019/MSSQL35 EVMSSQL 2019/MSSQL35 EVMSSQL 2019/MSSQL35 EVMSSQL 2019/MSSQL35	MSSCILSENHRAMSSCILLICOSI) MAAA_log Am MSSCILSENHRAMSSCILLICOSI) MAAA_Q Dadh MSSCILSENHRAMSSCILLICOSI (DATAQ Dadh QISSCILSENHRAMSSCILLICOSI (DATAQ Dadh MSSCILSENHRAMSSCILLICOSI (DATAQ DAdh)	
	Action		Duration
	😋 listarit Recovery start	ed at 1/10/2023 10:12:06 PM	
	Putrishing database		00:35
	🙁 Copying target files		08.28
	🙁 🕲 Database published a	1/10/2023 10:12:42 PM	
	Synchronizing files		
	Seady for switchover		
	Detechiner database		
	A neuronal damage		

For more detailed information on performing SQL Server restore operations with Veeam Explorer refer to the Microsoft SQL Server section in the Veeam Explorers User Guide.

### **Restore Oracle databases with Veeam Explorer**

Veeam Explorer for Oracle database provides the ability to perform a standard Oracle database restore or an uninterrupted restore using Instant Recovery. It also supports publishing databases for fast access, recovery of Data Guard databases and restores from RMAN backups.

For more detailed information on performing Oracle database restore operations with Veeam Explorer refer to the Oracle section in the Veeam Explorers User Guide.

In this section an Oracle database restore to a different server is covered using Veeam Explorer.

1. In the Veeam Backup and Replication console, navigate to the list of Oracle backups, right click on a server and select **Restore application items** and then **Oracle databases...**.

認 Backup Tools I Home Backup			Vee	am Backup and Replication
Instant Instant Disk Entire Virtual VM Guest Application Recovery VM Disks Files Files - Items - Restore	Amazon Microsoft Google EC2 Azure Itaas CE Restore to Cloud			
Home	Q. Type in an object name to search for	×		
· 後 Jobs 福 Backup Backups	Job Name 1 ▲ 환 Oracle Backups 타 ora_sry_01 특히 ora_sry_02	Creation Time 1/20/2023 2:24 PM 1/20/2023 6:00 PM 1/20/2023 6:02 PM	Restore Points 6 4	Repository Repository - Veeam Serve
Success Warning Failed	SQL Serve     Instant recovery     Instant disk recovery     Restore entire VM     Restore virtual disks     Restore VM files	10/2023 9:05 PM	5	Repository - Veeam Serve
	Restore guest files     Restore application item	<ul> <li>Fa Noracle databases</li> </ul>	6e (	
	Restore to Amazon EC2 Restore to Microsoft Azu Restore to Google CE Export backup Delete from disk	re		

2. In the Oracle Database Restore Wizard select a restore point from the list and click on **Next**.

Restore Point Reason Gummary	VM name: ora_srv_03 VM size: 38.5 GB Restore from the latest available backup Restore from this restore point:	Original ho	ost: vcenter.sddc-44-235-223-88.vm.
	Created	Туре	Backup
	<ul> <li>Iess than a day ago (6:01 PM Friday 1/</li> <li>Iess than a day ago (5:01 PM Friday 1/</li> <li>Iess than a day ago (4:02 PM Friday 1/</li> <li>Iess than a day ago (3:47 PM Friday 1/</li> <li>Iess than a day ago (2:47 PM Friday 1/</li> </ul>	Increment Increment Increment Full	Oracle Backups Oracle Backups Oracle Backups Oracle Backups Oracle Backups

3. Enter a **Restore reason** if desired and then, on the Summary page, click on the **Browse** button to launch Veeam Explorer for Oracle.

218

Oracle Database Restore	>
DRACLE' Summary	
Review the re to select data	estore point settings, and click Browse to exit the wizard and open Veeam Explorer for Oracle, where you will be able abases to restore.
Restore Point	Summary:
Reason	VM name: ora_srv_03 Restore point:
Summary	Current: ora_srv_03 less than a day ago (6:01 PM Friday 1/20/2023)
	< Previous Rest Provise Cancel
	~

4. In Veeam Explorer expand the list of database instances, click on the database to be restored and then from the **Restore Database** drop-down menu at the top select **Restore to another server...** 



5. In the Restore Wizard specify the restore point to restore from and click Next.

Specify restore point							
			10020424				
Specify point in time you w	ant to restore	the databas	se to:				
Restore to the point in t	ime of the sel	ected imag	je-level b	ackup			
Restore to a specific po	int in time (re	quires redo	log back	ups)			
5:01 PM 1/20/2023							6:01 PM 1/20/2023
	Fr	iday, Janua	iry 20, 202	23 6:01 PM			
Perform restore to 1	he specific tra	insaction					
Enables you to revi database to the mo	ew major data ment in time	base transa right befon	actions ar e the unv	ound the /anted cha	selected ange.	time, an	d restore the
🤼 To enable this f	unctionality, s	pecify the	staging (	)racle serv	er under	Menu >	Options.

6. Specify the target server the database will be restored to and the account credentials and click **Next**.

Account: oracle	Advanced.
Bernel ICC I have the second	
Password: [Click here to change the password]	- 4
Private key is required for this connection	1.7
Private key:	Browse
Passphrase:	

Cantral Eller		~
Control files	control01 ctl	
/ oracle/ app/ oraclasis/ oracleor/		
/oracle/app/recovery_area/ora	db01/control02.ctl	
Data files		
/oracie/app/oradata/oradbul/	systemul.dbf	
/oracle/app/oradata/oradb01/	sysaux01.dbf	
/oracle/app/oradata/oradb01/	undotbs01.dbf	
/oracle/app/oradata/oradb01/	pdbseed/system01.dbf	
/oracle/app/oradata/oradb01/	pdbseed/sysaux01.dbf	
/oracle/app/oradata/oradb01/	users01.dbf	

In this section a database is published to an alternate server for fast access without launching a full restore.

1. In the Veeam Backup and Replication console, navigate to the list of Oracle backups, right click on a server and select **Restore application items** and then **Oracle databases...**.



2. In the Oracle Database Restore Wizard select a restore point from the list and click on Next.

Restore Point Reason Summary	VM name: ora_srv_02 Original host: vcenter.sddc-44-235-22 VM size: 38.1 GB Restore from the latest available backup Restore from this restore point:					
	Created	Туре	Backup			
	🕑 less than a day ago (7:03 PM Friday 1/	Increment	Oracle Backups			
	🕒 less than a day ago (6:02 PM Friday 1/	Increment	Oracle Backups			
	I less than a day ago (5:02 PM Friday 1/	Increment	Oracle Backups			
	Ess than a day ago (4:03 PM Friday 1/ Ess than a day ago (3:49 PM Friday 1/	Full	Oracle Backups Oracle Backups			

- 3. Enter a **Restore reason** if desired and then, on the Summary page, click on the **Browse** button to launch Veeam Explorer for Oracle.
- 4. In Veeam Explorer expand the list of database instances, click on the database to be restored and then from the **Publish Database** drop-down menu at the top select **Publish to another server...**

<b>∃</b> • Databa	se					
Instant Recovery •	Publish Database •	Restore Database •	Export as RMAN back	; up▼ C	Export Patabase Files •	
Instant Recovery Publish to and			erver	Expor	t	
Databases	hire-110-2	1	Databa	ise Int	fo	
<ul> <li>ora_srv_02</li> <li>OraDB19Home1</li> <li>oradb01</li> </ul>		Name: Oracle Sl	D:		oradb01 oradb01	
		Log mod Backup t	le: ime:	1	ARCHIVELOG /20/2023 7:03 PM	
			Local list	ener:	L	ISTENER_ORADB01

- 5. In the Publish wizard, specify the restore point at which to publish the database from and click **Next**.
- 6. Finally, specify the target linux file system location and click on **Publish** to begin the restore process.

223

Browse

7. Once the publish has completed log into the target server and run the following commands to ensure the database is running:

oracle@ora\_srv\_01> sqlplus / as sysdba

SQL> select name, open\_mode from v\$database;



## Conclusion

VMware Cloud is a powerful platform for running business-critical applications and storing sensitive data. A secure data protection solution is essential for businesses that rely on VMware Cloud to ensure business continuity and help protect against cyber threats and data loss. By choosing a reliable and robust data protection solution, businesses can be confident that their critical data is safe and secure, no matter what.

The use case presented in this documentation focuses on proven data protection technologies that highlight the integration between NetApp, VMware, and Veeam. FSx for ONTAP is supported as supplemental NFS datastores for VMware Cloud in AWS and is used for all virtual machine and application data. Veeam Backup & Replication is a comprehensive data protection solution designed to help businesses improve, automate, and streamline their backup and recovery processes. Veeam is used in conjunction with iSCSI backup target volumes, hosted on FSx for ONTAP, to provide a secure and easy to manage data protection solution for application data residing in VMware Cloud.

## **Additional Information**

To learn more about the technologies presented in this solution refer to the following additional information.

- FSx for ONTAP User Guide
- Veeam Help Center Technical Documentation
- VMware Cloud on AWS Support. Considerations and Limitations

### TR-4955: Disaster Recovery with FSx for ONTAP and VMC (AWS VMware Cloud)

Disaster Recovery Orchestrator (DRO; a scripted solution with UI) can be used to seamlessly recover workloads replicated from on-premises to FSx for ONTAP. DRO automates the recovery from the SnapMirror level, through VM registration to VMC, to network mappings directly on NSX-T. This feature is included with all VMC environments.

# Overview

Disaster recovery to cloud is a resilient and cost-effective way of protecting the workloads against site outages and data corruption events (for example, ransomware). With NetApp SnapMirror technology, on-premises VMware workloads can be replicated to FSx for ONTAP running in AWS.

Disaster Recovery Orchestrator (DRO; a scripted solution with UI) can be used to seamlessly recover workloads replicated from on-premises to FSx for ONTAP. DRO automates the recovery from the SnapMirror level, through VM registration to VMC, to network mappings directly on NSX-T. This feature is included with all VMC environments.



# **Getting started**

## Deploy and configure VMware Cloud on AWS

VMware Cloud on AWS provides a cloud-native experience for VMware-based workloads in the AWS ecosystem. Each VMware Software-Defined Data Center (SDDC) runs in an Amazon Virtual Private Cloud (VPC) and provides a full VMware stack (including vCenter Server), NSX-T software-defined networking, vSAN software-defined storage, and one or more ESXi hosts that provide compute and storage resources to the workloads. To configure a VMC environment on AWS, follow the steps at this link. A pilot-light cluster can also be used for DR purposes.



In the initial release, DRO supports an existing pilot-light cluster. On-demand SDDC creation will be available in an upcoming release.

## Provision and configure FSx for ONTAP

Amazon FSx for NetApp ONTAP is a fully managed service that provides highly reliable, scalable, high-

performing, and feature-rich file storage built on the popular NetApp ONTAP file system. Follow the steps at this link to provision and configure FSx for ONTAP.

## Deploy and configure SnapMirror to FSx for ONTAP

The next step is to use NetApp BlueXP and discover the provisioned FSx for ONTAP on AWS instance and replicate the desired datastore volumes from an on-premises environment to FSx for ONTAP with the appropriate frequency and NetApp Snapshot copy retention:

NetApp BlueXP	Account Y Workspace nimolab nimolab	Connector AWSConnCtd
Canvas My Working Environments	My Opportunities New	🖽 Go to Tabular View
+ Add Working Environment	C Enable Services	(i)
nimfax F5x for ONTAP 7 13.01 Tile Valumers Capacity aws	ntaphci-a300e9u25 On-Premises ONTAP 131.27 TM Concetly	DETAILS On-Premises ONTAP
al		SERVICES
DemoFSiXN     PSix for ONTAP	ANF Azura NatApp Files	Backup and recovery Inable - 1
5 4.74 TIB Volumes Capacity aws	© Failed	Copy & sync 1,57 Till (1) • On Data Synced
		Loading
Azure Blob Storage	Amazon S3	Classification Enable 1
O Storiege Accounts	6 Buckets aws -+	Enter Working Environment

Follow the steps in this link to configure BlueXP. You can also use the NetApp ONTAP CLI to schedule replication following this link.



A SnapMirror relationship is a prerequisite and must be created beforehand.

### **DRO** installation

To get started with DRO, use the Ubuntu operating system on a designated EC2 instance or virtual machine to make sure you meet the prerequisites. Then install the package.

## Prerequisites

- Make sure that connectivity to the source and destination vCenter and storage systems exists.
- DNS resolution should be in place if you are using DNS names. Otherwise, you should use IP addresses for the vCenter and storage systems.
- Create a user with root permissions. You can also use sudo with an EC2 instance.

## **OS requirements**

- Ubuntu 20.04 (LTS) with minimum of 2GB and 4 vCPUs
- The following packages must be installed on the designated agent VM:

- Docker
- Docker-compose
- ∘ Jq

Change permissions on docker.sock: sudo chmod 666 /var/run/docker.sock.



The deploy.sh script executes all the required prerequisites.

### Install the package

1. Download the installation package on the designated virtual machine:

git clone https://github.com/NetApp/DRO-AWS.git



The agent can be installed on-premises or within an AWS VPC.

2. Unzip the package, run the deployment script, and enter the host IP (for example, 10.10.10.10).

tar xvf DRO-prereq.tar

3. Navigate to the directory and run the deploy script as follows:

```
sudo sh deploy.sh
```

4. Access the UI using:

```
https://<host-ip-address>
```

with the following default credentials:

Username: admin Password: admin

The password can be changed using the "Change Password" option.

NetApp	
Disaster Recovery Orchestrator     Surgerl Induses with DDI Username	FSX
Pasiword	
Login	

## **DRO** configuration

After FSx for ONTAP and VMC have been configured properly, you can begin configuring DRO to automate the recovery of on-premises workloads to VMC by using the read-only SnapMirror copies on FSx for ONTAP.

NetApp recommends deploying the DRO agent in AWS and also to the same VPC where FSx for ONTAP is deployed (it can be peer connected too), so that the DRO agent can communicate through the network with your on-premises components as well as with the FSx for ONTAP and VMC resources.

The first step is to discover and add the on-premises and cloud resources (both vCenter and storage) to DRO. Open DRO in a supported browser and use the default username and password (admin/admin) and Add Sites. Sites can also be added using the Discover option. Add the following platforms:

- On-premises
  - On-premises vCenter
  - ONTAP storage system
- Cloud
  - VMC vCenter
  - FSx for ONTAP



Q O Add New Site

Success

⊘ Success

•••

•••

• 44.235.223.88

• 172.21.253.160

View VM List

Once added, DRO performs automatic discovery and displays the VMs that have corresponding SnapMirror
replicas from the source storage to FSx for ONTAP. DRO automatically detects the networks and portgroups
used by the VMs and populates them.

2 Sites

Cloud

On Prem

Destination

Source

Cloud

On Prem

1

1

1

1

I NetApp	Disaster Recovery Orch	estrator 💊 Dashboard Di	icover Resource Groups Replic	ation Plans Job Monitoring	)	. 0	- ?- @
	Back		VM List				
			Site: On Prem   vCenter: 172.2	1253.160			
				VM Protection			
	<b>C</b>	10 latastores	219 Virtual Machines	S Protected	() Unprof	216 ected	
	38 vMs				۹	Create Resource Group	
	VM Name	C VM Status	🐨 VM State (1)	🗢 DetaStore	C CPU	C Memory (MB) C	
	a300-vcsa02	0 Not Protected	() Powered On	A300_NF5_D504	76	65538	
	PFSense	0 Not Protected	() Powered On	A300_NFS_D504	4	8192	
	PFSense260	0 Not Protected	() Pownred On	A300_NFS_DS04	4	16384	
	NimDC02	0 Not Protected	(1) Powered On	A300_NFS_DS04	4	8192	
	jhRBhoja-187	0 Not Protected	() Powered On	A300_NF5_D504	4	16384	
	JhNimo-187	9 Not Protected	(1) Powered On	A300_NFS_D504	4	16384	
	NimMSdesktop	0 Not Protected	() Powered On	A300_NFS_DS04	8	12288	

The next step is to group the required VMs into functional groups to serve as resource groups.

### **Resource groupings**

After the platforms have been added, you can group the VMs you want to recover into resource groups. DRO resource groups allow you to group a set of dependent VMs into logical groups that contain their boot orders, boot delays, and optional application validations that can be executed upon recovery.

To start creating resource groups, complete the following steps:

- 1. Access Resource Groups, and click Create New Resource Group.
- 2. Under **New resource group**, select the source site from the dropdown and click **Create**.
- 3. Provide Resource Group Details and click Continue.
- 4. Select the appropriate VMs using the search option.
- 5. Select the boot order and boot delay (secs) for the selected VMs. Set the order of the power-on sequence by selecting each VM and setting up the priority for it. Three is the default value for all VMs.

Options are as follows:

- 1 The first virtual machine to power on
- 3 Default
- 5 The last virtual machine to power on
- 6. Click Create Resource Group.

netApp	Disaster Recovery Orchestrator 💊	Dashboard Discover Resource Gr	oups Replication Plans Job Monitoring	
	Resource Group	C 1 Site	Center 1	3 Virtual Machines
	1 Resource Group			Q 🔿 Create New Resource Group
	Resource Group Name	C   Site Name	₩ Source vCenter	≂   VM List
	DemoRG1	On Prem	172.21.253.160	View VM List

## **Replication plans**

You need a plan to recover applications in the event of a disaster. Select the source and destination vCenter platforms from the drop down and pick the resource groups to be included in this plan, along with the grouping of how applications should be restored and powered on (for example, domain controllers, then tier-1, then tier-2, and so on). Such plans are sometimes also called blueprints. To define the recovery plan, navigate to the **Replication Plan** tab and click **New Replication Plan**.

To start creating a replication plan, complete the following steps:

1. Access Replication Plans, and click Create New Replication Plan.



2. Under **New Replication Plan**, provide a name for the plan and add recovery mappings by selecting the source site, associated vCenter, destination site, and associated vCenter.

NetApp Disaster Recovery Orchestra	ator 💊 Dashboard Discover Resource Groups	Replication Plans Job Monitoring	≜ ¢* ?* ©*
Create New Replication	Replication Plan and Site Details (2) Select Resource	Groups (3) Set Execution Order (4) Set VM Del	ails ×
	Replication	n Plan Details	
	Plan Name		0
	Recover	ry Mapping	
	Select Source Site +		
	Causes of Capitor	Destination (Center	
	Select Source vCenter ~	Select Destination vCenter +	
	Pre-requisite - You must configure SnapMirror rela create successful replication plan	tionships between the source site and target site to $$\times$$	
	co	Intinue	

3. After Recovery mapping is completed, select the cluster mapping.

Create New Replication Plan	Replication Plan and Site Details	2 Select Resource	Groups 3 Set Execution Orde	er (4) Set VM Details	
		Replication	Plan Details		
	Plan Name			0	
	DemoRP				
		Recover	Manning		
	Source Site	0	Destination Site	0	
	On Prem		Cloud	*	
	Source vCenter	0	Destination vCenter	0	
	172.21.253.160		44.235.223.88		
		Cluster	Mapping		
	Source Site Resource	O Destination	on Site Resource	0	
	TempCluster	× .	Cluster-1 -	Add	
	-				
	Source Resource	Destination	Resource		
	A300-Cluster01	Cluster-1		Delete	

- 4. Select **Resource Group Details** and click **Continue**.
- 5. Set the execution order for the resource group. This option enables you to select the sequence of operations when multiple resource groups exist.
- 6. After you are done, select the network mapping to the appropriate segment. The segments should already be provisioned within VMC, so select the appropriate segment to map the VM.
- 7. Based on the selection of VMs, datastore mappings are automatically selected.



SnapMirror is at the volume level. Therefore, all VMs are replicated to the replication destination. Make sure to select all VMs that are part of the datastore. If they are not selected, only the VMs that are part of the replication plan are processed.

NetApp Disaster Recovery Ord     Create New Replication     Pian	hestrator 💊 Dashboard Disco	ver Resource Groups Replication Plans Job Monitoring	t Details ×
		Replication Plan Details	
		Select Execution Order	
	Resource Group Name	Execution Order	
	DemoRG1	3	
	No more Source Resource	Source/Destination network resources available for mapping Destination Resource	
	VLAN 3375	sddc-cgw-network-1 Delete	
		DataStore Mapping	
	Source DataStore	Destination Volume	
	DRO_Mini	DRO_Mini_copy	
		Previous Continue	

8. Under the VM details, you can optionally resize the VM's CPU and RAM parameters; this can be very helpful when recovering large environments to smaller target clusters or for conducting DR tests without having to provision a one-to-one physical VMware infrastructure. Also, you can modify the boot order and boot delay (seconds) for all the selected VMs across the resource groups. There is an additional option to modify the boot order if there are any changes required from those selected during the resource-group boot-order selection. By default, the boot order selected during resource-group selection is used; however, any modifications can be performed at this stage.

		VM Details			
3 vms				Q	
VM Name	No. of CPUs	Memory (MB)	NIC/IP	Boot Order	
Resource Group	: DemoRG1				
Mini_Test01	1	2048	O Static O Dynamic	3	
Mini_Test02	1	2048	<ul><li>Static</li><li>Dynamic</li></ul>	2	
Mini_Test03	1	2048	O Static O Dynamic	1	

9. Click Create Replication Plan.

Source Details     Destination Details       Peplication Plans     Image: Application Plans     Image: Application Plans							
2   Image: Constraint of the second constraint o	_			Source Details		Destination Details	
Sites vCenters Sites vCenters	B 2	plication Plans	2 1 Resource Groups	<u>△</u> 1	🛃 1	<u> </u>	🔁 1
2 Replication Plans O O Create New Replication Plan	2 Replication	Plans				0.0	Create New Replication Plan
2 Replication Plans Q O Create New Replication Plan	2 Replication	Plans				۵ ۵	Create New Replication Plan
2 Replication Plans     Q O     Create New Replication Plan       Plan Name     ○   Active Site       Status       Compliance	2 Replication Plan Name	Plans \$   Active Site	Status	Compliance	Source Site =	Q O Destination Site 💠	Create New Replication Plan
2 Replication Plans     Q O     Create New Replication Plan       Plan Name     Create New Replication Plan	2 Replication	Plans	Status	Compliance	Source Site	Q O	Create New Replication Plan
Plan Name     Create New Replication Plan       Plan Name     Create New Replication Plan       Plan Name     Compliance       Source Site     Compliance       DemoRP     Source       O Source     Active       O Not Available     On Prem       Cloud     Resource Groups	2 Replication Plan Name DemoRP	Plans           Image: Active Site           Image: Source	Status O Active	Compliance () Not Available	Source Site 👳	Q O Destination Site 0	Create New Replication Pla
Plan Name     C Mattive Site     Status     Compliance     Source Site     Image: Compliance     Compli	2 Replication Plan Name DemoRP	Plans           C         Active Site           O Source	Status ⓒ Active	Compliance ① Not Available	Source Site 🛛 👻 On Prem	Q O Destination Site 0 Cloud	Create New Replication Plan

After the replication plan is created, the failover option, the test-failover option, or the migrate option can be exercised depending on the requirements. During the failover and test-failover options, the most recent SnapMirror Snapshot copy is used, or a specific Snapshot copy can be selected from a point-in-time Snapshot copy (per the retention policy of SnapMirror). The point-in-time option can be very helpful if you are facing a corruption event like ransomware, where the most recent replicas are already compromised or encrypted. DRO shows all available points in time. To trigger failover or test failover with the configuration specified in the replication plan, you can click **Failover** or **Test failover**.

			Source Details		Destination Det	tails
B 2	lication Plans	1 Resource Groups	C 1 Sites	Centers 1	Co 1 Sites	1 vCenters
2 Replication Pla	ans				Q 0	Create New Replication Plan
Plan Name	C Active Site	Status	Compliance	Source Site 👳	Destination Site 🗘 🗧	l I
DemoRP	⊘ Source	<ul> <li>Active</li> </ul>	Healthy	On Prem	Cloud	Resource Groups
DemoRP	⊘ Source	<ul> <li>Active</li> </ul>	Healthy	On Prem	Cloud	Plan Details Resource Edit Plan
						Failover
						Test Failover
						Migrate
						Run Compliance

Failover Details	×
Volume Snapshot Details <ul> <li>Use latest snapshot (i)</li> <li>Select specific snapshot (i)</li> </ul>	
Start Failover	

The replication plan can be monitored in the task menu:

🖬 NetApp	Disaster Rec	overy Orchestrator 🔌 Dashboard Discover Res	source Groups Replication Plans	Job Monitoring		۵	¢· ?· •·
	Back		Failover Steps Replication Plan: DemoRP				
	~	Breaking SnapMirror relationships (in parallel)		⊘ Suc	cess 11.3 Seconds 🛈		
	~	Mounting volumes and creating datastores (in parallel)		⊘ Suc	cess 34.7 Seconds 🛈		
	~	Registering VMs (in parallel)		⊙ Suc	cess 13.2 Seconds 🕕		
	~	Powering on VMs in protection group - DemoRG1 - in target		⊘ Suc	cess 95.8 Seconds 🕕		
	~	Updating replication status		⊘ Suc	cess 0.5 Seconds 🛈		

After failover is triggered, the recovered items can be seen in the VMC vCenter (VMs, networks, datastores). By default, the VMs are recovered to the Workload folder.

Instrument   Image: Contract	2	2 Support Course	1 Replacement Print	219	Enveranteed UNAs	· 218 Unprotected
etamer Averagy 3 Contract 2 2 2 Contract 2 2 2 Contract 2 2 2 Contract 2 2 2 Cont	Enderson and State	C Z Zaringe finishermanik	Topplage Camas			inneithe Ven ()
Constant and Associate Ass	etamerkenners	22	Or frame	Cent	5	
Constitue Labor Tana	Outline	Fathers	172 21 26A 249	44-236-222.8 49-49-6 view		
	Contraction Contraction	<ul> <li></li></ul>	172 A - 148 Mar 172 A - 164 249	44,256,201,8 46,46,6 Yes		

Failback can be triggered at the replication-plan level. For a test failover, the tear-down option can be used to roll back the changes and remove the FlexClone relationship. Failback related to failover is a two-step process. Select the replication plan and select **Reverse data sync**.

netApp	Disaster Recovery Orc	chestrator 💊 🛛 Da	ashboard Discover	Resource Groups	Replication Plans	Job Monitoring	1		¢* @* ®*
	Replication	Plans	1 Resource Groups	Source Details	2 1 vCenters		Destination Details	vCenters	
	2 Replication Plans						Q D 64	bate New Replication Plan	
	Plan Name 🗢 DemoRP	O Destination	Status     Running In Failover M	Compliance	On Prem	Cloud	on Site 🗢   Resol	urce Groups	
	DemoRP	<ul> <li>Source</li> </ul>	<ul> <li>Active</li> </ul>	Healthy	On Prem	Cloud	Resor	Plan Details Reverse Data Sync	
								Fallback	
■ NetApp	Disaster Recovery Orc	:hestrator 🗞 🛛 Di	ashboard Discover	Resource Groups	Replication Plans	Job Monitoring	ţ.	4	¢* 0* ©*
	Back			Reverse Data	a Sync Steps Plan: DemoRP				
	✓ Powerin	ig off VMs in protection g	group - DemoRG1 - in source				J In progress	- 🛈	
	∽ Reversir	ng SnapMirror relationshi	ips (in parallel)				✓ Initialized	- 🛈	

Once completed, you can trigger failback to move back to original production site.

n NetApp	Disaster Recovery Orchestrator 💊 Dashboard Discover	Resource Groups Replication Plans Job	o Monitoring	¢• ?• ®•
	Replication Plans 2 Resource Groups	Source Details	Destination Details	
	2 Replication Plans		Q 🕤 Create New Replication Plan	
	Plan Name     Clinical Active Site     Status       DemoRP     Operation     Active	Compliance Source Site ⊘ Healthy On Prem	Cloud Resource Groups	
	DemoRP	Healthy On Prem	Cloud Resource Failback	
■ NetApp	Disaster Recovery Orchestrator 💊 Dashboard Discover	Resource Groups Replication Plans Job	b Monitoring	¢• 9• 0•
	Back	Failback Steps		
	Powering off VMs in protection group - DemoRG1 - in target	Replication Plan: DemoRP	( in progress - O	

~	Powering off VMs in protection group - DemoRG1 - in target	C In progress	-0
~	Unregistering VMs in target (in parallel)	<ul> <li>Initialized</li> </ul>	- ①
~	Unmounting volumes in target (in parallel)	✓ Initialized	- ①
~	Breaking reverse SnapMirror relationships (in parallel)	<ul> <li>Initialized</li> </ul>	- ①
~	Updating VM networks (in parallel)	✓ Initialized	- 0
~	Powering on VMs in protection group - DemoRG1 - in source	✓ Initialized	- ①
~	Deleting reverse SnapMirror relationships (in parallel)	✓ Initialized	- 0
~	Resuming SnapMirror relationships to target (in parallel)	<ul> <li>Initialized</li> </ul>	-0

From NetApp BlueXP, we can see that replication health has broken off for the appropriate volumes (those that were mapped to VMC as read-write volumes). During test failover, DRO does not map the destination or replica volume. Instead, it makes a FlexClone copy of the required SnapMirror (or Snapshot) instance and exposes the FlexClone instance, which does not consume additional physical capacity for FSx for ONTAP. This process makes sure that the volume is not modified and replica jobs can continue even during DR tests or triage workflows. Additionally, this process makes sure that, if errors occur or corrupted data is recovered, the recovery can be cleaned up without the risk of the replica being destroyed.

III NetApp	Disaster Recovery Orches	trator 💊 Dashboard D	iscover Resource Groups Replic	etion Plans Job Monitoring			¢* @* \$*
	Constant 2 Sites	Resource Group	Plans 2	219 VMs	Protected VMs 3 Protected	0 216 Unprotected	
	Environments		Topology Canvas			Immersive View 👩	
	2 Virtual Environments	2 Storage Environments					
	vCenter Summary			)	aws		
	Custers	C 22 Folders	On Prem 172.21.253.160 172.21.254.210		Cloud 44,235,223,88 10,49,0,191		
	Datastores	9 45 Networks					
	Execution Jobs		Replication Plans				
	<b>Ø</b> 3	O	Replication Plan	Active Site	Status		
	Total Jobs	and staff and	DemoRP.	<ul> <li>Source</li> </ul>	@ Active	•).	

### Ransomware recovery

Recovering from ransomware can be a daunting task. Specifically, it can be hard for IT organizations to pinpoint where the safe point of return is and, once that is determined, to protect recovered workloads from reoccurring attacks from, for example, sleeping malware or vulnerable applications.

DRO addresses these concerns by enabling you to recover your system from any available point in time. You can also recover workloads to functional and yet isolated networks so that applications can function and communicate with each other in a location where they are not exposed to north-south traffic. This gives your security team a safe place to conduct forensics and make sure there is no hidden or sleeping malware.

### **Benefits**

- Use of the efficient and resilient SnapMirror replication.
- Recovery to any available point in time with Snapshot copy retention.
- Full automation of all required steps to recover hundreds to thousands of VMs from the storage, compute, network, and application validation steps.
- Workload recovery with ONTAP FlexClone technology using a method that doesn't change the replicated volume.
  - Avoids risk of data corruption for volumes or Snapshot copies.
  - Avoids replication interruptions during DR test workflows.
  - Potential use of DR data with cloud computing resources for workflows beyond DR such as DevTest, security testing, patch or upgrade testing, and remediation testing.
- CPU and RAM optimization to help lower cloud costs by allowing recovery to smaller compute clusters.

#### Using Veeam Replication and FSx for ONTAP for Disaster recovery to VMware Cloud on AWS

Amazon FSx for NetApp ONTAP integration with VMware Cloud on AWS is an AWS-

managed external NFS datastore built on NetApp's ONTAP file system that can be attached to a cluster in the SDDC. It provides customers with flexible, high-performance virtualized storage infrastructure that scales independently of compute resources.

Author: Niyaz Mohamed - NetApp Solutions Engineering

## Overview

For those customers looking to use VMware Cloud on AWS SDDC as the disaster recovery target, FSx for ONTAP datastores can be used to replicate data from on-premises using any validated third-party solution that provides VM replication capability. By adding FSx for ONTAP datastore, it will enable cost optimised deployment than building VMware cloud on AWS SDDC with enormous amount of ESXi hosts just to accommodate the storage.

This approach also helps customers to use pilot light cluster in VMC along with FSx for ONTAP datastores to host the VM replicas. The same process can also be extended as a migration option to VMware Cloud on AWS by gracefully failing over the replication plan.

## **Problem Statement**

This document describes how to use FSx for ONTAP datastore and Veeam Backup and replication to set up disaster recovery for on-premises VMware VMs to VMware Cloud on AWS using the VM replication functionality.

Veeam Backup & Replication allows onsite and remote replication for disaster recovery (DR). When virtual machines are replicated, Veeam Backup & Replication creates an exact copy of the VMs in the native VMware vSphere format on the target VMware Cloud on AWS SDDC cluster and keeps the copy synchronized with the original VM.

Replication provides the best recovery time objective (RTO) values as there is a copy of a VM in the ready-tostart state. This replication mechanism ensures that the workloads can quickly start in VMware Cloud on AWS SDDC in case of a disaster event. The Veeam Backup & Replication software also optimizes traffic transmission for replication over WAN and slow connections. In addition, it also filters out duplicate data blocks, zero data blocks, swap files and excluded VM guest OS files, and compresses the replica traffic.

To prevent replication jobs from consuming the entire network bandwidth, WAN accelerators and network throttling rules can be put in place. The replication process in Veeam Backup & Replication is job driven which means replication is performed by configuring replication jobs. In case of a disaster event, failover can be triggered to recover the VMs by failing over to its replica copy.

When failover is performed, a replicated VM takes over the role of the original VM. Fail over can be performed to the latest state of a replica or to any of its good known restore points. This enables ransomware recovery or isolated testing as needed. In Veeam Backup & Replication, failover and failback are temporary intermediate step that should be further finalized. Veeam Backup & Replication offers multiple options to handle different disaster recovery scenarios.



## **Solution Deployment**

### **High level steps**

- 1. Veeam Backup and Replication software is running in on-premises environment with appropriate network connectivity.
- Configure VMware Cloud on AWS, see the VMware Cloud Tech Zone article VMware Cloud on AWS integration with Amazon FSx for NetApp ONTAP Deployment Guide to deploy, configure VMware Cloud on AWS SDDC and FSx for ONTAP as NFS datastore. (A pilot-light environment set up with a minimal configuration can be used for DR purposes. VMs will fail over to this cluster in the event of an incident, and additional nodes can be added).
- 3. Set up replication jobs to create VM replicas using Veeam Backup and Replication.
- 4. Create failover plan and perform failover.
- 5. Switch back to production VMs once the disaster event is complete and primary site is Up.

### Pre-requisites for Veeam VM Replication to VMC and FSx for ONTAP datastores

- 1. Ensure Veeam Backup & Replication backup VM is connected to the source vCenter as well as the target VMware cloud on AWS SDDC clusters.
- 2. The backup server must be able to resolve short names and connect to source and target vCenters.
- 3. The target FSx for ONTAP datastore must have enough free space to store VMDKs of replicated VMs

For additional information, refer to "Considerations and Limitations" covered here.

### **Deployment Details**

Veeam Backup & Replication leverages VMware vSphere snapshot capabilities and during replication, Veeam Backup & Replication requests VMware vSphere to create a VM snapshot. The VM snapshot is the point-in-time copy of a VM that includes virtual disks, system state, configuration and so on. Veeam Backup & Replication uses the snapshot as a source of data for replication.

To replicate VMs, follow the below steps:

- 1. Open the Veeam Backup & Replication Console.
- 2. On the Home view, select Replication Job > Virtual machine > VMware vSphere.
- 3. Specify a job name and select the appropriate advanced control checkbox. Click Next.
  - Select the Replica seeding check box if connectivity between on-premises and AWS has restricted bandwidth.
  - Select the Network remapping (for AWS VMC sites with different networks) check box if segments on VMware Cloud on AWS SDDC do not match that of on-premises site networks.
  - If the IP addressing scheme in on-premises production site differs from the scheme in the AWS VMC site, select the Replica re-IP (for DR sites with different IP addressing scheme) check box.



4. Select the VMs that needs to be replicated to FSx for ONTAP datastore attached to VMware Cloud on AWS SDDC in the Virtual Machines step. The Virtual machines can be placed on vSAN to fill the available vSAN datastore capacity. In a pilot light cluster, the usable capacity of a 3-node cluster will be limited. The rest of the data can be replicated to FSx for ONTAP datastores. Click Add, then in the Add Object window select the necessary VMs or VM containers and click Add. Click Next.

### Virtual Machines

Select one or more VMs to replicate. Use exclusion settings to exclude specific VMs and virtual disks from replication.

	Name	Туре	Size	^	Add
Virtual Machines	TestVeeam21	Virtual Machine	873 MB		Remove
Destination	TestVeeam22	Virtual Machine	890 MB		nemme
o combaatin	TestVeeam23	Virtual Machine	883 MB		
Network	TestVeeam24	Virtual Machine	879 MB		Exclusions
	TestVeeam25	Virtual Machine	885 MB		Source
Job Settings	TestVeeam26	Virtual Machine	883 MB		
D. L. T	TestVeeam27	Virtual Machine	879 MB		
Data Iranster	TestVeeam28	Virtual Machine	880 MB		1 Up
Guest Processing	TestVeeam29	Virtual Machine	878 MB		+ Down
y	TestVeeam30	Virtual Machine	876 MB		
Schedule	TestVeeam31	Virtual Machine	888 MB		
	TestVeeam32	Virtual Machine	881 MB		
Summary	TestVeeam33	Virtual Machine	877 MB		
	TestVeeam34	Virtual Machine	875 MB		
	TestVeeam35	Virtual Machine	882 MB		Recalculate
	WinSQL401	Virtual Machine	20.3 GB		
	WinSQL405	Virtual Machine	24.2 GB		Total size:
	Pharman and			*	120 GB

5. After that, select the destination as VMware Cloud on AWS SDDC cluster / host and the appropriate resource pool, VM folder and FSx for ONTAP datastore for VM replicas. Then Click **Next**.

Name	Host or cluster:	
irtual Machines		Choose
estination	Resource pool:	
Jetwork	Resources	Choose
ob Settings	Pick resource pool for selected replicas VM folder:	
Data Transfer	vm	Choose
Suest Processing	Pick VM folder for selected replicas	-5.0
chedule	Datastore:	
, cricolore	Veeam [5.6 TB free]	Choose
ummary	Pick datastore for selected virtual disks	
		1

6. In the next step, create the mapping between source and destination virtual network as needed.

Name	Network mapping:		13.7/
Virtual Machines	Source network	Target network	Add
Destination	VM_3508 (vDS-Switch0)	SegmentTemp	Edit
Network			Remove
Job Settings			
Data Transfer			
Guest Processing			
Schedule			
Summary			

- 7. In the **Job Settings** step, specify the backup repository that will store metadata for VM replicas, retention policy and so on.
- 8. Update the **Source** and **Target** proxy servers in the **Data Transfer** step and leave **Automatic** selection (default) and keep **Direct** option selected and click **Next**.
- 9. At the **Guest Processing** step, select **Enable application-aware processing** option as needed. Click **Next**.

Name Virtual Machines	Enable application-aware processing Detects and prepares applications for consistent backup, performs transaction logs configures the OS to perform required application restore steps upon first boot.	processing, and
Destination	Customize application handling options for individual machines and applications Guest interaction proxy:	Applications
Network	Automatic selection	Choose
lob Settings	Guest QS credentials:	
Data Transfer		Add
	Manage accounts	
Suest Processing	Customize guest OS credentials for individual machines and operating systems	Credentials
ichedule	Verify network connectivity and credentials for each machine included in the job	Test Now

- 10. Choose the replication schedule to run the replication job to run on a regular basis.
- 11. At the **Summary** step of the wizard, review details of the replication job. To start the job right after the wizard is closed, select the **Run the job when I click Finish** check box, otherwise leave the check box unselected. Then click **Finish** to close the wizard.

Larkey Replacement CDP Alt - Alt - Soley Former 1991	with the state of								
Name	$Q_{\rm c}$ (species as adjust summarized as a second	(fer.:)	- 20	T Alpha					
r Sy Ma © Topicator - ⊴Thopica	Name 7 III ANR Anglobb?? III ANS Anglobb??	Non Application	Chiatta B 2	Slatun Stupped	2 days ago 21 days ago	Cart Result Failed	Next Run exct scheduled+	Dute-1 Dute-1	Control by VEAMERSR/00-Administration at 2/14/202 Created by VEAMERSR/00-Administration at 2/16/202
So Really So Really So Real House (2) Last 24 House (2) Last 24 House (2) Last (3) Last (4) Last (4) Last (4) Last (5) Last (	徽 79x41,18744,22230318 徽 75x41,9xy3x401,20230318	Volvene Replication Volvene Replication	н 3	Shapped Shapped	2 days ago 8 days ago	Success Success	rnot scheduled+ rnot scheduled+	172.33.160.68 172.30.160.60	Central by VIDAMEN950/05/Administrator at 1/10/202 Central by VIDAMEN950/05/Administrator at 1/10/202

Once the replication job starts, the VMs with the suffix specified will be populated on the destination VMC SDDC cluster / host.

E+ Home View Job					Veesm Beckup a	nd Replication					- 8		
State Supp Retry Statestics Report	Gove Distlik Delete												
Asene	Q. Type in an object name to a	earth fur		×									
5.Ws	Name	Tune	Objects	Status .	Levi Ruit	Last Recult	Next Run	Ternet	Description \$				
AD Replication	ML AVE Replayout	Whenese Realization		Stonet	10 days and	former	cost scheduled a	Output 1	Created by VERAM	Protection Anterior Interaction	A 116/2023 2:12 AL		
Ell Replicas	-18 ANF Repish01	WAware Replication	Minare Replication 2 Minare Replication 6		tion 2 stopp	Stopper	Stopped 39 osys age	Ja days ago Success 6 days ago Faled	<not scheduled=""></not>	Cluster-1	Created by VEEAME	CPSRV03\Administrator a	2/16/2023 T:27 AF
Seady	鉴 FSAN, Replich01, 20230313	Whivare Replication	n 5	Stopped	1 days ago	Says ago Success	<not scheduled=""></not>	172.30.160.66	Created by VEEAMBICPGRV05\Administrator at 3/11/2023 2		3/13/2023 2:53 AF		
Failover Plans	(00 F54N, 16V84, 20230316	Where Replication	s 16	Stopped	I days ago	Success	unot scheduled>	172.30.160.68	Created by VEFAMI	KPSAV05-Administrator a	1/16/2023 6-57 AM		
	SUMMARY Duration: 01 Processing inter 05	DA 21:27 Pro I MB/s Ree	TA second ed	256 GB (100%) 256 GB 38.0 MB (+996	STATUS Success Warnings d. Econs		16 O	UCHISUT (ALL TIME	5		Speed: 634 Mi		
	Name S	tatus Acti	ce.								Duration		
	Name S DistVecan01	tatus Action Success (1) p	un Hocessing Testi	/eesm05							Duration 08:13		
	Name S TestVeean01 S TestVeean02 O	tation Activ Soccess (0) ( Soccess (0) (	on Processing Test	/eeam05 /eeam06							Duration 08:13 07:09		
	Name S TestVecan01 TestVecan02 TestVecan03	tatus Actio Success O p Success O p Success O p	on Processing Test Processing Test	/eesm05 /eesm06 /eesm07							Duration 08:13 07:09 13:21		
	Name 5 Testvecan01 6 Testvecan02 6 Testvecan03 6 Testvecan04 6	tatus Actio 5 Success O s 9 Success O s 9 Success O s 9 Success O s	on Processing Test Processing Test Processing Test Processing Test	/eean/05 /eean/06 /eean/07 /eean/08							Duration 08:13 07:09 13:21 09:05		
	Name S Testvecan01 C Testvecan02 C Testvecan04 C Testvecan04 C Testvecan04 C	tatus Action 5 Success O s 9 Success O s 9 Success O s 9 Success O s 9 Success O s	en hocessing Test hocessing Test hocessing Test hocessing Test hocessing Test	/eeam05 /eeam06 /eeam08 /eeam08 /eeam09							Duration 08:13 07:09 13:21 09:05 14:59		
and Hanny	Name S TestVecan01 TestVecan02 TestVecan03 TestVecan04 TestVecan05 TestVecan06	tatun Artis Success O 5 Success O 5 Success O 5 Success O 5 Success O 5 Success O 5	on Processing Test Processing Test Processing Test Processing Test Processing Test Processing Test	/eeam05 /eeam06 /eeam07 /eeam08 /eeam09 /eeam10							Duration 08:13 07:09 13:21 09:05 14:59 06:53		
Filos	Name S FestVecan01 S TestVecan02 S TestVecan03 S TestVecan04 S TestVecan05 S TestVecan06 S TestVecan06 S	tatun Action Success O S Success O S Success O S Success O S Success O S Success O S Success O S	en Indeetsing Test Indeetsing Test Indeetsing Test Indeetsing Test Indeetsing Test Indeetsing Test Indeetsing Test	/eeam05 /eeam06 /eeam07 /eeam08 /eeam09 /eeam10 /eeam11							Duration 08:13 07:09 13:21 09:05 14:39 06:53 15:47		
Tione	Name S TestVean01 C TestVean02 C TestVean03 C TestVean04 C TestVean06 C TestVean06 C TestVean07 C TestVean07 C TestVean07 C TestVean07 C TestVean07 C TestVean07 C TestVean07 C TestVean07 C TestVean08 C TestVean08 C TestVean08 C TestVean07 C TestVean08 C	tatun Action Success O S Success O S	en hocessing Test hocessing Test hocessing Test hocessing Test hocessing Test hocessing Test hocessing Test	/eeam05 /eeam06 /eeam07 /eeam09 /eeam10 /eeam10 /eeam11 /eeam12							Duration 08:13 07:09 13:21 09:05 14:39 08:53 15:47 08:45		
None	Norma S TestVecan01 C TestVecan02 C TestVecan03 C TestVecan03 C TestVecan05 C TestVecan06 C TestVecan06 C TestVecan00 C TestVecan0 C TestVecan0 C TestVecan00 C TestVecan00 C	terum Activ Soccess O = Soccess O =	en hocessing Test hocessing Test hocessing Test hocessing Test hocessing Test hocessing Test hocessing Test hocessing Test	leesm05 leesm06 leesm07 leesm09 leesm10 leesm11 leesm11 leesm12 leesm13							Duration 08:13 07:09 13:21 09:05 14:39 06:53 15:47 08:45 09:24		
None Montoy	Name S an TestVecan01 C an TestVecan02 C an TestVecan03 C an TestVecan05 C an TestVecan06 C an TestVecan06 C an TestVecan06 C an TestVecan09 C	tatur Actó ) Succes 0 5 ) Suces 0 5 ) Succes 0 5 ) Suc	In the second se	leeam05 leeam06 leeam07 leeam08 leeam09 leeam10 leeam11 leeam13 leeam13							Duration 08:13 07:09 13:21 09:05 14:39 08:53 15:47 08:45 09:24 14:34		
Tione Newson To Secure Infrastructure	Name S TestVean01 C TestVean02 C TestVean02 C TestVean04 C TestVean06 C TestVean06 C TestVean07 C TestVean07 C TestVean08 C	Jores         Activity           Joress         0	Increasing Test Increasing Test Increasing Test Increasing Test Increasing Test Increasing Test Increasing Test Increasing Test Increasing Test	icean05 icean05 icean07 icean08 icean09 icean10 icean10 icean13 icean14 icean15							Duration 08:13 07:09 13:21 09:05 14:39 08:53 15:47 08:45 09:24 14:34 16:16		
Flatse Ventory Al Rectop Infrastructure	Name         S           TestVecan01         C           TestVecan02         TestVecan03           TestVecan03         C           TestVecan04         TestVecan05           TestVecan05         TestVecan06           TestVecan06         TestVecan06           TestVecan08         TestVecan06           TestVecan08         TestVecan08           TestVecan08         TestVecan09           TestVecan10         TestVecan11           TestVecan12         TestVecan12	tatus Actional Action	ter hocessing Test hocessing Test hocessing Test hocessing Test hocessing Test hocessing Test hocessing Test hocessing Test hocessing Test	leeam05 keam06 keam07 keam08 keam09 keam10 keam11 keam13 keam14 keam14 keam14							Dustion 08:13 07:09 13:21 09:05 14:59 08:53 15:47 09:24 14:34 16:16 17:21		
Kone     Mone     Mone     Mone     Mone     Society     Soci	Norme         Status           Test/sean01         Comparison           Test/sean02         Comparison           Test/sean03         Comparison           Test/sean04         Comparison           Test/sean05         Comparison           Test/sean06         Test/sean06           Test/sean07         Comparison           Test/sean08         Comparison           Test/sean09         Comparison           Test/sean10         Comparison           Test/sean11         Comparison           Test/sean12         Comparison	tation Action 5 Success Of 2 5 Success Of 2	In Accessing Test Inaccessing Test	leeam05 leeam06 leeam07 leeam09 leeam10 leeam10 leeam11 leeam12 leeam14 leeam14 leeam15 leeam16 moonaed for proving	2010						Dunation 02:53 07:09 13:21 09:05 14:59 06:53 15:47 06:45 09:24 14:34 16:16 17:21 00:00		
Tacky inflative	Name S TestVecan01 C TestVecan02 C TestVecan02 C TestVecan03 C TestVecan04 C TestVecan06 C TestVecan06 C TestVecan07 C TestVecan09 C TestVecan00 C TestVecan00 C TestVecan10 C TestVecan11 C TestVecan11 C TestVecan13 C TestVecan14 C	tatus Actional Construction of the second of	nocessing Text Inocessing Text	Iteram05 Iteram06 Iteram07 Iteram08 Iteram09 Iteram10 Iteram10 Iteram12 Iteram13 Iteram13 Iteram15 Iteram16 Iteram16 Iteram16 Iteram16	cessing Network 42% - T	n# 10%					Duration 02:13 07:09 13:21 09:05 14:39 06:53 15:47 06:45 09:24 14:34 14:34 14:34 14:34 14:34 14:34		
Hatte     Hatte     Hatte     Investory     Secure Inflastructure     Tope Inflastructure     Tope Inflastructure     Tope Inflastructure	Name S TestVecan01 C TestVecan02 C TestVecan03 C TestVecan04 C TestVecan05 C TestVecan06 C TestVecan09 C TestVecan09 C TestVecan09 C TestVecan09 C TestVecan10 C TestVecan10 C TestVecan11 C TestVecan12 C TestVecan12 C TestVecan13 C TestVecan14 C TestVecan14 C TestVecan14 C TestVecan14 C TestVecan14 C TestVecan14 C TestVecan14 C TestVecan15 C	tatus Actional Action	en hacessing Text hacessing Text hac	Veesm05 Veesm05 Veesm07 Veesm08 Veesm10 Veesm11 Veesm11 Veesm12 Veesm14 Veesm14 Veesm15 Veesm16 em queued for proc % > Proxy 36% > 1 whit Tonor	cessing Network 42% > Tee	get 30%					Dustion 0513 07.09 13.21 09.05 14.59 08.53 15.47 09.24 14.34 16.16 17.21 00.00		

For additional information for Veeam replication, refer to How Replication Works.

#### Step 2: Create a failover plan

When the initial replication or seeding is complete, create the failover plan. Failover plan helps in performing failover for dependent VMs one by one or as a group automatically. Failover plan is the blueprint for the order in which the VMs are processed including the boot delays. The failover plan also helps to ensure that critical dependant VMs are already running.

To create the plan, navigate to the new sub section called Replicas and select Failover Plan. Choose the appropriate VMs. Veeam Backup & Replication will look for the closest restore points to this point in time and use them to start VM replicas.



The failover plan can only be added once the initial replication is complete and the VM replicas are in Ready state.



The maximum number of VMs that can be started simultaneously when running a failover plan is 10.



During the failover process, the source VMs will not be powered off.

To create the Failover Plan, do the following:

- 1. On the Home view, select Failover Plan > VMware vSphere.
- 2. Next, provide a name and a description to the plan. Pre and Post-failover script can be added as required. For instance, run a script to shutdown VMs before starting the replicated VMs.

题 王· Home		Veeam Backup and Replication - 5
Backup Replication CDP Job * Jobs * Policy Primary Jobs Home	New Falover Plan General Type in name a	M description for this failover plan, and optionally specify scripts to trigger before and after the failover.
Success     Success     Success     Success     Success	General Virtual Machines Summary	Name:     [Failover,VMC_Demo]       Description:     Created by VEEAM8KPSRV05\Administrator at 8/15/2023 7:50 AM.       Pre-failover script:     Image: Created by VEEAM8KPSRV05\Administrator at 8/15/2023 7:50 AM.
Home  Home  Kernel  Ke		Next > Cancel

3. Add the VMs to the plan and modify the VM boot order and boot delays to meet the application dependencies.

Start Retry Unde	t Deleta		
Home	Q. Type in an object name	ta search tur 🛛 🗙	
Solos     September     September     Depication     Depication     Depication     Depication     Secoly     Secoly     Secoly     Secoly	Nome † Edit Fallover Plan (Fallover Virtual Machin Add virtual ma met	Plactions         Statut         Number of VMs           VMX_Demo)         X           es            thines to be failed over as a part of this plan. Use VM order and delays to ensure all application dependencies are	
Cill Last 24 Hours     S Success	C BAAM		
	General	Name Drine Renica state 6 Add VM	
	Virtual Machines	April 2 State Control Con	
		TertViewm04	
ff Harse		Charles and a second	
Inventory     Include Infrastructure     Tope Infrastructure     Tope Infrastructure     Tope Infrastructure     Tope Infrastructure		< Brevieus Apply Einich Cancel	

For additional information for creating replication jobs, refer Creating Replication Jobs.

During failover, the source VM in the production site is switched over to its replica at the disaster recovery site. As part of the failover process, Veeam Backup & Replication restores the VM replica to the required restore point and moves all I/O activities from the source VM to its replica. Replicas can be used not only in case of a disaster, but also to simulate DR drills. During failover simulation, the source VM remains running. Once all the necessary tests have been conducted, you can undo the failover and return to normal operations.



Make sure network segmentation is in place to avoid IP conflicts during DR drills.

To start the failover plan, simply click in **Failover Plans** tab and right click on the failover plan. Select **Start**. This will failover using the latest restore points of VM replicas. To fail over to specific restore points of VM replicas, select **Start to**.

ome	Q Type in an object name to	search for		×		
lobs	Name 1	Platform	Status	Number of VMs		
瘤 Replication	Failover_VMC_Demo	/Mware	Ready	21		
🗊 Replicas	10-0-0-0-0-0-0-0-0-0-0-0-0-0-0-0-0-0-0-					
Ready						
S Active (1)						
Failover Plans						
Last 24 Hours						
Success						
Failed						
and the second second	223					
lame: Failover_VMC_Demo lestore type: Failover Plan	Statu Start	s: In progress time: 3/23/2023 1	s 11:53:59 PM			
Vame: Failover_VMC_Demo Restore type: Failover Plan nitiated by: VEEAMBKPSRV05\Administrator	Statu Start	s: In progress time: 3/23/2023 1	s 11:53:59 PM		Cancel restore task	
lame: Failover_VMC_Demo lestore type: Failover Plan nitiated by: VEEAMBKPSRV05\Administrator	Statu Start	s: In progress time: 3/23/2023 1	s 11:53:59 PM		Cancel restore task	
Vame: Failover_VMC_Demo testore type: Failover Plan initiated by: VEEAMBKPSRV05\Administrator VM name Status TestVeeam02 () Processing	Statu Start	s: In progress time: 3/23/2023 1	s 11:53:59 PM		Cancel restore task	0
Hame:         Failover_VMC_Demo           testore type:         Failover Plan           initiated by:         VEEAMBKPSRV05\Administrator           /M name         Status           TestVeeam02         Processing           TestVeeam01         Processing	Statu Start Log Message () Processing VM: TestVeeam01	s: In progress time: 3/23/2023 1	s 11:53:59 PM		Cancel restore task Duratio 0:00:2	'n ,
Iame:         Failover_VMC_Demo           testore type:         Failover Plan           initiated by:         VEEAMBKPSRV05\Administrator           /M name         Status           TestVeeam02         Processing           TestVeeam03         Processing	Statu Start Message () Processing VM: TestVeeam01 (2) Waiting 0 sec before the nest V	s: In progress time: 3/23/2023 1	s 11 <mark>:</mark> 53:59 PM		Cancel restore task Duration 0:00:23	in '
lame: Failover_VMC_Demo lestore type: Failover Plan nitiated by: VEEAMBKPSRV05\Administrator /M name Status TestVeeam02 () Processing TestVeeam03 () Processing TestVeeam04 () Processing	Statu Start Log Message Processing VM: TestVeeam01 © Waiting 0 sec before the next V © Processing VM: TestVeeam03	s: In progress time: 3/23/2023 1 M	s 11:53:59 PM		Cancel restore task Duration 0.00.23 0:00-23	n '
Image:         Failover_VMC_Demo           lestore type:         Failover Plan           initiated by:         VEEAMBKPSRV05\Administrator           //M name         Status           TestVeeam02         Processing           TestVeeam03         Processing           TestVeeam04         Processing           TestVeeam05         Processing	Statu Start Log Message Processing VM: TestVeeam01 Processing VM: TestVeeam03 Processing VM: TestVeeam03 © Waiting 0 sec before the next V	s: <b>In progress</b> time: 3/23/2023 1 M	s 11:53:59 PM		Cancel restore task Duration 0.00.23 0:00.23	m * 3
Vame: Failover_VMC_Demo Restore type: Failover Plan initiated by: VEEAMBKPSRV05\Administrator VM name Status TestVeeam02 Processing TestVeeam04 Processing TestVeeam05 Processing TestVeeam05 Processing TestVeeam06 Processing	Statu Start Log Message Processing VM: TestVeeam01 Processing VM: TestVeeam03 Waiting 0 sec before the next V Processing VM: TestVeeam04	s: In progress time: 3/23/2023 1 M	s 11:53:59 PM		Cancel restore task Duration 0.00.23 0:00.23 0:00.23	in ' 3 3
Failover_VMC_Demo           lestore type:         Failover Plan           initiated by:         VEEAMBKPSRV05\Administrator           //M name         Status           TestVeeam02         Processing           TestVeeam03         Processing           TestVeeam04         Processing           TestVeeam05         Processing           TestVeeam06         Processing           TestVeeam07         Processing	Statu Start Log Message Processing VM: TestVeeam01 Waiting 0 sec before the next V Processing VM: TestVeeam03 Waiting 0 sec before the next V Processing VM: TestVeeam04 © Waiting 0 sec before the next V	s: In progress time: 3/23/2023 1 M M	s 11:53:59 PM		Cancel restore task Duratio 0.0023 0.0023 0.0023	ın ' 3 3
Name:     Failover_VMC_Demo       Restore type:     Failover Plan       initiated by:     VEEAMBKPSRV05\Administrator       VM name     Status       TestVeeam02     Processing       TestVeeam03     Processing       TestVeeam05     Processing       TestVeeam05     Processing       TestVeeam06     Processing       TestVeeam07     Processing       TestVeeam08     Processing	Statu Start Log Message Processing VM: TestVeeam01 Waiting 0 sec before the nest V Processing VM: TestVeeam03 Waiting 0 sec before the nest V Processing VM: TestVeeam04 Waiting 0 sec before the nest V Processing VM: TestVeeam04	s: In progress time: 3/23/2023 1 M M	s 11:53:59 PM		Cancel restore task Duratio 0.00;23 0.00;23 0:	n ' 3 2 2
Name:     Failover_VMC_Demo       Restore type:     Failover Plan       initiated by:     VEEAMBKPSRV05\Administrator       VM name     Status       TestVeeam02     Processing       TestVeeam04     Processing       TestVeeam05     Processing       TestVeeam06     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam09     Processing	Statu Start Log Message Processing VM: TestVeeam01 Waiting 0 sec before the nest V Processing VM: TestVeeam03 Waiting 0 sec before the nest V Processing VM: TestVeeam04 Waiting 0 sec before the nest V Processing VM: TestVeeam05 Waiting 0 sec before the nest V	s: In progress time: 3/23/2023 1 M M M	s 11:53:59 PM		Cancel restore task Duratio 0:00:23 0:00:23 0:00:23	2 2
Name:     Failover_VMC_Demo       Restore type:     Failover Plan       nitiated by:     VEEAMBKPSRV05\Administrator       VM name     Status       TestVeeam02     Processing       TestVeeam03     Processing       TestVeeam04     Processing       TestVeeam05     Processing       TestVeeam06     Processing       TestVeeam06     Processing       TestVeeam08     Processing       TestVeeam08     Processing       TestVeeam09     Processing       TestVeeam09     Processing       TestVeeam09     Processing       TestVeeam010     Processing	Statu Start Log Message Processing VM: TestVeeam01 Waiting 0 sec before the nest V Processing VM: TestVeeam03 Waiting 0 sec before the nest V Processing VM: TestVeeam04 Waiting 0 sec before the nest V Processing VM: TestVeeam05 Waiting 0 sec before the nest V Processing VM: TestVeeam05	s: In progress time: 3/23/2023 1 M M M	s 11:53:59 PM		Cancel restore task Duratio 0.00.23 0.	in ' 3 2 2 2
Name:     Failover_VMC_Demo       testore type:     Failover Plan       initiated by:     VEEAMBKPSRV05\Administrator       VM name     Status       TestVeeam02     Processing       TestVeeam03     Processing       TestVeeam04     Processing       TestVeeam05     Processing       TestVeeam06     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam09     Processing       TestVeeam09     Processing       TestVeeam10     Processing       TestVeeam11     Processing	Statu Start Log Message Processing VM: TestVeeam01 Waiting 0 sec before the nest V Processing VM: TestVeeam03 Waiting 0 sec before the nest V Processing VM: TestVeeam04 Waiting 0 sec before the nest V Processing VM: TestVeeam05 Waiting 0 sec before the nest V Processing VM: TestVeeam06 Waiting 0 sec before the nest V	s: In progress time: 3/23/2023 1 M M M M M	s 11:53:59 PM		Cancel restore task Duratio 0.00;2: 0:	n ' 3 2 2 2
Name:     Failover_VMC_Demo       Restore type:     Failover Plan       nitiated by:     VEEAMBKPSRV05\Administrator       VM name     Status       TestVeeam02     Processing       TestVeeam03     Processing       TestVeeam04     Processing       TestVeeam05     Processing       TestVeeam06     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam09     Processing       TestVeeam09     Processing       TestVeeam09     Processing       TestVeeam10     Processing       TestVeeam10     Processing       TestVeeam11     Pending       TestVeeam12     Pending	Statu Statu Start Message Processing VM: TestVeeam01 Waiting 0 sec before the nest V Processing VM: TestVeeam03 Waiting 0 sec before the nest V Processing VM: TestVeeam04 Waiting 0 sec before the nest V Processing VM: TestVeeam05 Waiting 0 sec before the nest V Processing VM: TestVeeam06 Waiting 0 sec before the nest V Processing VM: TestVeeam06 Waiting 0 sec before the nest V Processing VM: TestVeeam07	s: In progress time: 3/23/2023 1 M M M M M	s 11:53:59 PM		Cancel restore task Duratio 0:00:2 0:	n ' 3 2 2 2 2
Name:     Failover_VMC_Demo       Restore type:     Failover Plan       Initiated by:     VEEAMBKPSRV05\Administrator       VM name     Status       TestVeeam02     Processing       TestVeeam03     Processing       TestVeeam04     Processing       TestVeeam05     Processing       TestVeeam06     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam06     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam10     Processing       TestVeeam10     Processing       TestVeeam11     Pending       TestVeeam13     Pending       TestVeeam14     Pending	Statu Statu Start Message Processing VM: TestVeeam01 Waiting 0 sec before the next V Processing VM: TestVeeam03 Waiting 0 sec before the next V Processing VM: TestVeeam04 Waiting 0 sec before the next V Processing VM: TestVeeam05 Waiting 0 sec before the next V Processing VM: TestVeeam07 Waiting 0 sec before the next V Processing VM: TestVeeam07 Waiting 0 sec before the next V	s: In progress time: 3/23/2023 1 M M M M M M M	s 11:53:59 PM		Cancel restore task Duratio 0:00:2 0:	n ' 3 2 2 2 2
Name:     Failover_VMC_Demo       testore type:     Failover Plan       nitiated by:     VEEAMBKPSRV05\Administrator       VM name     Status       TestVeeam02     Processing       TestVeeam03     Processing       TestVeeam04     Processing       TestVeeam05     Processing       TestVeeam06     Processing       TestVeeam07     Processing       TestVeeam06     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam09     Processing       TestVeeam10     Processing       TestVeeam11     Pending       TestVeeam12     Pending       TestVeeam14     Pending       TestVeeam14     Pending	Statu Statu Start Message Processing VM: TestVeeam01 Waiting 0 sec before the next V Processing VM: TestVeeam03 Waiting 0 sec before the next V Processing VM: TestVeeam04 Waiting 0 sec before the next V Processing VM: TestVeeam05 Waiting 0 sec before the next V Processing VM: TestVeeam07 Waiting 0 sec before the next V Processing VM: TestVeeam08	s: In progress time: 3/23/2023 1 M M M M M M	s 11:53:59 PM		Cancel restore task Duratio 0:00:21 0:00:22 0:	n ' 3 2 2 2 2 2
Name:     Failover_VMC_Demo       Restore type:     Failover Plan       Initiated by:     VEEAMBKPSRV05\Administrator       VM name     Status       TestVeeam02     Processing       TestVeeam03     Processing       TestVeeam04     Processing       TestVeeam05     Processing       TestVeeam06     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam09     Processing       TestVeeam09     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam10     Processing       TestVeeam11     Pending       TestVeeam12     Pending       TestVeeam13     Pending       TestVeeam14     Pending       TestVeeam15     Pending	Statu Statu Start Message Processing VM: TestVeeam01 Waiting 0 sec before the next V Processing VM: TestVeeam03 Waiting 0 sec before the next V Processing VM: TestVeeam04 Waiting 0 sec before the next V Processing VM: TestVeeam05 Waiting 0 sec before the next V Processing VM: TestVeeam07 Waiting 0 sec before the next V Processing VM: TestVeeam07 Waiting 0 sec before the next V Processing VM: TestVeeam07 Waiting 0 sec before the next V Processing VM: TestVeeam08 Waiting 0 sec before the next V	s: In progress time: 3/23/2023 1 M M M M M M M M	s 11:53:59 PM		Cancel restore task         Duratio           0:00:21         0:00:21           0:00:22         0:00:22           0:00:21         0:00:22           0:00:22         0:00:22           0:00:22         0:00:22           0:00:22         0:00:22           0:00:22         0:00:22           0:00:22         0:00:22           0:00:22         0:00:22           0:00:22         0:00:22	n ' 3 2 2 2 2 2
Name:     Failover_VMC_Demo       testore type:     Failover Plan       initiated by:     VEEAMBKPSRV05\Administrator       VM name     Status       TestVeeam02     Processing       TestVeeam03     Processing       TestVeeam04     Processing       TestVeeam05     Processing       TestVeeam06     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam09     Processing       TestVeeam10     Processing       TestVeeam11     Processing       TestVeeam13     Pending       TestVeeam14     Pending       TestVeeam15     Pending       TestVeeam16     Pending	Statu Statu Start Message Processing VM: TestVeeam01 Waiting 0 sec before the next V Processing VM: TestVeeam03 Waiting 0 sec before the next V Processing VM: TestVeeam04 Waiting 0 sec before the next V Processing VM: TestVeeam05 Waiting 0 sec before the next V Processing VM: TestVeeam07 Waiting 0 sec before the next V Processing VM: TestVeeam07 Waiting 0 sec before the next V Processing VM: TestVeeam07 Waiting 0 sec before the next V Processing VM: TestVeeam08 Waiting 0 sec before the next V Processing VM: TestVeeam08 Waiting 0 sec before the next V	s: In progress time: 3/23/2023 1 M M M M M M M	s 11:53:59 PM		Cancel restore task         Duratio           0:00:21         0:00:21           0:00:21         0:00:21           0:00:21         0:00:21           0:00:21         0:00:21           0:00:21         0:00:21           0:00:21         0:00:21           0:00:21         0:00:21           0:00:22         0:00:21           0:00:21         0:00:21           0:00:22         0:00:21	n ' 3 2 2 2 2 2 2 2 2 2
Name:     Failover_VMC_Demo       Restore type:     Failover Plan       Initiated by:     VEEAMBKPSRV05\Administrator       VM name     Status       TestVeeam02     Processing       TestVeeam03     Processing       TestVeeam04     Processing       TestVeeam05     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam09     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam10     Processing       TestVeeam11     Pending       TestVeeam12     Pending       TestVeeam13     Pending       TestVeeam14     Pending       TestVeeam15     Pending       TestVeeam16     Pending       Weam17     Pending	Statu Start Start Start Message Processing VM: TestVeeam01 Waiting 0 sec before the next V Processing VM: TestVeeam03 Waiting 0 sec before the next V Processing VM: TestVeeam04 Waiting 0 sec before the next V Processing VM: TestVeeam05 Waiting 0 sec before the next V Processing VM: TestVeeam05 Waiting 0 sec before the next V Processing VM: TestVeeam08 Waiting 0 sec before the next V	s: In progress time: 3/23/2023 1 M M M M M M M M M M	s 11:53:59 PM		Cancel restore task	n ' 3 2 2 2 2 2 2 2 2 2
Name:     Failover_VMC_Demo       Restore type:     Failover Plan       Initiated by:     VEEAMBKPSRV05\Administrator       VM name     Status       TestVeeam02     Processing       TestVeeam03     Processing       TestVeeam04     Processing       TestVeeam05     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam09     Processing       TestVeeam10     Processing       TestVeeam11     Pending       TestVeeam13     Pending       TestVeeam14     Pending       TestVeeam15     Pending       TestVeeam16     Pending       TestVeeam16     Pending       TestVeeam17     Pending       TestVeeam18     Pending       TestVeeam14     Pending       TestVeeam15     Pending       TestVeeam16     Pending       TestVeeam16     Pending       WinSQL401     Pending       WinSQL401     Pending	Statu Start Start Message Processing VM: TestVeeam01 Processing VM: TestVeeam03 Waiting 0 sec before the next V Processing VM: TestVeeam04 Waiting 0 sec before the next V Processing VM: TestVeeam05 Waiting 0 sec before the next V Processing VM: TestVeeam06 Waiting 0 sec before the next V Processing VM: TestVeeam06 Waiting 0 sec before the next V Processing VM: TestVeeam07 Waiting 0 sec before the next V Processing VM: TestVeeam08 Waiting 0 sec before the next V Processing VM: TestVeeam08 Waiting 0 sec before the next V Processing VM: TestVeeam08 Waiting 0 sec before the next V Processing VM: TestVeeam09 Waiting 0 sec before the next V	s: In progress time: 3/23/2023 1 M M M M M M M M M	s 11:53:59 PM		Cancel restore task	n ' 3 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2
Name:     Failover_VMC_Demo       Restore type:     Failover Plan       Initiated by:     VEEAMBKPSRV05\Administrator       VM name     Status       TestVeeam02     Processing       TestVeeam03     Processing       TestVeeam04     Processing       TestVeeam05     Processing       TestVeeam06     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam09     Processing       TestVeeam09     Processing       TestVeeam01     Processing       TestVeeam03     Processing       TestVeeam04     Processing       TestVeeam05     Processing       TestVeeam10     Processing       TestVeeam11     Pending       TestVeeam12     Pending       TestVeeam13     Pending       TestVeeam14     Pending       TestVeeam15     Pending       WinSQL401     Pending       WinSQL402     Pending       WinSQL403     Pending	Statu Start Start Message Processing VM: TestVeeam01 Processing VM: TestVeeam03 Waiting 0 sec before the next V Processing VM: TestVeeam04 Waiting 0 sec before the next V Processing VM: TestVeeam05 Waiting 0 sec before the next V Processing VM: TestVeeam06 Waiting 0 sec before the next V Processing VM: TestVeeam06 Waiting 0 sec before the next V Processing VM: TestVeeam07 Waiting 0 sec before the next V Processing VM: TestVeeam08 Waiting 0 sec before the next V Processing VM: TestVeeam08 Waiting 0 sec before the next V Processing VM: TestVeeam09 Waiting 0 sec before the next V Processing VM: TestVeeam10 Waiting 0 sec before the next V	s: In progress time: 3/23/2023 1 M M M M M M M M M M M M	s 11:53:59 PM		Cancel restore task	n ' 3 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2
Name:     Failover_VMC_Demo       Restore type:     Failover Plan       initiated by:     VEEAMBKPSRV05\Administrator       VM name     Status       TestVeeam02     Processing       TestVeam03     Processing       TestVeam04     Processing       TestVeeam07     Processing       TestVeeam07     Processing       TestVeeam07     Processing       TestVeeam07     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam09     Processing       TestVeeam01     Processing       TestVeeam07     Processing       TestVeeam08     Processing       TestVeeam10     Processing       TestVeeam11     Pending       TestVeeam13     Pending       TestVeeam14     Pending       TestVeeam15     Pending       TestVeeam16     Pending       WinSQL401     Pending       WinSQL402     Pending       WinSQL403     Pending       WinSQL404     Pending	Statu Statu Statu Statu Statu Statu Message Processing VM: TestVeeam01 Waiting 0 sec before the nest V Processing VM: TestVeeam03 Waiting 0 sec before the nest V Processing VM: TestVeeam04 Waiting 0 sec before the nest V Processing VM: TestVeeam05 Waiting 0 sec before the nest V Processing VM: TestVeeam06 Waiting 0 sec before the nest V Processing VM: TestVeeam07 Waiting 0 sec before the nest V Processing VM: TestVeeam08 Waiting 0 sec before the nest V Processing VM: TestVeeam08 Waiting 0 sec before the nest V Processing VM: TestVeeam08 Waiting 0 sec before the nest V Processing VM: TestVeeam09 Waiting 0 sec before the nest V Processing VM: TestVeeam10	s: In progress time: 3/23/2023 1 M M M M M M M M M M M	s 11:53:59 PM		Cancel restore task	n 3 3 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2
Name:         Failover_VMC_Demo           Restore type:         Failover Plan           Initiated by:         VEEAMBKPSRV05\Administrator           VM name         Status           TestVeeam02         Processing           TestVeeam03         Processing           TestVeeam04         Processing           TestVeeam05         Processing           TestVeeam06         Processing           TestVeeam07         Processing           TestVeeam08         Processing           TestVeeam09         Processing           TestVeeam01         Processing           TestVeeam03         Processing           TestVeeam04         Processing           TestVeeam05         Processing           TestVeeam10         Processing           TestVeeam11         Pending           TestVeeam13         Pending           TestVeeam14         Pending           TestVeeam15         Pending           TestVeeam14         Pending           WinSQL401         Pending           WinSQL403         Pending           WinSQL4045         Pending	Statu Statu	s: In progress time: 3/23/2023 1 M M M M M M M M M M M M M M M	s 11:53:59 PM		Cancel restore task	n 3 3 2 2 2 2 2 2 2 2 2 2 2 1

The state of the VM replica changes from Ready to Failover and VMs will start on the destination VMware Cloud on AWS SDDC cluster / host.

vSphere Client Q, taget at the			ls failed over to C vCenter Server	0
8 8 8	Cluster-1 Incrition     Summary Monitor Configure Plemissions Heat	s VMs Datastores Networks Updates		
1000-Datacenter           102.00.08           102.00.08           0 Compute-HesourcePool           0 Mgmt-HesourcePool           0 TestVesem01           107 TestVesem02           107 TestVesem03           107 TestVesem06           107 TestVesem06	Cluster Details Total Processors 36 Tatar violation 0 Magnetice & 173	B Capacity and Usage Lat updated at 133 MB CHU STATE CHU READ AT 135 MB CHU READ AT 135 M	VSphere DRS  Cutier DRS VM DRS Score ()  Cutier DRS Score ()  Cutier DRS Score ()  Cutier DRS Score ()  Cutier DRS Score ()  VM DRS Score ()  Cutier DRS Score ()  VM DRS Score	11 3 Vin 3 Vin 3 Vin 3 Vin 8 Vin 6 Sin
(2) Test/Veeam10 (2) Test/Veeam11 (2) Test/Veeam13 (2) Test/Veeam13 (2) Test/Veeam15 (2) Test/Veeam15 (2) Test/Veeam15 (2) Win5GL401 (2) Win5GL403 (2) Win5GL403 (2) Win5GL403 (2) Win5GL404	Related Objects II VSphere HA Datacenter III SOCC-Datacenter Proctive HA Heat Meetings VM Koetking	Disabled Disabled Enabled sMi and Argingston Manhoring	VSAN VSAN UNITE CO. 7750 VSAN Insette VSAN Performance COPS. United the post 2 hours. DOPS @	II Details * Details Details

Once the failover is complete, the status of the VMs will change to "Failover".

Intervention           0017         005 feature           00017         105 feature           0017         105 feature	20230314 20230314 20230314 20230314 20230314 20230314	Type Reple Reple Reple	Datus Latinam Failunan Failunan	Constan Time 2/14/2023 2/15 AM 3/23/2023 11/15 AM 3/23/2023 11/15 AM	Ramon Pai 1 4	Original Latation a300-vice03.etval a200-vice03.etval	Explication 172.30.156.21(Outline 1 scanner adds: 37-165-115-210 universities claim).172.30.16068	Plathere Milliogra Milliogra
entry 105 been entry 115 been entry	N 20230314 20230314 20230314 20230314 20230314 20230314	Repúe Repúe Repúe	Fallow Fallow	2/16/2523 2/15 AM A/22/2623 11/13 PM A/23/2623 11/13 PM	4	alion-vestilaria. alion-vestilaria.	172.30.158.2/Cume-1 ecamaration_35-105-115-210.cms.artic.com/.172.30.16068	Videoper Videoper
seniit Platuitevi	20230314 20230314 20230314 20230314 20230314 20230314	Regular Regular Regular	Fallerer Fallerer	3/23/2023 11/13 PM 3/23/2023 11/13 PM	4	a200-vcae25.ehul	+carrier adds-33-765-115-210 ameanments carris (TZ.30.16068	Webware
een02 Florit, IEA een03 Florit, IEA een04 Florit, IEA een05 Florit, IEA een06 Florit, IEA	20230314 20230316 20230316 20230316	Regular Regular	Failurer	3/23/2623 11/15 PM				
eari03 FSAU181M eari04 FSAU181M eari03 FSAU181M eari03 FSAU181M	25230314 20230314 20230314	Repulse	A subscience of			\$200-ward1.ehut	volentes addc-35-145-113-210	Websen
eari03 4545,1858 eari03 4545,1858 eari03 4545,1858	20230314	Barris dama	100000	3/25/2023 13/13 PM	4	al00-result.abet.	scenter adds-33-165-115-210 ame are mit cam) 172.30.16068	WANTE
eardd Hick, 1974 eardd Hick, 1974	20230316	- all a second	fallow	3/21/2023 8-28 4M	3	a300-vcas01.evc8	vienterable-35-185-117-210 imeerencic cam/ 172.30.16088	Webser
ears06 FSeN_1EVM		hepdar	failurer	3/21/2025 8:31 AM	3	a200-xxad75.eho5	vcenter addc-35-185-113-218 -meanine.com/172.30.16048	Monare
and he are not	20230318	Regular	fallow	3/21/2023 8:32 AM	3	alloo-visadh.ehitt	vsantai adoc-33-185-115-210 vmvarpvnc.cum) 172.30.16088	Wednesse
earni? PSelf_TEVM	20230316	Regular	Fallover	3/21/2023 E32 AM	1	aboo-vesaits anex.	veranter salds-35-185-115-210 vmwantvmc.com/,172.30.16048	Viduare
earrold PSuN, 16144	20230316	Repulse	Fallecer	3/31/2023 E32 AM	3	\$300-cca#75-ehid	scanlar add: 23-165-115-210maann.mc.com/.172.30,160.66	Wheet
earth? Furt TEUM	20230316	Repulse	Failurer	3/21/2023 8:03 AM	3	a300-vcsal15 abust	vearties.addc-35-185-115-210.uma.arevine.com/.172.30.160/08	Welson
aam10 FS-N, HVM	20230316	Reputer	fallow.	3/21/2025 6-34 AM	3	allon-vessels-elect-	verantee adds-35-165-115-210 whereaveranc.com/172.30.140.68	Webware
went1 Fb-N, 19VM	20230316	Regular	Failurer	3/21/2023 & 34 AM	1	a300-vcsa01-ehot	volement adds-35-185-115-210 immeetering.com/172.30.16068	Weiners
ears12 FLACIERA	20230316	Repulse	Talipter	3/21/2023 8:34 AM	1	#200-vci#25.#vc8	rearise addi-33-185-115-210 cmaastering.com/.072.30.160.88	Websare
manil PLACIEVA	20230314	Angular	Palanet	3/21/2023 835 AM	1	allon-recalls.etest	vicentee addo-35-185-115-210-immediated, com/,172.30.16048	Viduare
eard4 FLN, 18VM	29230316	Regular	falser	3/21/2023 0.04 AM	8	a300-read/setes	scenteradds-35-165-115-210 mmanering card, 173.30,160.66	Wheee
wents its how	10030316	Anjular	Failurer	3/21/2023 0-36 AM	3	abob-waab5-ehult	volument adds-35-185-115-210meantries.com/.172.30.160.08	Western
mental Plant texts	20230316	Regular	failurer	3/21/2023 8:37 AM	3	aboo-wanthahad	vemmerador: 35-185-113-210	Videare
x401 Flatt Republ	61 20230315	Repulse	Fallow	3/17/2023 3:58 AM	4	aloo	vianter add:-35-185-115-310 omeaning cam/, 172.30,160.08	Webcare
E402 ISelf Replo	878 20220313	Regular	falmer	3/17/2023 3:58 AM		about wrants and	scenter addr-15-185-115-210 vmwartvmc.com/.172.30.16068	Videosta
E403 Flah Repla	d1 20230313	Repulse	Falmer	8/17/2023 4:05 AM	4	allon-results about	scanler adds-35-165-115-210	Whene
2404 Flatt Repla	01,20220313	Reputer	Enterner	3/117/2023 4:00 AM	*	allon-vessilt sheet	veneries adds-35-185-115-210 wms are inc. cum/ 172-30, 160/68	Vidwara
1405 Flah Septe	61 20230313	Angular	fallerer -	3/17/2021 4-02 AM		allos-erado etad.	volation adds-35-165-115-210 orthogenetics class 172.30 140/68	Magent
	aam10 Fisht, 1904, aam11 Fisht, 1904, aam11 Fisht, 1904, aam12 Fisht, 1904, aam13 Fisht, 1904, aam15 Fisht, 1904, aam15 Fisht, 1904, 640 Fisht, 5aplat 540 Fisht, 5aplat 2404 Fisht, 5aplat	sen13         Fib. (1964) 22220116           sen11         Fib. (1964) 22220116           sen12         Fib. (1964) 22220116           sen13         Fib. (1964) 22220116           sen14         Fib. (1964) 22220116           sen15         Fib. (1964) 22220116           sen16         Fib. (1964) 22220116           sen17         Fib. (1964) 22220116           sen16         Fib. (1964) 22220116           Sch1         Fib. (sequel) (22200116           Sch2         Fib. (sequel) (22200116           Sch2         Fib. (sequel) (22200117           Sch4         Fib. (sequel) (22200113           Sch9         Fib. (sequel) (22200113	Sentil         Flash (1964) 2020116         Regular Arguine           Minick (1964) 2020117         Regular Arguine           Minick (1964) 2020118         Regular Arguine           Minick (Regular) 2020118         Regular Arguine           Minick (Regular) 2020118         Regular Arguine           Minick (Regular) 2020118         Regular Arguine           Minick (Regular) 2020118         Regular           Minick (Regular) 2020118         Regular           Minick (Regular) 2020118         Regular           Minick (Regular) 2020118         Regular	sen11         Fib.(1).1964/22220116         Register         Follower           mm11         Fib.(1).1964/22220111         Register         Follower           SAP1         Fib.10.1.22220111         Register         Follower           SAP1         Fib.10.1.22220111	sam10         Fib.(L) FMA_20230116         Regular         Falsure         L/17 (2023 54.4 MJ)           sam11         Fib.(L) FMA_20230116         Regular         Falsure         L/17 (2023 54.4 MJ)           sam12         Fib.(L) FMA_20230116         Regular         Falsure         L/17 (2023 54.4 MJ)           sam14         Fib.(L) FMA_20230116         Regular         Falsure         L/17 (2023 54.3 MJ)           sam14         Fib.(L) FMA_20230116         Regular         Falsure         L/17 (2023 54.3 MJ)           sam15         Fib.(L) FMA_20230116         Regular         Falsure         L/17 (2023 54.3 MJ)           sam15         Fib.(L) FMA_20230116         Regular         Falsure         L/17 (2023 54.3 MJ)           Sk10         Fib.(L) FMA_20230116         Regular         Falsure         L/17 (2023 54.3 MJ)           Sk11         Fib.(L) FMA_20230111         Regular         Falsure         L/17 (2023 54.3 MJ)           Sk12         Fib.(L) Feasibility J. 20230111         Regular         Falsure         L/17 (2023 54.3 MJ)           Sk24         Fib.(L) Feasibility J. 20230111         Regular         Falsure         L/17 (2023 40.3 MJ)           Sk44         Fib.(L) Feasibility J. 20230113         Regular         Falsure         L/17 (2021 40.2 MJ)	sen11         Flack (1964) 22220116         Regime         Failure         L17 (2022 6.0.4.4.M.         S           sen11         Flack (1964) 22220116         Regime         Failure         L17 (2022 6.0.4.4.M.         S           sen12         Flack (1964) 22220116         Regime         Failure         L17 (2022 6.0.4.4.M.         S           sen13         Flack (1964) 22220116         Regime         Failure         L17 (2022 6.0.4.4.M.         S           sen14         Flack (1964) 22220116         Regime         Failure         L17 (2022 8.0.4.M.         S           sen14         Flack (1964) 22220116         Regime         Failure         L17 (2023 8.0.4.M.         S           sen15         Flack (1964) 22220111         Regime         Failure         L17 (2023 8.0.4.M.         S           Sk41         Flack (Regimed) 22220111         Regime         Failure         L17 (2023 8.0.4.M.         S           Sk41         Flack (Regimed) 22220113         Regime         Failure         L17 (2023 8.0.4.M.         S           Sk40         Flack (Regimed) 22210113         Regime         Failure         L17 (2023 8.0.4.M.         S           Sk40         Flack (Regimed) 201201313         Regime         Failure         L17 (2023 8.0.4.M.         S	Sam10         Fair (1944) 2020116         Regime         Fairmer         D27/2021 804 444         S         SD0=con20 4bdL.           sem11         Hold (1944) 2020116         Regime         Fairmer         D27/2021 804 444         S         SD0=con20 4bdL.           sem11         Hold (1944) 2020116         Regime         Fairmer         D27/2021 804 444         S         SD0=con20 4bdL.           sem12         Hold (1944) 2020116         Regime         Fairmer         D27/2021 803 444         S         SD0=con20 4bdL.           sem13         Hold (1944) 2020116         Regime         Fairmer         D27/2021 803 444         S         SD0=con20 4bdL.           sem15         Hold (1942) 2020116         Regime         Fairmer         D27/2021 808 444         S         SD0=con20 4bdL.           sem15         Hold (1942) 2020116         Regime         Fairmer         D27/2021 808 444         S         SD0=con20 4bdL.           SA01         Hold (1942) 2020118         Regime         Fairmer         D27/2021 808 444         S         SD0=con20 4bdL.           SA01         Hold (1944) 20201118         Regime         Fairmer         D27/2021 808 444         S         SD0=con20 4bdL.           SA01         Hold (1944) 202011118         Regime         Fairmer	sam10         Florit, 1940, 2020111         Regime         Fallow         3/17/2012 614 444         3         4200-cm25 Amula         scenario         5/15-17-210-mmergence.com/172.35 1608           amm11         Florit, 1940, 2020111         Regime         Fallow         3/17/2012 814 444         5         4200-cm25 Amula         scenario: 5/1-5/15-270-mmergence.com/172.35 1608           amm12         Florit, 1940, 2020111         Regime         Fallow         3/17/2021 814 44         5         4200-cm25 Amula         scenario: 5/1-5/15-270-mmergence.com/172.35 1608           amm12         Florit, 1940, 2020111         Regime         Fallow         3/17/2021 814 44         5         4200-cm25 Amula         scenario 3/16-115-115-1230-mmergence.com/172.35 1608           amm14         Fallow, 1940, 2020111         Regime         Fallow         3/17/2021 814 44         5         4200-cm25 Amula         scenario 3/16-115-115-1230-mmergence.com/172.35 1608           amm15         Fallow, 1940, 2020111         Regime         Fallow         3/17/2021 814 44         5         4200-cm25 Amula         5/16-115-115-1230-mmergence.com/172.35 1608           amm15         Fallow, 1940, 2020111         Regime         Fallow         3/17/2021 814 44         5         4200-cm25 Amula         5/16-115-112-120 ammergence.com/172.35 1608           Add1         Fallow



Veeam Backup & Replication stops all replication activities for the source VM until its replica is returned to the Ready state.

For detailed information about failover plans, refer to Failover Plans.
When the failover plan is running, it is considered as an intermediate step and needs to be finalized based on the requirement. The options include the following:

• Failback to production - switch back to the original VM and transfer all changes that took place while the VM replica was running to the original VM.



When you perform failback, changes are only transferred but not published. Choose **Commit failback** (once the original VM is confirmed to work as expected) or **Undo failback** to get back to the VM replica If the original VM is not working as expected.

- **Undo failover** switch back to the original VM and discard all changes made to the VM replica while it was running.
- **Permanent Failover** permanently switch from the original VM to a VM replica and use this replica as the original VM.

In this demo, Failback to production was chosen. Failback to the original VM was selected during the Destination step of the wizard and "Power on VM after restoring" check box was enabled.



E President					tinen Birbug an	Replication				
The States	Antonio International Internat	Corect Contraction Technic Contraction	Add to Find	×.			VM	s failing l	oack t	to rver
une .	Q Type in an object to	erne ha one with for	×			, v	ausue	in vesin		1000
Contractions of the second sec	New Y  Construction  Sectors	and Terror ArX, Depictor Hiski, (1444, 202001) Hiski, (1444, 2044, 2044, 2044, 2044, 2044, 2044, 2044, 2044, 2044, 2044, 2044, 2044, 2044, 2044, 2044, 204		References of the second secon	Cataline Time 2716-0221 315 AM 2027-02021 315 AM	Representations of the median	Conjunct Location allow crandid strutt, allow strutt allow strutt strutt for atlants for struttenes	Explor Location 172:33: 154: 2004;mm-1 senere and 2014;101:101:101;mmane senere and 2014;101:101;201;mmane senere and 2014;101:101;201;mmane senere and 2014;101:101;101:101;101 senere and 2014;101:101;101:101 and the two def foldware present if required.	A second 172 30 19924 and an 172 30 19924 A 172 30	Patron Minar Minar Minar Minar Minar Minar Minar Minar Minar Minar Minar Minar Minar Minar
a monthly							· Pening	Trady Can	Latin .	
Se faring internation								2		

Failback commit is one of the ways to finalize failback operation. When failback is committed, it confirms that the changes sent to the VM which is failed back (the production VM) are working as expected. After the commit operation, Veeam Backup & Replication resumes replication activities for the production VM.

For detailed information about the failback process, refer Veeam documentation for Failover and Failback for replication.

	a 5	2 2 2	e =	X		"(	Comn	nit Failback" to	o comple
e Talone Falson Falson Falson F	extense to Saddrane Dede Sectorion Time Talkada	Talley Ter Ben /	Add to Propert	ten farment Repet i				the Failback T	ask
	O tes a se destas	and in second line	×					the ranback r	usk
TO	led the second second	rea pa maren pa	~			100	a design of the second second	a the first of the	1622
ly here	Nerre	Ink Name	1740m	Tation	Deation Tree	Ramon Pol-	Original Sacation	Replica Locator	Patlem
ER Replication	EN DROTWINGT	425,5483(40)	heputer	Tationer	2/16/2023 213 44		#100-y15405 #N.B.	172.39.158.21D.00001	Approxim
( feplica.	Continuent of	FLAN TRUNK CONNETTS	Sec. in	Tallari	EGA/DOLT S.T.T. SH		and an and an a	supervised by the state of the second s	
To feely	(Classical)	PLAN 18144 JULY 18	Sec. 1	Table 1	1/14/02/1 1/11 444		all an an all and	Hartin also, FL 103, 113, 210, once on cont 172 16, 1604	Whene
Dis Active U.D.	-S lettiments	Fb/N,18VM,20230316	<b>Neplin</b>	Faillerb	1/24/2023 148 AM	AL.	allog-result about	reserved adult - 15-165-119-210. or server a current / 172.26.16048	Village
Last 24 March	EQ Testiment 65	PS/N_18VM_20200316	Reputer	Failura	1/54/2223 149 AM	4	a000-+++a05-mett-	viameratik 35-115-115-210	Villear
C farmers	Testimente .	Plick, 10584, 20230314	Fepular :	Failuria.	3/24/2023 1:24 AM		aligo vesets ends.	vcamaradok-35-165-115-210muaramic.com/772.5536568	VMean
(2 Warning	Testiment?	Fbin,18VM,20230318	Replice	falled.	3/34/2022 1/33 AM	41	a000-visa05.etv8	summer adds -23-065-213-210 umages/reg.aue/1172.30.16046	When
C Faled	Testveentill	F5-N, 16VM, 2020016	Reporter	<b>Juites</b>	1/34/2023 1/38 AM	30	a300-yeards enut.	scenteradok-35-185-119-218 umuseum.com/172.30.16066	Villege
Martin C	California (19	PSeN, 16VN8, 20230318	Repaire	Failurch	\$154/2005 1.37 AM	-AC	align reserves.	maniler salat: 25-165-115-210	Witness
	24 Terrissan 12	F6/N,16VM,25220316	Reprint	Faillack.	1/54/2023 1/27 AM		alto-coald met.	scanar add: 25-165-115-213 unu are mic cont.172.30.16008	Waren
	CQ Testiment I	FEAL HVAL BEDORTS	Fepler	Failure	1/24/2013 144 AM		allog-vysall5 aholi	marter adult 23-163-115-218-mare and com/172.56.16048	Vitere
	CG Testiment U	PS-N, 18V44,00230316	Repolar	faiture	B/04/0003 127 AM		#300-+11405 afutt.	100 million 25-765-715-210 - million and 2000, 172, 30 34006	Village .
	() Jenning ()	PEACHEVIA_DECOUCH	Reputer	Failinck.	1/3A/2013 1/3A AN	2	#300-cesa05-mell.	vcame.addi: 21-163-113-210-meanmic.com/UV2.30.25006	Villean
	CQ 1etTreen14	FLAG_18VM_20220318	Replat	Falbers	1/34/2013 1.30 AM		ASSO HEARTS after.	scantas add. 23-183-172-210.000-pre-me.com/172.30.16008	Villen
	25 Jertimen 15	754N,18VM,20200118	- Property	Taller's	1/24/2023 1:25 AM		#300 resets after	10 mm and 15-103-115-210.000 are not carry 172.10.16008	
	Cit restrations	Plant, 1998, Jacobille		Tarback.	1. 102W 2020 1132 AM		ADDO-HEARD APARL	10.0 million 201 102 112 210 million in 2010 172 30 18000	
	A MUSICIAN	The American Manual	1000	failers a	2/14/2013 112/ A00	- C	and a second state	10 PT 10 10 10 10 10 10 10 10 10 10 10 10 10	100 cm
	Che Westlands	Flats Receiption 20720113		famers.	TOTAL COLUMN THAT AND	-48	all states and	statements 15, 16, 115, 215, processing statements 2, 10, 16, 16, 16, 16, 16, 16, 16, 16, 16, 16	100 and 100
		Stati Respect 20230111	-	Reiller A	1/14/2023 1:53 AM		And Annual States	Hard 11, 185, 115, 210 and 111, 210 and 111, 210 and 111, 210 and 111, 210	Without
	100 wears and	First Resident's Statistics	-	fame.	1.114 (1003) V 10 444		all and the second second		1990-100
a 📲 😫 🏙 🗂 📤	an and an and								
A Construction COM Backing Copy - Rate - Anking Copy - Rate - Anking Copy - Rate - Anking Not -	Talawa Pan- Pan- Atom								
and Explosion CCP Market Copy Market Copy Market Copy Auditory Inter Auditory Inter Audi	Telawa Para- fare Q, type is an elipit to	erre is conti for .	×						
or Federation COP 1 Sta - Federation COP Protection COP Pro	Fallow Pres- fore date (C, 1)per in an eligit for a lange	anne bi aningth far i	×		dust true i	hatter			
a falter of the state of the st	Talawa Maran Ana Atalama Q, type is an algorit to be talama Atalama Atalama	arme (a' anise)h far '	X Second Log	Solar Same	Sectore 1	and fore			
a Alla Santa Congo	Adverse Factorse from a factor from a factor	ener in completion :	X Secon Yes Correct failures Correct failures	Solar Secons	30et Time 1 504/0023 207 #4 504/0023 207	Ind Ine Del Ine Del Del	1208 AM		
Antiney Note Sectors	Token To	nine (s march far :	X Second Fallback Control Fallback Control Fallback	Solar Secres Scient	Start Time & 504/3022 207 Ak 504/2022 207 Ak	bat fire 1/16/202 1/16/202 1/16/202	1200 AM 2200 AM 2200 AM		
a Forketore CFP Internet to CFP Intern	Telever Tel	eren (é seárch <i>fer</i> :	X Second Failures Connect Failures Connect Failures Connect Failures	Solar Secres Secres Secres	3641 Time 4 1654 (3023 2017 Ab 1654 (2023 2017 Ab 1654 (2023 2017 Ab 1654 (2023 2017 Ab	End Tone 3/34/2021 3/34/2021 3/24/2021 3/24/2021	206 AM 206 AM 206 AM 206 AM		
a Particular Control C	Advert Network Network Advert Adve	ane to samp for .	X Service Testes Connect Failback Connect Failback Connect Failback Connect Failback	Solar Secons Secons Secons Secons	3art Time 4 15442023 201 Ak 15442023 201 Ak 15442023 201 Ak 15442023 201 Ak 15442023 201 Ak	Led Tore 3:04/202 3:04/202 3:04/202 3:04/202 3:04/202 3:04/202	2.06 AM 2.08 AM 2.08 AM 2.08 AM 2.06 AM		
A starting car in fight car CD has been and have	Televier Televi	ana (i min) far -	X Service Trait Corrent Fallback Corrent Fallback Corrent Fallback Corrent Fallback Corrent Fallback	Statut Sectors Sectors Sectors Sectors Sectors	Diset Time 3 104-0013 207 AM 104-0013 207 AM 104-0012 307 AM 104-00213 207 AM 104-00213 207 AM	Led Tree 3/34/2022 3/24/2022 3/24/2022 3/24/2022 3/24/2022	2 206 AM 2 206 AM 2 206 AM 2 206 AM 3 206 AM 3 206 AM		
A region of the second	Adverse Adv	ane (case) for	X Senten Yeak Corrent Fallack Corrent Fallack Corrent Fallack Corrent Fallack Corrent Fallack Corrent Fallack	Secon Secon Secon Secon Secon Secon Secon	Start Time 4 1044/0013 201744 1044/0013 201744 1044/0013 201744 1044/0013 20174 1044/0013 20174 1044/0013 20174	Led Tree 3/04/2022 3/04/2022 3/04/2022 3/04/2022 3/04/2022 3/04/2022	206 AM 206 AM 206 AM 206 AM 206 AM 306 AM 306 AM		
an Refundor CDF Balance And Refundor CDF Balance And Refundor CDF Balance Status Market And Balance Status	Taliver Telever There's and the sector function The sector function	alma (is march for -	X Secon Tela Connet Fallback Connet Fallback Connet Fallback Connet Fallback Connet Fallback Connet Fallback Connet Fallback	Zoto Terms Socess Socess Socess Socess Socess Socess	Start Time 1 150400073 2014 1504200231 007 Ak 1504200231 007 Ak 150420023 007 Ak 150420023 007 Ak 150420023 007 Ak	End Time 5/34/2021 3/24/2021 3/24/2021 3/24/2021 3/24/2021 3/24/2021	200 AM 200 AM 200 AM 200 AM 200 AM 200 AM 200 AM		
a Patiento CO Statuto CO Sta	National Advancement Advanceme	ane (case) for	X Service Tableck Connet Fableck Connet Fableck Connet Fableck Connet Fableck Connet Fableck Connet Fableck Connet Fableck Connet Fableck	Zona Saran Sacan Sacan Sacan Sacan Sacan Sacan Sacan	5041 Time 3 1004/0003 2007 AN 1044/0013 2007 AN 1044/0013 2007 AN 1044/0023 2007 AN 1044/0023 2007 AN 1044/0023 2007 AN 1044/0023 2007 AN	List Time 2. 204/202 3. 204/202 3. 204/202 3. 204/202 3. 204/202 3. 204/202 3. 204/202 3. 204/202	2 206 AM 2 208 AM 2 208 AM 2 208 AM 2 208 AM 3 208 AM 3 208 AM 3 208 AM	Task succe	secful
a Reference Color New Production Color New Production Color Vision Product	Tallow Institute 1 fields Institute 1 fields	ning (g manyk fag :	X Sensor Fase Corrent Fallack Corrent Fallack Corrent Fallack Corrent Fallack Corrent Fallack Corrent Fallack Corrent Fallack Corrent Fallack Corrent Fallack	Solar Socan Socan Socan Socan Socan Socan Socan Socan Socan Socan	Spart Time 1 154-2022 2 6074 154-2022 1 607 Ab 154-2022 1 607 Ab 154-2022 1 607 Ab 154-2022 1 607 Ab 154-2022 2 607 Ab 154-2022 2 607 Ab 154-2022 2 607 Ab	Ind Tree 3/04/2021 3/04/2021 3/04/2021 3/04/2021 3/04/2021 3/04/2021 3/04/2021 3/04/2021 3/04/2021	2 208 AM 2 208 AM	Task succe	essful
a Performance CEP Internet Network State Performance State Performa	Notice Notice	ane (case) for	Senior Ivas Convet Fallact Convet Fallact Convet Fallact Convet Fallact Convet Fallact Convet Fallact Convet Fallact Convet Fallact	Secon Secon Secon Secon Secon Secon Secon Secon Secon Secon Secon	504-5023 207 AN 104-5023 207 AN	End Time 3.194/2022 3.094/2022 3.094/2022 3.094/2022 3.094/2022 3.094/2022 3.094/2022 3.094/2022 3.094/2022 3.094/2022	2 206 AM 2 208 AM 2 208 AM 2 208 AM 2 208 AM 3 208 AM 3 208 AM 2 208 AM 2 208 AM	Task succe	essful
e forganization norman units Annual units	Tallow Input Tapot	alma (o march far -	X Internative Control Fallmack Control Fallmack Control Fallmack Control Fallmack Control Fallmack Control Fallmack Control Fallmack Control Fallmack Control Fallmack Control Fallmack	Total Secons Secons Secons Secons Secons Secons Secons Secons Secons Secons	Seat Time 1 104-00023 cm Ak 104-00023 cm Ak	Lind Tore 3.154/2021 3.154/2021 3.154/2021 3.154/2021 3.154/2021 3.154/2021 3.154/2021 3.154/2021 3.154/2021 3.154/2021	2 256 AM 2 256 AM 2 506 AM 2 506 AM 2 506 AM 2 506 AM 2 508 AM 2 508 AM 2 508 AM 2 508 AM	Task succe	essful
e Performance Con- las A Failes Con- las A Failes Annuel Performance Con- las A Failes Performance Con- las A Failes Performance Participant Control Con- las Annuel Performance Control Con- Control Con- Con- Con- Con- Con- Con- Con- Con-	Non- ter States and S	ana (é nainte der :	Service Trans Connect Fallwack Connect Fallwack	Socan Socan Socan Socan Socan Socan Socan Socan Socan Socan Socan Socan Socan	Dest Time 1 104-00011 2014 104-00011 2014 104-00012 2014 104-0002 2014 104-00000 104-0000000000	End Trave 5 3/24/2022 3/24/2022 5/24/2022 5/24/2022 1/24/202	2 208 AM 2 258 AM 2 208 AM	Task succe Failback Cor	essful
a vice a vice vice for the formation vice for the formation vice for the formation vice f	Adarse Fallow Press	eine (o marifi far -	X Sector Yest Correct Fallows Correct Fallows	2000 Secon Secon Secon Secon Secon Secon Secon Secon Secon Secon Secon Secon Secon Secon	Sect Time 1 104/0022 2014 104/0023 2014	End Terre 5:054/2023 3:054/2023 3:054/2023 3:054/2023 3:054/2023 3:054/2023 3:054/2023 3:054/2023 3:054/2023 3:054/2023 3:054/2023	2209 AM 2259 AM 2509 AM 2508 AM 2508 AM 2508 AM 2508 AM 2508 AM 2508 AM 2508 AM 2508 AM 2508 AM	Task succe Failback Cor	essful mplete
a Persona Carlo International Carlo Internatio	Non- Non- ter Non- ter Non- No	ann (i march An -	Series Fore Connet Falses Connet Falses	Torio Sucres Suc	Start Time 1 104-0001 2014 104-0001 2014 104-0001 2014 104-0002 2014 104-0002000 104-00000 104-0000000000	Lind Tarre 5.04-2023 1.04-2023	2 208 AM 2 58 AM 2 58 AM 2 508 AM	Task succe Failback Cor	essful mplete
In Productions Pr	Adarse Program Sector S	eren (o march for -	Territor Territoria Connot Fallows, Connot Fallows,	2000 Tersine Socens Socens Socens Socens Socens Socens Socens Socens Socens Socens Socens Socens Socens Socens Socens	Sect Time 1 154-0003 2014 154-0003 2014	End Time 3/54/2021 3/54/2021 3/54/2021 3/54/2021 3/54/2021 3/54/2021 3/54/2021 3/54/2021 3/54/2021 3/54/2021 3/54/2021	2 209 AM 2 259 AM 2 206 AM 2 206 AM 2 206 AM 2 206 AM 2 208 AM 2 2	Task succe Failback Cor	essful mplete
a production COM task Factor Theorem 1 task Factor Production COM Task Factor Task Factor	Non- Non- ter Status School	aina (i marifi Arr -	Sector Tara Commit Fallwack Commit Fallwack	Torius Socras Socras Socras Socras Socras Socras Socras Socras Socras Socras Socras Socras Socras Socras Socras Socras Socras	Start Time 1 104-0003 2014 104-0003 2014 104-0000 2014 104-00000 104-0000000000	End Time 3.104/2021 3.044/2021 3.044/2021 3.044/2021 3.044/2021 3.044/2021 3.044/2021 3.044/2021 3.044/2021 3.044/2021 3.044/2021 3.044/2021	2 208 AM 2 206 AM 2 206 AM 2 206 AM 2 206 AM 2 208 AM	Task succe Failback Cor	essful mplete
n Persional Control Co	Adam Providence Provid	eren (o march for -	X Instance Treat Connect Fallows Connect Fallows	2000 Socon	Sect Time 1 154-00022 2014 154-00022 2014 154-00023 2014	End Tree 3/24/2022 3/24/20	2009 AM 2009 AM	Task succe Failback Cor	essful mplete
a production COM tas Packation The Packation Packati	New- New-	aina (i marif) far :	Sector Tara Commit Fallwack Commit Fallwack	Tota Socan	Start Time 1 10440003 2014 10440003 2014 10440000 2014 10440000 2014 10440000 2014 10440000 2014 104400000 2014 10440000 2014 10440000 1040000 1040000 1040000 1040000 1040000 104000 1040000 1040000 1040000 10400000 10400000 10400000 10400000 10400000000	Lind Taree 3.054/2023 3.044/2023	2 208 AM 2 206 AM 2 206 AM 2 206 AM 2 206 AM 2 208 AM	Task succe Failback Cor	essful mplete
And And And And And And And And And	Adare Adare Market	ane to early for .	Territor Terri Connot Fallows Connot Fallows	2000 Socean Socean Socean Socean Socean Socean Socean Socean Socean Socean Socean Socean Socean Socean Socean Socean Socean Socean Socean	Sect Time 1 154-0003 2014 154-0003 2014 154-0005 154-0005 154-0005 154-0005 154-0005 154-0005 154-0005 154-0005 15	End Tore 1.154/2021 1.154/2021 1.154/2021 1.154/2021 1.154/2021 1.154/2021 1.154/2021 1.154/2021 1.154/2021 1.154/2021 1.154/2021 1.154/2021 1.154/2021 1.154/2021	2 209 AM 2 259 AM 2 259 AM 2 206 AM 2 206 AM 2 206 AM 2 208 AM 2 2	Task succe Failback Cor	essful mplete
Projection Colle International	New- New-	alon (i morif) for .	Section Tops Commit Fallowsk Commit Fallowsk	Social Success Succes	Start Time 1 10440003 2014 10440003 2014 104400003 2014 10440003 2014 10440000 2014 10440000000000000000000000000000000	Ind Tree 3.04/2023	2 208 AM 2 208 AM	Task succe Failback Cor	essful mplete

After failback to production is successful, the VMs are all restored back to the original production site.

😑 vSphere Client 🛛 Q				∕ls r	recovered	C Steeneermatory	
. 8 . 9	¢	Cluster05 Actors Servery Monter Configure Permasons Hosts VMs Data	on On	Pre	m vCenter	-	
		Tota Processors 64 Tota Monter Migratiene 30					Page 1928 Capanity 1928 Page 2023 Capanity 1922 Capanity 1923 Capanity 1923
(自 Test/Vesam02) (首 Test/Vesam03)		Related Objects	~ ~	v5phere 0	CFR.		
(第 YestVeeam04 (第 YestVeeam05 (第 TestVeeam05		Datacenter IB: A300.0005		Cluster Se	lervices.		
(第 TestVesam0) (第 TestVesam0)		Outline Consumers	~	Custore A	Attributes		
(2. Test/VesarrO) (3. Test/VesarrO) (3. Test/VesarrO) (3. Test/VesarrO) (3. Test/VesarrO) (3. Test/VesarrO)		Taga August Tag Canyon S	entres v		www.committumitumitumitumi	-	
位 Test/vesamiti 位 Test/vesamiti 位 WenkQLACK 位 WenkQLACS 位 WenkQLACS		Aug. Down	No Anno So Davido No Anno So Davido	tin.			1 Martin
값 WestQL404 값 WestQL405		Outer Resources	÷				
<ul> <li>Recent Tasks Alarma</li> </ul>							
al None 🕴 1	Target 7	Status T Details T Inflatio	<ul> <li>Denet + Bart fire</li> </ul>	T	Completion Time * Server		
ower On whoat machine	db Wesce.401	Compared EHCDCCOM(Advan	radiator 4 mi 03/34/20	10961,19907 _	05/24/3023 139 07 . ab00-sta05 ehosion	2	

# Conclusion

FSx for ONTAP datastore capability enables Veeam or any validated third-party tool to provide low-cost DR solution using Pilot light cluster and without standing up large number of hosts in the cluster just to accommodate the VM replica copy. This provides a powerful solution to handle a tailored, customized disaster recovery plan and also allows to reuse existing backup products in house to meet the DR needs, thus enabling cloud-based disaster recovery by exiting DR datacentres on-premises. Failover can be done as planned failover or failover with a click of a button when disaster occurs, and decision is made to activate the DR site.

To learn more about this process, feel free to follow the detailed walkthrough video.

https://netapp.hosted.panopto.com/Panopto/Pages/Embed.aspx?id=15fed205-8614-4ef7-b2d0-b061015e925a

# Migrating Workloads on AWS / VMC

TR 4942: Migrate Workloads to FSx ONTAP datastore using VMware HCX

A common use case for VMware Cloud (VMC) on Amazon Web Services (AWS), with its supplemental NFS datastore on Amazon FSx for NetApp ONTAP, is the migration of VMware workloads. VMware HCX is a preferred option and provides various migration methods to move on-premises virtual machines (VMs) and their data, running on any VMware supported datastores, to VMC datastores, which includes supplemental NFS datastores on FSx for ONTAP.

Author(s): NetApp Solutions Engineering

# Overview: Migrating virtual machines with VMware HCX, FSx ONTAP supplemental datastores, and VMware Cloud

VMware HCX is primarily a mobility platform that is designed to simplify workload migration, workload rebalancing, and business continuity across clouds. It is included as part of VMware Cloud on AWS and offers many ways to migrate workloads and can be used for disaster recovery (DR) operations.

This document provides step-by-step guidance for deploying and configuring VMware HCX, including all its main components, on-premises and on the cloud data center side, which enables various VM migration mechanisms.

For more information, see Introduction to HCX Deployments and Install Checklist B - HCX with a VMware Cloud on AWS SDDC Destination Environment.

## **High-level steps**

This list provides the high-level steps to install and configure VMware HCX:

- 1. Activate HCX for the VMC software-defined data center (SDDC) through VMware Cloud Services Console.
- 2. Download and deploy the HCX Connector OVA installer in the on-premises vCenter Server.
- 3. Activate HCX with a license key.
- 4. Pair on-premises VMware HCX Connector with VMC HCX Cloud Manager.
- 5. Configure the network profile, compute profile, and service mesh.
- 6. (Optional) Perform Network Extension to extend the network and avoid re-IP.
- 7. Validate the appliance status and ensure that migration is possible.
- 8. Migrate the VM workloads.

## Prerequisites

Before you begin, make sure the following prerequisites are met. For more information, see Preparing for HCX Installation. After the prerequisites are in place, including connectivity, configure and activate HCX by generating a license key from the VMware HCX Console at VMC. After HCX is activated, the vCenter Plug- in is deployed and can be accessed by using the vCenter Console for management.

The following installation steps must be completed before proceeding with HCX activation and deployment:

- 1. Use an existing VMC SDDC or create a new SDDC following this NetApp link or this VMware link.
- 2. The network path from the on-premises vCenter environment to the VMC SDDC must support migration of VMs by using vMotion.
- 3. Make sure the required firewall rules and ports are allowed for vMotion traffic between the onpremises vCenter Server and the SDDC vCenter.
- 4. The FSx for ONTAP NFS volume should be mounted as a supplemental datastore in the VMC SDDC. To attach the NFS datastores to the appropriate cluster, follow the steps outlined in this NetApp link or this VMware link.

## **High Level Architecture**

For testing purposes, the on-premises lab environment used for this validation was connected through a site-to-site VPN to AWS VPC, which allowed on-premises connectivity to AWS and to VMware cloud SDDC through External transit gateway. HCX migration and network extension traffic flows over the internet between on-premises and VMware cloud destination SDDC. This architecture can be modified to use Direct Connect private virtual interfaces.

The following image depicts the high-level architecture.



# **Solution Deployment**

Follow the series of steps to complete the deployment of this solution:

To perform the installation, complete the following steps:

- 1. Log in to the VMC Console at vmc.vmware.com and access Inventory.
- 2. To select the appropriate SDDC and access Add- ons, click View Details on SDDC and select the Add Ons tab.
- 3. Click Activate for VMware HCX.



This step takes up to 25 minutes to complete.



- 4. After the deployment is complete, validate the deployment by confirming that HCX Manager and its associated plug-ins are available in vCenter Console.
- 5. Create the appropriate Management Gateway firewalls to open the ports necessary to access HCX Cloud Manager.HCX Cloud Manager is now ready for HCX operations.

For the on-premises Connector to communicate with the HCX Manager in VMC, make sure that the appropriate firewall ports are open in the on-premises environment.

- 1. From the VMC Console, navigate to the HCX Dashboard, go to Administration, and select the Systems Update tab. Click Request a Download Link for the HCX Connector OVA image.
- 2. With the HCX Connector downloaded, deploy the OVA in the on-premises vCenter Server. Right- click vSphere Cluster and select the Deploy OVF Template option.

😑 vSphere Client	Q	C & American Mercan co	⊪× © ©>
	K III A300-Cluste	r01 ( 3 Actions	_
	Deploy OVF Template	Select an OVF template ×	
• 11 AND CARACTER	Select an OVF template     Select a name and folder	Enter a URL to download and install the OVF package from the internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive. O URL	The
0 400-m	3 Select a compute resource 4 Janview details	Kito Little Direction and and an United Strategy of Little	
<ul> <li>G ANF_HC</li> <li>G ANFDAT</li> <li>G ANFDAT</li> </ul>	S. Select storage	UPLOAD FILES VMware-HCX-Connector-4: ova	
- G Kode 2010 2010 2010 2010 2010	<ul> <li>Beady to complete</li> </ul>		
10 HOL 10 HOL		CANCEL NEXT	
(2 110) (2 110)	Hotor_22	referios romand, o haitea 40.08 10.45.08 20.465. A 5.68 Fillet 11.58.08	* * * * * * *
A Recent Tasks	Alarma :		

- 3. Enter the required information in the Deploy OVF Template wizard, click Next and then Finish to deploy the VMware HCX Connector OVA.
- 4. Power on the virtual appliance manually.For step- by- step instructions, go to VMware HCX User Guide.

After you deploy the VMware HCX Connector OVA on-premises and start the appliance, complete the following steps to activate HCX Connector. Generate the license key from the VMware HCX Console at VMC and input the license during the VMware HCX Connector setup.

- 1. From the VMware Cloud Console, go to Inventory, select the SDDC, and click View Details. From the Add Ons tab, in the VMware HCX tile, click Open HCX.
- 2. From the Activation Keys tab, click Create Activation Key. Select the System Type as HCX Connector and click Confirm to generate the key. Copy the activation key.

vm	VMware HCX							D D		NetApp * 🗰
Sub	criptions Activation Keys	SODCE								C DARK
Act	ivation Keys								CRE	ATE ACTIVATION KEY
	Activation Key		Status y	Subscription	 System Type	Ŧ	System Id		τ.	Created
1	ABIEI	\$3	CONSUMED	VMware Cloud on AWS (	HCK Connector		205	73		9/19/22, 9:24 AM
1	92C1	75	CONSUMED	VMware Cloud on AWS (	HCK Cloud		201	15321		9/16/22, 9:56 AM
1	100	(846	DEACTIVATED	VMware Coud on AWS	HCX Cloud		202	26		8/11/22, 12:23 PM
										Showing 1 - 3 of 3 entries



A separate key is required for each HCX Connector deployed on-premises.

3. Log in to the on-premises VMware HCX Connector at "https://hcxconnectorIP:9443" using administrator credentials.



Use the password defined during the OVA deployment.

4. In the Licensing section, enter the activation key copied from step 2 and click Activate.



The on-premises HCX Connector must have internet access for the activation to complete successfully.

- 5. Under Datacenter Location, provide the desired location for installing the VMware HCX Manager onpremises. Click Continue.
- 6. Under System Name, update the name and click Continue.
- 7. Select Yes and then Continue.
- 8. Under Connect Your vCenter, provide the IP address or fully qualified domain name (FQDN) and the credentials for the vCenter Server and click Continue.



Use the FQDN to avoid communication issues later.

9. Under Configure SSO/PSC, provide the Platform Services Controller's FQDN or IP address and click Continue.



Enter the vCenter Server's IP address or FQDN.

- 10. Verify that the information is entered correctly and click Restart.
- 11. After complete, the vCenter Server is displayed as green. Both the vCenter Server and SSO must

have the correct configuration parameters, which should be the same as the previous page.



This process should take approximately 10–20 minutes and for the plug-in to be added to the vCenter Server.

m HCX Manager	Dashboard	Applance Summary	Configuration	Administration		77231354	K157 Version : 4.410 Type : Connector	admir
VMware-HCX	-440				0	CPU Used 1407 MHZ	Free 688 MHZ Capacity 2095 MHZ	67%
FGDN: IP Address: Version:	VMware-HCX-440 172.2 4.4.1.0 20 days 21 bours	9 minutes			()	Memory Used 9691 MB	Free 2316 MB Capacity 12008 MB	81%
Current Time:	Tuesday, 13 Septer	mber 2022 07:44:11 PM UTC			٢	Storage Used 29G	Free 98G Capacity 127G	23%
NSX			vCenter		s	so		
			https://a300-vcs	sa01.ehcdc.com	• ht	ttps://a300-vcsa01.ehcdc.com		_
MANAGE			MANAGE			ANAGE		

#### Step 4: Pair on-premises VMware HCX Connector with VMC HCX Cloud Manager

1. To create a site pair between the on-premises vCenter Server and the VMC SDDC, log in to the onpremises vCenter Server and access the HCX vSphere Web Client Plug- in.

noriculs											
(I)	ð	8	0			000	Π	(P)	tē.		1
Hosts and Clusters	VMs and Templates	Storage	Networking	Content Libraries	Global Inventory Lists	Workload Management	SnapCenter Plug-in for VMware vSphere	Cloud Provider Migration	Site Rec	overy HCX	
Ionitoring											
圇		æ	R		$\diamond$						
Task Console	Event Console	VM Customization Specifications	VM Storage Policies	Host Profiles	Lifecycle Manager	ONTAP tools					
dministratio	n										

2. Under Infrastructure, click Add a Site Pairing. To authenticate the remote site, enter the VMC HCX Cloud Manager URL or IP address and the credentials for the CloudAdmin role.

E vSphere Client Q			C S verentinetticoccom ~ (D (D ~
HCX Destboard intraktructure	si v	Site Pairing	C ADD A SITE PAIRING
Site Relation     Autoconnect     Services     Analysics     Services     Analysics Extension     Analysics     Analysics     Analysics     Constance Recovery	*	BACK-VCSa04-enterprise     Original, DV2 25254557.443     Original, DV2 2525457.443     Original, DV2 252547.443	
System	*		

HCX information can be retrieved from the SDDC Settings page.

(;)

<	C BACK			ÓFIN	NSX MANAGER OPEN VCF	NTER ACTIONS
	FSxNDemoSDDC	VMC on AWS SODC 👨 US West (Dregori)		10050		
Launchged	Remain Republics & Seculty	Anna Ant Con Habitananta	The second second second			
Inventory.	screening measuring a second	storage not this manierierte	manufacture account			
) suescriptions	SDDC					
Topin						
Developer Center	> Management Appliances (					
Montenance						
Notification Preferences	vcenter information					
	> Default vCenter User > Account					
	> vSphere Class (HTML5)	Ø				
	uCenter Server All					
	> Explorer					
	> PowerCLI Connect					
	> vCenter FGDN					
	HCX Information					
	HCK FOON					
	HCX F3DH		Amountain Antoineat	Part P	Province .	
	https://b	VIIIN BY BY THE COST	Public IP - resolvable from internet		172.30.361.215	ED
	NSX information					
	> Nox sanager button deraut aut	ens.				
0.011	a second while surgers broken					
- vtpriere Client - Q					C & annual sectors	un p
E - vSphare Gent - C		Site Paking	_	_	C (Similaria	inere 🛛 🖗
victore Over C		Sité Palring		-	C . S have been set and	
vipture Dient Q		Site Pairing			C & house and and a	ADD A SITE FAILUR
- vdphine Clent - Q connund : etholian Statigatur		Site Pairing	→ @mei		C . Li inner independent	ADD A SITE FAMILY
VSZPANA Ciant C		Site Pairing	→ Q DCI a demonstrations a demonstrations		C . Li konsek en ign inner	and a stre fairing
všazinie Cient – Cj. Internet I Antonio Statu Statunica I Statunica I Statunic	- - -	Site Pairing	→ Concel a desenant desenant		C . La incomé no la comercia de la c	abo a Sire famin
VSZPHINE Channel C. Series C. S.	* *	Site Pairing	a poci desenaria desenaria po Remote Site ×		C . Stanna and an and	abb a śre famin
VSZPANO Ciente Co Service di Service di		Site Pairing	a → ⊕ tocci de sense inter series de nomenane o Remote Site × Metorine ⊆ ⊕		C . Stanoord and state	ABD A SITE FAILUR
výszenne Cient – C. artiková ( materia materia dostavá dostavá dostavá ne velova materia postavá ne velova materia ne velovelove ne velove ne velove ne velove ne velove ne velove ne velovelove ne velovelove ne velovelove ne velovelove ne velovelovelovelovelovelovelovelovelovelo		Site Pairing	a → ⊕ bool desentations desentations or Remote Site × k			ABD A SITE FAILUR
vischer Ower C		Site Pairing	a → ⊕ boci dress and to boci dress and b Remote Site × http://www.ioca///www.ioca////www.ioca////////////////////////////////////			ABD A SITE FAILUR
výszeke Ower C overseel entered sentered Stanovel Atlantes Stanovel Atlantes	* * *	Site Pairing	a → ⊕ tocci de regulariza totato de regul			ADD A SIZE PANN
výszhike Ower – O artikard ( mitikard Stranuel Antonia Stranuel Antonia Stranuel Antonia Stranuel Antonia Stranuel Antonia Stranuel Antonia Stranuel Antonia	* *	Site Pairing				ADD A SITE FAMILY
výszekke Cient C entitisent sektosent <del>Statusent</del> Statusent Antikos Nelesis Coleman Nelesis Co	•	Site Pairing	→      →			ADD A SITE FAILUR
všeznika Okrat – Q articiara ( articiara) Articiara Arti	•	Site Pairing	a → → → hocki drima ANT220(561 drima ANT220(5			ADD A SITE FAIRIN
A VYSCHWE Clever Q		Site Pairing	→      →			ADO A SITE FAIRIN
Vestional Clean	-	Site Pairing				ADO A SITE YANN
Vysphirk Olext Q articlet articlet Statistica Stat		Site Pairing				ADO'A SITE FAMIL
And Andrewson Andrew Andrewson Andrewson Andre		Site Pairing				
Armont I and Armon		Site Pairing	a Annual 102201681 a Annual 102201681 a Annual Annual Inc. a Annual Annual Inc. CANCEL CONNECT			
Negative Own Q		Site Pairing		Anne T Kar Ina	<ul> <li>Answer 100</li> </ul>	

3. To initiate the site pairing, click Connect.



VMware HCX Connector must be able to communicate with the HCX Cloud Manager IP over port 443.

4. After the pairing is created, the newly configured site pairing is available on the HCX Dashboard.

The VMware HCX Interconnect (HCX-IX) appliance provides secure tunnel capabilities over the internet and private connections to the target site that enable replication and vMotion-based capabilities. The interconnect provides encryption, traffic engineering, and an SD-WAN. To create the HCI-IX Interconnect Appliance, complete the following steps:

1. Under Infrastructure, select Interconnect > Multi-Site Service Mesh > Compute Profiles > Create Compute Profile.



Compute profiles contain the compute, storage, and network deployment parameters required to deploy an interconnect virtual appliance. They also specify which portion of the VMware data center will be accessible to the HCX service.

For detailed instructions, see Creating a Compute Profile.

← → Q	A # • https://a300-vcsa01.ehcdc.com/ui/app/b	kugin/com.vmware.hybridity/com.vmware.ho	schybridConnect	台		0 1
$\equiv$ vSphere Client Q			С		٢	@ `
HCX Dashboard Infrastructure Ste Paining Sterconstect Erransport Analytics Services Network Extension Substant Recovery System System Support	<ul> <li>Interconnect         <ul> <li>Multi-Site Benice Mean</li> <li>Computer Hiofdia: Service Mean</li> <li>Computer Hiofdia: Service Mean</li> <li>Network</li> <li>hockdemo</li> <li>hockdemo</li> <li>Host ab00-exel(0 encic.com)</li> <li>Host ab00-exel(0 encic.com)</li> <li>Host ab00-exel(0 encic.com)</li> <li>Host ab00-exel(0 encic.com)</li> <li>Host ab00-clusterot</li> <li>Host ab00-exel(0 encic.com)</li> <li>Host aboveexel(0 encic.com)</li> <li>Host aboveexel(0 encic.com)</li> <li>Host aboveexel(0 encic.com)</li> <l< td=""><td>R Profiles Sentinel Management  20) is in ortical (red) state for service compute  20) is in ortical (red) state for service compute  20 and ortical (red) state for designment container  20 and ortical (red) state for designment container  20 and ortical com  21 and ortical com  22 and ortical com  23 and ortical com  24 and ortical com  24 and ortical com  24 and ortical com  25 an</td><td>Networks WM_3510 (Management) Network Container Ofersonik Enterna Retwork Container Ofersonik Enterna WD5-Switch0 (Unlamitted)</td><td>Q C CREATE COM</td><td>PUTE PP</td><td>IOFILE</td></l<></ul></li></ul>	R Profiles Sentinel Management  20) is in ortical (red) state for service compute  20) is in ortical (red) state for service compute  20 and ortical (red) state for designment container  20 and ortical (red) state for designment container  20 and ortical com  21 and ortical com  22 and ortical com  23 and ortical com  24 and ortical com  24 and ortical com  24 and ortical com  25 an	Networks WM_3510 (Management) Network Container Ofersonik Enterna Retwork Container Ofersonik Enterna WD5-Switch0 (Unlamitted)	Q C CREATE COM	PUTE PP	IOFILE

- After the compute profile is created, create the network profile by selecting Multi-Site Service Mesh > Network Profiles > Create Network Profile.
- 3. The network profile defines a range of IP address and networks that will be used by HCX for its virtual appliances.



This will require two or more IP address. These IP addresses will be assigned from the management network to virtual appliances.

$\equiv$ vSphere Client Q							Ca	Administrator@EHCDK	200M Y	3	0
HCX Dashboard Infrastructure Step Pairing Interconnect Et Transport Analytics	< ~	Interconnect Multi-Site Service Mesh Compute Profiles Service Mesh	Notwork Profiles	Sentinel Management	)			Q C CR	EATE NETWO	RK	DFILE
Services Network Extension Signation Disaster Recovery System Support	× ×	VM_3510	MTU 9000	IP Pools 172.21.25	IP Ranges 1.80 - 172.21.254.95	IP Usage(Used/Total) 4/ 16	Prefix Length 24	Gateway 172.21.254.230			
		EDIT DELETE									

For detailed instructions, see Creating a Network Profile.



(i)

If you are connecting with an SD-WAN over the internet, you have to reserve public IPs under the Networking and Security section.

4. To create a service mesh, select the Service Mesh tab within the Interconnect option and select onpremises and VMC SDDC sites.

The service mesh establishes a local and remote compute and network profile pair.

vm VMware HCX		@- administrator-
Dashboard     Infrastructure     Site Paring     Inferconnect	Interconnect Multi-Site Service Mesh Control C	
L: Transport Analytics - Services - Services - Setwork Extension - By Migration	Computer Holins Service Mean Inscission Holines Service Management	Q. C CREATE SERVICE MESH
Onsater Recovery     Administration     Administration     Administration     System Updates     Troubleshooting     Audit Logs	Ste Fung VMexare-HCC-440 Readp Readp Readp Readp Statistics Collection 2016 (2016) wether rope aver, above, com-cloud BARreadw Statistics Collection 2016 (2016)	,
<ul> <li>Activity Logs</li> <li>Alerts</li> <li>DICE</li> </ul>	VIEW APPLIANCES RESYNC EDIT DELETE HORE-	

Part of this process involves deploying HCX appliances that will be automatically configured on both the source and target sites, creating a secure transport fabric.

5. Select the source and remote compute profiles and click Continue.

Create Service Mesh	3 2 3 4 5	
Select Compute Profiles select one compute profile each in the source and remote pites for activating hybridity services. The select select Source Compute Profile (Crusteria )	bons will define the resources, where Virtual Machines will be able to consume HCX services. Select Bernote Compute Phothe ♥ (Computerhister)center))	
📩 Hint 3006 exc00 kinum zamljohost 2202 in in critical (vegi state for servar zongute.		•• CONTINU

6. Select the service to be activated and click Continue.

Edit Service Mesh		1 2 3 4	51		×
Select Services to be activated					
CS Assisted Migration Service can't be SRM totegration Service cannot be set	e selected as one or both the compute profiles as lected as they are not scenaed with this HCK with	ected in previous step doesn't have these services ad Nation.	tivated.		
Hybrid Intercented ( Research of the second	Wan Optimization () Porticities	Cost-cloud VMshowApraton () E2nchowtre Fisconheatty	Buck Magnation () Resolvemente PSchendel R. McBind	Repétation Assisted Wildown Migration () Etchelland.institutioner Picen Rheater	Network Extension @
DisasterRecovery () August RossenerSI Primativesce ADRIson	SIM Integration () Flore-Configuration Fund Integration (Contenting Ameri Participation (Contenting American Participation (Contenting American	(23 April 16 May an an Array of California)			
	Finalistic frame of the set activated				CONTINUE

An HCX Enterprise license is required for Replication Assisted vMotion Migration, SRM Integration, and OS Assisted Migration.

7. Create a name for the service mesh and click Finish to begin the creation process. The deployment should take approximately 30 minutes to complete. After the service mesh is configured, the virtual infrastructure and networking required to migrate the workload VMs has been created.

 $(\mathbf{i})$ 

≡ vichere Glent Q.									C 2		
HCK Or bertmant	f inte	rconr	ect								
C the Parries		-	n bezerten onertiden bezerten							COIT SERVICE	
Services v	* <u>- A</u>	Connected Determined D									
Ingestion     Constant Encourse			August Same	· · · Assessed Tape · · · ·	·* ******		fund Into	Derest Versee	Automa result		
System v A abreventure O Sugaret	* 1		CODITION of 2019/01-0126-071-0021 Ultimeterilities Assessed: A200-2018/01 Biorege #2001_01/1_01204	80 -Crosses			۲	4422	A418 100		
		1.35	COUT-68.0 an UT-69.0 Comman ADD-Cout-67 Temper ADD, VFL2020 Temper A	to contract	722.04k		۲	642.0	****		
			COOLS 405-6 er (ART749-766-6658 citcle ell'association fonnaux ART0-Cutrinol) deman ART0-Cutrinol) deman ART0-Cutrinol	Compact Manager				1444	N/A		
											100
	As	plano 	es on hox.Sebf3b067ddf4cc05e3HS.westeurope.av	ni lazura com-cloud	ne 749	e Moren 72 M N D (wengeness)				5arm 440	
		0007 44				1212 (2012) (2012) 1212 (2012) (2012) 1212 (2012) (2012) 1212 (2012) (2012) 1212 (2012) (				48.0	Q.
			17 m.	6.9	PCLAST D.T	1189H3 (III)				United and	
		000FW	P.M.	9	+Cimin OFI					738	

### Step 6: Migrating Workloads

HCX provides bidirectional migration services between two or more distinct environments such as onpremises and VMC SDDCs. Application workloads can be migrated to and from HCX activated sites using a variety of migration technologies such as HCX bulk migration, HCX vMotion, HCX Cold migration, HCX Replication Assisted vMotion (available with HCX Enterprise edition), and HCX OS Assisted Migration (available with HCX Enterprise edition).

To learn more about available HCX migration technologies, see VMware HCX Migration Types

The HCX-IX appliance uses the Mobility Agent service to perform vMotion, Cold, and Replication Assisted vMotion (RAV) migrations.



The HCX-IX appliance adds the Mobility Agent service as a host object in the vCenter Server. The processor, memory, storage and networking resources displayed on this object do not represent actual consumption on the physical hypervisor hosting the IX appliance.

$\equiv$ vSphere Client $$ Q	
<u>n</u> 89 8 Q	Image: Summary Monitor Configure Permissions VMs Resource Pools Datastores Networks Updates
✓      ✓      ✓      ✓      ✓      ✓      ▲300-ussa01.ehcdc.com     ✓     ▲300-DataCenter     ✓     ✓     ▲300-Duster01     ✓	Hypervisor:     VM ware ESXi, 7.0.3, 20305777       Model:     VM ware Mobility Platform       Processor Type:     VM ware Virtual Processor       Logical Processor:     768       NICs:     8       Virtual Machines:     0       State:     Connected       Uptime:     29 days

### VMware HCX vMotion

This section describes the HCX vMotion mechanism. This migration technology uses the VMware vMotion protocol to migrate a VM to VMC SDDC. The vMotion migration option is used for migrating the VM state of a single VM at a time. There is no service interruption during this migration method.



Network Extension should be in place (for the port group in which the VM is attached) in order to migrate the VM without the need to make an IP address change.

1. From the on-premises vSphere client, go to Inventory, right- click on the VM to be migrated, and select HCX Actions > Migrate to HCX Target Site.

$\equiv$ vSphere Client $O_{\rm c}$ to	Actions - HCX_Photos_14 Power Goest 05		noton_14	> 0 0	@ 62   ]actions	ų				c	8 Administration conco	on u i c	9 0
0 8 8	Corn Remote Console	2	ontor Cord	gure Permis	oons Datastores	Networks	Shipitots	lupdates					
- El A300-batademer	db Mgrate			Durnt OS	VMeans Photon OS (5	-bro							EW VEW
- III A300-Custarth	Clone	21	_	Compatibility VMaster Tools	ESH 6.7 and blar (VM Burning, version milds)	version Manute						U 01	łz
adoc-eskilltahodc.com	Pault Tolerance	7		Otio Name	word aird processor							25 20	MB
addo-execting encode core	VM Policies	,	INSCL.	P-ADDVELOPE	VEW ALL 2 IF ADDRES	163						10	
ANF_HCK_Dens     Ant/Dens	Template	,	I COMBOLE D	0 15	4000-41407.95c0: cor							7.10	101 MD
· G washing	Compacibility	1											
- @ HOKOHINE	Export System Logs						^	.Nobes					*
E HOLPHON,H	Jack de Carrison			1090	60			Construction of the					24
(2) HCR_Photos_36-865300045018	gran series.			82	08.00208 memory act	a		CUITORI APO DI	uner.				^
@ HCX_Photon_17-186305050438	Move to folder			2.08				Aminute			Ville		
(D. HOX, Process, 38	Penane.		Marter 1	VM.3	BO9 (convected)								
(2) HCK_PROTON_TR	Edit Notes		ten f	Dan	mailed								
(1) HCK, Steason, 30	Teos & Custom Attributes	1		71.4			- 20						
B HCK_Photon_21	Add Permission								_				_
(B. HCK, Photos, 24	Alarma	2	·	Devic	e on the virtual machine of for the virtual machine	PICI tous shat pr communicator	enterface					the bend	a to all soliday.
(B HCK_Photon_27				Asse	ional Hardware			FOL					
Y Recent Tasks Alarma	Delate from 2004												
Tech Name T Target			T Detett	2	. Nitatise	*	Sound T	. Vart Tree	, T	Campleton Title	T Server		*
Devers falory E3 HCK_Poston_34_X	NetApp SnapCenter	17			PHONE COMPARING	witrahur	3.00	09/13/2022, 4:45	17 P.	09/13/2022, 4 45 37 P	a300-vess01 ancos care		
Reconfigure Virtuel mach . (B. HCK_Poston_)4	th AS Side Recovery actions		Alignate to H	Target Site	DHCOC COMPAgnin	whatie:	2.mi	09/13/2022, 4:451	15 .	09/10/2022. 4:45-00 .	#300-vena01 encos com		
Move entities D VML mgrated to 0	HOX Actions			can marget side	COC COM A min	100 m	4.03	06/13/2022. 4.451	04 -	09/13/2022 # #5-04	#303-V[180].49108.009		1 at store

2. In the Migrate Virtual Machine wizard, select the Remote Site Connection (target VMC SDDC).

✓ Select Connection				C Recad C	onnecti
(there are 2 records found)					
Source: VMware-HCX-440 / VC a300-	vcsa01.ehcdc.com → Destination: (select)				
HCX Cloud - FSxNDemoSDDC / VC v https://cx.sdo:/54/100/6/08/vmwareumc.com	center.sddc-54-188-6-128.vmwarevmc.com				
0	/ VC 172.30.156.2				
https://172.30/156.8					
✓ Transfer and Placement:					
	(Mandatory: Storage)		(Migration Profile)		~
	Same format as source	-	(Optional: Switchover Schedule)		0
> Switchover:					
Enterned Continues					
Extended Options.					
Edit Extended Options					
0 selected				9	
VM for Migration	Disk / Memory / vCPU		Migration Info		
C Loading data					
() 1000-9 1000					

3. Add a group name and under Transfer and Placement, update the mandatory fields (Cluster, Storage, and Destination Network), Click Validate.

➤ Bource: VMware-HCX-440 / VC → Destination: HCX Cloud - FSxN https://dx.sode.54-184-610 vmwarevmc.com	a300-vcsa01.ehc DemoSDDC / V0	dc.com 2 vcenter.sddc-54-188-6-128.vmware	vmc.com		C Selasd Connect
roup Name: vMotion-vm14-2-vmc			Batch size: 1vw	/ 2.GB/ 2.GB/ 1vcPu	Select VMs for Migratio
<ul> <li>Transfer and Placement:</li> </ul>					*
Compute-ResourcePool		DemoDS01 (1644.08/19.18)		vMotion	~
Workloads		Same format as source	×	(Optional: Switchover Schedule)	0
Switchover:					
Force Power-off VM		Remove Snapshots			
Extended Options:     Edit Extended Options     Retain 10	AC)	Porte unnount SO images			
Extended Options:     Edit Extended Options     Reliais M	xc)	Disk / Memory / vCPU		Migration Info	a,
Extended Options:     Edst Extended Options     Retain 50	AC)	Disk / Memory / vCPU		Migration Info	Q.
	AC)	Disk / Memory / vCPU 2 GB / 2 GB / 1 vCPU (B) DemoDS01 (B64 GB / 1978)		Migration Info	¢.
Extended Options:     Edst Extended Options     Retain Ma  VM for Migration      + HCX_Photon_14      Compute-ResourcePool      Workloads	×C) 0 2 2	Disk / Memory / vCPU 2 GB / 2 GB / 1 vCPU G DemoDS01 (Bit 4 GB / 13 T0) Same format as source		Migration Info	¢. •
	AC)	Disk / Memory / vCPU 2 GB / 2 GB / 1 vCPU @ DemoDS01 (IBE4 GB / 1978) @ Same format as source	v.	Migration Info	á. •

4. After the validation checks are complete, click Go to initiate the migration.



The vMotion transfer captures the VM active memory, its execution state, its IP address, and its MAC address. For more information about the requirements and limitations of HCX vMotion, see Understanding VMware HCX vMotion and Cold Migration.

5. You can monitor the progress and completion of the vMotion from the HCX > Migration dashboard.

y5phere Client Q, hos				G	— Ханнализирскоссом ∨ (2) (2)
	4	Migration			
(CX					
Datroant		= Tracking E Management	D HURATE C		
Altrastructure	Ť				
C/ Site Paring		Auror	WMU Research Manuary/ CPUs	Programs Mart	End Status
C Transport Analytics		v a 100-vesa01 abode com	<ul> <li>vcenter sdoc 54-188-6-128 vmware</li> </ul>	vmc.com	
ervices					
E Network Extension		<ul> <li>vikimus entit 2 smc</li> </ul>	4. 208-208-1	1004 tool 201 Gas Ladiant	10 0 E
Migrature		0 / Tameled		THE CONCERNMENT OF COMPANY AND ADDRESS	
D Draainer Recovery		1 Vildi, have, 14	2 08 - 2 06 1	Swisning own	- Database states
lystem.	¥.				
		Mayalian Options (Reduce New) (	Annue IBS: ) (300: + Q L2E_VM_3509-3505-40041484	Service March Name VMC Service Conte - 1 min age Start Cohedia	i suarce details
		5 Cier22.25	4 808 808 4	S Magazian Complete	
		3 1986.00	4 100 100 4	O Migratus Complete	
		s value	1 206 208 1	Ci Vice and Complete	
		· ····	1 100 100 1	0	
				50 Mill 100 Contracts	
		> 2022-08-12-20-48-EFVPO			
<ul> <li>Becent Tarks Alarms</li> </ul>		> 2022 00 13 50 14 10 100	<ol> <li>2.08 - 2.08 - 1</li> </ol>	0.0	
v Recent Tasks Allerms	T Sates	> 2022 de 12 2048 EFVPO	1. 208 208 1	and T Statistics 5 T Completion Tree	bevel
V. Recent Tabla Alarms va Nane T Teast receite vitue machee B inCuDes	T Statue	2022 de raz zener provo-     Deners     Torners     Torners     Torners	<ul> <li>2.08 2.08 1</li> <li>7 мназа</li> <li>7 Бусор Сон-Аличиствани</li> </ul>	Ch University Constants And T Start Time & T Completion Time T rs 09/70/2022, 419-08	Benefi a000-vesalitatode com

#### VMware Replication Assisted vMotion

As you might have noticed from VMware documentation, VMware HCX Replication Assisted vMotion (RAV) combines the benefits of bulk migration and vMotion. Bulk migration uses vSphere Replication to migrate multiple VMs in parallel—the VM gets rebooted during switchover. HCX vMotion migrates with no downtime, but it is performed serially one VM at a time in a replication group. RAV replicates the VM in parallel and keeps it in sync until the switchover window. During the switchover process, it migrates one VM at a time with no downtime for the VM.

The following screenshot show the migration profile as Replication Assisted vMotion.



The duration of the replication might be longer compared to the vMotion of a small number of VMs. With RAV, only sync the deltas and include the memory contents. The following is a screenshot of the migration status—it shows how the start time of the migration is the same and the end time is different for each VM.

► GO @ VALIDATE M SAVE CLOSE

	<	Micratic	50																
HCK CR Deshtoerd Indeshterber		(E 1)	olog El Manager	oest _	( a wa	iate ]	2											(april 2	
O Ste Parry	×.	here				wear day	nyn/Ham	rry/ CPUs		Pre					liw)	Ð	4	linier	
Tis transport Analytics		·*	center addc 5410	6-128 V	mwarevir	ic.com		100 ves	a0tehcidic.com										
Annon Emerson	×	w.108	t <i>P</i>				8.149	3.02		0	hystor	Completi				1		<b>E</b>	0
Mgration			Q - n parameter																
*D Disaster Recovery		N G	3 HER, Phales, 12		60		2.08	1.08	4. L	0	theyation	Complete			83.29 PM		0	Meraturi completed	
ystem		10	> HCX, Photon, 52		9		2.08	3.08	1	0	Applier	Cemplen			E3 20 Hz	851	til Pe	Meater corpored	
C Lopert		1.0	> HEX,Photos,55		63		2.09	2.08	1			Cirtum	0	1	EX 20 me hep 31	43.	All Pres	Mention completed	
		10	> HDL/Hite, H	4	8		3.08	2.08	5. C	0		Constant			83.20 mil Ing 21	-	10 Pw 21	Mgrates completed	
		> 2023	49 22 15:24 NTTY				\$ 09	3 (3)	4	0	-	Compress	1						
		* av	centersiddc-54-18	1-6-128.v	mwarevii	ic.com	- 44	300 vcs	aO1.ehcdc.com										
		5 feet	1RTP			*	8.00	8.06	4	0	-	Comput		ĵ.		*			
V Recent Tasks Alarma									-										
as have T Target		,	Statue		Brom 1		,	better		۲	tume Fac	۰.	Stat Title		Compartue Terra	7	Sever		
eele istus nachre 🛛 🛱 HCK	Petrol, R. Stadney		Q-Completed					VMCU	CALLArmenthalour		2 ===		3WZM2022	4 01 04.	09/23/2022 4 191	i	voanter sodo 54 88	H GE umage unc com	
regiter value mattere 🛛 🖉 🗠 🔍	Phatan_H		@ completes					VINCL	Cachanneneter		2+1		29/25/2023	4 03 03	09/25/2022 4 03 0	۰.	-cantar \$450-54 m	S GLUMMARTIC CON	
eterhistus machines. 🛱 HCK.	Photo: N		@ Durighted					WHELE	CALM-Investmentar		4.00		08/21/2022	4.53:04.	08/35/2022 4 03 0	φ.	version adds 54 m	14 GB amore and Lore	
elicate visue mainre 🛛 🕀 HOL	Hutun_1		(i) Completing		Highting	Virtual Maio	hte ic	VHCU	CALLA development at an		4.0%		06/25/0022	e 00.55	00/25/0022 4:01/0	294	manter addo 54 m	6.6.421 considered, care	
veate vitual mácrime 🛛 🕆 5000	Calabertar		@-Companie					VMCU	CALMON INTO A		3.64		08/33/0022	3 50 4 7 .	00/35/0022.558.4	÷	viariter and did mit	14.01 masene cor	r
And and Acted allowing and the Physics of the	1001210		@ completes					VMCL	CRL\Administration		4		09/23/2022	A TRITE	00033/2022 85815	×.,	scandar adde-54 88		1.

For additional information about the HCX migration options and on how to migrate workloads from onpremises to VMware Cloud on AWS using HCX, see the VMware HCX User Guide.



VMware HCX vMotion requires 100Mbps or higher throughput capability.



# Conclusion

Whether you are targeting all-cloud or hybrid cloud and data residing on any type/vendor storage in onpremises, Amazon FSx for NetApp ONTAP along with HCX provide excellent options to deploy and migrate the workloads while reducing the TCO by making the data requirements seamless to the application layer. Whatever the use case, choose VMC along with FSx for ONTAP datastore for rapid realization of cloud benefits, consistent infrastructure, and operations across on-premises and multiple clouds, bidirectional portability of workloads, and enterprise-grade capacity and performance. It is the same familiar process and procedures used to connect the storage and migrate VMs using VMware vSphere replication, VMware vMotion or even NFC copy.

# Takeaways

The key points of this document include:

- You can now use Amazon FSx ONTAP as a datastore with VMC SDDC.
- You can easily migrate data from any on-premises datacenter to VMC running with FSx for ONTAP datastore
- You can easily grow and shrink the FSx ONTAP datastore to meet the capacity and performance requirements during migration activity.

# Where to find additional information

To learn more about the information described in this document, refer to the following website links:

• VMware Cloud documentation

https://docs.vmware.com/en/VMware-Cloud-on-AWS/

Amazon FSx for NetApp ONTAP documentation

https://docs.aws.amazon.com/fsx/latest/ONTAPGuide

VMware HCX User Guide

• https://docs.vmware.com/en/VMware-HCX/4.4/hcx-user-guide/GUID-BFD7E194-CFE5-4259-B74B-991B26A51758.html

# Region Availability – Supplemental NFS datastore for VMC

Learn more about the the Global Region support for AWS, VMC and FSx ONTAP.



NFS datastore will be available in regions where both services (VMC and FSx ONTAP) are available.

The availability of supplemental NFS datastores on AWS / VMC is defined by Amazon. First, you need to

determine if both VMC and FSxN are available in a specified region. Next, you need to determine if the FSxN supplemental NFS datastore is supported in that region.

- Check the availability of VMC here.
- Amazon's pricing guide offers information on where FSxN (FSx ONTAP) is available. You can find that information here.
- Availability of the FSxN supplemental NFS datastore for VMC is coming soon.

While information is still being released, the following chart identifies the current support for VMC, FSxN and FSxN as a supplemental NFS datastore.

# Americas

AWS Region	VMC Availability	FSx ONTAP Availability	NFS Datastore Availability
US East (Northern Virginia)	Yes	Yes	Yes
US East (Ohio)	Yes	Yes	Yes
US West (Northern California)	Yes	No	No
US West (Oregon)	Yes	Yes	Yes
GovCloud (US West)	Yes	Yes	Yes
Canada (Central)	Yes	Yes	Yes
South America (Sao Paulo)	Yes	Yes	Yes

Last updated on: June 2, 2022.

# EMEA

AWS Region	VMC Availability	FSx ONTAP Availability	NFS Datastore Availability
Europe (Ireland)	Yes	Yes	Yes
Europe (London)	Yes	Yes	Yes
Europe (Frankfurt)	Yes	Yes	Yes
Europe (Paris)	Yes	Yes	Yes
Europe (Milan)	Yes	Yes	Yes
Europe (Stockholm)	Yes	Yes	Yes

Last updated on: June 2, 2022.

## Asia Pacific

AWS Region	VMC Availability	FSx ONTAP Availability	NFS Datastore Availability
Asia Pacific (Sydney)	Yes	Yes	Yes
Asia Pacific (Tokyo)	Yes	Yes	Yes
Asia Pacific (Osaka)	Yes	No	No
Asia Pacific (Singapore)	Yes	Yes	Yes
Asia Pacific (Seoul)	Yes	Yes	Yes
Asia Pacific (Mumbai)	Yes	Yes	Yes
Asia Pacific (Jakarta)	No	No	No
Asia Pacific (Hong Kong)	Yes	Yes	Yes

# **NetApp Capabilities for Azure AVS**

Learn more about the capabilities that NetApp brings to the Azure VMware Solution (AVS) - from NetApp as a guest connected storage device or a supplemental NFS datastore to migrating workflows, extending/bursting to the cloud, backup/restore and disaster recovery.

Jump to the section for the desired content by selecting from the following options:

- Configuring AVS in Azure
- NetApp Storage Options for AVS
- NetApp / VMware Cloud Solutions

## **Configuring AVS in Azure**

As with on-premises, planning a cloud based virtualization environment is critical for a successful productionready environment for creating VMs and migration.

This section describes how to set up and manage Azure VMware Solution and use it in combination with the available options for connecting NetApp storage.



In-guest storage is the only supported method of connecting Cloud Volumes ONTAP to Azure VMware Solution.

The setup process can be broken down into the following steps:

- · Register the resource provider and create a private cloud
- · Connect to a new or existing ExpressRoute virtual network gateway
- · Validate the network connectivity and access the private cloud

View the detailed configuration steps for AVS.

## **NetApp Storage Options for AVS**

NetApp storage can be utilized in several ways - either as guess connected or as a supplemental NFS datastore - within Azure AVS.

Please visit Supported NetApp Storage Options for more information.

Azure supports NetApp storage in the following configurations:

- Azure NetApp Files (ANF) as guest connected storage
- · Cloud Volumes ONTAP (CVO) as guest connected storage
- Azure NetApp Files (ANF) as a supplemental NFS datastore

View the detailed guest connect storage options for AVS. View the detailed supplemental NFS datastore options for AVS.

# Solution Use Cases

With NetApp and VMware cloud solutions, many use cases are simple to deploy in Azure AVS. se cases are defined for each of the VMware defined cloud areas:

- Protect (includes both Disaster Recovery and Backup / Restore)
- Extend
- Migrate

Browse the NetApp solutions for Azure AVS

# Protecting Workloads on Azure / AVS

## Disaster Recovery with ANF and JetStream

Disaster recovery to cloud is a resilient and cost-effective way of protecting the workloads against site outages and data corruption events (for example, ransomware). Using the VMware VAIO framework, on-premises VMware workloads can be replicated to Azure Blob storage and recovered, enabling minimal or close to no data loss and near-zero RTO.

JetStream DR can be used to seamlessly recover the workloads replicated from on-premises to AVS and specifically to Azure NetApp Files. It enables cost-effective disaster recovery by using minimal resources at the DR site and cost-effective cloud storage. JetStream DR automates recovery to ANF datastores via Azure Blob Storage. JetStream DR recovers independent VMs or groups of related VMs into recovery site infrastructure according to network mapping and provides point-in-time recovery for ransomware protection.

This document provides an understanding of the JetStream DR principles of operations and its main components.

- 1. Install JetStream DR software in the on-premises data center.
  - a. Download the JetStream DR software bundle from Azure Marketplace (ZIP) and deploy the JetStream DR MSA (OVA) in the designated cluster.
  - b. Configure the cluster with the I/O filter package (install JetStream VIB).
  - c. Provision Azure Blob (Azure Storage Account) in the same region as the DR AVS cluster.
  - d. Deploy DRVA appliances and assign replication log volumes (VMDK from existing datastore or shared iSCSI storage).
  - e. Create protected domains (groups of related VMs) and assign DRVAs and Azure Blob Storage/ANF.
  - f. Start protection.
- 2. Install JetStream DR software in the Azure VMware Solution private cloud.
  - a. Use the Run command to install and configure JetStream DR.
  - b. Add the same Azure Blob container and discover domains using the Scan Domains option.
  - c. Deploy required DRVA appliances.
  - d. Create replication log volumes using available vSAN or ANF datastores.
  - e. Import protected domains and configure RocVA (recovery VA) to use ANF datastore for VM placements.
  - f. Select the appropriate failover option and start continuous rehydration for near-zero RTO domains or VMs.
- During a disaster event, trigger failover to Azure NetApp Files datastores in the designated AVS DR site.
- 4. Invoke failback to the protected site after the protected site has been recovered.Before starting, make sure that the prerequisites are met as indicated in this link and also run the Bandwidth Testing Tool (BWT) provided by JetStream Software to evaluate the potential performance of Azure Blob storage and its replication bandwidth when used with JetStream DR software. After the pre-requisites, including connectivity, are in place, set up and subscribe to JetStream DR for AVS from the Azure Marketplace. After the software bundle is downloaded, proceed with the installation process described above.

When planning and starting protection for a large number of VMs (for example, 100+), use the Capacity Planning Tool (CPT) from the JetStream DR Automation Toolkit. Provide a list of VMs to be protected together with their RTO and recovery group preferences, and then run CPT.

CPT performs the following functions:

- · Combining VMs into protection domains according to their RTO.
- Defining the optimal number of DRVAs and their resources.
- Estimating required replication bandwidth.
- Identifying replication log volume characteristics (capacity, bandwidth, and so on).
- Estimating required object storage capacity, and more.



The number and content of domains prescribed depend upon various VM characteristics such as average IOPS, total capacity, priority (which defines failover order), RTO, and others.

## Install JetStream DR in On-Premises Datacenter

JetStream DR software consists of three major components: JetStream DR Management Server Virtual Appliance (MSA), DR Virtual Appliance (DRVA), and host components (I/O Filter packages). MSA is used to install and configure host components on the compute cluster and then to administer JetStream DR software. The following list provides a high-level description of the installation process:

- 1. Check prerequisites.
- 2. Run the Capacity Planning Tool for resource and configuration recommendations (optional but recommended for proof-of-concept trials).
- 3. Deploy the JetStream DR MSA to a vSphere host in the designated cluster.
- 4. Launch the MSA using its DNS name in a browser.
- 5. Register the vCenter server with the MSA.To perform the installation, complete the following detailed steps:
- After JetStream DR MSA has been deployed and the vCenter Server has been registered, access the JetStream DR plug-in using the vSphere Web Client. This can be done by navigating to Datacenter > Configure > JetStream DR.

vm vSphere Client	Menu 🗸 🛛 🔍 Search in all	environments	C	Administrator@EHCDC.COM V	C
□         □         ●           ∨         ⊡         a300-vcsa.ehcdc.com         ^	A300-DataCent	er ACTIONS ~			
<ul> <li>A300-Datacenter</li> <li>A300-Cluster</li> <li>a300-esxi02.eh</li> </ul>	<ul> <li>More         Alarm Definitions         Scheduled Tasks     </li> </ul>	JebSbream DR Protected Domains Statistics Storage Sites Appliances Configurations Task	Log		Ē
a300-esti03.en.	Network Protocol Pr	Site Details		Alarm Sett	ings
a300-esxi05.eh	JetStream DR	vCenter Server Hostname 172.21.253.160			
ANFJSDR-MSA0		Management Appliance Hostname ANFJSDR-msa			
🖧 AuctionAppA0		Software Version 4.0.0.443			
🖧 AuctionAppA2		Subscription ID 0000000-0000-0000-0000-000000000001	onfigure		
🖧 AuctionAppA3		Tenant ID / Application ID - Configure			
AuctionAppB0		Application Secret - Configure			

7. From the JetStream DR interface, select the appropriate cluster.

ite Details				Alarm Setting
Center Server Hostname	172.21.253.160			
Aanagement Appliance Hostname loftware Version	Configure Clusters			
iubscription ID fenant ID / Application ID		Select All Cit	ear All Q	
opplication Secret	🗹 Cluster Name 🔺	Datacenter Name 🔺		
onfigured Clusters	A300-Cluster	A300-DataCenter	^	
Configure Cluster				c
Cluster Name 🔺			~	Host Details 🔺
No cluster contigured		Cancel	Configure	
		Tu		

8. Configure the cluster with the I/O filter package.

Instacted Domaine Statistics	Storage Sites Appliances Configurations Tack Log	
Storage Sites	Add Storage Site	
+ Add Storage Site @ Scan Domains	^	
	Storage Site Type.*	
Name 🛦	Azure Blob Storage	
No Storage Site configured.		
	Access Type *	
	Key Access	
	Storage Site Name (Provide a name to identify this Site) *	
	ANFDemoblobrepo	
Storage Site Details Alarms		
	Azure Blob Storage Account Name *	
	anfdrdemostor	
No storage site selected. Select a storage		
	Azure Blob Storage Account Key *	
	······································	
	Cancel Add Storage Site	

- 9. Add Azure Blob Storage located at the recovery site.
- 10. Deploy a DR Virtual Appliance (DRVA) from the Appliances tab.



DRVAs can be automatically created by CPT, but for POC trials we recommend configuring and running the DR cycle manually (start protection > failover > failback).

The JetStream DRVA is a virtual appliance that facilitates key functions in the data replication process. A protected cluster must contain at least one DRVA, and typically one DRVA is configured per host. Each DRVA can manage multiple protected domains.

JetStream DR Protected Domains Statistics	Deploy New DR Virtual Ap	pliance (DRVA)					
DRVAs (DR Virtual Appliances)	1. General	2. DRVA VM	3. DRVA Networ	rk 4. Su	nmary		
+ Deploy New DRVA	Name		ANFdemo001		^		Q
Name 🔺	Description (Optional)					Details 🔺	
No DR Virtual Appliance configured.	Datacenter		A300-DataCenter				
	Cluster		A300-Cluster				
	Resource Pool (Optional)		-				
	VM Folder (Optional)		-				$\checkmark$
Replication Log Volume	Datastore		A300_NFS_DS04				
	Number Of CPUs		8				-
+ New Replication Log Volume	Memory Size		32GB				Q
Disk Path Name 🔺	Management Network		VM_187			Details 🔺	
No DRVA selected. Select a DRVA to vi	Host(iofilter) to DRVA Data	Network	VM_187				
	Replication Network to O	oject Store	VM_187				
	Replication Log Network		VM_187		~		$\sim$
			Cancel	Back	Deploy		

In this example, four DRVA's were created for 80 virtual machines.

- 1. Create replication log volumes for each DRVA using VMDK from the datastores available or independent shared iSCSI storage pools.
- 2. From the Protected Domains tab, create the required number of protected domains using information

about the Azure Blob Storage site, DRVA instance, and replication log. A protected domain defines a specific VM or set of VMs within the cluster that are protected together and assigned a priority order for failover/failback operations.

ect Protected Domain:	Create Protected Domain			+ Create	
	1 Ceneral 2 Pr		3 Summary		
	i. General 2.11	linary site	o. Summary	^	
	Protected Domain Name	ANFPD001			
	Priority Level (Optional)	1			
	Total estimated data size to be protected	1000GB			
	DR Virtual Appliance	ANFdemo001			
	Compression	Yes			
	Compression Level	Default			
	Normal GC Storage Overhead	50%			
	Maximum GC Storage Overhead	300%			
	Replication Log Storage	/dev/sdb			
	Replication Log Size	94.31GB			
	Metadata Size	31.56GB		~	
		*****	1		

3. Select VMs you want to protect and start VM protection of the protected domain. This begins data replication to the designated Blob Store.

Verify that the same protection mode is used for all VMs in a protected domain.

Write- Back(VMDK) mode can offer higher performance.

	Start	Protection						-
Select Protected Domain: ANFPD001	o tai t	i rotocacii					reate Delete	
Recoverable / Total VMs	Protec	tion Mode for selected VMs	_			•		Edit Details
Replication Status	VVIIE	e-Back(VMDK)				ų	ANFDemoblobrepo	^
		VM Name 🔺		# of Disks	Protection Mode		AL ( 172.21.253.160 )	
Remaining Background Data	_	1	×				0-DataCenter \ A300-Cluster	
Current RPO		AuctionAppA1		1	Write-Back(VMDK) V	^	bled	~
		AuctionAppB1		1	Write-Back(VMDK) V			
Protected VMs Settings Ala		AuctionDB1		2	Write-Back(VMDK) V			
		AuctionLB1		1	Write-Back(VMDK) V			
+ Start Protection		AuctionMSQ1		1	Write-Back(VMDK) V			Q
		AuctionNoSQL1		2	Write-Back(VMDK) v			
U VM Name 🛦		AuctionWebA1		1	Write-Back(VMDK) v	(	ground Dat Details	
No VM is protected.		AuctionWebB1		1	Write-Back(VMDK) V			
		Client1		1	Write-Back(VMDK) V			
		D00D04		0	(	~		
					Cancel Start Pr	otection		

Verify that replication log volumes are placed on high performance storage.



i

Failover run books can be configured to group the VMs (called Recovery Group), set boot order sequence, and modify the CPU/memory settings along with IP configurations.

# Install JetStream DR for AVS in an Azure VMware Solution private cloud using the Run command

A best practice for a recovery site (AVS) is to create a three-node pilot-light cluster in advance. This allows the recovery site infrastructure to be preconfigured, including the following items:

- Destination networking segments, firewalls, services like DHCP and DNS, and so on.
- Installation of JetStream DR for AVS
- Configuration of ANF volumes as datastores, and moreJetStream DR supports near-zero RTO mode for mission- critical domains. For these domains, destination storage should be preinstalled. ANF is a recommended storage type in this case.



Network configuration including segment creation should be configured on the AVS cluster to match on-premises requirements.

Depending on the SLA and RTO requirements, continuous failover or regular (standard) failover mode can be used. For near-zero RTO, continuous rehydration should be started at the recovery site.

To install JetStream DR for AVS on an Azure VMware Solution private cloud, complete the following steps:

1. From the Azure portal, go to the Azure VMware solution, select the private cloud, and select Run command > Packages > JSDR.Configuration.



The default CloudAdmin user in Azure VMware Solution doesn't have sufficient privileges to install JetStream DR for AVS. Azure VMware Solution enables simplified and automated installation of JetStream DR by invoking the Azure VMware Solution Run command for JetStream DR.

The following screenshot shows installation using a DHCP-based IP address.

= Microsoft Azure	.P. Southmoor	ter, services, and door (5+7)	🖂 🖓 🖓 🛞 🖓 🖉 🕬 🕬 🕬 Ministérictupp.com
Home > ANFOataClus			Run command - Install-JetDRWithDHCP ×
ANFDataClus   Run	command		This turn level Crimited Described birth burdle from MMS crimites a new same analysis
P Seieth (Ch1+)	🔿 Refresh 🕂 Feedback		elevated privilleges to the user, deploys JetDr Management Server Appliance/UEIA), registers vCenter to the JetDr MSR, configures cluster.
Access control (AM)	Packages Run execution statue		Command parameters
♦ bgs			Register/WWNp ()
Diagnose and solve problems	<ul> <li>Name</li> </ul>	Description	The I
Settings	V KINConfectation 214 Instant	Maria to conference of address fullying on \$20 for address fullying in the control	ProtectedCluster * ()
	Duarde left)BlocChatter	This Creditt unconficiants a chatter but science uncostal LetDR completely to other clusters	Dute 1
A 1998		polos	Datastore * (j)
Manage	Inable-JetON or Outline	This Crediet configures an additional cluster for protection. It installs vibs to all hosts in the	vsarDatestore
S Connectivity	Install, Information	This true level Condist Downloads (arthe true ide from MMS) creates a new user assistmented	VMName* ()
B Chaten		registers vCenter to the tetDr MSA, configures chaster.	antjuval-msa
	Instati-JerDifdWendlamatik	This top level Cindlet Downloads JetDr bundle from MMS, creates a new user, assigns elev-	Duster * 🕓
us identity		registers vCenter to the JetDr MSA, configures cluster.	Outer-1
<ul> <li>Storage (preview)</li> </ul>	Invole PrefightietDSmital	This Cinclust checks and doplays current state of the system it checks whether the minimal	- Credential 🕕
Placement policies		4 bods, if the cluster details are correct, if there is already a VM with the same name pricin	Otername 50
+ Add-ons	Invoke Prefight/etDRUkimball	This Criticlet checks and dopleys current state of the system it checks whether the minimal 4 foots of the checks details are connect and if any VCenter is semicons to the MCA	rept
descent and the second	A DESCRIPTION OF A DESC	A root is the result of the set of the set of the root	Password 4
Workload Networking	Universities setting	the top level Undlet creates a new user, assigns elevated privileges to the user, uncontigu	
< Segments	3 Microsoft/WS/Management 4141	winter on debits administration well take in managing Abure VMMeet Southors	HistName 🕤
TT DHCF			anfpoal-maa
E Port merceling			Network* (j)
0.016			DRSep
			3877240
Operation			Details
📮 Run command 🛛 👻			Retain up to

2. After JetStream DR for AVS installation is complete, refresh the browser. To access the JetStream DR UI, go to SDDC Datacenter > Configure > JetStream DR.

Protected Domains St	atistics	Storage Site	es Ap	pliances	Configurations	Task Log		
Site Details							Alarm Sett	tin
vCenter Server Hostname		172.30.15	56.2					
Management Appliance Host	name	anfjsval-n	nsa					
Software Version		<mark>4</mark> .0.2.450						
Subscription ID		- Config	ure					
Tenant ID / Application ID		- Config	ure					
Application Secret		- Config	ure					
Configure Cluster	pgrade	Duconfigure	🛠 Resolv	/e Configure Is	sue			(
Cluster Name		Datacenter	Name 🔺	Status 🔺	Software V	ersion 🔺	Host Details	
Cluster-1		SDDC-Data	center	🕝 Ok	4.0.2.132		Details	

3. From the JetStream DR interface, add the Azure Blob Storage account that was used to protect the on-premises cluster as a storage site and then run the Scan Domains option.

	Protected Domain	Description	Recoverable V	VMs	Import				
orage Sit	ANFPD000	Protected Domain Tile0	20	20	Import		^	188	
Add Stora	ANFPD001	-	20	20	Import				0
ame 🔺	ANFPD002	Protected Domain 02	20	20	Import				
VFDemoble	ANFPD003	Protected Domain Tile 03	20	20	Import		~		
	<					>		100	
_								186	
torage Si									
torage Si								- 100	
itorage Si								~	

4. After the protected domains are imported, deploy DRVA appliances. In this example, continuous rehydration is started manually from the recovery site using the JetStream DR UI.



These steps can also be automated using CPT created plans.

- 5. Create replication log volumes using available vSAN or ANF datastores.
- 6. Import the protected domains and configure the Recovery VA to use the ANF datastore for VM placements.

ect Protected Domain:	Continuous Fa	ilover Protected Dom	ain				🕅 Delete		More
de	•	•	•	•	•	•			Det
overable / Total VMs	1. General	2a. Failover Settings	2b. VM Settings	3. Recovery	VA 4. DR Setting	gs 5. Summary	reporec		
	Protected Dr	main Name					253.160 )		
	Datacenter	Jinani Name		SDDC-Datacenter					
	Cluster			Cluster-1					
	Resource Po	ool (Optional)		-			1.00		
rotected VMs Setti	VM Folder (C	Optional)		-					
	Datastore			ANFRecoDSU002					
	Internal Netv	vork		DRSeg					C
VM Name 🔺	External Rep	lication Network		DRSeg				Details	
AuctionAppA2	Managemen	t Network		DRSeg				Details	1
AuctionAppB2	Storage Site			ANEDemoblobren	orec			Details	
AuctionDB2	DR Virtual Ar	pliance						Details	
AuctionLB2	Devintual Ap			/dou/odb			$\sim$	Details	
AuctionMSQ2					Cancel Back	Continuous Failo	ver	Details	
AuctionNoSQL2					Bank			Details	



Make sure that DHCP is enabled on the selected segment and enough IPs are available. Dynamic IPs are temporarily used while domains are recovering. Each recovering VM (including continuous rehydration) requires an individual dynamic IP. After recovery is complete, the IP is released and can be reused.

7. Select the appropriate failover option (continuous failover or failover). In this example, continuous rehydration (continuous failover) is selected.

Protected Domains Statistics	Storage Sites A	ppliances	Configuratio	ons	Task Log	-
Select Protected Domain: ANFPD00	0 - View all		+ Create	1	Delete	■More
Mode	Imported	Configura	itions		O Restore	
Recoverable / Total VMs	20 / 20	Storage Si	te	1	→ Failover	
		Owner Site	•	REN	→ Continuo	us Failover
					→ Test Faild	over
Protected VMs Settings Ala	arms					
						c
VM Name 🔺	Protecti	on Status 🔺	F	rotectio	n Mode 🛦	Details
AuctionAppA0	Recov	erable	1	Vrite-Bac	k(VMDK)	Details '
AuctionAppB0	@ Recov	erable	i.	Nrite-Bac		Dotaile

Performing Failover / Failback

1. After a disaster occurs in the protected cluster of the on-premises environment (partial or full failure), trigger the failover.



CPT can be used to execute the failover plan to recover the VMs from Azure Blob Storage into the AVS cluster recovery site.

After failover (for continuous or standard rehydration) when the protected VMs have been started in AVS, protection is automatically resumed and JetStream DR continues to replicate their data into the appropriate/original containers in Azure Blob Storage.

Jet/Stream DR Protected Domains Statts	Complete Continuous Failove	r for Protected Domain ANFPD003			8
Select Protected Domain: Al	VM Network Mapping			O Fallover	I More
Terrarian and the second se	Protected VM Network A	Recovery VM Network	^		
Atodon	VM_3510	DRSeg *	^		
Receive/able / Total VMs				nobiobroporec	(A)
That as The successful Connect Design				72.21.253 160 )	
Line (Processiver Nown Kinns	Force Failover			penter ) Cluster-1	
Current Ship	S States States States		-		*
Protoclud VMs Softings	Force Failov required Complete ou Are you sure	ver of Protected Domain requested. Administrator consent is wnership of this Protected Domain will be taken over by this Site. e you want to continue?	~ .		
	Other S	Cancel Confirm			4
Auctoriancea	Planama manaret		_	Jetails Writes	0
AurhonAnges	Force Failover			and a second sec	
Aurtonilies					
Authors B3				Martin Contraction	
Automasco				and a second	
Australia S/04.3					10
		Cancel Com	plete Failover		



The task bar shows progress of failover activities.

2. When the task is complete, access the recovered VMs and business continues as normal.

Protected Domains	Stati	Result	
elect Protected Dom	ain: Al 🕢 Task Completed Successfu	By	Toelete E More
vlode			Edt Detai
Recoverable / Total VI	Protected Domain	ANFPD003	nobiobreporec
	VMs Recovery status	© success	2 30 156 2 )
teplication Status	Total VMs Recovered	20	tenter \ Cluster-1
temaining Backgroun	T Data		the second se
Pre-script Execution Status		Not defined	
Current RPO	Runbook Execution Status	© Success	~
	Post-script Execution Status	Not defined	
Protected VMs	Settings		
+ Start Protection	Barry Pro		Q
UM Name A			t D Details
AuctionAppA3			Details
AuctionApp83			Detans
AuctionDB3			Details
AuctionLB3			Details
AuctionMSQ3			Dismiss Details
AuctionNoSQL3	U POSSARE AUTO	TOR INCOMPOSITION	Details V

After the primary site is up and running again, failback can be performed. VM protection is resumed and data consistency should be checked.

3. Restore the on-premises environment. Depending upon the type of disaster incident, it might be necessary to restore and/or verify the configuration of the protected cluster. If necessary, JetStream DR software might need to be reinstalled.



Note: The recovery\_utility\_prepare\_failback script provided in the Automation Toolkit can be used to help clean the original protected site of any obsolete VMs, domain information, and so on.

4. Access the restored on-premises environment, go to the Jetstream DR UI, and select the appropriate protected domain. After the protected site is ready for failback, select the Failback option in the UI.

Select Protected Domain: ANFPD003	<ul> <li>View all</li> </ul>		+ Create	Delete	■More
Mode	Running in Failover	nfigurations		O Restore	
Active Site	172.30.156.2 St	orage Site		ANE O Resume Continuous Rehydrati	
Recoverable / Total VMs	20 / 20	Owner Site		OT ← Failback	
Protected VMs Settings Alarms					
VM Name	Protection Status	Protection Mode 🔺		Details	
VM Name  AuctionAppA3	Protection Status	Write-Back(VMDK)		Details Details	
VM Name A AuctionAppA3 AuctionAppB3	Protection Status Protection Status Recoverable Recoverable	Write-Back(VMDK)		Details Details Details	, i
VM Name A AuctionAppA3 AuctionAppB3 AuctionDB3	Protection Status / © Recoverable © Recoverable © Recoverable	Protection Mode ▲ Write-Back(VMDK) Write-Back(VMDK) Write-Back(VMDK)		Details Details Details Details	
VM Name A AuctionAppA3 AuctionAppB3 AuctionDB3 AuctionLB3	Protection Status / © Recoverable © Recoverable © Recoverable © Recoverable	Protection Mode     Write-Back(VMDK)     Write-Back(VMDK)     Write-Back(VMDK)     Write-Back(VMDK)		Details Details Details Details Details	
VM Name A AuctionAppA3 AuctionAppB3 AuctionDB3 AuctionLB3 AuctionMSQ3	Protection Status / © Recoverable © Recoverable © Recoverable © Recoverable © Recoverable	Protection Mode     Write-Back(VMDK)     Write-Back(VMDK)     Write-Back(VMDK)     Write-Back(VMDK)     Write-Back(VMDK)     Write-Back(VMDK)		Details Details Details Details Details Details	


The CPT generated failback plan can also be used to initiate the return of the VMs and their data from the object store back to the original VMware environment.



Specify the maximum delay after pausing VMs in the recovery site and restarting in the protected site. This time includes completing replication after stopping failover VMs, the time to clean recovery site, and the time to recreate VMs in protected site. The NetApp recommended value is 10 minutes.

Complete the failback process, and then confirm the resumption of VM protection and data consistency.

# **Ransomeware Recovery**

Recovering from ransomware can be a daunting task. Specifically, it can be hard for IT organizations to determine the safe point of return and, once determined, how to ensure that recovered workloads are safeguarded from the attacks reoccurring (from sleeping malware or through vulnerable applications).

JetStream DR for AVS together with Azure NetApp Files datastores can address these concerns by allowing organizations to recover from available points in time, so that workloads are recovered to a functional, isolated network if required. Recovery allows applications to function and communicate with each other while not exposing them to north- south traffic, thereby giving security teams a safe place to perform forensics and other necessary remediation.

JetStream DR Protected Domains	latistics Sto	rade Sites Apolia	ves Configu	rations Task I	pa.			ц.	16
Select Protected Domai	Failback Prote	cted Domain						Delete	≡More
Mode	<ul> <li>1. General</li> </ul>	2a. Failback Settings	2b. VM Settings	3. Recovery VA	4. DR	Settings	5. Summary		East Details
Active Site								▲ \$P0	^
Recoverable / Total VMs	Protected De	omain Name		ANFPD003				00.2 )	
	Failback Dat	acenter		A300-DataCenter					
	Failback Clu	ster		A300-Cluster					Y.
Drainalad Mile Co	Failback Res	iource Pool		2				1000	
PIDIECIED VINS DE	VM Folder (C	Optional)						1.000	
	Failback Dat	astore		A300_NFS_DS02				- 200	0
VM Name	Maximum De	alay After Stopping		60 Minutes					~
AuctionAppA3	Internal Net	work		VM_187					^
AuctionAppB3	External Rep	olication Network		VM_187					
AuctionDB3	Managemen	t Network		VM_187				~	
AuctionLB3					Cancel	Back	Failback		
AuctionMSQ3	_		Recoverat	le	Write-Back	(VMDK)	Deta	105	
AuctionNoSQL3			Recoverat	le	Write-Back	(VMDK)	Deta	<u>ults</u>	~

## Disaster Recovery with CVO and AVS (guest-connected storage)

Disaster recovery to cloud is a resilient and cost-effective way of protecting workloads against site outages and data corruption events such as ransomware. With NetApp SnapMirror, on-premises VMware workloads that use guest-connected storage can be replicated to NetApp Cloud Volumes ONTAP running in Azure.

# Overview

Authors: Ravi BCB and Niyaz Mohamed, NetApp

This covers application data; however, what about the actual VMs themselves. Disaster recovery should cover all dependent components, including virtual machines, VMDKs, application data, and more. To accomplish this, SnapMirror along with Jetstream can be used to seamlessly recover workloads replicated from on-premises to Cloud Volumes ONTAP while using vSAN storage for VM VMDKs.

This document provides a step-by-step approach for setting up and performing disaster recovery that uses NetApp SnapMirror, JetStream, and the Azure VMware Solution (AVS).



# Assumptions

This document focuses on in-guest storage for application data (also known as guest connected), and we assume that the on-premises environment is using SnapCenter for application-consistent backups.



This document applies to any third-party backup or recovery solution. Depending on the solution used in the environment, follow best practices to create backup policies that meet organizational SLAs.

For connectivity between the on-premises environment and the Azure virtual network, use the express route global reach or a virtual WAN with a VPN gateway. Segments should be created based on the on-premises vLAN design.



There are multiple options for connecting on-premises datacenters to Azure, which prevents us from outlining a specific workflow in this document. Refer to the Azure documentation for the appropriate on-premises-to-Azure connectivity method.

## **Deploying the DR Solution**

## **Solution Deployment Overview**

- 1. Make sure that application data is backed up using SnapCenter with the necessary RPO requirements.
- 2. Provision Cloud Volumes ONTAP with the correct instance size using Cloud manager within the appropriate subscription and virtual network.
  - a. Configure SnapMirror for the relevant application volumes.
  - b. Update the backup policies in SnapCenter to trigger SnapMirror updates after the scheduled jobs.
- 3. Install the JetStream DR software in the on-premises data center and start protection for virtual machines.
- 4. Install JetStream DR software in the Azure VMware Solution private cloud.
- 5. During a disaster event, break the SnapMirror relationship using Cloud Manager and trigger failover of virtual machines to Azure NetApp Files or to vSAN datastores in the designated AVS DR site.
  - a. Reconnect the ISCSI LUNs and NFS mounts for the application VMs.
- 6. Invoke failback to the protected site by reverse resyncing SnapMirror after the primary site has been recovered.

## **Deployment Details**

### Configure CVO on Azure and replicate volumes to CVO

The first step is to configure Cloud Volumes ONTAP on Azure (Link) and replicate the desired volumes to Cloud Volumes ONTAP with the desired frequencies and snapshot retentions.

Health Status 🕴	Source Volume		Target Volume =	Total Transfer Time	Status	Mirror State	Last Successful Transfer	6
0	gcsdrsqldb_sc46 ntaphci-a300e9u25		gcsdrsqldb_sc46_copy ANFCVODRDemo	17 seconds	idle	snapmirrored	May 6, 2022, 11:43:18 AM 105.06 KiB	
0	gcsdrsqlhld_sc46_copy ANFCVODRDemo	,	gcsdrsqlhld_sc46 ntaphci-a300e9u25	7 seconds	idle	snapmirrored	May 6, 2022, 11:42:20 AN 7.22 MiB	
$\odot$	gcsdrsqilog_sc46 ntaphci-a300e9u25		gcsdrsqilog_sc46_copy ANFCVODRDemo	16 seconds	idle	snapmirrored	May 6, 2022, 11:43:52 AM 130.69 KiB	

#### Configure AVS hosts and CVO data access

Two important factors to consider when deploying the SDDC are the size of the SDDC cluster in the Azure VMware solution and how long to keep the SDDC in service. These two key considerations for a disaster recovery solution help reduce the overall operational costs. The SDDC can be as small as three hosts, all the way up to a multi-host cluster in a full-scale deployment.

The decision to deploy an AVS cluster is primarily based on the RPO/RTO requirements. With the Azure VMware solution, the SDDC can be provisioned just in time in preparation for either testing or an actual disaster event. An SDDC deployed just in time saves on ESXi host costs when you are not dealing with a disaster. However, this form of deployment affects the RTO by a few of hours while SDDC is being provisioned.

The most common deployed option is to have SDDC running in an always-on, pilot-light mode of operation. This option provides a small footprint of three hosts that are always available, and it also speeds up recovery operations by providing a running baseline for simulation activities and compliance checks, thus avoiding the risk of operational drift between the production and DR sites. The pilot-light cluster can be scaled up quickly to the desired level when needed to handle an actual DR event.

To configure AVS SDDC (be it on-demand or in pilot-light mode), see Deploy and configure the Virtualization Environment on Azure. As a prerequisite, verify that the guest VMs residing on the AVS hosts are able to consume data from Cloud Volumes ONTAP after connectivity has been established.

After Cloud Volumes ONTAP and AVS have been configured properly, begin configuring Jetstream to automate the recovery of on-premises workloads to AVS (VMs with application VMDKs and VMs with inguest storage) by using the VAIO mechanism and by leveraging SnapMirror for application volumes copies to Cloud Volumes ONTAP.

JetStream DR software consists of three major components: the JetStream DR Management Server Virtual Appliance (MSA), the DR Virtual Appliance (DRVA), and host components (I/O filter packages). The MSA is used to install and configure host components on the compute cluster and then to administer JetStream DR software. The installation process is as follows:

- 1. Check the prerequisites.
- 2. Run the Capacity Planning Tool for resource and configuration recommendations.
- 3. Deploy the JetStream DR MSA to each vSphere host in the designated cluster.
- 4. Launch the MSA using its DNS name in a browser.
- 5. Register the vCenter server with the MSA.
- After JetStream DR MSA has been deployed and the vCenter Server has been registered, navigate to the JetStream DR plug-in with the vSphere Web Client. This can be done by navigating to Datacenter > Configure > JetStream DR.



- 7. From the JetStream DR interface, complete the following tasks:
  - a. Configure the cluster with the I/O filter package.

JetStream DR				
Protected Domains Statistics Storage Sites A	oppliances Configurations	Task Log		
Site Details				Alarm Settings
vCenter Server Hostname	172.21.253.160			
Management Appliance Hostname	ANFJSDR-msa			
Software Version	4.0.0.443			
Subscription ID	00000000-0000-0000	-0000-00000000001 Configure		
Tenant ID / Application ID	- Configure			
Application Secret	- Configure			
Configured Clusters	Configure Clusters			
Configure Cluster 1 Upgrade 1 Unconfigure 1 Res		Select All Clear All	Q	٩
Cluster Name 🔺	Cluster Name	Datacenter Name	sion 🔺	Host Details 🔺
No cluster configured	A300-Cluster	A300-DataCenter	*	
		12		
		Cancel	onfigure	

b. Add the Azure Blob storage located at the recovery site.

Storage Sites	Add Storage Site	
+ Add Storage Site Scan Domains	- Otomas Sta Tuna *	
Name 🛦	Azure Blob Storage	
No Storage Site configured.		
-	Access Type *	
	Key Access	
	Storage Site Name (Provide a name to identify this Site) *	
	ANFDemoblobrepo	
Storage Site Details Alarms		
	Agare Blob Storage Account Name *	
	- AMINI AND	_
No storage sile selected, select a storage	Adure Bob Storage Account Key *	
	······································	
	Cancel Add Storage Site	

8. Deploy the required number of DR Virtual Appliances (DRVAs) from the Appliances tab.

(i)

Use the capacity planning tool to estimate the number of DRVAs required.

Protected Domains Statistics Storage Sites	Appliances Configurations	Task Log			E
DRVAs (DR Virtual Appliances)					
+ Deprov New DRVA TUpgrade TUpconfigure					٩
Name 🔺	Status 🔺	Child Alarm	Software Version	Details 🔺	
Replication Log Volume					
Replication Log Volume					٩
Replication Log Volume + New Replication Log Volume Disk Path Name	Status	Child Alarm 🔺	Size (available/total) 🔺	Details 🔺	۹.

1. General 2. DRVA VM	3. DRVA Network 4. Summary	
Name	GCSDRPD001	c
Description (Optional)	Protected Domain for VMs with ANF and JS	Details 🔺
Datacenter	A300-DataCenter	
Cluster	A300-Cluster	
Resource Pool (Optional)	(*)	
VM Folder (Optional)		
Datastore	A300_NFS_vMotion	
Number Of CPUs	8	
Memory Size	32GB	C
Management Network	VM_187	Details 🔺
Host(iofilter) to DRVA Data Network	VM_187	
Replication Network to Object Store	VM_187	
Replication Log Network	VM_187 .	
	1. General 2. DRVA VM Name Description (Optional) Datacenter Cluster Resource Pool (Optional) VM Folder (Optional) Datastore Number Of CPUs Memory Size Management Network Host(iofilier) to DRVA Data Network Replication Network to Object Store Replication Network	1. General     2. DRVA VM     3. DRVA Network     4. Summary       Name     GCSDRPD011       Description (Optional)     Protected Domain for VMs with ANF and JS       Datacenter     A300-DataCenter       Cluster     A300-Cluster       Resource Pool (Optional)     -       VM Folder (Optional)     -       Datascete     A300_NFS_vMotion       Number Of CPUs     8       Memory Size     32GB       Management Network     VM_187       Host(joiltin') to DRVA Data Network     VM_187       Replication Network to Object Store     VM_187       Replication Log Network     VM_187

9. Create replication log volumes for each DRVA using the VMDK from the datastores available or the independent shared iSCSI storage pool.

Jeusoream DH Protected Domains Statistics Storage Sites A	ppliances Configurations	Task Log			
DRVAs (DR Virtual Appliances)					
+ Deploy New DRVA TUrgrade Duconfigure					C
Name 🔺	Status 🔺	Child Alarm 🔺	Software Version	Details A	
GCSDRPD001	O Running	00	4.0.0.134	Detain	
+ New Replication Log Volume					(
Disk Path Name	Status.	Child Alarm	Size (available/total) 🔺	Details A	
idex/sdb	O Ok	<b>0</b> 0	179 88 GB / 200 GB	Detats	
Replication Log Volume Details					

10. From the Protected Domains tab, create the required number of protected domains using information about the Azure Blob Storage site, the DRVA instance, and the replication log. A protected domain defines a specific VM or set of application VMs within the cluster that are protected together and assigned a priority order for failover/failback operations.

JebSbream DR Protected Domains Statistics Storage Sites Select Protected Domain: - <u>Viewell</u>	Apoliances Confinurations Task Lor Create Protected Domain	0			+ Create = More
	1. General 2. Pr	rimary Site	J. Summary		
	Protected Domain Name	GCSDRPD_Demo01			
	Priority Level (Optional)				
	Description	Protection domain ANF			
	Total estimated data size to be protected	1000GB			
	DR Virtual Appliance	GCSDRPD001			
	Compression	Yes			
	Compression Level	Default		121	
	Normal GC Storage Overhead	50%			
	Maximum GC Storage Overhead	300%			
	Replication Log Storage	/dev/sdb			
	Paplication Lon Size	50GR			
		Cancel I	Back Create		

Select Protected Domain: * View all	Create Protected Domain				+ Create =
	1. General	2. Primary Site	3. Summary		
	Compression	Yes		*	
	Compression Level	Default			
	Normal GC Storage Overhead	50%			
	Maximum GC Storage Overhead	300%			
	Replication Log Storage	/dev/sdb			
	Replication Log Size	50GB			
	Metadata Size	31 56GB			
	Primary Site Datacenter	A300-DataCenter			
	Primary Site Cluster	A300-Cluster			
	Storage Site	ANFDRDemoFailoverSite			
	Enable PITR	No			
				*	
		Cancel Ba	ack Create		

11. Select the VMs to be protected and group the VMs into applications groups based on dependency. Application definitions allow you to group sets of VMs into logical groups that contain their boot orders, boot delays, and optional application validations that can be executed upon recovery.



Make sure that the same protection mode is used for all VMs in a protected domain.

eUSUream DR rotected Domains Statistics Storage Sites	Appliant	es Configurations Task Le	6					
elect Protected Domain: GCSDRPD_Demo01 -	Vic Star	t Protection				+ Create	Denta	I block
ecoverable / Total VMs								Est Deta
epication Status	Pitt	tection Mode for pelected VMp 🔹			٩		rSilte	
smaning Background Data	Ö	VM Name	# of Disks	Protection Mod		LOCAL ( 177.21.253 160		
	0	ElasticWebA2	1			A300-DataCenter \ A300-	Cutstar	
treat RPO	- Ö	ElasticWebA3	1		v	THRADIED		
NAME AND ADDRESS OF A DESCRIPTION OF A D	0	ElasticWebB0	.1		U			
relied vita Settinga oxiatma	Ö	ElasticWebB1	1.1		÷.,			
and the second se	Ó	EtasticWebB2	0.8	Vome Through				1
* Starl Protection	0	ElasticWebB3	. 1	Witte Through	w.,			
VM Name 🔺	2	GCS-DR-DC	1	Write-Through	* N	Background Data 🔺	Details	
No VM is protected	2	GCS-DR-LinVM01	1	Write-Through	× 48			
	2	OCS-DR-SCA	1	Write-Through	*			
		GCS-DR-SQL01	4	Write-Through	~			
	2	GCS-DR-WeVM01	.1	Write-Through	*			
		jss-drva-GCSDRPD001	2		¥.			
	0	PrimeClient	2		×			
	0	Standby0	1	Webs-Through	¥			
	0	Standby1	.1		Q.,			
	0	Standby2	12		1. V			
	0	Standby3			1 N S			
	0	VMmark-Template01	1		v			

12. Make sure that replication log volumes are placed on high- performance storage.

elect Protected Domain: GCSDRPD_Demo01 •	Start Protection				+ Create Doiete Elime
ecoverable / Total Wits				1	Lot Date
epication Status	Write-Back(VMDK) •			Q	ANFDRDemoFalloverSite
emaining Background Data	VM Name 🔺	# of Disks	Protection Mode	2.481	LOCAL (172 21 253 180 )
	ElasticWebA2	1	Write Through	- · ·	Tournad
	ElasticWebA3	1	White-Through	×.	
And a second second second	ElasticWebB0	1		¥	
PRARCIPO VARS Semogra Adarma	ElasticWebS1	1		V	
to a second second	ElasticWebB2	1	Write-Through	V	
	ElasticWeb63	1	Write Through	10 a	
UM Marne 🔺	GCS-DR-DC	1	Write-Back(VMDK)	~	Bockground Data A Details
No VM is protected.	GCS-DR-LivVM01	1	Write-Back(VMDK)	~	
	🖸 GCS-DR-SCA	1	Write-Back(VMDK)	*	
	GCS-DR-SQL01	1	Write-Back(VMDK)	~	
	GCS-DR-WeVM01	31	Write-Back(VMDK)	×	
	int-diva-GCSDRPD001	2		N.	
	PrimeClient	2		× .	
	C Standby0	1	Wide Trends	¥.,	
	Standby1	1		10 I.	
	Standby2	1	White Through:	93.	
	Standby3	1	Write Through:	85 L	
	VMmark-Template01	1		10.00	
			Cancel Star	Protection	
			i concio i sec	N	

13. After you are done, click Start Protection for the protected domain. This starts data replication for the selected VMs to the designated Blob store.

Protected Domains Statistics Storage	Sites Appliances Configuratio	ons Task Log			Running Tasks	
elect Protected Domain: GCSDRPD_Demo0	1 • <u>View all</u>			+ (	Start Protection (GCS-DR-SCA)	50%
ecoverable / Total VMs		0/5	Configurations		Start Protection (GCS-DR-Win	50%
eplication Status		ок	Storage Site	ANFDRDe	Start Protection (GCS-DR-Lin	50%
			Owner Site	LOCAL ( 172.2	Start Protection (GCS DR DC)	ENG
emaining Background Data		0 B	Datacenter \ Cluster	A300-DataCen	Sian Protection (GCS-DK-DC)	50%
urrent RPO			Point-in-time Recovery	Disabled	Start Protection (GCS-DR-SQ.	50%
Protected VMs Settings Alarms					Configure VMDK Re Complete	ed 🔽
Protected VMs Settings Alarms     + Start Protection					Configure VMDK Re Complete	ed 🔽
Settings Alarms     Start Protection     M Name	Protection Status	Replication St	stus 🔺 Protection Mode 🔺	Background Da	Configure VMDK Re Complete Close	<sup>10</sup>
Protected VMs         Settings         Alarms           + Start Protection         If Stop Protection           VM Name ▲         GCS-DR-DC	Protection Status ▲ <ul> <li>Initializing</li> </ul>	Replication St	atus ▲ Protection Mode ▲ Write-Back(VMDK)	Background Da	Close Close Ita A Details Details	ed 🔽
Protected VMs Settings Alarms  + Start Protection VM Name  GCS-DR-QC GCS-DR-LinVW01	Protection Status ▲ Initializing Initializing	Replication St	atus ▲ Protection Mode ▲ Write-Back(VMDK) Write-Back(VMDK)	Background Da	Configure VMDK Re Complete Close ta A Details Retails Retails	ed 🔽
Protected VMs Settings Alarms  + Start Protection  VM Name ▲  GCS-DR-UnVM01  GCS-DR-SCA	Protection Status    Protection Status  Initializing  Initializing  Initializing	Replication St - -	atus  Protection Mode  Write-Back(VMDK) Write-Back(VMDK) Write-Back(VMDK)	Background Da - -	Configure VMDK Re Complete Close ta Details Details Details Details Details	ed 🔽
Protected VMs Settings Alarms   Start Protection VM Name GCS-DR-DC GCS-DR-SCA GCS-DR-SOL01	Protection Status ▲  initializing  initializing  initializing  initializing  initializing	Replication St - -	atus A Protection Mode A Write-Back(VMDK) Write-Back(VMDK) Write-Back(VMDK) Write-Back(VMDK)	Background Da - - -	Configure VMDK Re Complete Close ta Details Details Details Details Details	

14. After replication is completed, the VM protection status is marked as Recoverable.

JEGSGREATE DH Protected Domains Statistics Stor	age Sites Appliances Configura	tions Task Log					E
Select Protected Domain: GCSDRPD_Der	mo01 👻 <u>View all</u>				+ Create	Delete	≡ More
Recoverable / Total VMs		5/5	Configuration	ns			Edit Details
Replication Status		ок	Storage Site		ANFDRDemoFallo	werSite	
			Owner Site		LOCAL ( 172.21.253.16	50)	
Remaining Background Data		0 B	Datacenter \ Cl	luster	A300-DataCenter \ A30	0-Cluster	
Current RPO		05	Point-in-time R	lecovery	Disabled		
Protected VMs Settings Alarms							
+ Start Protection	Protection Status	Replication St	atus 🔺	Protection Mode 🔺	Background Data 🔺	Details	م
Start Protection     VM Name ▲     GCS-DR-DC	Protection Status A	Replication St	atus 🔺	Protection Mode  Write-Back(VMDK)	Background Data 🔺	Details Details	٩
Start Protection     Stop Protection     VM Name     GCS-DR-DC     GCS-DR-LinVM01	Protection Status A © Recoverable © Recoverable	Replication St OK OK	atus 🛦	Protection Mode  Write-Back(VMDK) Write-Back(VMDK)	Background Data 🔺 0 B 0 B	Details Details Details	٩
Start Protection     VM Name      GCS-DR-DC     GCS-DR-LinVM01     GCS-DR-SCA	Protection Status ▲ © Recoverable © Recoverable @ Recoverable @ Recoverable	Replication St OK OK OK	atus 🔺	Protection Mode Write-Back(VMDK) Write-Back(VMDK) Write-Back(VMDK)	Background Data A 0 B 0 B 0 B	Details Details Details Details	٩
Start Protection     Stop Protection     OCS-DR-DC     OCS-DR-LinVM01     OCS-DR-SCA     OCS-DR-SCA     OCS-DR-SCA01	Protection Status  Protection Status Recoverable Recoverable Recoverable Recoverable	Replication St OK OK OK OK	atus 🔺	Protection Mode  Write-Back(VMDK) Write-Back(VMDK) Write-Back(VMDK) Write-Back(VMDK)	Background Date ▲ 08 08 08 08	Details Details Details Details Details	٩



Failover runbooks can be configured to group the VMs (called a recovery group), set the boot order sequence, and modify the CPU/memory settings along with the IP configurations.

15. Click Settings and then click the runbook Configure link to configure the runbook group.

elect Protected Domain: GCSDRPD_Demo01  View all			+ Cre	ate 🛢 Delete	≡ More
Recoverable / Total VMs	5/5	Configurations			Edit Deta
Replication Status	OK	Storage Site	ANFDRDemi	oFailoverSite	
	20	Owner Site	LOCAL ( 172.21.2	253.160 )	
Remaining Background Data	08	Datacenter \ Cluster	A300-DataCenter	\A300-Cluster	
Current RPO	0s	Point-in-time Recovery	Disabled		
Protected VMs Settings Alarms					
Failover Runbook Not Configured Configure					
Test Failover Runbook Not Configured Configure					
Faliback Runbook 😺 Not Configured Configure					
Memory Setting Not Configured Configure					
GC Settings Configured Configure					
Concurrency Settings Not Configured Configure					

16. Click the Create Group button to begin creating a new runbook group.

If needed, in the lower portion of the screen, apply custom pre-scripts and post-scripts to automatically run prior to and following operation of the runbook group. Make sure that the Runbook scripts are residing on the management server.

JebSbream DR Protected Domains Statistics Storage Sites A	Failover Runbook Settings			Ē
Select Protected Domain: GCSDRPD_Demo01 * Vi	+ Create Group	Delete Group		+ Create Delete E More
Recoverable / Total VMs	<ul> <li>Shoup Name</li> </ul>	# of VM Power Off	Retain MAC	Edit Detain
Replication Status	Independent VMs	5 -	·*·	ANFDRDemoFalloverSite
Remaining Background Data				LOCAL (172:21:253:160) A300-DataCenter \ A300-Cluster
Current RPO				Disabled
Protected VMs Settings Alarms	0 0			
Failover Runbook Configured Details				
Test Failover Runbook Configured Details				
Fallback Runbook Configured Details				
Memory Setting Not Configured Configured				
GC Settings Configured Configure				
Concurrency Settings Not Configured Configure				

17. Edit the VM settings as required. Specify the parameters for recovering the VMs, including the boot sequence, the boot delay (specified in seconds), the number of CPUs, and the amount of memory to allocate. Change the boot sequence of the VMs by clicking the up or down arrows. Options are also provided to Retain MAC.

JebSbream DR Protected Domains Statistics Storage	Create Runbook Group										
Select Protected Domain: GCSDRPD_Demo	1. General	2. Select VM	s	3. Edit V	M Settings		4. Summa	ry	- Cteate	E Celitte	≡ More
Recommine / Total VMs	Retain MAC		Power Of	f VMs			O Reset	1			
Replication Status	VM Name	Boot Sequence	Boot Delay	CPU	Memory	Script	NIC	UFDR (197			
Renhining Background Data	GCS-DR-WinVM01	17 4	0s	32	64 GB	Config	View	-			
Communitation	GCS-DR-SCA	21 4	0s	4	16 GB	Config	View				
Company and a second	GCS-DR-DC	3 11m+	0s	4	16 GB	Config	View				
Protected VUA Settings Aurors	GCS-DR-LinVM01	41 4	0s	2	4 GB	Config	View				
	GCS-DR-SQL01	5 † ↓	0s	4	8 G8	Config	View				
Allemony Setting Not Configured Contrains Git Settings Configured Configured Concurrency Settings Not Configured Configured				Ca	ncel	Back	1	iext			

18. Static IP addresses can be manually configured for the individual VMs of the group. Click the NIC View link of a VM to manually configure its IP address settings.

 $(\mathbf{i})$ 

	nain: GCSDRPD_Dem	1. General	2. Select VM:		3. Edit VI	M Settings		4. Summary		+ Creala	Delete	華み
		Retain MAC			VMs			0	-			
		TO BE THE REAL						O Reset		EDRD-IntoFalloy	ar Siller	
		VM Name	Boot Sequence	Boot Delay	CPU	Memory	Script	NIC				
Semaining Backgrout		GCS-DR-WinVM01	1↑↓	0s	32	64 GB	Config	Yigw 1		staCentarii A300	Cluster	
		GCS-DR-SCA	2↑↓	Qs	4	16 GB	Config	View				
		GCS-DR-LinVM01	31 1	0s	2	4 GB	Config	View				
		GCS-DR-SQLOT	41 1	05	4	8 GB	Canto	View				
		GCS-DR-DC	5† 4	0s.	4	16 GB	Config	View .	2			
Concurrency Settings												

19. Click the Configure button to save NIC settings for the respective VMs.

k adapter 1					
	VM_3510	4000	192	Configure	ary
Configure	e Static IP Address				-
IP Address *					
172.21.25	4.185			- 11	1
Subnet Mask					
255.255.2	55.0			- 11	
Gateway *	4.1				
				_	
DNS - 172.30.15	3.20				
DNS Suffi	6			_	
	Reset	Can	cel Cont	igure	
				Close	Ne
_		12			
Demo Create Runbo	pok Group	•	•	+ous )	
Group Name	rei 2. Soleci vies ie VMs	GCSRecovery 5	artungs 4. summ	- FDRDensFalkerD	
Retain VM's Power Off V	s MAC addresses /Ms	truë Taise		And an and a state of the state	
Group n	Success	be applied to both Fallover.	and Faikback runbook.		
	Configure P Address - 172,21,25 Subret Mask 255,255,2 DNS - 172,21,25 DNS - 172,30,15 DNS Suffe DNS Suffe Create Rube Group Nam Number of Retain With Power Office Group Nam Suffer Group Nam Number of Retain With Power Office Group Nam Suffer Group Nam Number of Retain With Power Office Group Nam Suffer Group Nam Suffer Group Nam Number of Retain With Power Office Group Nam Suffer Group Nam Suffer	Configure Static IP Addresss PAddress PAddress T72 21 254 185 Submet Mask * 255 255 255 0 Gataway * 172 21 254 1 DNS * 172 30 153 20 DNS Suffix DNS Suffix Reset	Configure Static IP Address PAddress 172.21.254.185 Subnet Mask = 255.255.255.0 Gataway = 172.21.254.1 DNS = 172.30.153.20 DNS Suffix Reset Can Contemporate Cont	PAddress         172 21 254 185         Submet Mask -         255 255 255 0         Gateway *         172 21 254 1         DNS -         172 30.153 20         DNS Suffix         Concel	Contigure Static IP Address   I72.21.254.185   Submet Mark •   255.255.255.0   Garways •   172.21.254.1   DNS •   172.30.153.20   DNS Suffix   Reset   Cancel   Configure   Configure   Configure

The status of both the failover and failback runbooks is now listed as Configured. Failover and failback runbook groups are created in pairs using the same initial group of VMs and settings. If necessary, the settings of any runbook group can be individually customized by clicking its respective Details link and making changes.

A best practice for a recovery site (AVS) is to create a three-node pilot-light cluster in advance. This allows the recovery site infrastructure to be preconfigured, including the following:

- Destination networking segments, firewalls, services like DHCP and DNS, and so on
- Installation of JetStream DR for AVS
- Configuration of ANF volumes as datastores and more

JetStream DR supports a near-zero RTO mode for mission-critical domains. For these domains, destination storage should be preinstalled. ANF is a recommended storage type in this case.



Network configuration including segment creation should be configured on the AVS cluster to match on-premises requirements.



Depending on the SLA and RTO requirements, you can use continuous failover or regular (standard) failover mode. For near-zero RTO, you should start continuous rehydration at the recovery site.

 To install JetStream DR for AVS on an Azure VMware Solution private cloud, use the Run command. From the Azure portal, go to Azure VMware solution, select the private cloud, and select Run command > Packages > JSDR.Configuration.



The default CloudAdmin user of the Azure VMware Solution doesn't have sufficient privileges to install JetStream DR for AVS. The Azure VMware Solution enables simplified and automated installation of JetStream DR by invoking the Azure VMware Solution Run command for JetStream DR.

The following screenshot shows installation using a DHCP-based IP address.

Microsoft Azure	P. Search resources, se	more, and docs (G+/)	🗉 😝 🖉 🗇 🗇 🕫 niyaz@netapp.com
Home > ANFDataClus ANFDataClus   Run Ar5 Private clood	command –		Run command - Install-JetDRWithDHCP × This top level Cmdler Downloads JetDr burdle from MAS, creates a new user, assopris
P Seatch (Cht+7)	🕐 Refresh 🖉 Feedback		elevated privileges to the user, deploys JetDr Management Server Apphance(MSA), registers vCenter to the JetDr MSA, configures duster.
Tigi	Packages Run execution status		Command parameters RegisterWithip: ①
<ul> <li>Diagnose and solve problems</li> <li>Settings</li> </ul>	Name     SOR Configuration 224     reserver items	Description	ProtectedCluster * ()
A Locks	Disable-JetDRForCluster	This Cindlet unconfigures a cluster but doesn't uninstall JetDR completely so other clusters policies.	Chuster-1 Detastore* ③
Manage Connectivity	Enable-SetORForChaner Install-JetDRWithDHCP	This Cindlet configures an additional cluster for protection. It installs vibs to all hosts in the This top level Cindlet Downloads JetDr bundle from MMS, crisates a new user, assigns elevi	isanDatatore
Custers dentity	Instan-JetDRWitt StaticIP	registers vicenter to the Jerla MSA, consigned cluster. This top level Crediter Downloads JerlDr bundle from MMS; creates a new user, assigns elev- registers vicenter to the Jerlar MSA; consigned cluster.	Outer 1
Storage (preview)	Invoke-PreflightJerDRoutall	This Conditit checks and displays current state of the system it checks whether the minimal 4 hosts, if the cluster details are correct, if there is already a VM with the same name provid	Crideital 🖸
+ Add-ons	Invoke-Prefight)=tDRUwnstall	This Cindlet checks and displays current state of the system it checks whether the minimal 4 hosts. If the cluster details are correct and if any VCenter is registered to the MGA	root Password *
Workload Networking     Segments	MicrosoftAVSManagement +147     www	The top levels (indeet creates a new user, assigns elevated privileges to the user, unconfigu- c ordets to echobate see bass in nanegra Asee Vintime Soutone	HostName 💿
DHCP     Port minoring			anfipul-ma Network*
O DNS			Drian
Run command			Ritain up to

2. After JetStream DR for AVS installation is complete, refresh the browser. To access the JetStream DR UI, go to SDDC Datacenter > Configure > JetStream DR.

MAR THE STORE								
Site Details							Alarm Set	tti
vCenter Server Hostna	me	172.30.15	56.2					
Management Appliance	Hostname	anfjsval-n	nsa					
Software Version		4.0.2.450						
Subscription ID		- Config	ure					
Tenant ID / Application	ID	- Config	ure					
Application Secret		- Config	ure					
Configure Cluster	1 Upgrade	Duconfigure	* Resolv	e Configure l	ssue			
Cluster Name		Datacenter	Name 🔺	Status	Softwar	e Version 🔺	Host Details	5
Cluster-1		SDDC-Data	center	🔮 Ok	4.0.2.13	2	Details	

- 3. From the JetStream DR interface, complete the following tasks:
  - a. Add the Azure Blob Storage account that was used to protect the on-premises cluster as a storage site and then run the Scan Domains option.
  - b. In the pop-up dialog window that appears, select the protected domain to import and then click its Import link.

JetStream DR Protected Domains Statistics Storage Sites	Available Protected Domain(s) For Import	<u>in.</u>
Storage Sites       + Add Storage Site     Scan Domains <ul> <li>Remove</li> <li>Name A</li> <li>NFDemotifobrepore:</li> </ul>	Protected Domain Description Recoverable V VMs Import GCSDRPD_Demo01 Protection domain ANF 5 5 Import	re Blob Storage ^

4. The domain is imported for recovery. Go to the Protected Domains tab and verify that the intended domain has been selected or choose the desired one from the Select Protected Domain menu. A list of the recoverable VMs in the protected domain is displayed.

elect Protected Domain: GCSDRPD_Demo01 👻 View all			+ Create	Delete	■More
lode	Imported	Configurations			Det
ecoverable / Total VMs	5/5	Storage Site	ANFDemoblob	reporec	
		Owner Site			
VM Name	Protection Status	*	Protection Mode	Details	(
VM Name A GCS-DR-DC	Protection Status	•	Protection Mode  Write-Back(VMDK)	Details Details	(
VM Name ▲ GCS-DR-DC GCS-DR-LiniVM01	Protection Status Recoverable Recoverable	•	Protection Mode  Write-Back(VMDK) Write-Back(VMDK)	Details Details Details	
VM Name ▲ GCS-DR-DC GCS-DR-Lin/M01 GCS-DR-SCA	Protection Status © Recoverable © Recoverable © Recoverable	•	Protection Mode  Write-Back(VMDK) Write-Back(VMDK) Write-Back(VMDK)	Details Details Details Details	
VM Name ▲ GCS-DR-DC GCS-DR-LinVM01 GCS-DR-SCA GCS-DR-SOL01	Protection Status © Recoverable © Recoverable © Recoverable © Recoverable		Protection Mode  Write-Back(VMDK) Write-Back(VMDK) Write-Back(VMDK) Write-Back(VMDK)	Details Details Details Details Details	

5. After the protected domains are imported, deploy DRVA appliances.

These steps can also be automated using CPT- created plans.

- 6. Create replication log volumes using available vSAN or ANF datastores.
- 7. Import the protected domains and configure the recovery VA to use an ANF datastore for VM placements.

ect Protected Domain:	Continuous Failover Protected Don	ain		_			Delete	=	More
ie -		Ob. 1/10. Comingo			•	•			Dela
overable / Total VMs	i. General 2a. Pallover settings	20. vm settings	a. neco	very vA	4. DR Settings	5. Summary	reporec		1
	Protected Domain Name		ANFPD002				253.160 )		
	Datacenter		SDDC-Datac	enter					
	Cluster		Cluster-1						
	Resource Pool (Optional)		•				and the second second		
otected VMs Seth	VM Folder (Optional)								
	Datastore		ANFRecoDS	J002					
	Internal Network		DRSeg						C
VM Name 🔺	External Replication Network		DRSeg					Details	
AuctionAppA2	Management Network		DRSeg					Details	1
AuctionAppB2	Storage Site		ANFDemobio	breporec				Details	
AuctionDB2	DR Virtual Appliance		ANFRecDRV	A003				<u>Details</u>	
AuctionLB2	Doplication Los Storano		Mouladh				~	Details	
AuctionLB2 AuctionMSQ2	Banliastian I an Starson		Marian	Cancel	Back	Continuous Failow	~ ~	Details Details	
AuctionNoSOL2				0.00			-	Details	

1

Make sure that DHCP is enabled on the selected segment and that enough IPs are available. Dynamic IPs are temporarily used while domains are recovering. Each recovering VM (including continuous rehydration) requires an individual dynamic IP. After recovery is complete, the IP is released and can be reused.

8. Select the appropriate failover option (continuous failover or failover). In this example, continuous rehydration (continuous failover) is selected.



Although Continuous Failover and Failover modes differ on when configuration is performed, both failover modes are configured using the same steps. Failover steps are configured and performed together in response to a disaster event. Continuous failover can be configured at any time and then allowed to run in the background during normal system operation. After a disaster event has occurred, continuous failover is completed to immediately transfer ownership of the protected VMs to the recovery site (near-zero RTO).

JebSbream DR Protected Domains Statistics Storage Sites Appliances C	onfigurations Task	Log			
elect Protected Domain: GCSDRPD_Demo01 👻 View all			+ Create	Delete	■More
Node	Imported	Configurations		O Restore	
Recoverable / Total VMs	5/5	Storage Site Owner Site	ANFDemoblobrepor REMOTE (172.21.253.1)	→ Failover  → Continuous Fai  → Test Failover	loves
Protected VMs Settings Alarms O O					
VM Name	Protection Status	•	Protection Mode A	Details	
GCS-DR-DC	@ Recoverable		Write-Back(VMDK)	Details	1
GCS-DR-LinVM01	O Recoverable		Write-Back(VMDK)	Details	
GCS-DR-SCA	C Recoverable		Write-Back(VMDK)	Details	
GCS-DR-SQL01	O Recoverable		Write-Back(VMDK)	Details	
GCS-DR-WinVM01	Recoverable		Write-Back(VMDK)	Details	

The continuous failover process begins, and its progress can be monitored from the UI. Clicking the blue icon in the Current Step section exposes a pop-up window showing details of the current step of the failover process.

#### **Failover and Failback**

1. After a disaster occurs in the protected cluster of the on-premises environment (partial or complete failure), you can trigger the failover for VMs using Jetstream after breaking the SnapMirror relationship for the respective application volumes.



 $(\mathbf{i})$ 

This step can easily be automated to facilitate the recovery process.

2. Access the Jetstream UI on AVS SDDC (destination side) and trigger the failover option to complete failover. The task bar shows progress for failover activities.

In the dialog window that appears when completing failover, the failover task can be specified as planned or assumed to be forced.

	<ul> <li>View all</li> </ul>			[	+ Create	O Fallover	■Mor
de	Continuous Rehydration in Progress	Configurations					Details
coverable / Total VMs	4/4	Storage Site			FDemoblobrep	orec	^
		Owner Site	REMOTE ( 172.21.253.160 )			160)	
a (Processed/Known Remaining)	329.01 GB / 6.19 GB	Datacenter \ Cluster		SDDC	-Datacenter \ Cli	uster-1	
rrent Step	Recover VMs' data from Storage Site	Point-in-time Recovery		Disabl	led		~
Protected VMs Settings Alarms	0 0						
							Q
VM Name	Protection Status		Protect	ion Mode 🔺		Details	^
GCS-DR-DC	Recoverable     Recoverable		Write-Back(VMDK) Details				
GCS-DR-SCA	Recoverable		Write-Back(VMDK) Details			Details	
GCS-DR-SOL01	Recoverable		Write-B:	ack(VMDK)		Details	
GCS-DR-WinVM01	@ Recoverable		Write-Ba	ack(VMDK)		Details	
VM 3510	DDStratebSag	2	~				
0 0			v .				
O O							
O O Other Settings Planned Failover Force Failover Some VMFs guest credential are	required because of network configuration:	Configure	~ .	•			

Forced failover assumes the primary site is no longer accessible and ownership of the protected domain should be directly assumed by the recovery site.

Ð	Force Failover of required!	of Protected Domain	requested. Adminis	trator consent is
	Complete owne Site.	ership of this Protect	ted Domain will <mark>b</mark> e ta	aken over <mark>b</mark> y this
	Are you sure yo	u want to continue?		
			Cancel	Confirm

rotected VM Network ▲ Recovery VM Network M_3510 DRStretchSeg ▼	^ ^
M_3510 DRStretchSeg -	A
0.0	~ ~
her Settings	_
Planned Failover	
Force Failover	
Some VM's guest credential are required because of network configuration: Configure	

3. After continuous failover is complete, a message appears confirming completion of the task. When the task is complete, access the recovered VMs to configure ISCSI or NFS sessions.

 $(\mathbf{i})$ 

**i** 

The failover mode changes to Running in Failover and the VM status is Recoverable. All the VMs of the protected domain are now running at the recovery site in the state specified by the failover runbook settings.

To verify the failover configuration and infrastructure, JetStream DR can be operated in test mode (Test Failover option) to observe the recovery of virtual machines and their data from the object store into a test recovery environment. When a failover procedure is executed in test mode, its operation resembles an actual failover process.

Protected Domains Statistics Stora	Continuous Rehydration Task Resu	iit.		
Select Protected Domain: GCSDRPD002	Task Completed Successfully with	warnings	+ Create Delate	= More
wode -				Ent Data
	Protected Domain	GCSDRPD002	ANECVODE	
oscoverative / voran vives	VMs Recovery Status	Success with warnings	DCAL / 172 30 156 2.)	
	Total VMs Recovered	4	DDC-Datacenter ( Custer, 1	
Saturd a law on the balance of Product	VM(s) with warning	2 <u>View</u>	esabled	
cemaining background bata	GCSRecovery03 Status:		100000	
	Pre-script Execution Status	Not defined		
	Runbook Execution Status	O Success		
Protected VMs Settings Atarms	Post-script Execution Status	Not defined		
+ Start Protection				c
VM Name A		N	ackground Data 🔺 Details	
GCS-DR-SC46		14	8 Details	3
GC8-DR-SQL03			B Dataits	
GCSDR-W2K16-01			8 Detaits	
UbuntuSw001			B Details	

4. After the virtual machines are recovered, use storage disaster recovery for in-guest storage. To

demonstrate this process, SQL server is used in this example.
5. Log into the recovered SnapCenter VM on AVS SDDC and enable DR mode.
a. Access the SnapCenter UI using the browserN.
Image: the star and the star and

- b. In the Settings page, navigate to Settings > Global Settings > Disaster Recovery.
- c. Select Enable Disaster Recovery.
- d. Click Apply.



e. Verify whether the DR job is enabled by clicking Monitor > Jobs.



NetApp SnapCenter 4.6 or later should be used for storage disaster recovery. For previous versions, application-consistent snapshots (replicated using SnapMirror) should be used and manual recovery should be executed in case previous backups must be recovered in the disaster recovery site.

6. Make sure that the SnapMirror relationship is broken.

Canvas	Replication	Backup & Restore	Data Sense	File Cache	Compute	Sync All Serv	rices (+9) ~			
Repline	cation									
		Volume Re	lationships	9 4.78 Replicated	GIB d Capacity	Currently To	ransferring	S Healthy	S 0 Failed	
	3	Volume Relationships								90
		Realth Status 🗧	Source Volume	s Target Vo	dume =	၀— စ Total Transfer Time	÷ Status		• Last Successful Transfer	0
		0	gcsdrsqldb_sc46 ntaphci-a300e9u25	gcsdrsqld ANFCVOC	lb_sc46_copy )RDemo >	6 minutes 41 seconds	idle	broken-off	May 5, 2022, 12:08:34 PN 33.66 KiB	
		$\odot$	gcsdrsqlhid_sc46 ntaphci-a300e9u25	gcsdrsqlh ANFCVOD	Ild_sc46_copy RDemo	4 minutes 56 seconds	idle	broken-off	May 5, 2022, 12:09:15 PN 69.84 KiB	
		$\odot$	gcsdrsqllog_sc46 ntaphci-a300e9u25	gcsdrsq8 ANFCVOE	og_sc46_copy )RDemo	10 minutes 18 seconds	idle	broken-aff	May 5. 2022, 12:08:34 PA 104.34 KiB	

7. Attach the LUN from Cloud Volumes ONTAP to the recovered SQL guest VM with same drive letters.

📅 Disk Manageme	nt						- <del></del>	×
File Action View	w Help							
🗢 🌩   🗰   📔	🖬   🗩 📝 (	<b>5</b> =1						
Volume	Layout	Туре	File System	Status	Capacity	Free Spa	% Free	
-	Simple	Basic		Healthy (R	450 MB	450 MB	100 %	
-	Simple	Basic		Healthy (E	99 MB	99 MB	100 %	
- (C:)	Simple	Basic	NTFS	Healthy (B	89.45 GB	67.03 GB	75 %	
BACKUP (G:)	Simple	Basic	NTFS	Healthy (P	9.97 GB	9.92 GB	99 %	
- DATA (E:)	Simple	Basic	NTFS	Healthy (P	24.88 GB	24.57 GB	99 %	
- LOG (F:)	Simple	Basic	NTFS	Healthy (P	9.97 GB	8.93 GB	90 %	
					0	0		

8. Open iSCSI Initiator, clear the previous disconnected session and add the new target along with multipath for the replicated Cloud Volumes ONTAP volumes.

l'argets	Discovery	Favorite Targets	Volumes and Devices	RADIUS	Configuration
Quick C To disc DNS na Target	Connect cover and log ame of the ta	on to a target usin arget and then dick	g a basic connection, to Quick Connect.	ype the IP	address or
Discove	ered targets				Refresh
Name				Status	
	92-08.com.	netapp:sn.547772c	cc47811ecbb62000	Connecte	d
ign. 19		and a shift to be a strength of the strength o		Deserves	and the

9. Make sure that all the disks are connected using the same drive letters that were used prior to DR.



Services (Local)					
SQL Server (MSSQLSERVER)	Name SQL Full-text Filter Daemon	Description Service to la	Status Running	Startup Type Manual	Log ' NT
Description: Provides storage, processing and controlled access of data, and rapid transaction processing.	SQL Server (MSSQLSERVER) SQL Server Agent (MSS) SQL Server Browser SQL Server CEIP service SQL Server Integration S SQL Server Integration S SQL Server VSS Writer SSDP Discovery State Repository Service Still Image Acquisition E Storage Service Storage Tiers Managem Superfetch Sync Host_df83a System Event Notification S	Provides etco Start Stop Pause <sup>O</sup> O Resume Refeant All Tasks Refresh Properties Help This service Monitors sy	Running hing hing hing hing hing hing hing	Automatic Automatic Automatic Automatic Automatic Automatic Manual Manual Manual Manual Manual Automatic (D Automatic	GCS Loc NT NT Loc Loc Loc Loc Loc Loc

11. Make sure that the SQL resources are back online.



In the case of NFS, attach the volumes using the mount command and update the /etc/fstab entries.

At this point, operations can be run and business continues normally.

i.



On the NSX-T end, a separate dedicated tier-1 gateway can be created for simulating failover scenarios. This ensures that all workloads can communicate with each other but that no traffic can route in or out of the environment, so that any triage, containment, or hardening tasks can be performed without risk of cross-contamination. This operation is outside of the scope of this document, but it can easily be achieved for simulating isolation.

After the primary site is up and running again, you can perform failback. VM protection is resumed by Jetstream and the SnapMirror relationship must be reversed.

- 1. Restore the on-premises environment. Depending on the type of disaster incident, it might be necessary to restore and/or verify the configuration of the protected cluster. If necessary, JetStream DR software might need to be reinstalled.
- 2. Access the restored on-premises environment, go to the Jetstream DR UI, and select the appropriate protected domain. After the protected site is ready for failback, select the Failback option in the UI.



The CPT-generated failback plan can also be used to initiate the return of the VMs and their data from the object store back to the original VMware environment.

JebStream DR Protected Domains Statistics Storage Sites Appliances Configu	rations Task Log		
Select Protected Domain: GCSDRPD_Demo01 🔻 View all		+ Create	Delete E More
Mode Rui	ning in Failover Configur	ations	O Restore
Active Site	172.30.156.2 Storage S	ite 🔨 ANFCVODR	O Resume Continuous Rehydratio
Recoverable / Total VMs	4/4 Owner Sit	e REMOTE ( 172.3	← Failback
Protected VMs Settings Alarms O O			c
VM Name 🔺	Protection Status	Protection Mode	Details
GCS-DR-DC	Recoverable	Write-Back(VMDK)	Details
GCS-DR-LinVM01	Recoverable	Write-Back(VMDK)	Details
GCS-DR-SCA	Recoverable	Write-Back(VMDK)	Details
GCS-DR-SQL01	Ø Recoverable	Write-Back(VMDK)	Details
GCS-DR-WinVM01	Recoverable	Write-Back(VMDK)	Details



Specify the maximum delay after pausing the VMs in the recovery site and restarting them in the protected site. The time need to complete this process includes the completion of replication after stopping failover VMs, the time needed to clean the recovery site, and the time needed to recreate VMs in the protected site. NetApp recommends 10 minutes.

•	•	•	•			0
. General	2a. Failback Settings	2b. VM Settings	3. Recovery VA	4. DR	Settings	5. Summary
Failback Da	tacenter		A300-DataCenter			
Failback Clu	uster		A300-Cluster			
Failback Re	source Pool					
VM Folder (	Optional)					
Failback Da	tastore		A300_NFS_vMotion			
Maximum D	elay After Stopping		10 Minutes			
Internal Net	wo9k		VM_187			
External Re	plication Network		VM_187			
Managemen	it Network		VM_187			
Storage Site	9		ANFCVODR			
DR Virtual A	ppliance		GCSDRVA002			
Replication	Log Storage		/dev/sdb			

3. Complete the failback process and then confirm the resumption of VM protection and data consistency.

JetStream DR Protected Domains Statistics Storage Si	Failback Task Result	
Select Protected Domain: GCSDRPD002 *	Task Completed Successfully	
Recoverable / Total VMs	Protected Domain	GCSDRPD002
Replication Status	VMs Recovery Status	O Success
Remaining Background Data	Total VMs Recovered	4
remaining beinground bein	GCSRecovery03 Status:	
Current RPO	Pre-script Execution Status	Not defined
	Runbook Execution Status	O Success
Protected VMs Settings Alarms	Post-script Execution Status	Not defined

4. After the VMs are recovered, disconnect the secondary storage from the host and connect to the primary storage.

-	gredrenidh er46		gcsdrsqldb_sc46_copy				Mar. 5, 2022, 12:00:24 PM
$\odot$	ntaphci-a300e9u25	14	ANFCVODRDemo	6 minutes 41 seconds	idle	broken-off	33.66 KiB
0	gcsdrsqlhld_sc46		gcsdrsqlhld_sc46_copy ANFCVODRDemo	4 minutes 56 seconds	idle	broken-off	Information
-	arsdrsollog sr46	19	<pre>gcsdrsqllog_sc46_copy</pre>	10 minuter 10			Resync
9	ntaphci-a300e9u25	3	ANFCVODRDemo	seconds	idle	broken-off	Reverse Resync
							Edit Schedule
							Edit Max Transfer Rate
							Delete

J Volume Re	elationships	6.54 GIB Replicated Capacity	O Currently Trans	ferring	3 Healthy	⊗ 0 <sub>Failed</sub>	
Volume Relationships			0 0				۹ (
Health Status 🔅	Source Volume 🔹	Target Volume 🕴	Total Transfer Time 🕴	Status		E Last Successful Transfer	Ð
$\odot$	gcsdrsqldb_sc46 ntaphcl-a300e9u25	gcsdrsqldb_sc46_copy ANFCVODRDemo	19 seconds	idle	snapmirrored	May 6, 2022, 11:03:000A 5.73 MiB	
$\odot$	gcsdrsqlhld_sc46_copy ANFCVODRDemo	gcsdrsqlhld_sc46 ntaphci-a300e9u25	1 minute 46 seconds	idle	snapmirrored	May 6, 2022, 11:01:39 AN 800.76 MIB	
$\odot$	gcsdrsqllog_sc46 ntaphci-a300e9u25	gcsdrsqllog_sc46_copy ANFCVODRDemo	51 seconds	idle	snapmirrored	May 6, 2022, 11:03:15 AN 785.8 MiB	

- 5. Restart the MSSQL server service.
- 6. Verify that the SQL resources are back online.



To failback to the primary storage, make sure that the relationship direction remains the same as it was before the failover by performing a reverse resync operation.

To retain the roles of primary and secondary storage after the reverse resync operation, perform the reverse resync operation again.

This process is applicable to other applications like Oracle, similar database flavors, and any other applications using guest-connected storage.

i.

÷.

As always, test the steps involved for recovering the critical workloads before porting them into production.

# Benefits of this solution

- Uses the efficient and resilient replication of SnapMirror.
- Recovers to any available points in time with ONTAP snapshot retention.
- Full automation is available for all required steps to recover hundreds to thousands of VMs, from the storage, compute, network, and application validation steps.
- SnapCenter uses cloning mechanisms that do not change the replicated volume.
  - This avoids the risk of data corruption for volumes and snapshots.
  - Avoids replication interruptions during DR test workflows.
  - Leverages the DR data for workflows beyond DR, such as dev/test, security testing, patch and upgrade testing, and remediation testing.
- CPU and RAM optimization can help lower cloud costs by enabling recovery to smaller compute clusters.

### TR-4955: Disaster Recovery with Azure NetApp Files (ANF) and Azure VMware Solution (AVS)

Disaster recovery using block-level replication between regions within the cloud is a resilient and cost-effective way of protecting the workloads against site outages and data corruption events (for example, ransomware).

Author(s): Niyaz Mohamed, NetApp Solutions Engineering

## Overview

With Azure NetApp files (ANF) cross-region volume replication, VMware workloads running on an Azure VMware Solution (AVS) SDDC site using Azure NetApp files volumes as an NFS datastore on the primary AVS site can be replicated to a designated secondary AVS site in the target recovery region.

Disaster Recovery Orchestrator (DRO) (a scripted solution with a UI) can be used to seamlessly recover workloads replicated from one AVS SDDC to another. DRO automates recovery by breaking replication peering and then mounting the destination volume as a datastore, through VM registration to AVS, to network mappings directly on NSX-T (included with all AVS private clouds).



### Prerequisites and general recommendations

- Verify that you have enabled cross-region replication by creating replication peering. See Create volume replication for Azure NetApp Files.
- You must configure ExpressRoute Global Reach between the source and target Azure VMware Solution private clouds.
- You must have a service principal that can access resources.
- The following topology is supported: primary AVS site to secondary AVS site.
- Configure the replication schedule for each volume appropriately based on business needs and the datachange rate.



Cascading and fan- in and fan- out topologies are not supported.

## **Getting started**

#### **Deploy Azure VMware Solution**

The Azure VMware Solution (AVS) is a hybrid cloud service that provides fully functional VMware SDDCs within a Microsoft Azure public cloud. AVS is a first-party solution fully managed and supported by Microsoft and verified by VMware that uses Azure infrastructure. Therefore, customers get VMware ESXi for compute virtualization, vSAN for hyper-converged storage, and NSX for networking and security, all while taking advantage of Microsoft Azure's global presence, class-leading data- center facilities, and proximity to the rich ecosystem of native Azure services and solutions. A combination of Azure VMware Solution SDDC and Azure NetApp Files provides the best performance with minimal network latency.

To configure an AVS private cloud on Azure, follow the steps in this link for NetApp documentation and in this link for Microsoft documentation. A pilot- light environment set up with a minimal configuration can be used for DR purposes. This setup only contains core components to support critical applications, and it can scale out and spawn more hosts to take the bulk of the load if a failover occurs.



In the initial release, DRO supports an existing AVS SDDC cluster. On-demand SDDC creation will be available in an upcoming release.

### Provision and configure Azure NetApp Files

Azure NetApp Files is a high-performance, enterprise-class, metered file- storage service. Follow the steps in this link to provision and configure Azure NetApp Files as a NFS datastore to optimize AVS private cloud deployments.

#### Create volume replication for Azure NetApp Files-powered datastore volumes

The first step is to set up cross- region replication for the desired datastore volumes from the AVS primary site to the AVS secondary site with the appropriate frequencies and retentions.



Follow the steps in this link to set up cross-region replication by creating replication peering. The service level for the destination capacity pool can match that of the source capacity pool. However, for this specific use case, you can select the standard service level and then modify the service level in the event of a real disaster or DR simulations.

A cross- region replication relationship is a prerequisite and must be created beforehand.

### **DRO** installation

To get started with DRO, use the Ubuntu operating system on the designated Azure virtual machine and make sure you meet the prerequisites. Then install the package.

### **Prerequisites:**

- · Service principal that can access resources.
- Make sure that appropriate connectivity exists to the source and destination SDDC and Azure NetApp Files instances.
- DNS resolution should be in place if you are using DNS names. Otherwise, use IP addresses for vCenter.

### OS requirements:

- Ubuntu Focal 20.04 (LTS)The following packages must be installed on the designated agent virtual machine:
- Docker
- Docker- compose
- JqChange docker.sock to this new permission: sudo chmod 666 /var/run/docker.sock.



The deploy.sh script executes all required prerequisites.

The steps are as follows:

1. Download the installation package on the designated virtual machine:

```
git clone https://github.com/NetApp/DRO-Azure.git
```



The agent must be installed in the secondary AVS site region or in the primary AVS site region in a separate AZ than the SDDC.

2. Unzip the package, run the deployment script, and enter the host IP (for example, 10.10.10.10).

```
tar xvf draas_package.tar
Navigate to the directory and run the deploy script as below:
sudo sh deploy.sh
```

- 3. Access the UI using the following credentials:
  - Username: admin
  - Password: admin

	NetApp	
	Disaster Recovery Orchestrator English Existence on GU     Vername	
	Passeerd	
	- English	_

### **DRO** configuration

After Azure NetApp Files and AVS have been configured properly, you can begin configuring DRO to automate the recovery of workloads from the primary AVS site to the secondary AVS site. NetApp recommends deploying the DRO agent in the secondary AVS site and configuring the ExpressRoute gateway connection so that the DRO agent can communicate via the network with the appropriate AVS and Azure NetApp Files components.

The first step is to Add credentials. DRO requires permission to discover Azure NetApp Files and the Azure VMware Solution. You can grant the required permissions to an Azure account by creating and setting up an Azure Active Directory (AD) application and by obtaining the Azure credentials that DRO needs. You must bind

the service principal to your Azure subscription and assign it a custom role that has the relevant required permissions. When you add source and destination environments, you are prompted to select the credentials associated with the service principal. You need to add these credentials to DRO before you can click Add New Site.

To perform this operation, complete the following steps:

- 1. Open DRO in a supported browser and use the default username and password (admin/admin). The password can be reset after the first login using the Change Password option.
- 2. In the upper right of the DRO console, click the **Settings** icon, and select **Credentials**.
- 3. Click Add New Credential and follow the steps in the wizard.
- 4. To define the credentials, enter information about the Azure Active Directory service principal that grants the required permissions:
  - · Credential name
  - Tenant ID
  - Client ID
  - · Client secret
  - Subscription ID

You should have captured this information when you created the AD application.

5. Confirm the details about the new credentials and click Add Credential.

NetApp Disaster Recovery Orchestrator 🂊   Dashboard   Discover	Resource Groups Replication Plans Job	b Monitoring	<b>≜ ⇔ </b> 0 ©
Add New Credential	1 Credentials Details		×
	Enter Credentials D	etails	
	Credential Name	•	
	Tenant Id		
	Client Id	•	
	Client Secret		
	Subscription Id		
	<u></u>		
		-	
	Add Credential		

After you add the credentials, it's time to discover and add the primary and secondary AVS sites (both vCenter and the Azure NetApp files storage account) to DRO. To add the source and destination site, complete the following steps:

- 6. Go to the **Discover** tab.
- 7. Click Add New Site.
- 8. Add the following primary AVS site (designated as **Source** in the console).

- SDDC vCenter
- Azure NetApp Files storage account
- 9. Add the following secondary AVS site (designated as **Destination** in the console).
  - SDDC vCenter
  - Azure NetApp Files storage account

NetApp Disaster Recovery Orchestrator A Dashboard Discover Resource Groups Replication Plans Job Monitoring	ė	¢	?	۹
Add New Site 10 Site Type 3 Site Details 3 vCenter Details 4 Storage Details				×
Site Type				
Source Destination				
Continue				

10. Add site details by clicking **Source**, entering a friendly site name, and select the connector. Then click **Continue**.



For demonstration purposes, adding a source site is covered in this document.

- 11. Update the vCenter details. To do this, select the credentials, Azure region, and resource group from the dropdown for the primary AVS SDDC.
- 12. DRO lists all the available SDDCs within the region. Select the designated private cloud URL from the dropdown.
- 13. Enter the cloudadmin@vsphere.local user credentials. This can be accessed from Azure Portal. Follow the steps mentioned in this link. Once done, click **Continue**.

Add New Site	(🕑 Site Type 🕞	) Site Details (3) vCenter Details (4)	) Storage Details	
		Source AVS Private Cloud		
	Select Credentials	Azure Region	Azure Resource Group	
	DemoCred -	West Europe 👻	ANFAVSVal2 v	
	Add New Credential [5			
		AVS Details		
		Ava octalia		
	Web Clien	URL	0	
		ANFDataClus	*	
	licename			
	clouded	min@venhere local		
	Password		0	
		•••••		
		Accept self-signed certificates		

14. Select the Source Storge details (ANF) by selecting the Azure Resource group and NetApp account.

## 15. Click Create Site.

pp	Disaster Recovery Orchestrator 💊 🕴 Dashi	ooard Discover Resource Gro	oups Replication Plans Job Mor	hitoring	4	•
	C 2 Sites 2 vCentr	rs 2 Storages	Site Type	Site Location	(a) 2 Cloud	
	2 sites				Q O Add New Site	i.
	Site Name	©   Site Type = 〒   Locati	on 🖙   vCenter 🗘   Storage 🗘	VM List Discovery Status	Û.	
	DemoDest	Destination Cloud	1 1	<ul> <li>https://10.75.0.2/</li> </ul>	⊘ Success	

Once added, DRO performs automatic discovery and displays the VMs that have corresponding cross- region replicas from the source site to the destination site. DRO automatically detects the networks and segments used by the VMs and populates them.

Disaster Recovery Orchestrator	Sector R	escurce Groups   Replication Plans   Job	Monitoring		<b>A</b> -
Back					
		VM List Site: DemoSRC   vCenter: https://172	30.1562/		
421		1.000	VM Protection		
CP 7	stores	128 Virtual Machines	🤣 2. Protected	0 126 Unprotected	
128 vm				0	Create Resource Group
VM Name	C     VM Stitus	VM State	DutaStore	: I (00	2 Memory (MD) 2
HDSench,2.5.1	0 Not Protected	() Powered Gn	vianDatastore	8	8192
nci-fio-datastore-13984-0-1	0 Not Protected	D Powered Off	HCRstDS	12	65536
ICCA2005-WID-R1	O Not Protected	() Powered On	vianDataitose		14336
ICCA2005-NE-R1	0 Not Protected	() Powered On	vianDataitore		3072
10CA2005-01.81	0 Not Protected	() Powered Gn	vianDatastiore	8	3072
HCX_Demo_05	9 Not Protected	D Powered Off	Demo002	1	2048
between defaulture 19844-0.1	10 Had Scalasted	(D) Encount Off	W/RMOK	14	254173

The next step is to group the required VMs into their functional groups as resource groups.

### **Resource groupings**

After the platforms have been added, group the VMs you want to recover into resource groups. DRO resource groups allow you to group a set of dependent VMs into logical groups that contain their boot orders, boot delays, and optional application validations that can be executed upon recovery.

To start creating resource groups, click the **Create New Resource Group** menu item.

1. Access Resource Grou\*ps and click \*Create New Resource Group.

netApp	Disaster Recovery Orchestrator 💊 📔 D	ashboard   Discover   Resource Groups   Re	eplication Plans   Job Monitoring	٨	¢? 2
	-				
	Resource Group	☐ 1 Site	Center 1	<mark>같</mark> 2 Virtual Machines	
	1 Resource Group			Q O Create New Resource Group	
	Resource Group Name	‡   Site Name	⇒   Source vCenter		
	DemoRG	DemoSRC	https://172.30.156.2/	View VM List	

- 2. Under New Resource Group, select the source site from the dropdown and click Create.
- 3. Provide the resource group details and click **Continue**.
- 4. Select appropriate VMs using the search option.
- 5. Select the Boot Order and Boot Delay (secs) for all the selected VMs. Set the order of the power- on sequence by selecting each virtual machine and setting up the priority for it. The default value for all virtual machines is 3. The options are as follows:
  - The first virtual machine to power on
  - Default

• The last virtual machine to power on

letApp	Disaster Recovery Orchestrator 💊 🛛	Dashboard Discov		Replication Plans	Job Monitoring		
	Edit Resource Group		Resource Group Deta	ails 🕢 Select VMs	Boot order	and Delay	
				Boot order and	l Delay		
		VM Name	Boot Order 💿		Boot Delay	(secs)	
		QALin1	3	[2]	0	10	
		QALin	3	10	0	[0]	

### 6. Click Create Resource Group.

Disaster Recovery Orchestrator 💊	Dashboard   Discover   Resource Groups   Re	plication Plans   Job Monitoring	<b>A</b> :
2 1 Resource Group	C 1 Ste	Center	Virtual Machines
1 Resource Group			Q O Create New Resource Group
Resource Group Name			포   VM List
DemoRG	DemoSBC	https://172.30.156.2/	View VM List

## **Replication plans**

You must have a plan to recover applications in the event of a disaster. Select the source and destination vCenter platforms from the drop down, pick the resource groups to be included in this plan, and also include the grouping of how applications should be restored and powered on (for example, domain controllers, tier-1, tier-2, and so on). Plans are often called blueprints as well. To define the recovery plan, navigate to the Replication Plan tab, and click **New Replication Plan**.

To start creating a replication plan, complete the following steps:

1. Navigate to Replication Plans and click Create New Replication Plan.

n NetApp	Disaster Recovery C	Drchestrator 💊 🛛 Da	shboard   Discover   Resource G	iroups   Replication Plans   Job Monitoring					• •	0	
	B 1 Rep	lication Plans	2 1 Resource Groups	Source Details	vCenters	Destination	n Details	1 iers			
	1 Replication Plan						Q D Creat	e New Replication Pi	in T		
	Plan Name DemoRP	Active Site     Source	Status ② Active	Compliance	Source Site		Resource Grou	ups			

2. On the **New Replication Plan**, provide a name for the plan and add recovery mappings by selecting the Source Site, associated vCenter, Destination Site, and associated vCenter.

NetApp Disaster Recovery Orchestrator	Dashboard Discover Resource Grou	ups   Replication P	ans   Job Monitoring		<b>A</b> 4	<b>8</b>	9
Create New Replication Plan	() Replication Plan and Site Details	2 Select Resource	Groups ③ Set Execution Order ④	Set VM Details			×
		Replication	Plan Details				
	Plan Name			0			
	DemoRP						
		Recovery	Mapping				
	Source Site	0	Destination Site	0			
	DemoSRC	*	DemoDest	*			
	Source vCenter	0	Destination vCenter	0			
	https://172.30.156.2/	*	https://10.75.0.2/	*			
		Cluster I	Mapping				
	Source Site Resource	O Destination	Site Resource 🕖				
	Cluster-1	-	Cluster-1 +	Add			
	Source Resource	Destinat	on Resource				
		No Mappi	ngs added!				
		6	ntinue				

3. After recovery mapping is complete, select the **Cluster Mapping**.

NetApp Disaster Recovery Orchestrator	Dashboard Discover Resource Gro	ups   Replication Pla	ns Job Monitoring		۵	٠	?	9
Create New Replication Plan	Replication Plan and Site Details	2 Select Resource G	roups (3) Set Execution Order (4) Set VM Details					×
	Plan Name			0				
	DemoRP							
Recovery Mapping								
	Source Site	0	Destination Site	0				
	DemoSRC	. *	DemoDest	*				
	Source vCenter	0	Destination vCenter	0				
	https://172.30.156.2/	*	https://10.75.0.2/	*				
	No more Source/Destination cluster resources available for mapping							
	Source Resource Destination Resource							
	Cluster-1	Cluster-1	Delete					
		Cont	tinue					

- 4. Select Resource Group Details and click Continue.
- 5. Set the execution order for the resource group. This option enables you to select the sequence of operations when multiple resource groups exist.
- 6. Once done, set network mapping to the appropriate segment. The segments should already be provisioned on the secondary AVS cluster, and, to map the VMs to those, select the appropriate segment.
- 7. Datastore mappings are automatically selected based on the selection of VMs.



Cross- region replication (CRR) is at the volume level. Therefore, all VMs residing on the respective volume are replicated to the CRR destination. Make sure to select all VMs that are part of the datastore, because only virtual machines that are part of the replication plan are processed.
NetApp Disaster Recovery Orchestrator 💊	Dashboard Discover Resource G	iroups   Replication Plans   Job N	Ionitoring	4	۵ (	? 9	
Create New Replication Plan	Replication Plan and Site Details	Select Resource Groups 3	Set Execution Order (4) Set VM Details			×	
		Replication Plan Det	ails				
		Select Execution Orde	r				
	Resource Group Name		Execution Order 🜒				
	DemoRG		3 [3]				
		Network Mapping					
	No more	Source/Destination network resources	available for mapping				
	Source Resource	Destination Resource					
	SepSeg	SegDR	Delete				
		DataStore Mapping					
	Source DataStore	Destination Volume					
	TestSrc01	gwc_ntap_acct/gwc_DRO_cp/testsrci	О1сору				
		Previous	iue				

8. Under VM details, you can optionally resize the VMs CPU and RAM parameters. This can be very helpful when you are recovering large environments to smaller target clusters or when you are conducting DR tests without having to provision a one-to-one physical VMware infrastructure. Also, modify the boot order and boot delay (secs) for all the selected VMs across the resource groups. There is an additional option to modify the boot order if any changes are required from what you selected during resource- group boot-order selection. By default, the boot order selected during resource- group selection is used, however any modifications can be performed at this stage.

NetApp Disaster Recovery Orchestrat	tor 💊   Dashboard   Discover	Resource Groups   Replication	on Plans   Job Monitoring				-	•	?	۹
Create New Replication Plan	<ul> <li>Replication Plan and Sit</li> </ul>	e Details 🕢 Select Resour	rce Groups 🕜 Set Exect	ution Order	Set VM Details					×
		V	M Details							
	2 vms					۹				
	VM Name	No. of CPUs	Memory (MB)	NIC/IP	Boot Order 💿 🔳 Override					
	Resource Group : DemoRG									
	QALin1	1 8	1024 [2]	O Static O Dynamic	3	0				
	QALin	4	1024 [0]	<ul> <li>Static</li> <li>Dynamic</li> </ul>	3	0				
		Previous	Create Replication Plan							

9. Click **Create Replication Plan**. After the replication plan is created, you can exercise the failover, test failover, or migrate options depending on your requirements.

🗖 NetApp	Disaster Recovery C	Orchestrator 💊 📔 Dashboar	rd   Discover   Resource G	Sroups   Replication Plans	Job Monitoring				•	3	9
	B 1 Rep	lication Plans	7 1 Resource Groups	Source Details	1 vCenters	Destinatio	n Details	2 1 vCenters			
	1 Replication Plan						۵ ۵	Create New Replication Plan			
	Plan Name	C Active Site	Status	Compliance	Source Site	⇒ Destination Site	÷1	Ĵ.			
	DemoRP	Source	G Active	A Partially Healthy	DemoSRC	DemoDest	Re	source Groups			
								Pan Details Edit Pan Falover Test Falover Migrate			
								Delete Plan			

During the failover and test failover options, the most recent snapshot is used, or a specific snapshot can be selected from a point-in-time snapshot. The point-in-time option can be very beneficial if you are facing a corruption event like ransomware, where the most recent replicas are already compromised or encrypted. DRO shows all available time points.

		_	Source Details		Destination Details		
B 1 Replica	tion Plans	Resource Groups	Sites	t vCenters	Sites 1	vCenters	
		Testfailover Details			×		
1 Replication Plan		<ul> <li>Use latest snapshot </li> <li>Select specific snapshot</li> </ul>	o		^ Q	O Create New Replication	on Plan
Plan Name		Volume	-	Snapshot	¢.	Ъ.	
DemoRP	⊘ Source	WEANFAVSacct/testcap/te	stsrc01	Select Snapshot	. (	Resource Groups	211
				2023-04-28			
				2023-04-28711:31:55.000Z - gwc_ntap	has		
				2023-04-28T11:21:54.000Z - gwc_ntap			

To trigger failover or test failover with the configuration specified in the replication plan, you can click **Failover** or **Test Failover**. You can monitor the replication plan in the task menu.

Back			
	Test Failover Steps Replication Plan: DemoRP		
~	Cloning volumes for test (in parallel)	⊙ Success	0.7 Seconds 🕕
~	Mounting cloned volumes and creating datastores (in parallet)	<ul> <li>Success</li> </ul>	0.9 Seconds 🗿
~	Registering VMs (in parallel)	⊙ Success	0.1 Seconds 🕕
~	Powering on VMs in protection group - DemoRG - in target (in parallel)	() Success	0.1 Seconds ()

After failover is triggered, the recovered items can be seen in the secondary site AVS SDDC vCenter (VMs, networks, and datastores). By default, the VMs are recovered to Workload folder.

C 2 Sites	1 Resource Group	Replication	Plan	128 VMs Potected	VMs 0 127 Ungestanded	
Environments 2 Vortual Environments	2. Auff Storage Accounts	Topology Canvas			Immensive View (7	
SDDC Summary	D 14	DV Mays. P	MisRC 17236 19629	Demiciesi heys/118/26/2		
12 Datastores						
Execution John		Replication Plans			1	
🕑 1 Tettal Solar	C 1 In Program	Reportation Flam	Active Site	Status	Mode	

Failback can be triggered at the replication plan level. In case of test failover, the tear down option can be used to roll back the changes and remove the newly created volume. Failbacks related to failover are a two- step process. Select the replication plan and select **Reverse Data sync**.

🗖 NetApp	Disaster Recovery Orchestrator 💊   Dashboard   Discover	Resource Groups Replication Plans	Job Monitoring		A 🌣 🕄 🛎
	Replication Plans	Groups	vCenters	Destination Details	vCenters
	1 Replication Plan Plan Name	Compliance	Source Site 🛛 😤	Q D	Create New Replication Plan
	DemoRP	ning In Failover Mode 🕝 Healthy	DemoSRC	DemoDest Res	ource Groups
					Plan Details
					Reverse Data Sync
					Fallante

After this step is complete, trigger failback to move back to the primary AVS site.

netApp	Disaster Recovery Orchestrator 💊   Dashboard   Discover	Resource Groups Replication Plans Job Monitoring	4 ¢ @ ©
	B 1 Replication Plans 1 Resource	e Groups Source Details	Destination Details I I I VCenters
	1 Replication Plan		Q O Create New Replication Plan
	Plan Name     Clinical Active Site     Status       DemoRP     Operation     Operation	Compliance   Source Site ♥   Desti we ⓒ Healthy DemoSRC Demo	antion Site 0       SDest Resource Groups (33)
			Pan Deals Faiback
II NetApp	Disaster Recovery Orchestrator 🔌   Dashimmed   Discuss	er   Resource Groups   Replication Plans   Job Monitoring	A O O O
	2 1	<b>a</b> 1 <b>b</b> 128	Protected VMs
	Resource G	ogi 🧐 Repication Plan	Protected Unprotected
	2     virtual Environments     All Storage Accounts	sopology Lanvas	minimize view D
	SDDC Summary		
	Chatters 2 14 Folders	DemoSRC	DemoCent pps://10.75.0.2/
	12     Datastores     2     12     Networks		
	Execution Jobs	Replication Plans	
	3     Total Jobs     m Progens	Ingelonition Pain Active Site	Status (7) Addive

From the Azure portal, we can see that the replication health has been broken off for the appropriate volumes that were mapped to the secondary site AVS SDDC as read/write volumes. During test failover, DRO does not map the destination or replica volume. Instead, it creates a new volume of the required cross- region replication snapshot and exposes the volume as a datastore, which consumes additional physical capacity from the capacity pool and ensures that the source volume is not modified. Notably, replication jobs can continue during DR tests or triage workflows. Additionally, this process makes sure that the recovery can be cleaned up without the risk of the replica being destroyed if errors occur or corrupted data is recovered.

# **Ransomware recovery**

Recovering from ransomware can be a daunting task. Specifically, it can be difficult for IT organizations to pinpoint what the safe point of return is, and, once that's determined, how to ensure that recovered workloads are safeguarded from the attacks reoccurring (for example, from sleeping malware or through vulnerable applications).

DRO addresses these concerns by allowing organizations to recover from any available point-in-time. Workloads are then recovered to functional and yet isolated networks, so that applications can function and communicate with each other but are not exposed to any north- south traffic. This process gives security teams a safe place to conduct forensics and identify any hidden or sleeping malware.

# Conclusion

The Azure NetApp Files and Azure VMware disaster recovery solution provide you with the following benefits:

- Leverage efficient and resilient Azure NetApp Files cross- region replication.
- Recover to any available point-in-time with snapshot retention.
- Fully automate all required steps to recover hundreds to thousands of VMs from the storage, compute, network, and application validation steps.
- Workload recovery leverages the "Create new volumes from the most recent snapshots" process, which doesn't manipulate the replicated volume.
- Avoid any risk of data corruption on the volumes or snapshots.
- Avoid replication interruptions during DR test workflows.
- Leverage DR data and cloud compute resources for workflows beyond DR, such as dev/test, security testing, patch and upgrade testing, and remediation testing.
- CPU and RAM optimization can help lower cloud costs by allowing recovery to smaller compute clusters.

# Where to find additional information

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

· Create volume replication for Azure NetApp Files

https://learn.microsoft.com/en-us/azure/azure-netapp-files/cross-region-replication-create-peering

Cross-region replication of Azure NetApp Files volumes

https://learn.microsoft.com/en-us/azure/azure-netapp-files/cross-region-replication-introduction#service-level-objectives

Azure VMware Solution

https://learn.microsoft.com/en-us/azure/azure-vmware/introduction

• Deploy and configure the Virtualization Environment on Azure

# Setup AVS on Azure

• Deploy and configure Azure VMware Solution

https://learn.microsoft.com/en-us/azure/azure-vmware/deploy-azure-vmware-solution?tabs=azure-portal

## Using Veeam Replication and Azure NetApp Files datastore for disaster recovery to Azure VMware Solution

Azure NetApp Files (ANF) datastores decouples storage from compute and unlocks the flexibility needed for any organisation to take their workloads to the cloud. It provides customers with flexible, high-performance storage infrastructure that scales independently of compute resources. Azure NetApp Files datastore's simplifies and optimizes the deployment alongside Azure VMware Solution (AVS) as a disaster recovery site for on premises VMWare environments.

# Overview

Azure NetApp Files (ANF) volume based NFS datastores can be used to replicate data from on-premises using any validated third-party solution that provides VM replication capability. By adding Azure NetApp Files datastores, it will enable cost optimised deployment vs building an Azure VMware Solution SDDC with enormous amount of ESXi hosts to accommodate the storage. This approach is called a "Pilot Light Cluster". A pilot light cluster is a minimal AVS host configuration (3 x AVS nodes) along with Azure NetApp Files Datastore capacity.

The objective is to maintain a low-cost infrastructure with all the core components to handle a failover. A pilot light cluster can scale out and provision more AVS hosts if a failover does occur. And once the failover is complete and normal operations are restored, the pilot light cluster can scale back down to low-cost mode of operations.

# Purposes of this document

This article describes how to use Azure NetApp Files datastore with Veeam Backup and replication to set up disaster recovery for on-premises VMware VMs to (AVS) using the Veeam VM replication software functionality.

Veeam Backup & Replication is a backup and replication application for virtual environments. When virtual machines are replicated, Veeam Backup & Replication is replicated from on AVS, the software will create an exact copy of the VMs in the native VMware vSphere format on the target AVS SDDC cluster. Veeam Backup & Replication will keep the copy synchronized with the original VM. Replication provides the best recovery time objective (RTO) as there is a mounted copy of a VM at the DR site in a ready-to-start state.

This replication mechanism ensures that the workloads can quickly start in a AVS SDDC in the case of a disaster event. The Veeam Backup & Replication software also optimizes traffic transmission for replication over WAN and slow connections. In addition, it also filters out duplicate data blocks, zero data blocks, swap files, and "excluded VM guest OS files". The software will also compress the replica traffic. To prevent replication jobs from consuming the entire network bandwidth, WAN accelerators and network throttling rules can be utilized.

The replication process in Veeam Backup & Replication is job driven which means replication is performed by configuring replication jobs. In the case of a disaster event, failover can be triggered to recover the VMs by failing over to its replica copy. When failover is performed, a replicated VM takes over the role of the original VM. Failover can be performed to the latest state of a replica or to any of its good known restore points. This enables ransomware recovery or isolated testing as needed. Veeam Backup & Replication offers multiple options to handle different disaster recovery scenarios.



# **Solution Deployment**

# **High level steps**

- 1. Veeam Backup and Replication software is running in an on-premises environment with appropriate network connectivity.
- 2. Deploy Azure VMware Solution (AVS) private cloud and attach Azure NetApp Files datastores to Azure VMware Solution hosts.

A pilot-light environment set up with a minimal configuration can be used for DR purposes. VMs will fail over to this cluster in the event of an incident, and additional nodes can be added).

- 3. Set up replication job to create VM replicas using Veeam Backup and Replication.
- 4. Create failover plan and perform failover.
- 5. Switch back to production VMs once the disaster event is complete and primary site is Up.

## Pre-requisites for Veeam VM Replication to AVS and ANF datastores

- 1. Ensure the Veeam Backup & Replication backup VM is connected to the source as well as the target AVS SDDC clusters.
- 2. The backup server must be able to resolve short names and connect to source and target vCenters.
- 3. The target Azure NetApp Files datastore must have enough free space to store VMDKs of replicated VMs.

For additional information, refer to "Considerations and Limitations" covered here.

## **Deployment Details**

Veeam Backup & Replication leverages VMware vSphere snapshot capabilities/During replication, Veeam Backup & Replication requests VMware vSphere to create a VM snapshot. The VM snapshot is the pointin-time copy of a VM that includes virtual disks, system state, configuration and metadata. Veeam Backup & Replication uses the snapshot as a source of data for replication.

To replicate VMs, follow the below steps:

- 1. Open the Veeam Backup & Replication Console.
- 2. On the Home view. Right click the jobs node and select Replication Job > Virtual machine.
- 3. Specify a job name and select the appropriate advanced control checkbox. Click Next.
  - Select the Replica seeding check box if connectivity between on-premises and Azure has restricted bandwidth.

\*Select the Network remapping (for AVS SDDC sites with different networks) check box if segments on Azure VMware Solution SDDC do not match that of on-premises site networks.

• If the IP addressing scheme in on-premises production site differs from the scheme in the target AVS site, select the Replica re-IP (for DR sites with different IP addressing scheme) check box.

Name	Name:
(intual Machines	AVS_20230522_RepJob01
incoar infactinites	Description:
estination	Created by VEEAMBKPSRV05\Administrator at 5/21/2023 10:52 PM.
Network	
ob Settings	Show advanced controls:
	Replica seeding (for low bandwidth DR sites)
Data Transfer	Network remapping (for DR sites with different virtual networks)
Guest Processing	Replica re-IP (for DR sites with different IP addressing scheme)
Schedule	
Summary	
	High priority Backup infrastructure resources are offered to high priority jobs first. Use this option for jobs sensitive to the start time, or jobs with strict RPO requirements.

4. Select the VMs to be replicated to Azure NetApp Files datastore attached to a Azure VMware Solution SDDC in the Virtual Machines\* step. The Virtual machines can be placed on vSAN to fill the available vSAN datastore capacity. In a pilot light cluster, the usable capacity of a 3-node cluster will be limited. The rest of the data can be easily placed on Azure NetApp Files datastores so that the VMs can recovered, and cluster can be expanded to meet the CPU/mem requirements. Click Add, then in the Add Object window select the necessary VMs or VM containers and click Add. Click Next.

## Virtual Machines

Select one or more VMs to replicate. Use exclusion settings to exclude specific VMs and virtual disks from replication.

	Name	Туре	Size	^	Add
Virtual Machines	TestVeeam21	Virtual Machine	873 MB		Ramour
Destination	TestVeeam22	Virtual Machine	890 MB		THE FILL
	TestVeeam23	Virtual Machine	883 MB		
Network	TestVeeam24	Virtual Machine	879 MB		Exclusions
	TestVeeam25	Virtual Machine	885 MB		Source
Job Settings	TestVeeam26	Virtual Machine	883 MB		
Data Tanafar	TestVeeam27	Virtual Machine	879 MB		
Data Iransfer	TestVeeam28	Virtual Machine	880 MB		T Up
Guest Processing	TestVeeam29	Virtual Machine	878 MB		+ Down
	TestVeeam30	Virtual Machine	876 MB		
Schedule	TestVeeam31	Virtual Machine	888 MB		
	TestVeeam32	Virtual Machine	881 MB		
Summary	TestVeeam33	Virtual Machine	877 MB		
	TestVeeam34	Virtual Machine	875 MB		-
	TestVeeam35	Virtual Machine	882 MB		Recalculate
	WinSQL401	Virtual Machine	20.3 GB		
	WinSQL405	Virtual Machine	24.2 GB		Total size:
	Phone con tor			*	120 GB

5. After that, select the destination as Azure VMware Solution SDDC cluster / host and the appropriate resource pool, VM folder and FSx for ONTAP datastore for VM replicas. Then click **Next**.

teres -	Host or cluster	1
vame /irtual Machines	Cluster-1	Choose
Destination	Resource pool:	
Vetwork	Resources	Choose
ob Settings	Pick resource pool for selected replicas VM folder:	
Data Transfer	vm	Choose
Suest Processing	Pick VM folder for selected replicas	
	Datastore:	
chedule	ds001 [152.6 GB free] ds001 is an ANF Datastore	Choose
Summary	Pick datastore for selected virtual disks	

6. In the next step, create the mapping between source and destination virtual network as needed.

Name	Network mapping:		
Virtual Machines	Source network SVM_3508 (vDS-Switch0)	Target network	Add
Destination	VM_3510 (vDS-Switch0)	SegmentTemp	2
Data Transfer Guest Processing Schedule Summary			

- 7. In the **Job Settings** step, specify the backup repository that will store metadata for VM replicas, retention policy and so on.
- 8. Update the **Source** and **Target** proxy servers in the **Data Transfer** step and leave **Automatic** selection (default) and keep **Direct** option selected and click **Next**.
- 9. At the **Guest Processing** step, select **Enable application-aware processing** option as needed. Click **Next**.

ame irtual Machines	Enable application-aware processing Detects and prepares applications for consistent backup, performs transaction logs configures the OS to perform required application restore steps upon first boot.	processing, and
estination	Customize application handling options for individual machines and applications Guest interaction proxy:	Applications
etwork	Automatic selection	Choose
b Settings	Guest QS credentials:	
ata Transfer		Add
	Manage accounts	
uest Processing	Customize guest OS credentials for individual machines and operating systems	Credentials
chedule	Verify network connectivity and credentials for each machine included in the job	Test Now
ummary		

10. Choose the replication schedule to run the replication job to run on a regular basis.

Name	Run the job automatically	r -					
Virtual Machines	Daily at this time:	10:00 PM	•	Everyday		~	Days
	<ul> <li>Monthly at this time:</li> </ul>	10:00 PM	*	Fourth $\vee$	Saturday	. v	Months
Destination	O Periodically every:	1	se.	Hours		~	Schedule
Network	<ul> <li>After this job:</li> </ul>	Replication	Job 2	2 (Created by VEEA	MBKPSRV05\A	dministr	ator at 6/6/
Data Transfer Guest Processing Schedule Summary	Backup window  Backup window  Terminate job if it exce If the job does not cor terminated to prevent	essing: attempt for: eeds allowed t mplete within snapshot con	3 10 acki alloc	times minutes up window ated backup wind during production	low, it will be n hours.		Window,

11. At the **Summary** step of the wizard, review details of the replication job. To start the job right after the

wizard is closed, select the **Run the job when I click Finish** check box, otherwise leave the check box unselected. Then click **Finish** to close the wizard.

10	_	-
	4	-
- 2	Æ	ΞI

# Summary

The job's settings have been saved successfully. Click Finish to exit the wizard.

Name	Summary:	
	Name: AVS 20230522 RepJob01	^
Virtual Machines	Type: VMware Replication	
	Source items:	
Destination	TestVeeam21 (a300-vcsa05.ehcdc.com)	
Sestinotion	TestVeeam22 (a300-vcsa05.ehcdc.com)	
Mathuards	TestVeeam23 (a300-vcsa05.ehcdc.com)	
INELWOIK	TestVeeam24 (a300-vcsa05.ehcdc.com)	
	TestVeeam25 (a300-vcsa05.ehcdc.com)	
Job Settings	TestVeeam26 (a300-vcsa05.ehcdc.com)	
	TestVeeam27 (a300-vcsa05.ehcdc.com)	
Data Transfer	TestVeeam28 (a300-vcsa05.ehcdc.com)	
	TestVeeam29 (a300-vcsa05.ehcdc.com)	
Guest Processing	TestVeeam30 (a300-vcsa05.ehcdc.com)	
	TestVeeam31 (a300-vcsa05.ehcdc.com)	
Schedule	TestVeeam32 (a300-vcsa05.ehcdc.com)	
	TestVeeam33 (a300-vcsa05.ehcdc.com)	
Summany	TestVeeam34 (a300-vcsa05.ehcdc.com)	
Summary	TestVeeam35 (a300-vcsa05.ehcdc.com)	
	WinSQL401 (a300-vcsa05.ehcdc.com)	
	WinSQL405 (a300-vcsa05.ehcdc.com)	
	WinSQL404 (a300-vcsa05.ehcdc.com)	~
	1 146-COL 402 (-2000E -L-J)	
	Run the job when I click Finish	
		1
	< Previous Next > Finish	Cancel

Once the replication job starts, the VMs with the suffix specified will be populated on the destination AVS SDDC cluster / host.

C 48 660	×								
. (nm.	S.c.m			T. 1 mg					
tran Sharan Chila Alta Sharan Sharan Sharan	States	-	-		12	131	No. of Concession, Name	111	
		1							
	-		-		1	1			
	-	-	-						1000
		1.000	C manufat C helitypic C helit						1
angle and									
1.0									

For additional information for Veeam replication, refer How Replication Works

### Step 2: Create a failover plan

When the initial replication or seeding is complete, create the failover plan. Failover plan helps in performing failover for dependent VMs one by one or as a group automatically. Failover plan is the blueprint for the order in which the VMs are processed including the boot delays. The failover plan also helps to ensure that critical dependant VMs are already running.

To create the plan, navigate to the new sub section called **Replicas** and select **Failover Plan**. Choose the appropriate VMs. Veeam Backup & Replication will look for the closest restore points to this point in time and use them to start VM replicas.



The failover plan can only be added once the initial replication is complete and the VM replicas are in Ready state.



The maximum number of VMs that can be started simultaneously when running a failover plan is 10



During the failover process, the source VMs will not be powered off

To create the Failover Plan, do the following:

1. On the Home view. Right click the Replicas node and select Failover Plans > Failover Plan > VMware vSphere.

Backup Replication CDP Job + Job + Policy Primary Jobs	Restore Failover Plan - Restore Actions	ort kup					
Home	Q. Type in an object name to search for X All jobs						
Jobs Image: Replication Replicas Ready Active (4) Failover Plans	Name↓	Type VMware Replication VMware Replication	Objects 1 20	Status Stopped Stopped	Last Run 63 days ago 8 days ago		

2. Next provide a name and a description to the plan. Pre and Post-failover script can be added as required. For instance, run a script to shutdown VMs before starting the replicated VMs.

ieneral	Name:	
irtual Machines	Description:	
ummary	Created by VEEAMBKPSRV05\Administrator at 5/24/2023 9:08 AM.	
	Re-failover regint	
		Browseur
	Post-failover script:	
		Browsen

3. Add the VMs to the plan and modify the VM boot order and boot delays to meet the application dependencies.

ame  TestVeeam21  TestVeeam23	Delay 2 sec	Replica state 63 days ago (5:52 AM T	^	Add VM
TestVeeam21 TestVeeam23	2 sec	63 days ago (5:52 AM T	6	-
TestVeeam23	2000	and a second s		naccost
T 11/ 24	c sec	7 days ago (10:12 AM T	i i i	Kemove
TestVeeam24	2 sec	7 days ago (10:20 AM T		Set Delay
TestVeeam22	2 sec	7 days ago (10:10 AM T		-
WinSQL401	2 sec	7 days ago (3:52 AM Tu		
WinSQL405	2 sec	8 days ago (4:05 PM Mo		
TestVeeam25	2 sec	7 days ago (10:14 AM T		
TestVeeam26	2 sec	7 days ago (10:17 AM T		
TestVeeam27	2 sec	7 days ago (10:18 AM T		
TestVeeam28	2 sec	7 days ago (10:14 AM T		
TestVeeam29	2 sec	7 days ago (10:18 AM T		
TestVeeam30	2 sec	7 days ago (10:15 AM T		
TestVeeam31	2 sec	7 days ago (10:21 AM T		
TestVeeam32	2 sec	7 days ago (10:13 AM T		
TestVeeam33	2 sec	7 days ago (10:15 AM T		
TestVeeam34	2 sec	7 days ago (10:14 AM T		
TestVeeam34	2 sec	7 days ago (10:14 AM T		✤ Up
TestVeeam34	2 sec	7 days ago (10:14 AM T		✤ Up
TestVeeam34	2 sec	7 days ago (10:14 AM T		
The sector was shown in the sector of the se		t abje ege (territerin		
	I TestVeeam22 I WinSQL401 I WinSQL405 I TestVeeam25 I TestVeeam26 I TestVeeam27 I TestVeeam28 I TestVeeam29 I TestVeeam30 I TestVeeam31 I TestVeeam32 I TestVeeam32 I TestVeeam33	IVestVeeam222 secIVinSQL4012 secIVinSQL4052 secITestVeeam252 secITestVeeam262 secITestVeeam272 secITestVeeam282 secITestVeeam292 secITestVeeam302 secITestVeeam312 secITestVeeam322 secITestVeeam332 sec	Inestveeam22       2 sec       7 days ago (10:10 AM T         IWinSQL401       2 sec       7 days ago (3:52 AM Tu         IWinSQL405       2 sec       8 days ago (4:05 PM Mo         ITestVeeam25       2 sec       7 days ago (10:14 AM T         ITestVeeam26       2 sec       7 days ago (10:17 AM T         ITestVeeam27       2 sec       7 days ago (10:18 AM T         ITestVeeam28       2 sec       7 days ago (10:18 AM T         ITestVeeam29       2 sec       7 days ago (10:18 AM T         ITestVeeam30       2 sec       7 days ago (10:15 AM T         ITestVeeam31       2 sec       7 days ago (10:21 AM T         ITestVeeam32       2 sec       7 days ago (10:13 AM T         ITestVeeam33       2 sec       7 days ago (10:15 AM T	Inestveeam22       2 sec       7 days ago (10:10 AM T         IWinSQL401       2 sec       7 days ago (3:52 AM Tu         IWinSQL405       2 sec       8 days ago (4:05 PM Mo         ITestVeeam25       2 sec       7 days ago (10:14 AM T         ITestVeeam26       2 sec       7 days ago (10:17 AM T         ITestVeeam27       2 sec       7 days ago (10:18 AM T         ITestVeeam28       2 sec       7 days ago (10:18 AM T         ITestVeeam29       2 sec       7 days ago (10:18 AM T         ITestVeeam30       2 sec       7 days ago (10:18 AM T         ITestVeeam31       2 sec       7 days ago (10:15 AM T         ITestVeeam32       2 sec       7 days ago (10:21 AM T         ITestVeeam33       2 sec       7 days ago (10:15 AM T

During failover, the source VM in the production site is switched over to its replica at the disaster recovery site. As part of the failover process, Veeam Backup & Replication restores the VM replica to the required restore point and moves all I/O activities from the source VM to its replica. Replicas can be used not only in case of a disaster, but also to simulate DR drills. During failover simulation, the source VM remains running. Once all the necessary tests have been conducted, you can undo the failover and return to normal operations.



Make sure network segmentation is in place to avoid IP conflicts during failover.

To start the failover plan, simply click in **Failover Plans** tab and right click on your failover plan. Select **\*Start**. This will failover using the latest restore points of VM replicas. To fail over to specific restore points of VM replicas, select **Start to**.

Start Start Retry Undo to Actions Details	Edit Delete			
Home	Q Type in an object	name to search for	A.	×
🛚 🆓 Jobs	Name 1	Platform	Status	Number of VMs
籀 Replication	ANF_AVS_FP01	VMware	Completed	20
■ E Replicas	(e	Start		
Ready	6	Start to		
Active (4)	CE CE	S Undo		
Failover Plans	h.	II Statistics		
I Last 24 Hours		Dalata	-	
Success		Edit		

Name: ANF_AVS_FP01 Restore type: Failover Plan nitiated by: VEEAMBKPSRV0	5\Admi	Status: In progress Start time: 8/9/2023 3:37:41 AM nistrator <u>Cancel</u>	restore task
/M name Status	^	Log	
TestVeeam21 🕑 Proce		Message	Duration ^
TestVeeam23 🜔 Proce		🙄 Waiting 2 sec before the next VM	0:00:02
TestVeeam24 🕑 Proce		Processing VM: TestVeeam22	0:00:13
TestVeeam22 🕑 Proce		Waiting 2 sec before the next VM	0:00:02
WinSQL401 () Proce		Processing VM: WinSQL401	0:00:10
WinSQL405 (D) Proce		Waiting 2 sec before the next VM	0:00:02
TestVeeam25 🜔 Proce		Processing VM: WinSQL405	0:00:08
TestVeeam26 🕑 Proce		Waiting 2 sec before the next VM	0:00:02
🗇 TestVeeam27 🜔 Pendi		Processing VM: TestVeeam25	0:00:06
🔁 TestVeeam28 🜔 Pendi		Waiting 2 sec before the next VM	0:00:02
TestVeeam29 🜔 Pendi		Processing VM: TestVeeam26	0:00:04
TestVeeam30 🜔 Pendi		Waiting 2 sec before the next VM	0:00:02
TestVeeam31 🜔 Pendi		Processing VM: TestVeeam27	0:00:02
TestVeeam32 🕑 Pendi		Waiting 2 sec before the next VM	0:00:02
🔁 TestVeeam33 🜔 Pendi	~		v

The state of the VM replica changes from Ready to Failover and VMs will start on the destination Azure VMware Solution (AVS) SDDC cluster / host.

· Sup	· Thinks input in	-baser i A late	landar khannit kan k 🖉 china kann hajina	to a Bings	* Øiren	- 19 k					8
+ = 0 (	0 & ++	LIN2	and the second	Sand Service States of the	ditional of						+ 0 U A D
= elaborations Q,											
		1 (A V	teart Appleat Leter	1/1		word on /	1/0	CDDC			
0 77 17 9		Pullers	ty Monte Lothers' Permanen Ben	viation Ma	vis powe	red on A	115	SUDC			
D is herbelicher sittentit in B 2000 Gelaame B Classert			a manufact and transmission of the second		at	ter failov	er				4.000
C web return and earlier	the fifth and a strength of	on a contraction of the contract	Barke 1 - Bask		ne huturiti	er dettere	-	water a	-	1	
C and a state of the	Collection (18) second of	C.	B tenend Powerth		nend BUSS	10.68	4.4				
the state party of the state	Contraction and the	0	If the work in the second to		minute 19.85.28	07229-00	4-6	221.948		1	
a th dates		0	· # Norwards Room) (h		100714 101718	1022-07100	-2-6	10149	12.04	se -	
1 - 12 familie		0	- # Technology Research		mina Millia	878.37.66	-216	127/148	11101	ME .	
A CO Mart Plan		0	- @ Interestion Research		100mg 8.0108	475.26 w0	-1-1	20140			
El Marris Hamas			- 2 Second Report In		mma #9138	070.21 493	1.4	00.00			
C arrival		0	Strategic Parents		ment would	000-07-00	2.46	22.48			
@ waterspirit		0	· 2 Telever(4 Restrict)		100 March 10 March 10	0751346	1.00	10.44			
Of Chickware		0	- 2 Testiments Street		married \$5,251.00	472-9-148	1.00	2214			
(E-n)(en)(_)+1		0	· # Indianality Provention		- monut - 16-01-08	675.68.949		11.48			
(§ 40, Poin, II		0	2 Interest Received in		1000 N 10 10	10101146	1-10				
B NAMES		0	E Submarial Parentia		- Horna - M-M-10	100,04,040	1.4	31-48			
Ch. Hardenberg a		0.	Distanti Puterilia		mond NUMBER	072.0148	1.4				
In succession of the		0	· 2 formulat moved in		mana matta	01232-00	1.4				
Of an insert, we an		0	2 Incoments		100mg 10.01.08	810.01.08	1.0				
() design and the			and an and a second sec		(1000 - 100		1.00				10.0 m 2 4 7
and the second stationers		- Render									
<ul> <li>Recent Table — Alaries</li> </ul>											
ur tere	Too Served	and a second	T i beat	T other	(T.) 2017	the big	17	Company Inte		Non Contraction of Contraction	
teres in struct has been	d weakers	Q-Companied	Dreaming Life Sea Strong Matters	1944041.00 A.10 A.10 A	ame ( ) - 14	11/10/2012 0-14-14		annengis, brief be we		an Annal San San Anna anna anna anna anna anna a	
territoria de la contra de	D WHERE	C-Louise	Recording only Minute in Antibula	ted visibility (CALIERS	1999 ( 1998	1110203-014-04		including to be ad as	-	of Ann and Annual An	
and the second sec	D weaking	(Classifier	Incodinguing (1944 Institute of perchased	ee roestication	100	2010/02/04 2020 04		streaming where as		to be the contract of the cont	
Number of the Number	B webber	() ( research	Interligency 2014 Names in Annual	na cerete cocum	100	01702503, 8:00 40-04		division in the state		scherteleinen othetti wonerge es ave om	
New 2r stud Laries	B weater	O Estamot	Proverting are the first attractive to	signation a COA Lines	atta 244	Differences a database		University \$ 40.55 Au		a before the control of the second of a second	
hereigen of all machine	-D WINDAM	O Deserved	Rearrighting Struk Review of American	ind vitretti 100,500	100 Line	21/962620, F-175, B-44		CONTRACTOR DATA	e .	a period for the child and a second statement	
time in struk harme	- worker	42 Enricement	Proverting but the rate fortuge that them	VEREILLAUNE	100	2119(303), 9-(8.3) or		100000333 \$100-45 Av		a period provide the second second second second	
europere visual machine	O weetlack	4) Cargonie	Recentparing infrue Nacione on decimation	INE VEHENELOOGUNA	249	01/02/2010 02012 04		CONTRACTOR NOTING IN CONTRACTOR OF CONTRACTO		a personale concernance as now our	
enterior a contra enteriore	@ W68405	CO Conserved	Receipting that here in beington	AND ADDREED OF ANY	alter 2.44	01/90203, 8:04:24 on		010010123 2-09-23 54		<ul> <li>Jest Mikrischer Stellen einer und ans ander sein</li> </ul>	
entropy while manage	8 W62494	Clamana	Benefitivity (modified to contractor	ve vestically	100 C	11/10/2011 \$125,01 am				a ben bis natur (de 181 activities an altre of	
Number of the market	-B WY65425	C) Conjune	Recompany (1944 Marcel or Solitation	we version of	1899) - 2796	11112203.0003.00				a permitta and a second second on any or	
And the state of t	-B WARACO	O Companie	Bearing on a barren or being an	CONTRACTOR STATE	200 549	and the second second second				A DECKED AND A DEC	
and a spinor	a merel	C Caracterios		YEP-KIR LOOA She		- TROUG BOR IS AN					
and its and ranks	in recovering	Contractions.	making in the last since therein	COMPANY OF A DAMAGE	2.00	CONTRACTOR OF THE OWNER.		the second second second		a second s	
	a reference	Concession of the local division of the loca	summing on the local birth decision					summer and the second second	-		
		Co Companyo	Transmistration of the second second	CONTRACTOR OF THE OWNER		and solution in the second second		summer have been			Activity Workson
the state of the later	a house of	Contraction of the local division of the loc	And the second se	and weekend the local		and the second s		increased in the later of	-		
and the second second	and the second second	ALCONGRAM .	second the part of the second second			**************************************		and the second s	-	of her state when the state of	

Once the failover is complete, the status of the VMs will change to "Failover".

E- Honer Teplics				11	Visian Riccop and Explicat				14 A
Tenner Falcor New Falcor Falcor Falcor	Farback to Sectoryour to Sectory Production Production T		Application Renn - Restore	Add to Falcorer Plan - Many	Projector Annos Sector				
lome	Q Type in an adjust nor	taé ita Generite few	×						
The lates	Alaria 2	30ki hiama	Type	Status	Crashon Time	Rastore Founda	O-symul Location	Faplica Location	Itation
All Servicemen	19 TestVetam22	AVS.20220522, Replet01	Rogular	Failovor	0/0/2028 5:58 AM	3	a300 vcsa05 chick:.com/(Duster05	172.20.156.2\Cluster 1	VM/ware
it d Realizes	TestVeesm23	AVS_20230522_Replob01	Regular	Failover	6/6/2023 5/52 AM	4	a300-vcsa05 ehodo.comi/Cluster05	172.30.156.2\Cluster-1	Vistware
CL Rents	Test/reem?4	AV5_20230522_Replab01	Septer	Fellover	6/6/2025-3-32 AM	3	a300-vcsa05 ehcdc.com/.Cluster05	172.30.156.2\Guster-1	Shiwer.
III Active (19)	Testivecani25	AVE 20210522 Replet01	Rogetar	Fallovat	6/6/2023 7:48 AM	4	all00-ycsa05.ehedr.som/cluster05	172.30.156.2\Chuster-1	VMware
(iii) Falser Flass	TestVecari26	AVS_20230522_Replot01	Regular	Failover	6/6/2023 10:44 AM	1	a300-vcsa05-eheds.com/Cluster05	172.30.156.2\Guster 1	VMware
Cit Let 24 Marca	TextVeeem27	AV5_20230522_Replob01	Regular	Fallover	6/6/2023 12:19 PM	3	a300-vesa05.ebcdc.coml@uster05	172.30.156.2\Guster-1	V5-twore
(Te Surreys	Tan/Vasiarin28	AV5 20210522 Replote1	Reputer	Falevat	6/6/7023 1.07 PM	3	a509-wise05-ehido.com/Dister03	172.30.156.2\flusher-1	Where
C Whomas	Testveesin20	AV5_20220522_Replo601	Rogular	Failovet	6/6/2023 1:16 PM	1	a300-vcsa05 ehodo com/Cluster05	172.30.156.2\Ouster-1	Whyare
C. Andready	TestVeeam00	AV\$_20230522_Replot01	Beguler:	Tellover	6/6/2023 2:24 PM	1	a300-vcsa05.ehodo.com/iCluster05	172.30.156.2\Quster-1	VMvpre
	Tent/Verentil1	AV5_20230522_Replot01	Reputer	Fallows	7/31/2023-4-43-AM	3	a500-vicsa05 eticik convUDuiter01	172.30.156.25Gaster-1	Videore
	Tectivicanit2	AVE JUJ10522, Replote1	Rogular	Ealover	6/6/2028 #31 PM	1	ali00-vcsi05 ehode confiduator03	172.20.155.2\duster 1	VMaaro
	TestVeesm23	AV\$ 20230522 Replot01	Regular	Failover	6/6/2023 3:31 PM	4	a300-vcsa05.ebcdc.com/Cluster05	172.20.155.2\Chuster-1	19-twore
	TestVeeem34	AV5_20230522_Replot01	Repular	Failover	6/6/2023 4.31 PM	4	a300-vcsa05.ekcdc.comt/Quster05	172.30.156.2\Cluster-1	Whene
	Testvesen13	Av5 20250522 Republic1	Regular	Failbest	6/6/2028 5-30 PM	8	add0 vicsa05 effects com (Distar03	172.30.156.2\Gustin-1	VMmark
	WinSQL401	AV5,20230522_Replat01	Regular	Failover	6/6/2023 6-52 AM	5	a300-vcsa05 ehodc.com/Cluster05	172.30.156.2\Cluster-1	VMware
	WASOL NO2	AV5_20230522_Replot01	Segular	Failover	6/7/2023 6:11 AM	3	a300-vcsa05 ehcdc.com/Guster05	172.30.156.2\Guster-1	VMware
lione	WVSC/401	AVS 20210522 Replicid 1	Septiar	Failows	6/7/2028 12:28 454	1	a200-example of the constitue of the	172.30.156.2\Cluster-1	Whenew
	Win906404	AVS 20230522 Replot01	Regular	Fallover	6/6/2023 6:29 PM	3	a300 visso05 shode com//Ouster05	172.20.156.2\Guster-1	staware
E Inventory	WinSOL405	AV\$_20230522_Replot01	Regular	Fallover	6/6/2023 7:50 AM	3	a300-vcsa05 ehcdc.com//Quster05	172.30.156.21Cluster-1	Shhware
	1		11.53.53			- 36			
2 Dackup left act ucture									
an Storage Infrastructure									
N THE REAL PROPERTY AND									
tape initiaequelure									
Cft inter									
-									
De									



Veeam Backup & Replication stops all replication activities for the source VM until its replica is returned to the Ready state.

For detailed information about failover plans, refer Failover Plans.

When the failover plan is running, it is considered as an intermediate step and needs to be finalized based on the requirement. The options include the following:

 Failback to production - switch back to the original VM and transfer all changes that took place while the VM replica was running to the original VM.



When you perform failback, changes are only transferred but not published. Choose **Commit failback** (once the original VM is confirmed to work as expected) or Undo failback to get back to the VM replica If the original VM is not working as expected.

- **Undo failover** switch back to the original VM and discard all changes made to the VM replica while it was running.
- **Permanent Failover** permanently switch from the original VM to a VM replica and use this replica as the original VM.

In this demo, Failback to production was chosen. Failback to the original VM was selected during the Destination step of the wizard and "Power on VM after restoring" check box was enabled.



Perce	eduation illiack Mode mmary	Name References				
eduation inhack Mute inhack	edination illiack Mode eremany	E-Stretteen3	Six	Crisical Investiga		Select 2
Allack Mude ommary allack Mude ommary allack Mude ammary allack Mude a	illack Mode	COLUMN ADD ADD ADD ADD	11110	hall would shake a	TIM	
Interference of the second se	remary	Contraction 71	00.7 AM	fallfournal about o	ami (T	Clear A
Primary Prima	ermany.	Contractioner 22	754149	lable-result atests of	emi IT_	Popula
Summay       Summay         Pictor       Summay         Summay       Summay         Pictor       Summay         Pictor       Summay         Summay       Summay         Pictor       Summay <td< td=""><td></td><td>El Territeranito</td><td>01.0 Mg</td><td>(wild-vesal5.ehodc.o</td><td>am][T_</td><td></td></td<>		El Territeranito	01.0 Mg	(wild-vesal5.ehodc.o	am][T_	
Sommary         Pick         Summary         Mode Control         Mark         Set Mark         Mare         Mare <td></td> <td>E Testieum 12</td> <td>857 MB</td> <td>(#300-vesatilis ehode o</td> <td>om]{7</td> <td></td>		E Testieum 12	857 MB	(#300-vesatilis ehode o	om]{7	
Summary     Fastimeenili     BL488     [A000-cca85.abcdx.com] [T		E Windows	15.68	(al00-vesal5.ehedc.o	am] (#5	
Sammary       Pics       Summary       Pics		12 Testimon 25	\$1,2 MB	(a00-scalif-shotic o	am] [T_	
Beck     Summary       Pice     Summary		Tertiesende	\$1.7 MB	(a)00-vesab5.aheste.o	am]][T	
Sommary     Next Set Sector (PL)       Sommary     Bit No       Sommary     Bit No       Statemary:     Next Set Sector (PL)       Statemary:     Bit No       Statemary:     Bit No       Market Notice     Market Notice		E Testiesamil	96.43/8	(a300-vesal)5.ehodc.o	am]]7_	
back     Summary       Seconary       Rower failback settings, and click Freich to start failback agreention. Now will be able to under failback process if regulations.       Back       Seconary       Rower failback settings, and click Freich to start failback agreention. Now will be able to under failback process if regulations.       Back       Seconary       Rower failback settings, and click Freich to start failback agreention. Now will be able to under failback process if regulations.       Back       Seconary       Min name. TestTiveemID       Wind Cliff       Wind name. TestTiveemID       Wind name. TestTiveemID <td></td> <td>MindQL403</td> <td>1.6.08</td> <td>(x000-veza05.ehutic.c</td> <td>am]]P%_</td> <td></td>		MindQL403	1.6.08	(x000-veza05.ehutic.c	am]]P%_	
Pice     IRLING		boos lettiesamili	81116	JASEP VISAES ethelic o	om){1	
Sommary     Sommary       Pice     Sommary       Roke     Sommary       Bit Americ     Test/insem121       Marce     Test/insem121       Marce     Test/insem121       Marce     Test/insem121       Marce     Test/insem121		C2 Chryston and	10.00	patrie-results anote o	am) (Pr	
Summary     Mass -       Back     Summary       Sommary     Mass -       Mode     Mass -       Mark     Series       Series     Series       Series		C. Contraction and	11.08	Jalifformali abole o	omj (Pa.	
Summary       71.8 M0       (a)20		Plan Testimon 3	81.2 MB	(a305-sesal), about o	ami IT.	
		El Tettiseen27	77.6 MB	(a)00-veset5 ahods o	om)(T_	
		13 a Testiveanity	\$14 MB	(a300-vezal/5-ehode.o	-T1[ma	
pica Summary: Minane TettVeen20 bick Mode VM name TettVeen20 WM name TettVeen20 WM name TettVeen20 VM name TettVeen20	Review failbac	a settings, and click Finish to start failba	ch operation. You v	vill be able to undo fail	tack process #	Third
PLS     P		Summer .				
shination Back Mode Wit name: TertYexam23 Wit name: TertYexam23 Wit name: TertYexam23 Wit name: TertYexam32 Wit name: Wick28.402 Wit name: TertYexam33	paca	Ald same Tellingen 7				
Back Mode VM name TetTissam23 VM name TetTissam22 VM name TetTissam22 VM name Visit20.422 VM name TetTissam34 VM name TetTissam34 VM name TetTissam31	stination					
Back Mode VM name: TettVesen23 HTMIY VM name: TettVesen32 VM name: VettVesen32 VM name: VettVesen32 VM name: TettVesen34 VM name: TettVesen34 VM name: TettVesen31		MA Auron Techiesen 21				
Inmany UM name: TestTineam30 VM name: TestTineam32 VM name: UBICS.A02 VM name: TestTineam35 VM name: TestTineam34 VM name: TestTineam31	Aback Mode	VM name TestVesen22				
VM name: Testilasamitz VM name: WinSCS.AC VM name: Testilasamit3 VM name: Testilasamit4 VM name: Testilasamit4	200707	and and a second s				
VM name: Tertileanni2 VM name: Wick28.40 VM name: Tertileanni25 VM name: Tertileanni34 VM name: Tertileanni31	onnous a second	- Pro Autre representati				
VM name: WinSCB.AD VM name: TestVesem21 VM name: TestVesem34 VM name: TestVesem31	on many 2	1954 name: Testilessen12				
VM name TestVesen25 VM name TestVesen34 VM name TestVesen31	mmey /					
VM name: Tettissen23 VM name: Tettissen34 VM name: Tettissen31	mmary .	Minama Mini Ch. 87				
VM.name.Testimam34 VMI.name.Testimam31	niney /	VM name WinSQL422				
(Vid name: Test/Insem31	innay ;	VM name: WindCLAD2 VM name: TestVesam23				
(Vit name retriesens)	ninaey .	154 name Wold2,402 VM name Testileam25 VM name Testileam34				
	minary .	HM name WolCLAC VM name Testiveam21 VM name Testiveam34				
Whene Wir52,40	array .	MA name: Work28.422 MA name: TestVoaan23 MA name: TestVoaan34 MA name: TestVoaan31				
	nnay.	VM name: Wold2842 VM name: Tectivean23 VM name: Tectivean34 VM name: Tectivean31 VM name: Wol22403				
[7] Proof on heard this day participan	mmay	VM name: Wold2842 VM name: TestVeam25 VM name: TestVeam34 VM name: TestVeam31 VM name: Wold2443				
C Prose on local Mildle sections	mmay	VM name: Wolc2,452 VM name: TestVeam25 VM name: TestVeam34 VM name: TestVeam31 VM name: Wolc2,453				

Failback commit is one of the ways to finalize failback operation. When failback is committed, it confirms that the changes sent to the VM which is failed back (the production VM) are working as expected. After the commit operation, Veeam Backup & Replication resumes replication activities for the production VM.

For detailed information about the failback process, refer Veeam documentation for Failover and Failback for replication.

Image:         Image:<	a diama Tapita					Same Print	and Replication						5
Constraints     Constrain			States Tables	and a state	MATRA TRANSPORT								
Odd         Tame         Table         Ta	+	Q but in an eligibility	arms by page do for	×									
Operation         Operation <t< th=""><th></th><th>Hame</th><th>July Name</th><th>Time</th><th>Status 7</th><th>Creation Time</th><th>Tamos Pairm</th><th>Original lacation</th><th>Replica Location</th><th>Patient</th><th></th><th></th><th></th></t<>		Hame	July Name	Time	Status 7	Creation Time	Tamos Pairm	Original lacation	Replica Location	Patient			
Replan         Contraction         Product WC2202322, April Margin         April Mark         Ref 2010 101 2011 002 and 101 more Contraction         Product WC220322, April Mark         April Mark         Ref 2010 101 2011 002 and 101 more Contraction         Product WC220322, April Mark         April Mark         Ref 2010 101 2011 002 and 101 more Contraction         Product WC22032000 000 more Contraction         Product WC2203200 000 more Contraction         Product WC22030000 more Contraction	Illi Bastrator	FG TerrivesordA	AV5_20332522_Rapin67	Septer	Faiture	BE1/25251014 XM	1	a300-ica01ahusi spiri-Ourier03	172.50.138.2\Dume-1	Whene			
Sector         Operation         O	faction	Tanti/assars23	W/5_20230522_Feetine01	Aspile:	Fallinik	BITS 2023 YESTZ AM		SherrorD.month.incele 20aper 000a	172.20.126.25Oumer I	Mage			
Anter (15)         (a) Performent         (b) Performent         (c) Perform	The Reads	Test/www.02	AV5_20232522_AuplinD1	Arpolar	Fallinth	8/1/25253013 AM	4	aboo-icia05ahosk.com/Chatar05	172.20.154.2-Cume-1	(Make)			
Beller Mere         Community         Follower/d         Add/Sci200212, Separation         Name         Environment         Environment         TT2.55 (18.4 2 Counter)         Weisser           Link Let Kimmer         Territower/d         Add/Sci200212, Separation         Name         Environment         Environment         TT2.55 (18.4 2 Counter)         Weisser           Science         Territower/d         Add/Sci200222, Separation         Name         Environment         Environment         TT2.55 (18.4 2 Counter)         Weisser           Weisser         Territower/d         Add/Sci200222, Separation         Name         Environment         Environment         TT2.55 (18.4 2 Counter)         Weisser           Weisser         Add/Sci200222, Separation         Name         Environment         Environment         TT2.55 (18.4 2 Counter)         Weisser           Weisser         Add/Sci20022, Separation         Name         Environment         Environment         TT2.55 (18.2 Counter)         Weisser           C         Territower/d         Add/Sci20022, Separation         Name         Environment         Environment         TT2.55 (18.2 Counter)         Weisser           C         Territower/d         Add/Sci20022, Separation         Name         Environment         Environment         TT2.55 (18.2 Counter)         Weisser	Active (19)	Test//exerc2?	Av5_20200522_Repint(1)	Septer	Autorix	8/1/2022 10:18 AM	14	ablo-scale interaction (Seaso-Oblia	172.20.136.2\Dume 1	West			
Line & Horsen         Implementation         Article 2002 (2000)         Register 2000         Register 2000 <thregister 2000<="" th="">         Register 2000         Re</thregister>	2 Failo-er Plant	Hig TertVeneni24	AV520230522_fepixed1	Angelar	Falbaik	B/1/2023 1017 AM	- 14	10mm/Orms intel Same Olda	172.85.156.25Outer 1	Mean			
Openanie	Lett 24 Hours	Tex/Veceni28	AV5_20130523_Aepix801	Angelar	Falback	8/1/2023 10/14 AM	4	400-usa01.etoal.com/Durad5	172.33.136.2 Durn+1	Were			
Operation         Operation <t< td=""><td>(Se Surgers</td><td>(3) Sentimental 1</td><td>AVS_20230523_Rep30601</td><td>Repda</td><td>Fallack</td><td>8/1/2023 10:21 AM</td><td>- 4</td><td>60eeu-O/mus.siste.f6ezor-00ta</td><td>172.30.136.21/Duster-1</td><td>Moan</td><td>De.</td><td>Falterny lotter.</td><td></td></t<>	(Se Surgers	(3) Sentimental 1	AVS_20230523_Rep30601	Repda	Fallack	8/1/2023 10:21 AM	- 4	60eeu-O/mus.siste.f6ezor-00ta	172.30.136.21/Duster-1	Moan	De.	Falterny lotter.	
Op/Entries         Op/Entries         An1/2021/1014/A         Source/Link	( Werning	WWSGLACE	AV5_20230522_Repixe01	Rep/ier	Faibark	#/1/2023 XB/17 AM	14	a000-scalit ehods.asmiChumer11	0230.158.24Dune-1	UMage.	10×.	Panned talmen.	
Colstance-Id         Adv. 2023/0222, Registriol         Register         Ref/100021         Filter         Ref/100021         Filter         Adv. 2023/2023         Register         Adv. 2023/2023         Ref/100021         Ref/100021 <thref 100021<="" th=""> <thref 100021<="" th=""></thref></thref>	C Tales	Cil Testiveserri25	Avg_20130522_Replat01	Repoler	Palbalk	8/1/2023 101# AM	1	60emuQ/mess store control of the	172.30.194.2:Outlet-1	April	22	Add to failover plan	•
Construction         Av12,022322,24paint         Register         Av12,02232,151,44         Av12,02232,24paint         Register         Av12,0223,022,04paint         Register         Av12,0223,022,04paint         Register         Av12,0223,022,04paint         Register         Av12,0223,022,04paint         Register         Av12,022,022,04paint         Register         Av12,022,022,04paint         Register         Av12,022,022,04paint         Register         Av12,022,022,04paint         Register         Av12,022,022,04paint         Register         Revise Av12,022,04paint         Revise Av12,022,04paint <threvise av12,022,04pa<="" td=""><td></td><td>Ci Testiveeen2.6</td><td>4V5_20230522_Feptide01</td><td>Pepiler</td><td>Feitherk</td><td>E/1/2023 Y0/29 AM</td><td>4</td><td>elite scall study cost.Quand5</td><td>172.20.154.25Champ-1</td><td>White</td><td>10</td><td>fuller's to production .</td><td></td></threvise>		Ci Testiveeen2.6	4V5_20230522_Feptide01	Pepiler	Feitherk	E/1/2023 Y0/29 AM	4	elite scall study cost.Quand5	172.20.154.25Champ-1	White	10	fuller's to production .	
Conferences AV_20120222_20ealul Region Refere AV_2022110114.04 4 AV00-codd actual anti-anti-anti-anti-anti-anti-anti-anti-		Cil Testiverani)0	AV5_20130522_Reptiled (	Replet	Failtain	A/1/2023 10 13 AM	-	able-scale) and a service and changes	172.30 []#.2\Oyater [	Village.	- 54	Corrorat failback	
Construction of the second se		C4 Techvesers29	AV1_20132522_Report1	Septor .	Pattern	8/1/22/28 10:18 AM	11	allog-coatt anus, com/closed/5	172.30,154.25Oume-1	(1)7+FT	5	Underfailleach	
Conference 2 AV2/20131222 Avanues Avance and Avance Avanc		() Tev//essen13	WV0_20230522_Reprinter	Arpiw	Failback	R/1/2023 TR/15 AM	10	a100-scsa01 attents contributients	172.30.134.25Quiner I.	Wege		and its states	-
Bin wold Adv. 2012 Japanio Magne Magne Fallow 71/02/2013 Vol2 Adv All Once State County Volant II Advance Advance Volant II Advance Advance Volant II Advance Adva		Of remember	AV1_20130322_Maph001	deport.	( ALL ALL ALL ALL ALL ALL ALL ALL ALL AL	ACT/2013 TO 19 AM		alto-scapt anode tam charact	172 20 198 210 400001		19	Return govid files	1
a metodeta antigene regione regione regione regione regione a sub-regione regione regi		The interaction of the	AND DESCRIPTION AND ADDRESS	Magazar	Parties.	Retronte Sauce And	-	entry which and a second second	ATT IN THE PLANET	UNTERN.	15	Names from configuration	
Distance and Annual		Distantia atta	AND DOLDONE MEMORY	- Andrew	Tangoar	ALCONDED TO DE ANN	1	and search and an in the second	172 20 126 2 Water 1	Total and	14	Debeta from due	
		The second and	Ave. 200 00 00 00 00000	and the second	faire an	1/11/10/11 2/14 MM		allow status decisi, contracting a	VITA SALESA MALANCE		100	Accession in the second se	
		Car menadore	1.4 1010000 Augusta	report	- and the	1000 (0000 4 400 PM	12	and weaks draw composition	112 20-126 2 000000-1			TOP TOP TO	
Wold(2404 AVC_20201012_feptime11 Regular Fallowr 7/11/20214.06.PM 6 e800-csal3 and card(Card(Card(Card(Card(Card(Card(Card(C		WHATLADA	Av6_20230522_Argiliti01	Reputer	Falover	7/\$1/2023 4:06 PM		s100-cos03 atop: com/clume03	172.30.134.21Dunier-1	shiere	100		
	Press.												
	a Tatica Mamutan												
Janes	1 Damage inflationation												
Tennes       Speaking       Laskag inflammednes       Speaking inflammednes													
Paramati ( )       Streaming ( )       Backage inflammation ( )       Disrege inflammation ( )       There inflammatic ( )       There inflammatic ( )       There inflammatic ( )	Two inflationships												

After failback to production is successful, the VMs are all restored back to the original production site.

								<u> </u>	
	C VeeamTest in	1068							
80 60 69	Summary Monitor Con	four Permission Resour	a Pools - VMs -						
9 a000-acu/01-ahrst.com		and I see 1							
18 A300 OCC5	Without Machines With Tarr	TITLE CONTRACT							
<ul> <li>ID: QuarterD9.</li> </ul>								- 10 m	
a300-exx09 efcdc.com	2011 Land	Contractor	1. Weber	1. million million	1 Viet from	1 mile real	1 Intelligence 1		
alloo existo encarcismi	Child And Andrews	Descent Ab	A Report	10.00	NAME & ADD	A last	474 140		
<ul> <li>i (3) HOBerth_0001</li> </ul>	Child At Summer	Downey Co.	- Anna -	14.5.8	und lief Anti-	1000	1011408		
<ul> <li>O HOBERS_0002</li> </ul>		Prompt and Col.		10.00	1000 000 PM	10.04	24240		
> Q HOBerch_0503	Li in reconnuts	POWERQ OF		N GB	1111111	- Q PE	241 140		
<ul> <li>O VesamTest</li> </ul>		-owered on	·	10.00	101414 1411	Ung .	200 Ma		
(if Textonianizt	D I IP HELWERTO	Powersed Cos	V. 942718	NO GE	890.21 Mg	014	STL ME		
/唐 TextVelam22	C I ID TellWearDC	Plimerad Dh	V North	10 GB	876.8.145	0.46	226,140		
62 Performant23	C = CT Tectybeam37	Powersd On	- Normal	16.08	874.45 MB	.010	242.949		
(3 Testivear:34	C + (2 Testweet2)	Powered On	V Nortal	16.08	875.12 HB	0.40	202 MB		
GE Textileon/25	C (# testWeam20	Proversid Dis	- Named	16.08	873 54 MB	0.H2	340 MS		
(2) Technolecula	A technemical	Powered On	V Nortal	16 GB	87129 MB	0.142	300 MB		
10 Termeanum	🗋 = 🖨 Testyleartill	Powersd On	V Normal	16:08	882,92.949	-0 HQ	343.66		
if Testviesamili	C / C testVerant2	Provend OI	<ul> <li>Normal</li> </ul>	10.08	675.34 xell	0.H2	333 MB		
(2) Textybeam20	C A 25 terlyleards	Powered On	V Normal	10 GB	872 07 ME	0.40	337 MB		
(2 performantal)	C / // test/reamile	Powered On	V. Normal	16.08	\$72.39 Mill	0.142	300 HB		
(B Testoieeentit	C / A terrylecamob	Powered (M	v North	16:08	877.62 MB	0.00	228.149		
of Textmeaning	C + O WINSCRADI	Provenation:	- Nortal	308 22 58	20.32.08	ONE	08		
(Z. Testimearit)	C) A (\$ WYDOLAGO	Poweres Off	- Normal	308.22.08	20,63.08	0.0162	.08		
(2 Termineam)4	C) (2 Wester 403	Powered On	V Nortal	300.01.68	20.51 GB	851410	2.04.00		
(p terminearcil)	CT (1) (D YERSON ADA	Powered Off	- tarma	318.21.08	224266	0.40	0.8		
(IP WHIDGE 400	() : () WING AT	Provinati OT	V Nacimia	208.22.08	24.24.00	010	0.0		
(B) WebGL402	C + (2 1000 C)	Powered Cri	No. No.	90.08	19 9 O.E	214 346-02	W 100 (c)		
CP WHIGH 400	Gi - G - Martin								
(B) MARKAGARAN									
(B. ANAROCHOD									
CP WHISE HIS									
da pecterente									
di nantanin									
18 interested in									
Characterization from									
18 Landberry Die 1978									
18 1 10 10 10 10 10 10									
78 Landshamm	IT COMPANY							Bandra and All	
and the second second second	CIT IN COORD							and the field	1

# Conclusion

Azure NetApp Files datastore capability enables Veeam or any validated third-party tool to provide a low-cost DR solution by leveraging Pilot light clusters instead of standing up a large cluster only to accommodate VM replicas. This provides an efficacious way to handle a tailored, customized disaster recovery plan and to reuse existing backup products in house for DR, enabling cloud-based disaster recovery by exiting on-premises DR datacenters. It is possible to failover by clicking a button in case of disaster or to failover automatically if a

disaster occurs.

To learn more about this process, feel free to follow the detailed walkthrough video.

https://netapp.hosted.panopto.com/Panopto/Pages/Embed.aspx?id=2855e0d5-97e7-430f-944a-b061015e9278

# Migrating Workloads on Azure / AVS

## TR-4940: Migrate workloads to Azure NetApp Files datastore using VMware HCX - Quickstart guide

One of the most common use cases for the Azure VMware Solution and Azure NetApp Files datastore is the migration of VMware workloads. VMware HCX is a preferred option and provides various migration mechanisms to move on-premises virtual machines (VMs) and its data to Azure NetApp Files datastores.

Author(s): NetApp Solutions Engineering

# Overview: Migrating virtual machines with VMware HCX, Azure NetApp Files datastores, and Azure VMware solution

VMware HCX is primarily a migration platform that is designed to simplify application migration, workload rebalancing, and even business continuity across clouds. It is included as part of Azure VMware Solution Private Cloud and offers many ways to migrate workloads and can be used for disaster recovery (DR) operations.

This document provides step-by-step guidance for provisioning Azure NetApp Files datastore followed by downloading, deploying, and configuring VMware HCX, including all its main components in on-premises and the Azure VMware Solution side including Interconnect, Network Extension, and WAN optimization for enabling various VM migration mechanisms.



VMware HCX works with any datastore type as the migration is at the VM level. Hence this document is applicable to existing NetApp customers and non-NetApp customers who are planning to deploy Azure NetApp Files with Azure VMware Solution for a cost-effective VMware cloud deployment.

## **High-level steps**

This list provides the high-level steps necessary to install and configure HCX Cloud Manager on the Azure cloud side and install HCX Connector on-premises:

- 1. Install HCX through the Azure portal.
- 2. Download and deploy the HCX Connector Open Virtualization Appliance (OVA) installer in the onpremises VMware vCenter Server.
- 3. Activate HCX with the license key.
- 4. Pair the on-premises VMware HCX Connector with Azure VMware Solution HCX Cloud Manager.
- 5. Configure the network profile, compute profile, and service mesh.
- 6. (Optional) Perform network extension to avoid re-IP during migrations.
- 7. Validate the appliance status and ensure that migration is possible.
- 8. Migrate the VM workloads.

## Prerequisites

Before you begin, make sure the following prerequisites are met. For more information, see this link. After the prerequisites, including connectivity, are in place, configure and activate HCX by generating the license key from the Azure VMware Solution portal. After the OVA installer is downloaded, proceed with the installation process as described below.



HCX advanced is the default option and VMware HCX Enterprise edition is also available through a support ticket and supported at no additional cost.

- Use an existing Azure VMware solution software-defined data center (SDDC) or create a private cloud by using this NetApp link or this Microsoft link.
- Migration of VMs and associated data from the on-premises VMware vSphere- enabled data center requires network connectivity from the data center to the SDDC environment. Before migrating workloads, set up a site-to-site VPN or Express route global reach connection between the on-premises environment and the respective private cloud.
- The network path from on-premises VMware vCenter Server environment to the Azure VMware Solution private cloud must support the migration of VMs by using vMotion.
- Make sure the required firewall rules and ports are allowed for vMotion traffic between the onpremises vCenter Server and SDDC vCenter. On the private cloud, routing on the vMotion network is configured by default.
- Azure NetApp Files NFS volume should be mounted as a datastore in Azure VMware Solution. Follow the steps detailed in this link to attach Azure NetApp Files datastores to Azure VMware Solutions hosts.

## **High Level Architecture**



# **Solution Deployment**

Follow the series of steps to complete the deployment of this solution:

To perform the installation, complete the following steps:

- 1. Log in to the Azure Portal and access the Azure VMware Solution private cloud.
- 2. Select the appropriate private cloud and access Add-ons. This can be done by navigating to **Manage** > **Add-ons**.
- 3. In the HCX Workload Mobility section, click Get Started.



4. Select the I Agree with Terms and Conditions option and click Enable and Deploy.



The default deployment is HCX Advanced. Open a support request to enable the Enterprise edition.



The deployment takes approximately 25 to 30 minutes.



For the on-premises Connector to connect to the HCX Manager in Azure VMware Solution, make sure the appropriate firewall ports are open in the on-premises environment.

To download and install HCX Connector in the on-premises vCenter Server, complete the following steps:

From the Azure portal, go to the Azure VMware Solution, select the private cloud, and select Manage > Add-ons > Migration using HCX and copy the HCX Cloud Manager portal to download the OVA file.



Use the default CloudAdmin user credentials to access the HCX portal.

Home > Asure VMware Sol          Appendixed could             Appendixed could                  Create @ Manage view @ ****                  Create @ Manage view @ ****                  Create @ Manage view @ ****                 Create @ Manage view @ ****                 Create @ Manage view @ ****                 Create @ Manage view @ ****                 Create @ Manage view @ ****                 Create @ Manage view @ ****                 Create @ Manage view @ ****                 Create @ Manage view @ ****                 Create @ Manage view @ ****                 Create @ Manage view @ ****                 Mane ?                 Avide                 Create @ Manage view @ ****                 Create @ Manage view @ ****                 Create @ Manage view @ ***********************************	≡ Microsoft Azure	arch resources, services, and docs (G+/)	📄 🖸 🖗 🖉 🚳 Ø R	niyaz@netapp.com
	Microsoft Axure Solution > A       Home > Azure VMware Solution > A       Azure VMware Solution > A       Hybrid Cloud TME       + Create ③ Manage view < ····       Filter for any field_       Name 1            • ANF       • AVS	arch resources, services, and docs (G+/)                 ANFDataClus   Add-ons	tion migration, workload rebalancing, and noce (OVA file) from Adminstration page and more.	niyat@netapp.com

2. After you access the HCX portal with cloudadmin@vsphere.local using the jumphost, navigate to Administration > System Updates and click Request Download Link.



Either download or copy the link to the OVA and paste it into a browser to begin the download process of the VMware HCX Connector OVA file to deploy on the onpremises vCenter Server.

vm VMware HCX	ent ( renter and damaged) (m.	Con-Scienced & Constraint Constraint of Constraints of	and the second has				c	ଲି* =*	cloudadmir
Dashboard      Infrastructure     C Site Paining     Sinterconnect      Services      Compute     Network Extension     Migration     Disaster Recovery	System Upd Pair your remote da Resultant por Local HCX Glack Resultant	ates							
Administration     Administration     System Updates     Troubleshooting     Audit Logs	Current Version	System Name	vs.azure.co.	Status	info	System Type HCX Ooud	NSX Version 7 312.0.017883600	VC Version v 7.0.3.19234570	Copy To Clipboard
Activity Logs     DICE     Support	7 22							Numb	er of Applanc
	Remote HCX	System Name	τ Status	info		System Type	· *	Copy To Clipboard	
				8					
								Number	of Appliances

3. After the OVA is downloaded, deploy it on to the on-premises VMware vSphere environment by using the **Deploy OVF Template** option.

VSphere Client			
	* El A300-Cluste	erOI Account	
	Deploy OVF Template	Select an OVF template ×	
- B ADD-ONISCH	1 Select an OVF template	Enter a URL to download and install the OVF package from the internet, or browse to a location accessible from your computer, such as a locat hard drive, a network share, or a CD/DVD drive.	The.
0 4300-es 0 4300-es	<ol> <li>Select a name and folder</li> <li>Select a compute resource</li> </ol>	New Collection Press Dealer and and a Collection of Collec	
© #100-00 → © #1#2_Hc	4 (Perview details.)	Local file     UPLOAD FILES     VMwareHCX-Connector-4:     owa	
1 G ANFON 1 G AVSAN	<ol> <li>Select storage</li> <li>Ready to complete</li> </ol>		
- G Houpe B Hol			
(3 HO). (3 HO).			
(g) HOL (g) HOL			
B HOL			
la nou	Photon 34		* 1 1 1 1 2 2

4. Enter all the required information for the OVA deployment, click **Next**, and then click **Finish** to deploy the VMware HCX connector OVA.



Power on the virtual appliance manually.

For step-by-step instructions, see the VMware HCX User Guide.

After you deploy the VMware HCX Connector OVA on-premises and start the appliance, complete the following steps to activate HCX Connector. Generate the license key from the Azure VMware Solution portal and activate it in VMware HCX Manager.

- 1. From the Azure portal, go to the Azure VMware Solution, select the private cloud, and select **Manage** > Add-ons > Migration using HCX.
- 2. Under Connect with on-premise Using HCX keys, click Add and copy the activation key.





A separate key is required for each on-premises HCX Connector that is deployed.

3. Log into the on-premises VMware HCX Manager at "https://hcxmanagerIP:9443" using administrator credentials.



Use the password defined during the OVA deployment.

4. In the licensing, enter the key copied from step 3 and click Activate.



The on-premises HCX Connector should have internet access.

- 5. Under **Datacenter Location**, provide the nearest location for installing the VMware HCX Manager onpremises. Click **Continue**.
- 6. Under System Name, update the name and click Continue.
- 7. Click Yes, Continue.
- 8. Under **Connect your vCenter**, provide the fully qualified domain name (FQDN) or IP address of vCenter Server and the appropriate credentials and click **Continue**.



Use the FQDN to avoid connectivity issues later.

9. Under Configure SSO/PSC, provide the Platform Services Controller's FQDN or IP address and click

# Continue.



Enter the VMware vCenter Server FQDN or IP address.

- 10. Verify that the information entered is correct and click Restart.
- 11. After the services restart, vCenter Server is displayed as green on the page that appears. Both vCenter Server and SSO must have the appropriate configuration parameters, which should be the same as the previous page.



This process should take approximately 10 to 20 minutes and for the plug-in to be added to the vCenter Server.

m HCX Manager	Dashboard	Appliance Summary	Configuration	Administration			172.21.254.157	Version: 4.410 Type : Connector	admi
VMware-HCX     FODN:     IP Address:     Version:     Uptime:     Current Time:	-440 VMware-HCX-440 172.2 4.4.1.0 20 days, 21 hours Tuesday, 13 Septe	).ehcdc.com , 9 minutes mber 2022 07:44:11 PM UT	c		0 0	CPU Used 1407 MHZ Memory Used 9691 MB Storage Used 29G		Free 688 MHZ Capacity 2095 MHZ Free 2316 MB Capacity 12008 MB Free 98G Capacity 1276	679 819 239
NSX			vCenter https://a300-vc	ia01.ehcdc.com	• h	ISO https://a300-vcsa01.ehc	dc.com		
MANAGE			MANAGE			ANAGE			

## Step 4: Pair on-premises VMware HCX Connector with Azure VMware Solution HCX Cloud Manager

After HCX Connector is installed in both on-premises and Azure VMware Solution, configure the onpremises VMware HCX Connector for Azure VMware Solution private cloud by adding the pairing. To configure the site pairing, complete the following steps:

 To create a site pair between the on-premises vCenter environment and Azure VMware Solution SDDC, log in to the on-premises vCenter Server and access the new HCX vSphere Web Client plugin.

ihortcuts									1 84 1-24			
oventories										2		1
([]]	ē.		Ø	11	8	000		6		U.	۲	
Hosts and Clusters	VMs and Templates	Storage	Networking	Content Libraries	Global Inventory Lists	Workload Management	SnapCenter Plug-in for VMware VSphere	Cloud Provider Migration		Site Recovery	нсх	1
Ionitoring												
會		æ	8	32	$\diamond$							
Task Console	Event Console	VM Costomization Specifications	VM Storage Policies	Host Profiles	Lifecycle Manager	ONTAP tools						
dministratic	n											
Q												
-												

1. Under Infrastructure, click Add a Site Pairing.



Enter the Azure VMware Solution HCX Cloud Manager URL or IP address and the credentials for CloudAdmin role for accessing the private cloud.

$\epsilon \rightarrow \sigma$	0 & ≓	• • https://a300-vcsa01	ehcdc.com/ui/app/blugir	r/com.vmware.hybridity/com.vmware.hor	sitePairing		☆	9 8
vSphere Client ()								9 0
HCX © Destrocans Intrastructure ClistResconnect Tomport Analytics Services © Notinivis Extension © Migration © Disaster litecowry System © Administration © Secont	« «	Site Pairing	Connect to R Remote HCK URL Username Password	Remote Site https://72. cloudsdmin@vsphere.local	Сомиест	9V3.43300.000-		TE DANNIG

1. Click Connect.

i.

VMware HCX Connector must be able to route to HCX Cloud Manager IP over port 443.

1. After the pairing is created, the newly configured site pairing is available on the HCX Dashboard.

	<	Site Dairi	ing				
cx 3 Dashboard hfrastructure	~	Site Pair				A SITE PAI	RING
Site Paring Shterconnect Transport Analytics ervices Network Extension Migration Dispate Decoupting	×	Ø	WMware-HCX-440 Phttps://12.21254.157.443 Raleigh Linterconnect(s)	<b>→</b>	<pre> where the state of the st</pre>		
vstem	~	EDIT CO	NNECTION DISCONNECT				
Administration Support		Ø	VMware-HCX-440 v https://172.21254.157.443 Raieigh 1 Interconnect(s)	$\rightarrow$	HCX     P https     @ US W		

## Step 5: Configure the network profile, compute profile, and service mesh

The VMware HCX Interconnect service appliance provides replication and vMotion-based migration capabilities over the internet and private connections to the target site. The interconnect provides encryption, traffic engineering, and VM mobility. To create an Interconnect service appliance, complete the followings steps:

1. Under Infrastructure, select Interconnect > Multi-Site Service Mesh > Compute Profiles > Create Compute Profile.



The compute profiles define the deployment parameters including the appliances that are deployed and which portion of the VMware data center are accessible to HCX service.

$\Theta ~ \in ~ \rightarrow$	0&≠	# 0+ https://a300-vcsa01.ehodc.com/ui/app/plug	in/com.vmware.hybridity/com.vmware.ho	x/hybridConnect			0
$\equiv$ vSphere Client $$ $$ $$	2			C &	loministrator@EHCDC.COM >	٢	0
HCX Dashboard Infrastructure Ste Paining Interconnect E Transport Analytics.	< ~	Interconnect Multi-Site Device Mesh Compute Profiles Service Mesh Network Pr	rofiles Senticel Management		Q. C CREATE COM	PUTE PRO	DFILE
Services Network Extension Megration Disaster Recovery System & Administration © Support	*	hcxdemo     hcxtaeno     horr s000-exs01 shock conto host 0202)     hort s000-exs01 shock conto host 0202)     hort s000-exs01 shock conto     sorter Reservices     ala00-exsa01.shock.com     Ala00-luxter01 HCx Services     for inces     for inces     for inces     for inces     for inces     for inces	s in critical (red) state for service compute is in critical (red) state for deployment container Orpskyment Container (Ca300-Vesta01.ehrdic.com (Ca300-Custer01 Catastore (Ca300,NFS_0504 Catastore (Ca300,NFS_0504 Catastore (Ca300,NFS_0504) Catastore (Ca300,NFS_0504) Catastore (Ca300,NFS_0504)	Networks WM_3510 (Manaptiment) (Vigonere III Network Container (Network Entersion Applia III vDS-Switch0 (Unlimited)	epication) ((Issina) (vincisce) (Q new Limit)	EDIT	
		This Compute Pratter is being used in 2 Service  EDIT DELETE REVIEW CONNECTION PL	e Mesheti				

After the compute profile is created, create the network profiles by selecting Multi-Site Service Mesh
 Network Profiles > Create Network Profile.

The network profile defines a range of IP address and networks that are used by HCX for its virtual appliances.



This step requires two or more IP addresses. These IP addresses are assigned from the management network to the Interconnect Appliances.

							- And an		. 0
ICX Dashboard nfrastructure Site Paring Interconnect Transport Analytics	v	Interconnect Multi-Site Service Mesh Compute Profiles Service Mesh	Notwork Profiles Sentin	het Management			Q C CPE	ATE NETWOR	K PROFILI
Transport Analytics ervices Network Extension Migration Disaster Recovery ystem Administration O Support	* *	VM_3510 Network Details Backing VM_3510 show more	ыти 9000	IP Pools IP Ranges 172.21.254.80 - 172.21.254.95	IP Usage(Used/Total) 4/ 16	Prefix Length 24	Gateway 172.21.254.230		
		EDIT DELETE							

- 1. At this time, the compute and network profiles have been successfully created.
- 2. Create the Service Mesh by selecting the **Service Mesh** tab within the **Interconnect** option and select the on-premises and Azure SDDC sites.
- 3. The Service Mesh specifies a local and remote compute and network profile pair.



As part of this process, the HCX appliances are deployed and automatically configured on both the source and target sites in order to create a secure transport fabric.

$\leftarrow \  \   \rightarrow \  \   G$	0 8 =	• Pr. https://s300-vcsa01.ehodc.com/ui/app/plugin/com.vmware.hybridity/com.vmware.hcc.hybridConnect		\$	⊚ ≡
= vSphere Client C	2		C	Administrator@EHCDC.COM ≻	© 0~
HCX Dashboard Infrastructure Site Pairing Interconnect E: Transport Analytics Services	~	Interconnect Multi-Site Service Mesh Compute Profiles Service Mesh Network Profiles Service Management ICC007			E MESH
<ul> <li>Network Extension</li> <li>Migration</li> <li>Disaster Recovery</li> </ul> System & Administration Support	•	New version for service meth applances is available: Click on Update Applances to Upgrade to latest.           Size Paring           VMWare-           HCX-440           © Amsterdam           Nockerno           THT93-HCX-COMPUTE_PROFILE	(		×

1. This is the final step of configuration. This should take close to 30 minutes to complete the deployment. After the service mesh is configured, the environment is ready with the IPsec tunnels successfully created to migrate the workload VMs.

viphere Clent Q										C 2**			
	1	nterconr	rect										
Phone 1		10.0 00 041	na Mari										
inutive													
As Parrie	-	Completion of	TA CONTRACTOR ( LEGENSON ( )									_	
amport Arwytee	-	+ )	CC007 V								dout stand	ar eesse	
**		Artemp Bapteres Green											
where a parameter													
Nation Employees			Andrew 1		Assessed Type: 1	7.449.046		funnet instea	Darrente Versione	Average result.			
ippine An Administration C Engenet	×	0.1	00007-0-0 40 2019-91-020-979-0202 00255660000 Алануын A300-0200-0 Винның A300-02010-0		NO -Crimeria	(11) (1) (1) (1) (1) (1) (1) (1) (1) (1)		۲	4422				
		0.2	COUT HE II as UTHAT ROBATING TO THE RECORD Consume ADD, VEL, 2014 Through ADD, VEL, 2014 Western Konsumer (ST), 201701 Western Konsumer (ST), 201701 Marcana Konsumer (ST), 201701		() control	722(B+R <mark>(</mark>	annan ann ann ann ann ann ann ann ann a	۲	6428	****			
			COULT-WO-H we ANTTHIN POIN-MERIE WITH WEAK-STOOL Entering ACCOUNTLASSE Service ACCOUNTLASSE		(C) HOLE MAN OF				1200	N/A			
						-							
		Applano	es on hox.Sebt3b0b7ddf4cc08e3f85.westeurope.a 	5 westeurope ani azum com-cloud			e Adore				5a-		
		10000714	14			CONTRA	1770 2014 (modeleta) 1770 2014 (modeleta) 1770 2014 (modeleta) 1770 2014 (modeleta)				48	0.0	
		100001 W	0#	0	C HOLE MARK OF				73	60.			

#### Step 6: Migrate workloads

Workloads can be migrated bidirectionally between on-premises and Azure SDDCs using various VMware HCX migration technologies. VMs can be moved to and from VMware HCX-activated entities using multiple migration technologies such as HCX bulk migration, HCX vMotion, HCX Cold migration, HCX Replication Assisted vMotion (available with HCX Enterprise edition), and HCX OS Assisted Migration (available with the HCX Enterprise edition).

To learn more about various HCX migration mechanisms, see VMware HCX Migration Types.

## **Bulk migration**

This section details the bulk migration mechanism. During a bulk migration, the bulk migration capability of HCX uses vSphere Replication to migrate disk files while recreating the VM on the destination vSphere HCX instance.

To initiate bulk VM migrations, complete the following steps:

1. Access the Migrate tab under Services > Migration.

$\leftrightarrow \rightarrow \ C$	08=	https://x300-vcsa01.ehodc.com/ui/a	pp/plugin/o	onsonware	hýbridity	/com.vin	vare, ho	xmigrationold		ŵ		5 4	1	
$\equiv$ vSphere Client	Q								C	온 Administrator@EHCD		© (	@ ·	
HCX Dashboard Infrastructure	× ~	Migration	merit	ET MIG	RATE	<u>c</u> (	e			Search				
Ste Paring Sinterconnect		Name	VMs/ Storage/ Memory/ CPUs Progress					Start	End	Statua		4		
Services	v													
Migration		Y a300-vcsa01.ehcdd	.com →	<b>6</b> 1724	30.156.2	2								
System		> 2022-09-26 09:00 FLJVU > 2022-09-26 08:35 8XMTM		1	2 08 2 08	2 GB 2 GB	1	Migration Complete Migration Complete	9 5	*				
@ Support		> 2022-09-18 16:21 ERCZO	(Strail)	2	4 GB	4 08	2	@ Draft	*	×.				
		> M5-18cbce94 / Sep 16 > M6-04abdee8 / Sep 16	(100000	1	2 68	10 GB	1	Migration Complete	12.44 AU Sec 10 12.25 Au					
		> MG-ef7374dd / Sep 16	(TELE22)	1	2.GB	2 G8	1	S Migration Complete	540 10 12:11 AX 540 10	4				
		> MG-d2of93ef / Sep 14	0000	5	10 68	10 08	5	O lagration Complete	02.05 mm Sep 14	4				
		> MG-99fecac8 / Sep 14	(1111)	3	2.08/	2.08	*	Migration Comptela	11.02 AM Sep 14	×				
		> MG-548618cb / Sep 14	(feer Colg)	1	2.08	2.08	1	Stagration Complete	10.04 AM Seg 14	*				
		> MG-dd475274 / Sep 12	(Harrow)	- 2	.4.697	4.68	2	Migration Complete	12.25 Ptd					

- 1. Under **Remote Site Connection**, select the remote site connection and select the source and destination. In this example, the destination is Azure VMware Solution SDDC HCX endpoint.
- 2. Click **Select VMs for Migration**. This provides a list of all the on-premises VMs. Select the VMs based on the match:value expression and click **Add**.
- 3. In the **Transfer and Placement** section, update the mandatory fields (**Cluster**, **Storage**, **Destination**, and **Network**), including the migration profile, and click **Validate**.
|   |   |    |  | Select VMs for Migration |
|---|---|----|--|--------------------------|
| Transfer and Placement:   | This is migrating to ANF datastore  |    |  |                          |
| Duster-1  | Datastore02 ((115(276))   | 9  | Buk Mgradion   | 4                        |
| C3 (Soecity Destination Folder)   | Same foritiat as source   | ¥. | (Optional Switchover Schedule)   | Ö                        |
| Switchover:   |   |    |  |                          |
| Extended Options:   |   |    |  |                          |
|   |   |    |  |                          |
|   |   |    |  |                          |
| (Trus MAC)  |   |    |  |                          |
| Carton WAC  |   |    |  | 0                        |
| (Better MAC)  | Disk / Memory / vCPU  |    | Migration linfo  | <u>a</u>                 |
| (Set of Hold County)<br>(Set us: MAC)<br>M for High abov.   | 00k / Memory / vCPU<br>0 2.08.: 2.08.: 1vCPU  |    | Migration Mito   | <u>a</u> .               |
| (British MAC)<br>(British MAC)<br>M for Highelion<br>3 Demo_HCX_BUIK01<br>3 Demo_HCX_BUIK01   | 0046 / Memory / vCPU<br>02 2081 2085 1 vCPU<br>02 2082 2085 1 vCPU  |    | Migration Mro<br>Bulk Migration<br>Bulk Migration                                      | <u>s</u> .               |
| Centro MAC     C | 0148 / Memory / vCPU<br>0 2.08 2.08 1.1vCPU<br>0 2.08 2.08 1.1vCPU<br>0 2.08 2.08 1.1vCPU   |    | Migration Into-<br>Dull Migration<br>Dull Migration<br>Dull Migration                  | Q.                       |
| Centrol (Store)     Centrol (MAC)  Million Higheration     Demo_HCX_Buildot     Demo_HCX_Buildot     Demo_HCX_Buildot     Demo_HCX_Buildot     Demo_HCX_Buildot   | Ots./ Memory / vCPU           Ø         2.08         2.08         1VCPU           Ø         2.08         2.08         1VCPU |    | Migration Nrto<br>Bulk Migration<br>Bulk Migration<br>Bulk Migration<br>Outh Migration | 9                        |

1. After the validation checks are complete, click **Go** to initiate the migration.

(i)

+ + C	O & # musical	such wheelcham / vi	Spillipeters	-	manho	perchaser in									\$		0	2 11
😑 vSphere Clie	nt Q													C	& territory processo		۵	<b>0</b> -
HCX Deshboard Infrastructure Infrastructure		Migration	: Maragement M	(2) MISAATS	Storeget H		AND	Program			1	Start	. Ent		Status ;	8	i.	
Semiconer L: transport Analytic Services P testwork Extension Magazitak																		
Disaster Recovery System		V = alloo vo	sa01 ehodo co Juliilis	m → 🖕 172.30	156.2 29	0 250	.,	Hgaleg.				29144.407			Şwitchcoor stated			
Se Administration Q <sup>2</sup> Support		D > Dama, HCX	Jub04 percection	-	30	0 268	.*	Hipstei				2.01 Per 601 Des 14			Belichter statet			
		Dens_HCR	(Reb03 percise sizer	-	- 29	0 2.08	e t	37% Save Surv				205.444.807 Teap 14			Transfer Started			
		Dens_HCK	\$6662 (6140010181	-	10	8 2.68	. *	Marshee.	******		aneres.	2:05 PM 607 Bag 74			Delt/haver atlanted			
		> Deno_HCK	Byb01 birecol.com	-	10	8 2.68		Nursing.				2:05 PALEDY Deg 14			Detchove/stated			
		D > Detto_HCK	Motue	-	79	8 2.68		O Migration Col	ekig			1102 Av1-827 Tex 14	11.10 AM-857 Bap-14	3 100	Migration completed			
v Recent Tasks	Alarma	1					-											-
Tax Nate	Y Target	9 Itela	- ×	Getalu	٠	ninte		٠	Donund For	*	Start Time	٠	Completion Time		Y Server			•
Bulk Migration	(B. Dema, HDC, Buildo		245.0	Traisler Started		America	00 D	HODC COM	0.01		08/14/202	7.2.05 TLP54			Hybrioty-manager			-
Bulk Migration	(D Deve, HCK, Build)		01-0	Transfer Starting		America	0.00	HEDGOOM	0.ms		08/54/202	2.2:03 11 (994			hybrioty-manager			
Bak Mgration	D Denie HCK_BURD2		0.40	Transfer Starten		American	0.00	HEDCION	0.me		08/14/202	7,20511994			every warmer			
Bulk Migration	(2 Dens, JCK, BARS		0.42	Transfer Martine		Administrat	is (pic)	HOD COM	0.946		09/14/202	2,2:0511244			trypratty-manager			
Buik Migration	(D. Dens, HCK, Builde	1000	225 0	Traister Started		American	w@D	HODC DOM	0.ms		08/14/202	2.2051199			hybrioty-manager			

During this migration, a placeholder disk is created on the specified Azure NetApp Files datastore within the target vCenter to enable replication of the source VM disk's data to the placeholder disks. HBR is triggered for a full sync to the target, and after the baseline is complete, an incremental sync is performed based on the recovery point objective (RPO) cycle. After the full/incremental sync is complete, switchover is triggered automatically unless a specific schedule is set.

1. After the migration is complete, validate the same by accessing the destination SDDC vCenter.

e + 0 (	0 A # moult	enerar garperari Daatti rakahi	(+-2543-873+564)-1000-434	0401-0164052400	Charles (Series			\$	0 2 6
$\equiv$ vSphere Client Q							C & southeterer		0 0
	DatastoreO2 Annews Summary Menter Configure	Permissions Files Hosts	VMs						
<ul> <li>         Vicinie/social/europeanerse     </li> <li>         B_ scoc-betacenter     </li> </ul>	Victual Mathines. VM Templates								
Vic Beersook Sochie cobest     B SODC Outscenter	Virtual Machines VM Templates VMs on AMP datastore after migra	Bon	1. Buckeyed Back	- childhear	Contract Prints 1	1 Markhan		* Deho	uqut
Construction from a construction     B SCOC-Detacenter     Construction     Construction     Construction     Construction	Virtual Machines VM Templates VMs on ANF datastore after migra We on ANF datastore after migra	Bon   State T State	Provinced Spain	Unit Space	mut OV	Head March		* Date	usul
<ul> <li>Overheitsbezigen verheiten verhei</li></ul>	Victore Matchines VMs on ANF-datastore after migra Imare I	Bon blue t bole Presental On v 1 Presental On v 1	Protocol Second Normal 2:08	Ved See. 240.07 MB 751.24 MB	9944 OPV 0 HQ 0 HQ	Next Mass 264 Mill 200 Mill		.* Dente	ucul
<ul> <li>Social Statement</li> <li>Bissocial Statement</li> <li>Bissocial Statement</li> <li>Bissocial Statement</li> <li>Bissocial Statement</li> <li>Bissocial Statement</li> <li>Bissocial Statement</li> </ul>	Victure Matchines VMI Templates VMIs on ANF datastore after migra  I more I = @ Demo_PCC_Bandit I = @ Demo_PCC_BANDI I = @ Demo_PCC_BANDI	Son Nove f Solve Powerki On 71 Powerki On 71 Powerki On 71	Produced Spec Namu 200 Namu 200 Namu 200	044 April 246,87 MB 745,87 MB 751,24 MB 755,46 MB	0 HQ 0 HQ 0 HQ 0 HQ	Head Mass 264 MB 280 MB 264 MB		v Deite	unut
<ul> <li>Social Social Social Social Control Social So</li></ul>	Victure International VM Temperates VMs on ANF datasetore after migra  U to an ANF datasetore after mi	Son Nee f Nee Poserel On C I Poserel On C I Poserel On C I	Produced Spece Namual 208 Namual 208 Namual 208 Namual 208	Viel Space 248,87 MB 75124 MB 755,46 MB 76122 MB	0 He 0 He 0 He 0 He	Heal Mass 204.448 200.948 200.948 205.948 205.948		. s <u>Den</u>	wool

For additional and detailed information about various migration options and on how to migrate workloads from on-premises to Azure VMware Solution using HCX, see VMware HCX User Guide.

To learn more about this process, feel free to watch the following video:

# Workload Migration using HCX

Here is a screenshot of HCX vMotion option.

+ + 0	O & at mount	0.000	(abduem)(injo)	i girdan e	ويتراد وارخه وهو	errien bo	Newspapers.	0							\$	0	\$
😑 VSphere Client	à													C	А монитерсосори -		0
HCX C best-count initiastructures C base Parrieg	~	Mig	macking 121 mil	ragement	27 MERCEN	) 🐨 (	et 💿	Alloria				Start	. Del		Tarka	85	
E Transport Andatico invicas Distanció Esteració	÷	F															
P Disaster Recovery witers	Ŷ	-	er alloo vesat bene jick vier terreneljice	ehcdc co	n → † 172.00	16.2	08 2.08	1.2	Its free Serv.	8	) •101 ad bra	112 or 10 <sup>1</sup> but to			abbellert Pranafer Int Programm		
Q <sup>4</sup> tomport		0	<ul> <li>Dens, BCK, Coll sciences (21) C MCX, Photor, 47 Ammonia (21) C</li> </ul>	digration E-cone C-cone			08 2.08		C Maratum Con	Aunte Konte		1206 Mil 827 Dag 14 1225 Pec 807 Dag 13	12.15 Hardon Des 14 12.57 Hardon Des 15	Stars	Migration comprised		
		0	<ul> <li>HOLPHERING</li> <li>HOLPHERING</li> <li>HOLPHERING</li> <li>HOLPHERING</li> </ul>	C.1544	_		CB 2 CB	1 6	S Mangan Chu	upiele .		1225 Av. 801 846 13 1217 Av. 921 946 10	12,31,443,57 Ber 10 12,34,94 (20) Ber 10	t sie	Nigotus contribut		
		0	> HOLDene 25	1.1241		2	GB 2.09	1 6	O Migration Oper	-Crime		12((1)(mc.42)) (mar.7)	1217 earstyr Sec 19	7.400	Wgales cripted		
lati Rame T	faget		sunia .	+ 10	prim.	*	adutor.		*	Garant . W	Start Type	•	Comparison Type		* bron		
ecodigue votal hach. Voiver ortual itactica (). Ielesi votal methine s.	Demo, HCK, VMcRom     Demo, HCK, VMcRom     Demo, HCK, VMcRom		Comparised Comparised Completies		Beamlyung Vrb	ual Marits.	EHEDC.COM EHEDC.COM	Polanen Polanen Polanen	etraku- etrakur etrakur	5.ms 4.ms 6.ms	0904/30	12, 10:53 17 . 12, 10:59 58 . 10, 10:59 58 .	09/4/2022 101 29/4/2022 101 09/4/2022 102	10.55 AM 10.58 AM 12.33 AM	mass standes Clauser + 6004 mass, standes Clauser + 6006 mass, standes Clauser, + 6006		
when not divide in.	0 01221254.00		© Compacted				0+00000	15	distor.	3.94	09/14/30	2.110433	0604/2022,110	4.33.34	6300-vession article com		

To learn more about this process, feel free to watch the following video:

# HCX vMotion

(;

i.

Make sure sufficient bandwidth is available to handle the migration.

The target ANF datastore should have sufficient space to handle the migration.

# Conclusion

Whether you're targeting all-cloud or hybrid cloud and data residing on any type/vendor storage in onpremises, Azure NetApp Files and HCX provide excellent options to deploy and migrate the application workloads while reducing the TCO by making the data requirements seamless to the application layer. Whatever the use case, choose Azure VMware Solution along with Azure NetApp Files for rapid realization of cloud benefits, consistent infrastructure, and operations across on-premises and multiple clouds, bidirectional portability of workloads, and enterprise-grade capacity and performance. It is the same familiar process and procedures used to connect the storage and migrate VMs using VMware vSphere Replication, VMware vMotion, or even network file copy (NFC).

# Takeaways

The key points of this document include:

- You can now use Azure NetApp Files as a datastore on Azure VMware Solution SDDC.
- You can easily migrate data from on-premises to Azure NetApp Files datastore.
- You can easily grow and shrink the Azure NetApp Files datastore to meet the capacity and performance requirements during migration activity.

# Where to find additional information

To learn more about the information described in this document, refer to the following website links:

Azure VMware Solution documentation

https://docs.microsoft.com/en-us/azure/azure-vmware/

Azure NetApp Files documentation

https://docs.microsoft.com/en-us/azure/azure-netapp-files/

• VMware HCX User Guide

https://docs.vmware.com/en/VMware-HCX/4.4/hcx-user-guide/GUID-BFD7E194-CFE5-4259-B74B-991B26A51758.html

# Region Availability – Supplemental NFS datastore for ANF

Learn more about the the Global Region support for Azure, AVS and ANF.



NFS datastore will be available in regions where both services (AVS and ANF) are available.

The availability of supplemental NFS datastores on Azure / AVS is defined by Microsoft. First, you need to determine if both AVS and ANF are available in a specific region. Next, you need to determine if the ANF supplemental NFS datastore is supported in that region.

- Check the availability of AVS and ANF here.
- Check the availability of the ANF supplemental NFS datastore here.

# NetApp Capabilities for Google Cloud Platform GCVE

Learn more about the capabilities that NetApp brings to the Google Cloud Platform (GCP) Google Cloud VMware Engine (GCVE) - from NetApp as a guest connected storage device or a supplemental NFS datastore to migrating workflows, extending/bursting to the cloud, backup/restore and disaster recovery.

Jump to the section for the desired content by selecting from the following options:

- Configuring GCVE in GCP
- NetApp Storage Options for GCVE
- NetApp / VMware Cloud Solutions

## **Configuring GCVE in GCP**

As with on-premises, planning a cloud based virtualization environment is critical for a successful productionready environment for creating VMs and migration.

This section describes how to set up and manage GCVE and use it in combination with the available options for connecting NetApp storage.



In-guest storage is the only supported method of connecting Cloud Volumes ONTAP and Cloud Volumes Services to GCVE.

The setup process can be broken down into the following steps:

- Deploy and Configure GCVE
- Enable Private Access to GCVE

View the detailed configuration steps for GCVE.

#### NetApp Storage Options for GCVE

NetApp storage can be utilized in several ways - either as guess connected or as a supplemental NFS datastore - within GCP GCVE.

Please visit Supported NetApp Storage Options for more information.

Google Cloud supports NetApp storage in the following configurations:

- · Cloud Volumes ONTAP (CVO) as guest connected storage
- · Cloud Volumes Service (CVS) as guest connected storage
- · Cloud Volumes Service (CVS) as a supplemental NFS datastore

View the detailed guest connect storage options for GCVE.

Read more about NetApp Cloud Volumes Service datastore support for Google Cloud VMware Engine (NetApp blog) or How to use NetApp CVS as datastores for Google Cloud VMware Engine (Google blog)

#### **Solution Use Cases**

With NetApp and VMware cloud solutions, many use cases are simple to deploy in Azure AVS. se cases are defined for each of the VMware defined cloud areas:

- Protect (includes both Disaster Recovery and Backup / Restore)
- Extend
- Migrate

Browse the NetApp solutions for Google Cloud GCVE

# Protecting Workloads on GCP / GCVE

#### Application Consistent Disaster Recovery with NetApp SnapCenter and Veeam Replication

Disaster recovery to cloud is a resilient and cost-effective way of protecting workloads against site outages and data corruption events such as ransomware. With NetApp SnapMirror, on-premises VMware workloads that use guest-connected storage can be replicated to NetApp Cloud Volumes ONTAP running in Google Cloud.

Authors: Suresh Thoppay, NetApp

# Overview

Many customers are looking for an effective disaster recovery solution for their application VMs hosted on VMware vSphere. Many of them use their existing backup solution to perform recovery during diaster. Many times that solution increase the RTO and doesn't meet their expectations. To reduce the RPO and RTO, Veeam VM replication can be utilized even from on-prem to GCVE as long as network connectivity and environemnt with appropriate permissions are available.

NOTE: Veeam VM Replication doesn't protect VM guest connected storage devices like iSCSI or NFS mounts inside the guest VM. Need to protect those seperately.

For application consistent replication for SQL VM and to reduce the RTO, we used SnapCenter to orchestrate snapmirror operations of SQL database and log volumes.

This document provides a step-by-step approach for setting up and performing disaster recovery that uses NetApp SnapMirror, Veeam, and the Google Cloud VMware Engine (GCVE).



# Assumptions

This document focuses on in-guest storage for application data (also known as guest connected), and we assume that the on-premises environment is using SnapCenter for application-consistent backups.



This document applies to any third-party backup or recovery solution. Depending on the solution used in the environment, follow best practices to create backup policies that meet organizational SLAs.

For connectivity between the on-premises environment and the Google Cloud network, use the connectivity options like dedicated interconnect or Cloud VPN. Segments should be created based on the on-premises VLAN design.



There are multiple options for connecting on-premises datacenters to Google Cloud, which prevents us from outlining a specific workflow in this document. Refer to the Google Cloud documentation for the appropriate on-premises-to-Google connectivity method.

# Deploying the DR Solution

#### **Solution Deployment Overview**

- 1. Make sure that application data is backed up using SnapCenter with the necessary RPO requirements.
- 2. Provision Cloud Volumes ONTAP with the correct instance size using BlueXP within the appropriate subscription and virtual network.
  - a. Configure SnapMirror for the relevant application volumes.
  - b. Update the backup policies in SnapCenter to trigger SnapMirror updates after the scheduled jobs.
- 3. Install the Veeam software and start replicating virtual machines to Google Cloud VMware Engine instance.
- 4. During a disaster event, break the SnapMirror relationship using BlueXP and trigger failover of virtual machines with Veeam.
  - a. Reconnect the ISCSI LUNs and NFS mounts for the application VMs.
  - b. Bring up applications online.
- 5. Invoke failback to the protected site by reverse resyncing SnapMirror after the primary site has been recovered.

# **Deployment Details**

The first step is to configure Cloud Volumes ONTAP on Google Cloud (cvo) and replicate the desired volumes to Cloud Volumes ONTAP with the desired frequencies and snapshot retentions.

ogle Cloud		1200 - 100120-200		
VMware Engine		Customer VPC		
		Cloud Vo	olumes AP	
-	Private Connection	•	Compute engine	
VSAN		E Pe	rsistent disks	Cloud storage
<b>.</b>		Hot	data	Cold data
vSphere NSX		L		

For sample step-by-step instructions on setting up SnapCenter and replicating the data, Refer to Setup Replication with SnapCenter

Review of SQL VM protection with SnapCenter

#### Configure GCVE hosts and CVO data access

Two important factors to consider when deploying the SDDC are the size of the SDDC cluster in the GCVE solution and how long to keep the SDDC in service. These two key considerations for a disaster recovery solution help reduce the overall operational costs. The SDDC can be as small as three hosts, all the way up to a multi-host cluster in a full-scale deployment.

NetApp Cloud Volume Service for NFS Datastore and Cloud Volumes ONTAP for SQL databases and log can be deployed to any VPC and GCVE should have private connection to that VPC to mount NFS datastore and have VM connect to iSCSI LUNs.

To configure GCVE SDDC, see Deploy and configure the Virtualization Environment on Google Cloud Platform (GCP). As a prerequisite, verify that the guest VMs residing on the GCVE hosts are able to consume data from Cloud Volumes ONTAP after connectivity has been established.

After Cloud Volumes ONTAP and GCVE have been configured properly, begin configuring Veeam to automate the recovery of on-premises workloads to GCVE (VMs with application VMDKs and VMs with in-guest storage) by using the Veeam Replication feature and by leveraging SnapMirror for application volumes copies to Cloud Volumes ONTAP.

Based on deployment scenario, the Veeam backup server, backup repository and backup proxy that needs to be deployed. For this use case, there is no need to deploy object store for Veeam and Scale-out repository also not required.

Refer to the Veeam documentation for the installation procedure For additional information, please refer Migration with Veeam Replication

#### Setup VM Replication with Veeam

Both on-premises vCenter and GCVE vCenter needs to be registered with Veeam. Setup vSphere VM Replication Job At the Guest Processing step of wizard, select disable application processing as we will be utilizing SnapCenter for application aware backup and recovery.

https://netapp.hosted.panopto.com/Panopto/Pages/Embed.aspx?id=8b7e4a9b-7de1-4d48-a8e2-b01200f00692

#### Failover of Microsoft SQL Server VM

https://netapp.hosted.panopto.com/Panopto/Pages/Embed.aspx?id=9762dc99-081b-41a2-ac68-b01200f00ac0

#### Benefits of this solution

- Uses the efficient and resilient replication of SnapMirror.
- Recovers to any available points in time with ONTAP snapshot retention.
- Full automation is available for all required steps to recover hundreds to thousands of VMs, from the storage, compute, network, and application validation steps.
- SnapCenter uses cloning mechanisms that do not change the replicated volume.
  - This avoids the risk of data corruption for volumes and snapshots.
  - Avoids replication interruptions during DR test workflows.
  - Leverages the DR data for workflows beyond DR, such as dev/test, security testing, patch and upgrade testing, and remediation testing.
- Veeam Replication allows changing VM IP addresses on DR site.

#### Application Disaster Recovery with SnapCenter, Cloud Volumes ONTAP and Veeam Replication

Disaster recovery to cloud is a resilient and cost-effective way of protecting workloads against site outages and data corruption events such as ransomware. With NetApp SnapMirror, on-premises VMware workloads that use guest-connected storage can be replicated to NetApp Cloud Volumes ONTAP running in Google Cloud.

Authors: Suresh Thoppay, NetApp

# Overview

This covers application data; however, what about the actual VMs themselves. Disaster recovery should cover

all dependent components, including virtual machines, VMDKs, application data, and more. To accomplish this, SnapMirror along with Veeam can be used to seamlessly recover workloads replicated from on-premises to Cloud Volumes ONTAP while using vSAN storage for VM VMDKs.

This document provides a step-by-step approach for setting up and performing disaster recovery that uses NetApp SnapMirror, Veeam, and the Google Cloud VMware Engine (GCVE).



# Assumptions

This document focuses on in-guest storage for application data (also known as guest connected), and we assume that the on-premises environment is using SnapCenter for application-consistent backups.



This document applies to any third-party backup or recovery solution. Depending on the solution used in the environment, follow best practices to create backup policies that meet organizational SLAs.

For connectivity between the on-premises environment and the Google Cloud network, use the connectivity options like dedicated interconnect or Cloud VPN. Segments should be created based on the on-premises VLAN design.



There are multiple options for connecting on-premises datacenters to Google Cloud, which prevents us from outlining a specific workflow in this document. Refer to the Google Cloud documentation for the appropriate on-premises-to-Google connectivity method.

# **Deploying the DR Solution**

# **Solution Deployment Overview**

- 1. Make sure that application data is backed up using SnapCenter with the necessary RPO requirements.
- 2. Provision Cloud Volumes ONTAP with the correct instance size using Cloud manager within the appropriate subscription and virtual network.

- a. Configure SnapMirror for the relevant application volumes.
- b. Update the backup policies in SnapCenter to trigger SnapMirror updates after the scheduled jobs.
- 3. Install the Veeam software and start replicating virtual machines to Google Cloud VMware Engine instance.
- 4. During a disaster event, break the SnapMirror relationship using Cloud Manager and trigger failover of virtual machines with Veeam.
  - a. Reconnect the ISCSI LUNs and NFS mounts for the application VMs.
  - b. Bring up applications online.
- 5. Invoke failback to the protected site by reverse resyncing SnapMirror after the primary site has been recovered.

# **Deployment Details**

#### Configure CVO on Google Cloud and replicate volumes to CVO

The first step is to configure Cloud Volumes ONTAP on Google Cloud (cvo) and replicate the desired volumes to Cloud Volumes ONTAP with the desired frequencies and snapshot retentions.

Architectur	e: Cloud Volumes ONTA	P communication to GCVE
oogle Cloud		
VMware Engine		Customer VPC
	Private Connection	Cloud Volumes ONTAP
Datastore vSAN		Persistent Cloud disks storage
		Hot data Cold data

For sample step-by-step instructions on setting up SnapCenter and replicating the data, Refer to Setup Replication with SnapCenter

Setup Replication with SnapCenter

Two important factors to consider when deploying the SDDC are the size of the SDDC cluster in the GCVE solution and how long to keep the SDDC in service. These two key considerations for a disaster recovery solution help reduce the overall operational costs. The SDDC can be as small as three hosts, all the way up to a multi-host cluster in a full-scale deployment.

Cloud Volumes ONTAP can be deployed to any VPC and GCVE should have private connection to that VPC to have VM connect to iSCSI LUNs.

To configure GCVE SDDC, see Deploy and configure the Virtualization Environment on Google Cloud Platform (GCP). As a prerequisite, verify that the guest VMs residing on the GCVE hosts are able to consume data from Cloud Volumes ONTAP after connectivity has been established.

After Cloud Volumes ONTAP and GCVE have been configured properly, begin configuring Veeam to automate the recovery of on-premises workloads to GCVE (VMs with application VMDKs and VMs with in-guest storage) by using the Veeam Replication feature and by leveraging SnapMirror for application volumes copies to Cloud Volumes ONTAP.

#### **Install Veeam Components**

Based on deployment scenario, the Veeam backup server, backup repository and backup proxy that needs to be deployed. For this use case, there is no need to deploy object store for Veeam and Scale-out repository also not required.

Refer to the Veeam documentation for the installation procedure

#### Setup VM Replication with Veeam

Both on-premises vCenter and GCVE vCenter needs to be registered with Veeam. Setup vSphere VM Replication Job At the Guest Processing step of wizard, select disable application processing as we will be utilizing SnapCenter for application aware backup and recovery.

Setup vSphere VM Replication Job

#### Failover of Microsoft SQL Server VM

Failover of Microsoft SQL Server VM

#### Benefits of this solution

- Uses the efficient and resilient replication of SnapMirror.
- Recovers to any available points in time with ONTAP snapshot retention.
- Full automation is available for all required steps to recover hundreds to thousands of VMs, from the storage, compute, network, and application validation steps.
- SnapCenter uses cloning mechanisms that do not change the replicated volume.
  - This avoids the risk of data corruption for volumes and snapshots.
  - Avoids replication interruptions during DR test workflows.

- Leverages the DR data for workflows beyond DR, such as dev/test, security testing, patch and upgrade testing, and remediation testing.
- Veeam Replication allows changing VM IP addresses on DR site.

# Migrating Workloads on GCP / GCVE

Migrate workloads to NetApp Cloud Volume Service datastore on Google Cloud VMware Engine using VMware HCX - Quickstart guide

One of the most common use cases for the Google Cloud VMware Engine and Cloud Volume Service datastore is the migration of VMware workloads. VMware HCX is a preferred option and provides various migration mechanisms to move on-premises virtual machines (VMs) and its data to Cloud Volume Service NFS datastores.

Author(s): NetApp Solutions Engineering

# Overview: Migrating virtual machines with VMware HCX, NetApp Cloud Volume Service datastores, and Google Cloud VMware Engine (GCVE)

VMware HCX is primarily a migration platform that is designed to simplify application migration, workload rebalancing, and even business continuity across clouds. It is included as part of Google Cloud VMware Engine Private Cloud and offers many ways to migrate workloads and can be used for disaster recovery (DR) operations.

This document provides step-by-step guidance for provisioning Cloud Volume Service datastore followed by downloading, deploying, and configuring VMware HCX, including all its main components in on-premises and the Google Cloud VMware Engine side including Interconnect, Network Extension, and WAN optimization for enabling various VM migration mechanisms.



VMware HCX works with any datastore type as the migration is at the VM level. Hence this document is applicable to existing NetApp customers and non-NetApp customers who are planning to deploy Cloud Volume Service with Google Cloud VMware Engine for a cost-effective VMware cloud deployment.

#### **High-level steps**

This list provides the high-level steps necessary to pair & Migrate the VMs to HCX Cloud Manager on the Google Cloud VMware Engine side from HCX Connector on-premises:

- 1. Prepare HCX through the Google VMware Engine portal.
- 2. Download and deploy the HCX Connector Open Virtualization Appliance (OVA) installer in the onpremises VMware vCenter Server.
- 3. Activate HCX with the license key.
- 4. Pair the on-premises VMware HCX Connector with Google Cloud VMware Engine HCX Cloud Manager.
- 5. Configure the network profile, compute profile, and service mesh.
- 6. (Optional) Perform network extension to avoid re-IP during migrations.
- 7. Validate the appliance status and ensure that migration is possible.
- 8. Migrate the VM workloads.

#### Prerequisites

Before you begin, make sure the following prerequisites are met. For more information, see this link. After the prerequisites, including connectivity, are in place, download HCX license key from the Google Cloud VMware Engine portal. After the OVA installer is downloaded, proceed with the installation process as described below.



HCX advanced is the default option and VMware HCX Enterprise edition is also available through a support ticket and supported at no additional cost. Refer this link

- Use an existing Google Cloud VMware Engine software-defined data center (SDDC) or create a private cloud by using this NetApp link or this Google link.
- Migration of VMs and associated data from the on-premises VMware vSphere- enabled data center requires network connectivity from the data center to the SDDC environment. Before migrating workloads, set up a Cloud VPN or Cloud Interconnect connection between the on-premises environment and the respective private cloud.
- The network path from on-premises VMware vCenter Server environment to the Google Cloud VMware Engine private cloud must support the migration of VMs by using vMotion.
- Make sure the required firewall rules and ports are allowed for vMotion traffic between the onpremises vCenter Server and SDDC vCenter.
- Cloud Volume Service NFS volume should be mounted as a datastore in Google Cloud VMware Engine. Follow the steps detailed in this link to attach Cloud Volume Service datastores to Google Cloud VMware Engines hosts.

#### **High Level Architecture**

For testing purposes, the lab environment from on-premises used for this validation was connected through a Cloud VPN, which allows on-premises connectivity to Google Cloud VPC.



# **Solution Deployment**

Follow the series of steps to complete the deployment of this solution:

#### Step 1: Prepare HCX through the Google VMware Engine Portal

HCX Cloud Manager component automatically gets installed as you provision private cloud with VMware Engine. To prepare for site pairing, complete the following steps:

1. Log in to the Google VMware Engine Portal and sign-in to the HCX Cloud Manager.

You can login to HCX Console either by clicking on the HCX version link image::gcpd-hcx-image2.png[HCX Console access with link on GCVE resource] or clicking on HCX FQDN under vSphere Management Network tab. image::gcpd-hcx-image3.png[HCX Console access with FQDN link]

- 2. In HCX Cloud Manager, go to Administration > System Updates.
- Click Request download link and download the OVA file. image::gcpd-hcx-image4.png[Request download link]
- 4. Update HCX Cloud Manager to the latest version available from the HCX Cloud Manager UI.

For the on-premises Connector to connect to the HCX Manager in Google Cloud VMware Engine, make sure the appropriate firewall ports are open in the on-premises environment.

To download and install HCX Connector in the on-premises vCenter Server, complete the following steps:

- 1. Have the ova downloaded from the HCX Console on Google Cloud VMware Engine as stated in previous step.
- 2. After the OVA is downloaded, deploy it on to the on-premises VMware vSphere environment by using the **Deploy OVF Template** option.

Deploy OVF Template Select an OVF template	×
Select an OVF template from remote URL or local file system Enter a URL to download and install the OVF package from the internet, or browse to a location accessible from your co	mputer,
such as a local hard drive, a network share, or a CD/DVD drive.	1.1916.1912.1
2 Select a name and folder O URL	
3 Select a compute resource	
A Review details	
UPLOAD FILES VMware-HCX-Connector-4.5.2.0-20914338.ova	
6 Ready to complete	
CANCEL	EXT

3. Enter all the required information for the OVA deployment, click **Next**, and then click **Finish** to deploy the VMware HCX connector OVA.



Power on the virtual appliance manually.

For step-by-step instructions, see the VMware HCX User Guide.

After you deploy the VMware HCX Connector OVA on-premises and start the appliance, complete the following steps to activate HCX Connector. Generate the license key from the Google Cloud VMware Engine portal and activate it in VMware HCX Manager.

- From the VMware Engine portal, Click on Resources, select the private cloud, and click on download icon under HCX Manager Cloud Version. image::gcpd-hcx-image6.png[Download HCX License]
   Open Downloaded file and copy the License Key String.
- 2. Log into the on-premises VMware HCX Manager at "https://hcxmanagerIP:9443" using administrator credentials.



Use the hcxmanagerIP and password defined during the OVA deployment.

3. In the licensing, enter the key copied from step 3 and click Activate.



The on-premises HCX Connector should have internet access.

- 4. Under **Datacenter Location**, provide the nearest location for installing the VMware HCX Manager onpremises. Click **Continue**.
- 5. Under System Name, update the name and click Continue.
- 6. Click Yes, Continue.
- 7. Under **Connect your vCenter**, provide the fully qualified domain name (FQDN) or IP address of vCenter Server and the appropriate credentials and click **Continue**.



Use the FQDN to avoid connectivity issues later.

8. Under **Configure SSO/PSC**, provide the Platform Services Controller's(PSC) FQDN or IP address and click **Continue**.



For Embedded PSC, Enter the VMware vCenter Server FQDN or IP address.

- 9. Verify that the information entered is correct and click Restart.
- 10. After the services restart, vCenter Server is displayed as green on the page that appears. Both vCenter Server and SSO must have the appropriate configuration parameters, which should be the same as the previous page.



This process should take approximately 10 to 20 minutes and for the plug-in to be added to the vCenter Server.

O HCX-RI	P		(!)	CPU	Free 1543 MHZ	26
IP Address:	172.21.254.155			Used 552 MHZ	Capacity 2095 MHZ	
Version:	4.5.2.0		0	Memory	Free 2472 MB	70
Current Time:	Thursday, 16 February 2023	05:59:00 PM UTC	U	Used 9535 MB	Capacity 12008 MB	13
			~	Storage	Free 76	G
			0	Used 7.7G	Capacity 84	9 G
			1.1			
NSX		vCenter	s	SO		-
		https://a300-vcsa01.ehcdc.com	• h	ttps://a300-vcsa01.el	icdc.com	
MANAGE		MANAGE		ANAGE		

#### Step 4: Pair on-premises VMware HCX Connector with Google Cloud VMware Engine HCX Cloud Manager

After HCX Connector is deployed and configured on on-premises vCenter, establish connection to Cloud Manager by adding the pairing. To configure the site pairing, complete the following steps:

1. To create a site pair between the on-premises vCenter environment and Google Cloud VMware Engine SDDC, log in to the on-premises vCenter Server and access the new HCX vSphere Web Client plug-in.

≡ vSphere 0	Client Q							C		0	0×
Shortcuts											
Inventories											
([])	图		Ø	II		000		<b>(</b>	٢		
Hosts and Clusters	VMs and Templates	Storage	Networking	Content Libraries	Global Inventory Lists	Workload Management	SnapCenter Plug-in for VMware vSphere	Site Recovery	нсх		
Monitoring											
		ćą.	R		п	$\Leftrightarrow$					
Task Console	Event Console	VM Customization Specifications	VM Storage Policies	Host Profiles	ONTAP tools	Lifecycle Manager					
Administratio	'n										
Q											
Licenting											

2. Under Infrastructure, click Add a Site Pairing.



Enter the Google Cloud VMware Engine HCX Cloud Manager URL or IP address and the credentials for user with Cloud-Owner-Role privileges for accessing the private cloud.

Remote HCX URI	https://bcy-58042.f7458c8f.europe-west3.c	
2	https://ficx-56642.1/456c61.edrope-west5.g	
Username	cloudowner@gve.local	í
Password		
	CANCEL	CONNECT

# 3. Click Connect.



VMware HCX Connector must be able to route to HCX Cloud Manager IP over port 443.

4. After the pairing is created, the newly configured site pairing is available on the HCX Dashboard.



#### Step 5: Configure the network profile, compute profile, and service mesh

The VMware HCX Interconnect service appliance provides replication and vMotion-based migration capabilities over the internet and private connections to the target site. The interconnect provides encryption, traffic engineering, and VM mobility. To create an Interconnect service appliance, complete the followings steps:

1. Under Infrastructure, select Interconnect > Multi-Site Service Mesh > Compute Profiles > Create Compute Profile.



The compute profiles define the deployment parameters including the appliances that are deployed and which portion of the VMware data center are accessible to HCX service.

😑 vSphere Client	c	i -		C & Administrator@EHCDC.COM ~ ③ ⑦ ~
HCX Dashboard Infrastructure Ste Pairing Inferconnect Entransport Analytics Converse	×	Interconnect Multi-Site Service Mesh Concule Ptoffers Service Mesh Network P	rofiles Sentinel Management	Q C CREATE COMPUTE PROFILE
Services  Network Extension  Nigration  Disaster Recovery  System  Andministration  Support	~	MCX-CP Service Resources Service Resources Ca300-vesa01.ebddc.com A300-Cluster01 HCX Services (a)	Deployment Container Ca300-Vesa01.ehcdc.com Casaotore Casastore CourMemory Reservations CourMemory Reservations CourMemory Reservations CourMemory Reservations CourMemory Reservations CourMemory Reservations CourMemory Reservations	Networks <b>©VM_3510</b> (Warkagement) (vSphere Beplication) (Uplink) (vMotion) (): Ebit Network Cootarer (Network Extension Applance Limit) @ vDS-Switch0 (Unlimited)
		EDIT OBLETE REVIEW CONNECTION R	ULES	

2. After the compute profile is created, create the network profiles by selecting Multi-Site Service Mesh
 > Network Profiles > Create Network Profile.

The network profile defines a range of IP address and networks that are used by HCX for its virtual appliances.



This step requires two or more IP addresses. These IP addresses are assigned from the management network to the Interconnect Appliances.

HCX (Deshboard Infrastructure	Interconnect Muti-Ste Service Mesh							
C Site Paring Interconnect E Transport Analytics	Compute Profiles Service Mech	Notwork Profiles Settine	Management			Q C CRE	ATE NETWORK P	ROFILE
Services × Network Extension Migration Disatter Recovery System × Co Administration Co Subport	C) VM_3510 Networy Detais Backing: VM_3510 show more EDIT_DELETE	MTU 1350	0 Pools 17 Ranges 172.21.254.81 - 172.21.254.95	P Usege(Used/Total) 2/15	Prefix Leegm 24	Gateway 172,21,254,231		

- 3. At this time, the compute and network profiles have been successfully created.
- 4. Create the Service Mesh by selecting the **Service Mesh** tab within the **Interconnect** option and select the on-premises and GCVE SDDC sites.
- 5. The Service Mesh specifies a local and remote compute and network profile pair.

As part of this process, the HCX appliances are deployed and automatically configured on both the source and target sites in order to create a secure transport fabric.

	- <	Interconnect		
cx Dashboard		Multi-Site Service Mesh		
Site Pairing	×	Compute Profiles Sensice Mesh Network Profiles Sentine/ Management		
Transport Analytics				MESH
ervices	~	RTP-GCVE		
<ul> <li>Network Extension</li> <li>Migration</li> <li>Disaster Recovery</li> </ul>		Ste Fanng HCX-RTP  ODurfam OF makfurt		
System	×	HCX-CP HCX-GCVE Uplinis (Overnidden) Vplinis (Overnidden) Vplinis (Overnidden) Vplinis (Overnidden)	~~~~~	>
		VIEW APPLIANCES RESYNC EDIT DELETE MORE -		

(i)

6. This is the final step of configuration. This should take close to 30 minutes to complete the deployment. After the service mesh is configured, the environment is ready with the IPsec tunnels successfully created to migrate the workload VMs.

	- 6	Internet	neet							
		1000000	Inclusion							
handdoard		Multi-Star Sar	reve them							
ofracture		Danaka Pro	Here Service Mark							
Interconnect		÷	ятя-осує и						EDIT SERVICE MESH	
149		Lines	77 BApplaners Etails							
<ul> <li>Antonya Tytheraan</li> <li>Megapoor</li> <li>Disapter Recovery</li> </ul>		Appliance	tes on MCX-817P						e	
a grout all a four	- 10		Applanes Name	2.3	Applanet Type - 1	PADED	Turnet Trans-	Content Version		
(Chispor	1.8	J.	0 >	ITT-SCVE (4.8) is beatty-acc-autometaboutines; Energies Also Quantiti Barryot Also Quantiti Barryot Also QUEL		HCK WARKS	(122.048) Annumer (Calmo Balcala)	۲	4128	
				D B BTTN-00 Ke an Comp Stores Bases Paties	ITTP-OC/EX.46-5 Mic 011523 RevLates A/TO-excEllation000 Compare 2000 Comm/01 Revenue 2000 Artic 2000 Revenue Container (CD-SouthOL External Revenue X/ID		HOX HET LEAT	00000000000000000000000000000000000000		4520
			HTP-COM/WD-0 wr/2014703-079-470-4988-4988443acoa# Campion A3Co-Quarrell Interage A3Co-Q470_C0501		-CH WARK CIT			1280		
									(Apple 199	
		Applanc	es on nex-58042.17458c81.europe-west3.gve	9009-clou	Ø2					
		Approve	Name	Assesse	Tape # Address				Carriell Western	
		079-0291	5 or Ar	NG	Notice Rotation	(Respected Colores Issues), man (Sp			43.20	
		#UP-OCM	¥-WQ-81	Δ.	or www.expert				7382	

#### Step 6: Migrate workloads

Workloads can be migrated bidirectionally between on-premises and GCVE SDDCs using various VMware HCX migration technologies. VMs can be moved to and from VMware HCX-activated entities using multiple migration technologies such as HCX bulk migration, HCX vMotion, HCX Cold migration, HCX Replication Assisted vMotion (available with HCX Enterprise edition), and HCX OS Assisted Migration (available with the HCX Enterprise edition).

To learn more about various HCX migration mechanisms, see VMware HCX Migration Types.

The HCX-IX appliance uses the Mobility Agent service to perform vMotion, Cold, and Replication Assisted vMotion (RAV) migrations.



The HCX-IX appliance adds the Mobility Agent service as a host object in the vCenter Server. The processor, memory, storage and networking resources displayed on this object do not represent actual consumption on the physical hypervisor hosting the IX appliance.

#### **HCX vMotion**

This section describes the HCX vMotion mechanism. This migration technology uses the VMware vMotion protocol to migrate a VM to GCVE. The vMotion migration option is used for migrating the VM state of a single VM at a time. There is no service interruption during this migration method.



Network Extension should be in place (for the port group in which the VM is attached) in order to migrate the VM without the need to make an IP address change.

1. From the on-premises vSphere client, go to Inventory, right- click on the VM to be migrated, and select HCX Actions > Migrate to HCX Target Site.

	6 (A 54) (100)	es i co				
	of Move2GCV	E	C G & B (11000			
	Summary Monito	C. Can	figure Permisions Datastores Netwools Sri	apshors Lipdates		
帝 (CC000-WG-6						Switch to Marker with
(i) (CL3/OCH-H II)	the second s		Gunst Oli VMmare (Rotun Oli (54-bit) Consectables ED) 6.1 and later (VM access M)			CI CIVULAU
# KCANNERE			VMware Tomo: Romming, remain (1222 (Sum) Managert)			0 Hz
(B) ICCAROCE-WO-R			DNS frame proto-of			C O B
(B IS-ena	IN A REAL PROPERTY.		PARMEN VI21212			10 100AUL1A
(B) (an draw Dynamic)	CD Actions - Move2DCVE	30	wear all'O' soull'antest com			<sup>Ⅲ</sup> 731,45 м
60 Move25CVE	Guest DS	5 W.	Ara			
gi invatic	Snepshots	2	- 17.			
Charmen C	C Open Remote Comple			6	A Notes	
段 nma500kli			SPREZELL :			
(b anvaloo	(a Morstell		E contra			
(D rection/anti-root)	CON	×	The start is the second states		Contam Attributes.	2.6
() in coverves	Fault Tolerance	5	2.08		T defeas	
(\$ 1045 Source Works			VM_IBCO (convertant)			
·益 att=ocvs.ocm	VM PORDES	4	Dispression			
E ATE-OCVENIER	Template	÷	a best			
di amiocia wold	Companishty	2	Device in the origin machine PD not th	et provides autoort for the		
CB Scapability	Expert System Loop		vital subsequences and interface			
12 BridgCermervMexal	a short of some to do		Additional Hardware		10	of seat 1 lines
(B) BRING	@ bak Sellings		\$300-6.7 and later (MM average VI)			
EP Techania	NAME OF COMPANY				3 Agis.	
(D. Tell-Leven	Destation				designed flag Company	Description
(D) Test-Terrini2	Edit Notes.			2	*	
(2) Yestika-min2	Tags & Custon Attributes	6	Alton Chilthirdf			
(p Test_De			Addd-exact enablisms			
di Testinei	Add Pertresidort		IB 59(2000)			
(# Tellbel00)	ALATIN	4	Accessive access			The series to displace
fiecent Tasks - Alarma						
an Name * Tarrat	Contraction Contraction		* Details * Indiana	* Outer * ]	RatTree 1 * Completion Tree * Leven	
met Crisistust martine /B. M.			Presenting on the new Wittak Dadam	4.04	NUMBER TRENT NUMBER TO DO IN ADD-CADARTS ON	
time preserve on the A	Therefore Contail team		EHEDE COMME	initiation 3 mil	00/W/2021, 2:30 10 . 03/W/2021 2:30 10 . 4800-esa01 Misds.com	
	Annapo SnapCemer		Highle to HCK Earph Buccc contactor	Milliatur 3 mil	00/6/0023 3 30 33 P. 03/6/0023 3 30 33 P. add0-caed/white com	
econopore entrainenti da la	SALSAE Recovery actions	. ~	Decor Company	natiator 6.99	00/M/2001, 210:39 PM 00/M/2001, 2:30:30 . #300-ctudy #tudy.com	
	NOV A Direct	100	D			

2. In the Migrate Virtual Machine wizard, select the Remote Site Connection (target GCVE).

HCX: Migrate Virtual Machine × Remote Site Connection: Source: HCX-RTP / VC: a300-vcsa01.ehcdc.com → Ø Destination: hcx-58042.f7458c8f.europe-west3.gve.goog-cloud / VC vcsa-57901.f7458c8f.europe-west3.gve.goog ntips #0.0 % 0 Transfer and Placement: (Mandatory: Compute Container) (Mandatory: Storage) (Migration Profile) 20 2 ÷ Same format as source (Specify Destination Folder) -\* (Optional: Switchover Schedule) 0 > Switchover: ✓ Extended Options: Edit Extended Options VM for Migration Disk / Memory / vCPU Migration Info C 2 GB / 2 GB / 1 vCPU (Migration profile is not specified!) > Move2GCVE

► GO

Ø VALIDATE

CLOSE

3. Update the mandatory fields (Cluster, Storage, and Destination Network), Click Validate.

Source: HCX-RTP / V⊂ a300-vc → ② Destination: hcx-58042.17458c https://10.036.13	sa01.ehcdc.com 8f.europe-west3.g	ve.goog-cloud / VC vcsa-57901.17458c8	3f.europe-west3.	gve.goog	C Related Conne
<ul> <li>Transfer and Placement:</li> </ul>					
🔗 Workload	9	gcp-ve-4 (107.0408./178)		vMotion	4
(Specify Destination Folder)	2	Same format as source	Ŷ	(Optional: Switchover Schedule)	G
> Switchover:					
✓ Extended Options:					
Colt Extended Ontions					
					15
		and a subscription of the same to		and control of the second	C
VM for Migration		Disk / Memory / VCPU		Migration into	
Move2GCVE	C	2 GB / 2 GB / 1 VCPU			
Workload	1	gcp-ve-4 (007/6/GB/178)	<b>1</b>	vMotion	8
(Specify Destination Folder)	1	Same format as source	×		
E Force Power-off VM					
Enable Seed Checkpoint					
Edit Extended Options Resent	AC #				
>	Network	adapter 1 (VM 3509) → L2E VM 350	09-3509-a0041a8	ki i	
				O VALIDATE	C.L.



The vMotion transfer captures the VM active memory, its execution state, its IP address, and its MAC address. For more information about the requirements and limitations of HCX vMotion, see Understanding VMware HCX vMotion and Cold Migration.

5. You can monitor the progress and completion of the vMotion from the HCX > Migration dashboard.

ece Gitarbiant	¢	Migration	(5 Martin 1 (17 1 (17 1	( wattown )						
C bie Ferry		D Hereine VM	. Stars	par Warrany LONG	Poges		flat for	Data		
To transport Analytics		■ s300-vcsa01ebcdc.com →	xcsa-6790117458c8	europe-west3 ave apor	N.					
Services	100	C MINANCAE		2.00	O Mpakai Carpina		241/wat 200	war in Appleton	impated.	
Initian Schemme		Constanting of Consta		Same III and a f			Page 16 Page 10	Which the star and a second started	460	
D Quality Barmaty	-	perinates Interactor D Data	etter	Disfand G fande	THE AS SMUTCH :		Miguiter Once 11	Within a site acts hist juritude	418	
Supton.		Manager Contract of Barray	-				Abjustus Profile	is status		
A Americanteration							Manhood Street	Discharger .		
Chapert			OWDR + O O	2.00 ton sub-sources			Service de la reine	and woman in the second s		
						1000		26000		
						Iveau	200	Collecting resurs delate		
						1 percept		Reading O applance		
						1 - Sinte apr.	-46	Collecting Largel debets		
						2.000.000	19	Recordports source Resulty Agent		
						3 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1	1.00	WARD COLLEGE AND ADDRESS		
						2 100 400	- 10	Creating placetoolese VM for eMotion	at two or sills	
						Tree op	100	Martery tencols late or larger one		
						<ol> <li>Costupe</li> </ol>	120	Stading valuatie late on taken with		
								(These local)		
		C > WHOR ID MICE	-	49.00 + 00 2	O Mpater Corpore		10.5444.621 11.31	AN EST ST Migration o	unyinter.	
· Aecent Tasks Alarm	i i				<ul> <li>A state of the sta</li></ul>	1	Commissions Press	Samuel		
v <u>Aecest Tasks</u> Alam task Name * Te	ni nir	17   mag (		Indiator	Re .	Bartine L	- Competition Long			
<ul> <li>Recent Tasks Alarm</li> <li>Test Name</li> <li>Tree patient virtue macrone</li> </ul>	ni niet B. Miniszoczał	T Inna T I B Conpeter	••••	Inflator Encluid COM-Administration	Tes	Bartine 1 PM	02/05/0112 3 34 11 PM	4302-mails shop, can		
<ul> <li>Becest tasks Allers</li> <li>Mark * Te</li> <li>Tespiter vitue nucleus Allers</li> </ul>	ne net B. Howarocvit B. Howarocvit	Initia     Concretese     Of Concretese	•	NUMBER FINCTIC CONTAGENEERS and DACIDC CONTAGENEERS and	2 mg 5 mg	Bark/2021, 2 54 21 PM 62/96/2023, 2 54 21 PM	02/16/2013 234 11/94 02/16/2013 234 11/94	allog -scale reaction 2008 allog -scale forces com		

# Conclusion

Whether you're targeting all-cloud or hybrid cloud and data residing on any type/vendor storage in onpremises, Cloud Volume Service and HCX provide excellent options to deploy and migrate the application workloads while reducing the TCO by making the data requirements seamless to the application layer. Whatever the use case, choose Google Cloud VMware Engine along with Cloud Volume Service for rapid realization of cloud benefits, consistent infrastructure, and operations across on-premises and multiple clouds, bidirectional portability of workloads, and enterprise-grade capacity and performance. It is the same familiar process and procedures used to connect the storage and migrate VMs using VMware vSphere Replication, VMware vMotion, or even network file copy (NFC).

#### Takeaways

The key points of this document include:

- You can now use Cloud Volume Service as a datastore on Google Cloud VMware Engine SDDC.
- You can easily migrate data from on-premises to Cloud Volume Service datastore.
- You can easily grow and shrink the Cloud Volume Service datastore to meet the capacity and performance requirements during migration activity.

# Videos from Google and VMware for reference

## From Google

- Deploy HCX Connector with GCVE
- Configure HCX ServiceMesh with GCVE
- Migrate VM with HCX to GCVE

#### From VMware

- HCX Connector deployment for GCVE
- HCX ServiceMesh configuration for GCVE
- HCX Workload Migration to GCVE

#### Where to find additional information

To learn more about the information described in this document, refer to the following website links:

· Google Cloud VMware Engine documentation

https://cloud.google.com/vmware-engine/docs/overview

Cloud Volume Service documentation

https://cloud.google.com/architecture/partners/netapp-cloud-volumes

• VMware HCX User Guide

https://docs.vmware.com/en/VMware-HCX/index.html

# VM Migration to NetApp Cloud Volume Service NFS Datastore on Google Cloud VMware Engine using Veeam Replication feature

Customers who currently use Veeam for their data protection requirements continue using that solution to migrate the workloads to GCVE and enjoy the benefits of NetApp Cloud Volume Service NFS Datastores.

# Overview

Authors: Suresh Thoppay, NetApp

VM Workloads running on VMware vSphere can be migrated to Google Cloud VMware Engine (GCVE) utilizing Veeam Replication feature.

This document provides a step-by-step approach for setting up and performing VM migration that uses NetApp Cloud Volume Service, Veeam, and the Google Cloud VMware Engine (GCVE).



#### Assumptions

This document assumes you have either Google Cloud VPN or Cloud Interconnect or other networking option in place to establish network connectivity from existing vSphere servers to Google Cloud VMware Engine.



There are multiple options for connecting on-premises datacenters to Google Cloud, which prevents us from outlining a specific workflow in this document. Refer to the Google Cloud documentation for the appropriate on-premises-to-Google connectivity method.

# **Deploying the Migration Solution**

#### **Solution Deployment Overview**

- 1. Make sure NFS datastore from NetApp Cloud Volume Service is mounted on GCVE vCenter.
- 2. Ensure Veeam Backup Recovery is deployed on existing VMware vSphere environment
- 3. Create Replication Job to start replicating virtual machines to Google Cloud VMware Engine instance.
- 4. Perform Failover of Veeam Replication Job.
- 5. Perform Permanent Failover on Veeam.

#### **Deployment Details**

#### Make sure NFS datastore from NetApp Cloud Volume Service is mounted on GCVE vCenter

Login to GCVE vCenter and ensure NFS datastore with sufficient space is available. If not, Please refer Mount NetApp CVS as NFS datastore on GCVE

#### Ensure Veeam Backup Recovery is deployed on existing VMware vSphere environment

Please refer Veeam Replication Components documentation to install required components.

# Create Replication Job to start replicating virtual machines to Google Cloud VMware Engine instance.

Both on-premises vCenter and GCVE vCenter needs to be registered with Veeam. Setup vSphere VM Replication Job Here is a video explaining how to Configure Replication Job.



Replica VM can have different IP from the source VM and can also be connected to different port group. For more details, check the video above.

#### Perform Failover of Veeam Replication Job

To Migrate VMs, perform Perform Failover

#### Perform Permanent Failover on Veeam.

To treat GCVE as your new source environment, perform Permanent Failover

#### Benefits of this solution

- Existing Veeam backup infrastructure can be utilized for migration.
- Veeam Replication allows changing VM IP addresses on target site.
- Has ability to remap existing data replicated outside of Veeam (like replicated data from BlueXP)
- Has ability to specify different network portgroup on target site.
- Can specify the order of VMs to power on.
- Utilizes VMware Change Block Tracking to minimize the amount of data to send across WAN.
- Capability to execute pre and post scripts for replication.
- · Capability to execute pre and post scripts for snapshots.

# Region Availability – Supplemental NFS datastore for Google Cloud Platform (GCP)

Learn more about the the Global Region support for GCP, GCVE and CVS.



NFS datastore will be available in regions where both services (GCVE and CVS Performance) are available.

Supplemental NFS datastore for GCVE is supported with NetApp Cloud Volume Service.



Only CVS-Performance volumes can be used for GCVE NFS Datastore. For the available location, refer Global Region Map

Google Cloud VMware Engine is available at following locations image::gcve\_regions\_Mar2023.png[] To minimize latency, NetApp CVS Volume and GCVE where you intent to mount the volume should be in same availability zone.

Work with Google and NetApp Solution Architects for availability and TCO optimizations.

# Security overview - NetApp Cloud Volumes Service (CVS) in Google Cloud

Oliver Krause, Justin Parisi, NetApp

Security, particularly in the cloud where infrastructure is outside of the control of storage administrators, is paramount to trusting your data to service offerings provided by cloud providers. This document is an overview of the security offerings that NetApp Cloud Volumes Service provides in Google Cloud.

## Intended audience

This document's intended audience includes, but is not limited to, the following roles:

- Cloud providers
- Storage administrators
- Storage architects
- · Field resources
- Business decision makers

If you have questions about the content of this technical report, see the section "Contact us."

Abbreviation	Definition
CVS-SW	Cloud Volumes Service, Service Type CVS
CVS-Performance	Cloud Volume Service, Service Type CVS- Performance
PSA	

#### How Cloud Volumes Service in Google Cloud secures your data

Cloud Volumes Service in Google Cloud provides a multitude of ways to natively secure your data.

#### Secure architecture and tenancy model

Cloud Volumes Service provides a secure architecture in Google Cloud by segmenting the service management (control plane) and the data access (data plane) across different endpoints so that neither can impact the other (see the section "Cloud Volumes Service architecture"). It uses Google's private services access (PSA) framework to provide the service. This framework distinguishes between the service producer, which is provided and operated by NetApp, and the service consumer, which is a Virtual Private Cloud (VPC) in a customer project, hosting the clients that want to access Cloud Volumes Service file shares.

In this architecture, tenants (see the section "Tenancy model") are defined as Google Cloud projects that are completely isolated from each other unless explicitly connected by the user. Tenants allow complete isolation of data volumes, external name services, and other essential pieces of the solution from other tenants using the Cloud Volumes Service volume platform. Because the Cloud Volumes Service platform is connected through VPC peering, that isolation applies to it also. You can enable sharing of Cloud Volumes Service volumes between multiple projects by using a shared-VPC (see the section "Shared VPCs"). You can apply access controls to SMB shares and NFS exports to limit who or what can view or modify datasets.

# Strong identity management for the control plane

In the control plane where Cloud Volumes Service configuration takes place, identity management is managed by using Identity Access Management (IAM). IAM is a standard service that enables you to control authentication (logins) and authorization (permissions) to Google Cloud project instances. All configuration is performed with Cloud Volumes Service APIs over a secure HTTPS transport using TLS 1.2 encryption, and authentication is performed by using JWT tokens for added security. The Google console UI for Cloud Volumes Service translates user input into Cloud Volumes Service API calls.

# Security hardening - Limiting attack surfaces

Part of effective security is limiting the number of attack surfaces available in a service. Attack surfaces can include a variety of things, including data at-rest, in-flight transfers, logins, and the datasets themselves.

A managed service removes some of the attack surfaces inherently in its design. Infrastructure management, as described in the section "Service operation," is handled by a dedicated team and is automated to reduce the number of times a human actually touches configurations, which helps reduce the number of intentional and unintentional errors. Networking is fenced off so that only necessary services can access one another. Encryption is baked into the data storage and only the data plane needs security attention from Cloud Volumes Service administrators. By hiding most of the management behind an API interface, security is achieved by limiting the attack surfaces.

# Zero Trust model

Historically, IT security philosophy has been to trust but verify, and manifested as relying solely on external mechanisms (such as firewalls and intrusion detection systems) to mitigate threats. However, attacks and breaches evolved to bypass the verification in environments through phishing, social engineering, insider threats and other methods that provide the verification to enter networks and wreak havoc.

Zero Trust has become a new methodology in security, with the current mantra being "trust nothing while still verifying everything." Therefore, nothing is allowed access by default. This mantra is enforced in a variety of ways, including standard firewalls and intrusion detection systems (IDS) and also with the following methods:

- Strong authentication methods (such as AES-encrypted Kerberos or JWT tokens)
- Single strong sources of identities (such as Windows Active Directory, Lightweight Directory Access Protocol (LDAP), and Google IAM)
- Network segmentation and secure multitenancy (only tenants are allowed access by default)
- · Granular access controls with Least Privileged Access policies
- · Small exclusive lists of dedicated, trusted administrators with digital audit and paper trails

Cloud Volumes Service running in Google Cloud adheres to the Zero Trust model by implementing the "trust nothing, verify everything" stance.

# Encryption

Encrypt data at-rest (see the section "Data encryption at rest") by using XTS-AES-256 ciphers with NetApp Volume Encryption (NVE) and in-flight with "SMB encryption" or NFS Kerberos 5p support. Rest easy knowing cross-region replication transfers are protected by TLS 1.2 encryption (see the section "Cross-region replication"). In addition, Google networking also provides encrypted communications (see the section "Data encryption in transit") for an added layer of protection against attacks. For more information about transport encryption, see the section "Google Cloud network".

# Data protection and backups

Security isn't just about the prevention of attacks. It is also about how we recover from attacks if or when they occur. This strategy includes data protection and backups. Cloud Volumes Service provides methods to replicate to other regions in case of outages (see the section "Cross-region replication") or if a dataset is affected by a ransomware attack. It can also perform asynchronous backups of data to locations outside of the Cloud Volumes Service instance by using Cloud Volumes Service backup. With regular backups, mitigation of security events can take less time and save money and angst for administrators.

## Fast ransomware mitigation with industry leading Snapshot copies

In addition to data protection and backups, Cloud Volumes Service provides support for immutable Snapshot copies (see the section "Immutable Snapshot copies") of volumes that allow recovery from ransomware attacks (see the section "Service operation") within seconds of discovering the issue and with minimal disruption. Recovery time and effects depend on the Snapshot schedule, but you can create Snapshot copies that provide as little as one-hour deltas in ransomware attacks. Snapshot copies have a negligible effect on performance and capacity usage and are a low-risk, high-reward approach to protecting your datasets.

#### Security considerations and attack surfaces

The first step in understanding how to secure your data is identifying the risks and potential attack surfaces.

These include (but are not limited to) the following:

- · Administration and logins
- Data at rest
- Data in flight
- Network and firewalls
- · Ransomware, malware, and viruses

Understanding attack surfaces can help you to better secure your environments. Cloud Volumes Service in Google Cloud already considers many of these topics and implements security functionality by default, without any administrative interaction.

# **Ensuring secure logins**

When securing your critical infrastructure components, it is imperative to make sure that only approved users can log in and manage your environments. If bad actors breach your administrative credentials, then they have the keys to the castle and can do anything they want—change configurations, delete volumes and backups, create backdoors, or disable Snapshot schedules.

Cloud Volumes Service for Google Cloud provides protection against unauthorized administrative logins through the obfuscation of storage as a service (StaaS). Cloud Volumes Service is completely maintained by the cloud provider with no availability to login externally. All setup and configuration operations are fully automated, so a human administrator never has to interact with the systems except in very rare circumstances.

If login is required, Cloud Volumes Service in Google Cloud secures logins by maintaining a very short list of trusted administrators that have access to log in to the systems. This gatekeeping helps reduce the number of potential bad actors with access. Additionally, the Google Cloud networking hides the systems behind layers of network security and exposes only what is needed to the outside world. For information about the Google Cloud, Cloud Volumes Service architecture, see the section "Cloud Volumes Service architecture."

## Cluster administration and upgrades

Two areas with potential security risks include cluster administration (what happens if a bad actor has admin access) and upgrades (what happens if a software image is compromised).

## Storage administration protection

Storage provided as a service removes the added risk of exposure to administrators by removing that access to end users outside of the cloud data center. Instead, the only configuration done is for the data access plane by customers. Each tenant manages their own volumes, and no tenant can reach other Cloud Volumes Service instances. The service is managed by automation, with a very small list of trusted administrators given access to the systems through the processes covered in the section "Service operation."

The CVS-Performance service type offers cross-region replication as an option to provide data protection to a different region in the event of a region failure. In those cases, Cloud Volumes Service can be failed over to the unaffected region to maintain data access.

#### Service upgrades

Updates help protect vulnerable systems. Each update provides security enhancements and bug fixes that minimize attack surfaces. Software updates are downloaded from centralized repositories and are validated before the updates are allowed to verify that official images are used and that the upgrades are not compromised by bad actors.

With Cloud Volumes Service, updates are handled by the cloud provider teams, which removes risk exposure for administrator teams by providing experts well versed in configuration and upgrades that have automated and fully tested the process. Upgrades are nondisruptive, and Cloud Volumes Service maintains the latest updates for best overall results.

For information about the administrator team that performs these service upgrades, see the section "Service operation."

# Securing data at-rest

Data-at-rest encryption is important to protect sensitive data in the event of a disk that is stolen, returned, or repurposed. Data in Cloud Volumes Service is protected at rest by using software-based encryption.

- · Google-generated keys are used for CVS-SW.
- For CVS-Performance, the per-volume keys are stored in a key manager built into Cloud Volumes Service, which uses NetApp ONTAP CryptoMod to generate AES-256 encryption keys. CryptoMod is listed on the CMVP FIPS 140-2 validated modules list. See FIPS 140-2 Cert #4144.

Starting in November 2021, preview Customer-managed Encryption (CMEK) functionality was made available for CVS-Performance. This functionality allows you to encrypt the per-volume keys with per-project, per-region master-keys that are hosted in Google Key Management Service (KMS). KMS enables you to attach external key managers.

For details about how to configure KMS for CVS-Performance, see the Cloud Volumes Service documentation.

For more information about architecture, see the section "Cloud Volumes Service architecture."

# Securing data in-flight

In addition to securing data at rest, you must also be able to secure data when it is in flight between the Cloud Volumes Service instance and a client or replication target. Cloud Volumes Service provides encryption for in-

flight data over NAS protocols by using encryption methods such as SMB encryption using Kerberos, the signing/sealing of packets, and NFS Kerberos 5p for end-to-end encryption of data transfers.

Replication of Cloud Volumes Service volumes uses TLS 1.2, which takes advantage of AES-GCM encryption methods.

Most insecure in-flight protocols such as telnet, NDMP, and so on are disabled by default. DNS, however, is not encrypted by Cloud Volumes Service (no DNS Sec support) and should be encrypted by using external network encryption when possible. See the section "Data encryption in transit" for more information about securing data in-flight.

For information about NAS protocol encryption, see the section "NAS protocols."

# Users and groups for NAS permissions

Part of securing your data in the cloud involves proper user and group authentication, where the users accessing the data are verified as real users in the environment and the groups contain valid users. These users and groups provide initial share and export access, as well as permission validation for files and folders in the storage system.

Cloud Volumes Service uses standard Active Directory-based Windows user and group authentication for SMB shares and Windows-style permissions. The service can also leverage UNIX identity providers such as LDAP for UNIX users and groups for NFS exports, NFSv4 ID validation, Kerberos authentication, and NFSv4 ACLs.



Currently only Active Directory LDAP is supported with Cloud Volumes Service for LDAP functionality.

#### Detection, prevention and mitigation of ransomware, malware, and viruses

Ransomware, malware, and viruses are a persistent threat to administrators, and detection, prevention, and mitigation of those threats are always top of mind for enterprise organizations. A single ransomware event on a critical dataset can potentially cost millions of dollars, so it is beneficial to do what you can to minimize the risk.

Although Cloud Volumes Service currently doesn't include native detection or prevention measures, such as antivirus protection or automatic ransomware detection, there are ways to quickly recover from a ransomware event by enabling regular Snapshot schedules. Snapshot copies are immutable and read only pointers to changed blocks in the file system, are near instantaneous, have minimal impact on performance, and only use up space when data is changed or deleted. You can set schedules for Snapshot copies to match your desired acceptable recovery point objective (RPO)/recovery time objective (RTO) and can keep up to 1,024 Snapshot copies per volume.

Snapshot support is included at no additional cost (beyond data storage charges for changed blocks/data retained by Snapshot copies) with Cloud Volumes Service and, in the event of a ransomware attack, can be used to roll back to a Snapshot copy before the attack occurred. Snapshot restores take just seconds to complete, and you then can get back to serving data as normal. For more information, see The NetApp Solution for Ransomware.

Preventing ransomware from affecting your business requires a multilayered approach that includes one or more of the following:

- Endpoint protection
- · Protection against external threats through network firewalls
- · Detection of data anomalies

- Multiple backups (onsite and offsite) of critical datasets
- · Regular restore tests of backups
- Immutable read-only NetApp Snapshot copies
- Multifactor authentication for critical infrastructure
- · Security audits of system logins

This list is far from exhaustive but is a good blueprint to follow when dealing with the potential of ransomware attacks. Cloud Volumes Service in Google Cloud provides several ways to protect against ransomware events and reduce their effects.

#### Immutable Snapshot copies

Cloud Volumes Service natively provides immutable read-only Snapshot copies that are taken on a customizable schedule for quick point-in-time recovery in the event of data deletion or if an entire volume has been victimized by a ransomware attack. Snapshot restores to previous good Snapshot copies are fast and minimize data loss based on the retention period of your Snapshot schedules and RTO/RPO. The performance effect with Snapshot technology is negligible.

Because Snapshot copies in Cloud Volumes Service are read-only, they cannot be infected by ransomware unless the ransomware has proliferated into the dataset unnoticed and Snapshot copies have been taken of the data infected by ransomware. This is why you must also consider ransomware detection based on data anomalies. Cloud Volumes Service does not currently provide detection natively, but you can use external monitoring software.

#### **Backups and restores**

Cloud Volumes Service provides standard NAS client backup capabilities (such as backups over NFS or SMB).

- CVS-Performance offers cross-region volume replication to other CVS-Performance volumes. For more information, see volume replication in the Cloud Volumes Service documentation.
- CVS-SW offers service-native volume backup/restore capabilities. For more information, see cloud backup in the Cloud Volumes Service documentation.

Volume replication provides an exact copy of the source volume for fast failover in the case of a disaster, including ransomware events.

# **Cross-region replication**

CVS-Performance enables you to securely replicate volumes across Google Cloud regions for data protection and archive use cases by using TLS1.2 AES 256 GCM encryption on a NetApp-controlled backend service network using specific interfaces used for replication running on Google's network. A primary (source) volume contains the active production data and replicates to a secondary (destination) volume to provide an exact replica of the primary dataset.

Initial replication transfers all blocks, but updates only transmit the changed blocks in a primary volume. For instance, if a 1TB database that resides on a primary volume is replicated to the secondary volume, then 1TB of space is transferred on the initial replication. If that database has a few hundred rows (hypothetically, a few MB) that change between the initialization and the next update, only the blocks with the changed rows are replicated to the secondary (a few MB). This helps to make sure that the transfer times remain low and keeps replication charges down.

All permissions on files and folders are replicated to the secondary volume, but share access permissions

(such as export policies and rules or SMB shares and share ACLs) must be handled separately. In the case of a site failover, the destination site should leverage the same name services and Active Directory domain connections to provide consistent handling of user and group identities and permissions. You can use a secondary volume as a failover target in the event of a disaster by breaking the replication relationship, which converts the secondary volume to read-write.

Volume replicas are read-only, which provides an immutable copy of data offsite for quick recovery of data in instances where a virus has infected data or ransomware has encrypted the primary dataset. Read-only data won't be encrypted, but, if the primary volume is affected and replication occurs, the infected blocks also replicate. You can use older, non-affected Snapshot copies to recover, but SLAs might fall out of range of the promised RTO/RPO depending on how quickly an attack is detected.

In addition, you can prevent malicious administrative actions, such as volume deletions, Snapshot deletions, or Snapshot schedule changes, with cross-region replication (CRR) management in Google Cloud. This is done by creating custom roles that separate volume administrators, who can delete source volumes but not break mirrors and therefore cannot delete destination volumes, from CRR administrators, who cannot perform any volume operations. See Security Considerations in the Cloud Volumes Service documentation for permissions allowed by each administrator group.

# **Cloud Volumes Service backup**

Although Cloud Volumes Service provides high data durability, external events can cause data loss. In the event of a security event such as a virus or ransomware, backups and restores become critical for resumption of data access in a timely manner. An administrator might accidentally delete a Cloud Volumes Service volume. Or users simply want to retain backup versions of their data for many months and keeping the extra Snapshot copy space inside the volume becomes a cost challenge. Although Snapshot copies should be the preferred way to keep backup versions for the last few weeks to restore lost data from them, they are sitting inside the volume and are lost if the volume goes away.

For all these reasons, NetApp Cloud Volumes Service offers backup services through Cloud Volumes Service backup.

Cloud Volumes Service backup generates a copy of the volume on Google Cloud Storage (GCS). It only backs up the actual data stored within the volume, not the free space. It works as incremental forever, meaning it transfers the volume content once and from there on continues backing up changed data only. Compared to classical backup concepts with multiple full backups, it saves large amounts of backup storage, reducing cost. Because the monthly price of backup space is lower compared to a volume, it is an ideal place to keep backup versions longer.

Users can use a Cloud Volumes Service backup to restore any backup version to the same or a different volume within the same region. If the source volume is deleted, the backup data is retained and needs to be managed (for example, deleted) independently.

Cloud Volumes Service backup is built into Cloud Volumes Service as option. Users can decide which volumes to protect by activating Cloud Volumes Service backup on a per-volume basis. See the Cloud Volumes Service backup documentation for information about backups, the number of maximum backup versions supported, scheduling, and pricing.

All backup data of a project is stored within a GCS bucket, which is managed by the service and not visible to the user. Each project uses a different bucket. Currently, the buckets are in same region as the Cloud Volumes Service volumes, but more options are being discussed. Consult the documentation for the latest status.

Data transport from a Cloud Volumes Service bucket to GCS uses service-internal Google networks with HTTPS and TLS1.2. Data is encrypted at-rest with Google-managed keys.
To manage Cloud Volumes Service backup (creating, deleting, and restoring backups), a user must have the roles/netappcloudvolumes.admin role.

# Architecture

# Overview

Part of trusting a cloud solution is understanding the architecture and how it is secured. This section calls out different aspects of the Cloud Volumes Service architecture in Google to help alleviate potential concerns about how data is secured, as well as call out areas where additional configuration steps might be required to obtain the most secure deployment.

The general architecture of Cloud Volumes Service can be broken down into two main components: the control plane and the data plane.

# **Control plane**

The control plane in Cloud Volumes Service is the backend infrastructure managed by Cloud Volumes Service administrators and NetApp native automation software. This plane is completely transparent to end users and includes networking, storage hardware, software updates, and so on to help deliver value to a cloud-resident solution such as Cloud Volumes Service.

# Data plane

The data plane in Cloud Volumes Service includes the actual data volumes and the overall Cloud Volumes Service configuration (such as access control, Kerberos authentication, and so on). The data plane is entirely under the control of the end users and the consumers of the Cloud Volumes Service platform.

There are distinct differences in how each plane is secured and managed. The following sections cover these differences, starting with a Cloud Volumes Service architecture overview.

# **Cloud Volumes Service architecture**

In a manner similar to other Google Cloud native services such as CloudSQL, Google Cloud VMware Engine (GCVE), and FileStore, Cloud Volumes Service uses Google PSA to deliver the service. In PSA, services are built inside a service producer project, which uses VPC network peering to connect to the service consumer. The service producer is provided and operated by NetApp, and the service consumer is a VPC in a customer project, hosting the clients that want to access Cloud Volumes Service file shares.

The following figure, referenced from the architecture section of the Cloud Volumes Service documentation, shows a high-level view.



The part above the dotted line shows the control plane of the service, which controls the volume lifecycle. The part below the dotted line shows the data plane. The left blue box depicts the user VPC (service consumer), the right blue box is the service producer provided by NetApp. Both are connected through VPC peering.

# Tenancy model

In Cloud Volumes Service, individual projects are considered unique tenants. This means that manipulation of volumes, Snapshot copies, and so on are performed on a per- project basis. In other words, all volumes are owned by the project that they were created in and only that project can manage and access the data inside of them by default. This is considered the control plane view of the service.

# **Shared VPCs**

On the data plane view, Cloud Volumes Service can connect to a shared VPC. You can create volumes in the hosting project or in one of the service projects connected to the shared VPC. All projects (host or service) connected to that shared VPC are able to reach the volumes at the network layer (TCP/IP). Because all clients with network connectivity on the shared- VPC can potentially access the data through NAS protocols, access control on the individual volume (such as user/group access control lists (ACLs) and hostnames/IP addresses for NFS exports) must be used to control who can access the data.

You can connect Cloud Volumes Service to up to five VPCs per customer project. On the control plane, the project enables you to manage all created volumes, no matter which VPC they are connected to. On the data plane, VPCs are isolated from one another, and each volume can only be connected to one VPC.

Access to the individual volumes is controlled by protocol specific (NFS/SMB) access control mechanisms.

In other words, on the network layer, all project s connected to the shared VPC are able to see the volume, while, on the management side, the control plane only allows the owner project to see the volume.

# **VPC Service Controls**

VPC Service Controls establish an access control perimeter around Google Cloud services that are attached to

the internet and are accessible worldwide. These services provide access control through user identities but cannot restrict which network location requests originate from. VPC Service Controls close that gap by introducing the capabilities to restrict access to defined networks.

The Cloud Volumes Service data plane is not connected to the external internet but to private VPCs with welldefined network boundaries (perimeters). Within that network, each volume uses protocol-specific access control. Any external network connectivity is explicitly created by Google Cloud project administrators. The control plane, however, does not provide the same protections as the data plane and can be accessed by anyone from anywhere with valid credentials (JWT tokens).

In short, the Cloud Volumes Service data plane provides the capability of network access control, without the requirement to support VPC Service Controls and does not explicitly use VPC Service Controls.

# Packet sniffing/trace considerations

Packet captures can be useful for troubleshooting network issues or other problems (such as NAS permissions, LDAP connectivity, and so on), but can also be used maliciously to gain information about network IP addresses, MAC addresses, user and group names, and what level of security is being used on endpoints. Because of the way Google Cloud networking, VPCs, and firewall rules are configured, unwanted access to network packets should be difficult to obtain without user login credentials or JWT tokens into the cloud instances. Packet captures are only possible on endpoints (such as virtual machines (VMs)) and only possible on endpoints internal to the VPC unless a shared VPC and/or external network tunnel/IP forwarding is in use to explicitly allow external traffic to endpoints. There is no way to sniff traffic outside of the clients.

When shared VPCs are used, in-flight encryption with NFS Kerberos and/or SMB encryption can mask much of the information gleaned from traces. However, some traffic is still sent in plaintext, such as DNS and LDAP queries. The following figure shows a packet capture from a plaintext LDAP query originating from Cloud Volumes Service and the potential identifying information that is exposed. LDAP queries in Cloud Volumes Service currently do not support encryption or LDAP over SSL. CVS-Performance support LDAP signing, if requested by Active Directory. CVS-SW does not support LDAP signing.





unixUserPassword is queried by LDAP and is not sent in plaintext but instead in a salted hash. By default, Windows LDAP does not populate the unixUserPassword fields. This field is only required if you need to leverage Windows LDAP for interactive logins through LDAP to clients. Cloud Volumes Service does not support interactive LDAP logins to the instances.

The following figure shows a packet capture from an NFS Kerberos conversation next to a capture of NFS over AUTH\_SYS. Note how the information available in a trace differs between the two and how enabling in-flight encryption offers greater overall security for NAS traffic.

		IP addresses of the NFS client an	nd CVS instance	Genericized NFS call/reply	
No	Time	Source	Destination	Protocol Length Info	
	380 9,218014	10,193,67,225	10,193,67,219	NES 346 V4 Call (Reply In 381)	
Ť	381 9,218480	10,193,67,219	10.193.67.225	NFS 426 V4 Reply (Call In 380)	
	382 9 218641	10.193.67.225	10.193.67.219	NFS 370 V4 Call (Renly In 397)	
	207 0 260025	10.103.67.225	10.193.07.215	NES 458 V4 Call (Keply in 397)	
	597 9.509055	10.195.07.219	10.195.07.225	NFS 450 V4 Reply (Call III 562)	
>	Frame 381: 426 bytes	on wire (3408 bits), 4	26 bytes captured (	(3408 bits)	
>	Ethernet II, Src: Int	elCor_7f:da:bc (90:e2:	ba:7f:da:bc), Dst:	VMware_a0:2c:2d (00:50:56:a0:2c:2d)	
>	Internet Protocol Ver	sion 4, Src: 10.193.67	.219, Dst: 10.193.0	67.225	
>	Transmission Control	Protocol, Src Port: 20	49, Dst Port: 738,	Seq: 6305, Ack: 6569, Len: 360	
>	Remote Procedure Call	, Type:Reply XID:0xef5	e998d		
~	GSS-Wrap			GSS wrapped NFS calls/replies with no other identifying information	
	Length: 300				
	GSS Data: 050407ff	3000000000000000259134	51ee1d43d298cf3031		
	krb5_blob: 050407f	F00000000000000000025913	451ee1d43d298cf3031	1	
~	Network File System				
	[Program Version: 4	41			
	[V4 Procedure: COM	POUND (1)]			
	33 0.958480 34 0.958784 35 0.958784 35 0.959284 ) Opcode: PUTFH (22) ) Opcode: SETATTR (34) ♥ Opcode: GETATTR (9) Status: NFS4_OK (0) ♥ Attr mask[0]: 0x00100112 ) reqd_attr: Type (1) ) reqd_attr: Size (4) ) reqd_attr: Size (4) ) reqd_attr: FSID (8) ♥ reco_attr: FIIeId (20) fileid: 9232254136 ♥ Attr mask[1]: 0x00100254136 ♥ Attr mask[1]: 0x00100254136	.67.201 10.193.67.204 .67.204 10.193.67.201 .67.201 10.193.67.204 .67.201 10.193.67.204 .67.201 File ID .597092620 . (Mode Numlinks, Owner, Owner,	NFS 458 V4 Rep NFS 306 V4 Cal NFS 358 V4 Rep leId)	Access. Time Metadata. Time Modify. Mounted on FileId)	
	v reco_attr: Mode (33)	(node, numerics, owner, owner	_droup; space_osed; rime_	Petroistion informatic	m
	> mode: 0644, Name:	Unknown, Read permission for o	wner, Write permission for	r owner, Read permission for group, Read permission for other	rs.
	<pre>&gt; reco_attr: NumLinks (     reco_attr: Owner (36)     &gt; fattr4_owner: root     reco_attr: Owner_Grou     fattr4_owner_group     reco_attr: Space_Used     reco_attr: Time_Acces     reco_attr: Time_Metad     reco_attr: Time_Modif</pre>	Owner and group ID strings           @NTAP.LOCAL           p (37)           : root@NTAP.LOCAL           (45)           s (47)           ata (52)           y (53)			
	> reco_attr: Mounted_on	_LITETO (22)			

## VM network interfaces

One trick attackers might attempt is to add a new network interface card (NIC) to a VM in promiscuous mode (port mirroring) or enable promiscuous mode on an existing NIC in order to sniff all traffic. In Google Cloud, adding a new NIC requires a VM to be shut down entirely, which creates alerts, so attackers cannot do this unnoticed.

In addition, NICs cannot be set to promiscuous mode at all and will trigger alerts in Google Cloud.

# **Control plane architecture**

All management actions to Cloud Volumes Service are done through API. Cloud Volumes Service management integrated into the GCP Cloud Console also uses the Cloud Volumes Service API.

# **Identity and Access Management**

Identity and Access Management (IAM) is a standard service that enables you to control authentication (logins) and authorization (permissions) to Google Cloud project instances. Google IAM provides a full audit trail of permissions authorization and removal. Currently Cloud Volumes Service does not provide control plane auditing.

# Authorization/permission overview

IAM offers built-in, granular permissions for Cloud Volumes Service. You can find a complete list of granular permissions here.

IAM also offers two predefined roles called netappcloudvolumes.admin and netappcloudvolumes.viewer. These roles can be assigned to specific users or service accounts.

Assign appropriate roles and permission to allow IAM users to manage Cloud Volumes Service.

Examples for using granular permissions include the following:

- Build a custom role with only get/list/create/update permissions so that users cannot delete volumes.
- Use a custom role with only snapshot.\* permissions to create a service account that is used to build application- consistent Snapshot integration.
- Build a custom role to delegate volumereplication.\* to specific users.

# Service accounts

To make Cloud Volumes Service API calls through scripts or Terraform, you must create a service account with the roles/netappcloudvolumes.admin role. You can use this service account to generate the JWT tokens required to authenticate Cloud Volumes Service API requests in two different ways:

- Generate a JSON key and use Google APIs to derive a JWT token from it. This is the simplest approach, but it involves manual secrets (the JSON key) management.
- Use Service account impersonation with roles/iam.serviceAccountTokenCreator. The code (script, Terraform, and so on.) runs with Application Default Credentials and impersonates the service account to gain its permissions. This approach reflects Google security best practices.

See Creating your service account and private key in the Google cloud documentation for more information.

# **Cloud Volumes Service API**

Cloud Volumes Service API uses a REST-based API by using HTTPS (TLSv1.2) as the underlying network transport. You can find the latest API definition here and information about how to use the API at Cloud Volumes APIs in the Google cloud documentation.

The API endpoint is operated and secured by NetApp using standard HTTPS (TLSv1.2) functionality.

# JWT tokens

Authentication to the API is performed with JWT bearer tokens (RFC-7519). Valid JWT tokens must be obtained by using Google Cloud IAM authentication. This must be done by fetching a token from IAM by providing a service account JSON key.

# Audit logging

Currently, no user-accessible control plane audit logs are available.

# Data plane architecture

Cloud Volumes Service for Google Cloud leverages the Google Cloud private services access framework. In this framework, users can connect to the Cloud Volumes Service. This framework uses Service Networking and VPC peering constructs like other Google Cloud services, ensuring complete isolation between tenants.

For an architecture overview of Cloud Volumes Service for Google Cloud, see Architecture for Cloud Volumes Service.

User VPCs (standalone or shared) are peered to VPCs within Cloud Volumes Service managed tenant projects, which hosts the volumes.



The preceding figure shows a project (the CVS consumer project in the middle) with three VPC networks connected to Cloud Volumes Service and multiple Compute Engine VMs (GCE1-7) sharing volumes:

- VPC1 allows GCE1 to access volumes A and B.
- VPC2 allows GCE2 and GCE4 to access volume C.

• The third VPC network is a shared VPC, shared with two service projects. It allows GCE3, GCE4, GCE5, and GCE6 to access volumes D and E. Shared VPC networks are only supported for volumes of the CVS-Performance service type.



GCE7 cannot access any volume.

Data can be encrypted both in-transit (using Kerberos and/or SMB encryption) and at-rest in Cloud Volumes Service.

# Data encryption in transit

Data in transit can be encrypted at the NAS protocol layer, and the Google Cloud network itself is encrypted, as described in the following sections.

# **Google Cloud network**

Google Cloud encrypts traffic on the network level as described in Encryption in transit in the Google documentation. As mentioned in the section "Cloud Volumes Services architecture," Cloud Volumes Service is delivered out of a NetApp-controlled PSA producer project.

In case of CVS-SW, the producer tenant runs Google VMs to provide the service. Traffic between user VMs and Cloud Volumes Service VMs is encrypted automatically by Google.

Although the data path for CVS-Performance isn't fully encrypted on the network layer, NetApp and Google use a combination of IEEE 802.1AE encryption (MACSec), encapsulation (data encryption), and physically restricted networks to protect data in transit between the Cloud Volumes Service CVS-Performance service type and Google Cloud.

# **NAS** protocols

NFS and SMB NAS protocols provide optional transport encryption at the protocol layer.

# **SMB** encryption

SMB encryption provides end-to-end encryption of SMB data and protects data from eavesdropping occurrences on untrusted networks. You can enable encryption for both the client/server data connection (only available to SMB3.x capable clients) and the server/domain controller authentication.

When SMB encryption is enabled, clients that do not support encryption cannot access the share.

Cloud Volumes Service supports RC4-HMAC, AES-128-CTS-HMAC-SHA1, and AES-256-CTS-HMAC-SHA1 security ciphers for SMB encryption. SMB negotiates to the highest supported encryption type by the server.

# NFSv4.1 Kerberos

For NFSv4.1, CVS-Performance offers Kerberos authentication as described in RFC7530. You can enable Kerberos on a per-volume basis.

The current strongest available encryption type for Kerberos is AES-256-CTS-HMAC-SHA1. NetApp Cloud Volumes Service supports AES-256-CTS-HMAC-SHA1, AES-128-CTS-HMAC-SHA1, DES3, and DES for NFS. It also supports ARCFOUR-HMAC (RC4) for CIFS/SMB traffic, but not for NFS.

Kerberos provides three different security levels for NFS mounts that offer choices for how strong the Kerberos security should be.

As per RedHat's Common Mount Options documentation:

sec=krb5 uses Kerberos V5 instead of local UNIX UIDs and GIDs to authenticate users. sec=krb5i uses Kerberos V5 for user authentication and performs integrity checking of NFS operations using secure checksums to prevent data tampering. sec=krb5p uses Kerberos V5 for user authentication, integrity checking, and encrypts NFS traffic to prevent traffic sniffing. This is the most secure setting, but it also involves the most performance overhead.

As a general rule, the more the Kerberos security level has to do, the worse the performance is, as the client and server spend time encrypting and decrypting NFS operations for each packet sent. Many clients and NFS servers provide support for AES-NI offloading to the CPUs for a better overall experience, but the performance impact of Kerberos 5p (full end-to-end encryption) is significantly greater than the impact of Kerberos 5 (user authentication).

The following table shows differences in what each level does for security and performance.

Security level	Security	Performance
NFSv3—sys	<ul> <li>Least secure; plain text with numeric user IDs/group IDs</li> </ul>	<ul> <li>Best for most cases</li> </ul>
	<ul> <li>Able to view UID, GID, client IP addresses, export paths, file names, permissions in packet captures</li> </ul>	
NFSv4.x—sys	<ul> <li>More secure than NFSv3 (client IDs, name string/domain string matching) but still plain text</li> </ul>	<ul> <li>Good for sequential workloads (such as VMs, databases, large files)</li> </ul>
	<ul> <li>Able to view UID, GID, client IP addresses, name strings, domain IDs, export paths, file names, permissions in packet captures</li> </ul>	<ul> <li>Bad with high file count/high metadata (30-50% worse)</li> </ul>

Security level	Security	Performance
NFS—krb5	<ul> <li>Kerberos encryption for credentials in every NFS packet—wraps UID/GID of users/groups in RPC calls in GSS wrapper</li> <li>User requesting access to mount needs a valid Kerberos ticket (either through username/password or manual key tab exchange); ticket expires after a specified time period and user must reauthenticate for access</li> <li>No encryption for NFS operations or ancillary protocols like mount/portmapper/nlm (can see export paths, IP addresses, file handles, permissions, file names, atime/mtime in packet captures)</li> </ul>	<ul> <li>Best in most cases for Kerberos; worse than AUTH_SYS</li> </ul>
NFS—krb5i	<ul> <li>Kerberos encryption for credentials in every NFS packet—wraps UID/GID of users/groups in RPC calls in GSS wrapper</li> <li>User requesting access to mount needs a valid Kerberos ticket (either via username/password or manual key tab exchange); ticket expires after a specified time period and user must reauthenticate for access</li> <li>No encryption for NFS operations or ancillary protocols like mount/portmapper/nlm (can see export paths, IP addresses, file handles, permissions, file names, atime/mtime in packet captures)</li> <li>Kerberos GSS checksum is added to every packet to ensure nothing intercepts the packets. If checksums match, conversation is allowed.</li> </ul>	<ul> <li>Better than krb5p because the NFS payload is not encrypted; only added overhead compared to krb5 is the integrity checksum. Performance of krb5i won't be much worse than krb5 but will see some degradation.</li> </ul>

Security level	Security	Performance
NFS – krb5p	<ul> <li>Kerberos encryption for credentials in every NFS packet—wraps UID/GID of users/groups in RPC calls in GSS wrapper</li> <li>User requesting access to mount needs a valid Kerberos ticket (either via username/password or manual keytab exchange); ticket expires after specified time period and user must reauthenticate for access</li> </ul>	<ul> <li>Worst performance of the security levels; krb5p has to encrypt/decrypt more.</li> <li>Better performance than krb5p with NFSv4.x for high file count workloads.</li> </ul>
	<ul> <li>All of the NFS packet payloads are encrypted with the GSS wrapper (cannot see file handles, permissions, file names, atime/mtime in packet captures).</li> </ul>	
	<ul> <li>Includes integrity check.</li> </ul>	
	<ul> <li>NFS operation type is visible (FSINFO, ACCESS, GETATTR, and so on).</li> </ul>	
	<ul> <li>Ancillary protocols (mount, portmap, nlm, and so on) are not encrypted - (can see export paths, IP addresses)</li> </ul>	

In Cloud Volumes Service, a configured Active Directory server is used as Kerberos server and LDAP server (to lookup user identities from an RFC2307 compatible schema). No other Kerberos or LDAP servers are supported. NetApp highly recommends that you use LDAP for identity management in Cloud Volumes Service. For information on how NFS Kerberos is shown in packet captures, see the section "Packet sniffing/trace considerations."

# Data encryption at rest

All volumes in Cloud Volumes Service are encrypted-at-rest using AES-256 encryption, which means all user data written to media is encrypted and can only be decrypted with a per-volume key.

- For CVS-SW, Google-generated keys are used.
- For CVS-Performance, the per-volume keys are stored in a key manager built into the Cloud Volumes Service.

Starting in November 2021, preview customer-managed encryption keys (CMEK) functionality was made available. This enables you to encrypt the per-volume keys with a per-project, per-region master key that is hosted in Google Key Management Service (KMS). KMS enables you to attach external key managers.

For information about configuring KMS for CVS-Performance, see Setting up customer-managed encryption keys.

# Firewall

Cloud Volumes Service exposes multiple TCP ports to serve NFS and SMB shares:

- Ports required for NFS access
- Ports required for SMB access

Additionally, SMB, NFS with LDAP including Kerberos, and dual-protocol configurations require access to a Windows Active Directory domain. Active Directory connections must be configured on a per-region basis. Active Directory Domain controllers (DC) are identified by using DNS-based DC discovery using the specified DNS servers. Any of the DCs returned are used. The list of eligible DCs can be limited by specifying an Active Directory site.

Cloud Volumes Service reaches out with IP addresses from the CIDR range allocated with the gcloud compute address command while on-boarding the Cloud Volumes Service. You can use this CIDR as source addresses to configure inbound firewalls to your Active Directory domain controllers.

Active Directory Domain Controllers must expose ports to the Cloud Volumes Service CIDRs as mentioned here.

# NAS protocols

# NAS protocols overview

NAS protocols include NFS (v3 and v4.1) and SMB/CIFS (2.x and 3.x). These protocols are how CVS allows shared access to data across multiple NAS clients. In addition, Cloud Volumes Service can provide access to NFS and SMB/CIFS clients simultaneously (dual-protocol) while honoring all of the identity and permission settings on files and folders in the NAS shares. To maintain the highest possible data transfer security, Cloud Volumes Service supports protocol encryption in flight using SMB encryption and NFS Kerberos 5p.



Dual-protocol is available with CVS-Performance only.

# **Basics of NAS protocols**

NAS protocols are ways for multiple clients on a network to access the same data on a storage system, such as Cloud Volumes Service on GCP. NFS and SMB are the defined NAS protocols and operate on a client/server basis where Cloud Volumes Service acts as the server. Clients send access, read, and write requests to the server, and the server is responsible for coordinating the locking mechanisms for files, storing permissions and handling identity and authentication requests.

For example, the following general process is followed if a NAS client wants to create a new file in a folder.

1. The client asks the server for information about the directory (permissions, owner, group, file ID, available space, and so on); the server responds with the information if the requesting client and user have the

necessary permissions on the parent folder.

- 2. If the permissions on the directory allow access, the client then asks the server if the file name being created already exists in the file system. If the file name is already in use, creation fails. If the file name does not exist, the server lets the client know it can proceed.
- 3. The client issues a call to the server to create the file with the directory handle and file name and sets the access and modified times. The server issues a unique file ID to the file to make sure that no other files are created with the same file ID.
- 4. The client sends a call to check file attributes before the WRITE operation. If permissions allow it, the client then writes the new file. If locking is used by the protocol/application, the client asks the server for a lock to prevent other clients from accessing the file while locked to prevent data corruption.

# NFS

NFS is a distributed file system protocol that is an open IETF standard defined in Request for Comments (RFC) that allows anyone to implement the protocol.

Volumes in Cloud Volumes Service are shared out to NFS clients by exporting a path that is accessible to a client or set of clients. Permissions to mount these exports are defined by export policies and rules, which are configurable by Cloud Volumes Service administrators.

The NetApp NFS implementation is considered a gold standard for the protocol and is used in countless enterprise NAS environments. The following sections cover NFS and specific security features available in Cloud Volumes Service and how they are implemented.

# Default local UNIX users and groups

Cloud Volumes Service contains several default UNIX users and groups for various basic functionalities. These users and groups cannot currently be modified or deleted. New local users and groups cannot currently be added to Cloud Volumes Service. UNIX users and groups outside of the default users and groups need to be provided by an external LDAP name service.

The following table shows the default users and groups and their corresponding numeric IDs. NetApp recommends not creating new users or groups in LDAP or on the local clients that re-use these numeric IDs.

Default users: numeric IDs	Default groups: numeric IDs
• root:0	• root:0
• pcuser:65534	• daemon:1
• nobody:65535	• pcuser:65534
	• nobody:65535



When using NFSv4.1, the root user might display as nobody when running directory listing commands on NFS clients. This is due to the client's ID domain mapping configuration. See the section called NFSv4.1 and the nobody user/group for details on this issue and how to resolve it.

# The root user

In Linux, the root account has access to all commands, files, and folders in a Linux-based file system. Because of the power of this account, security best practices often require the root user to be disabled or restricted in some fashion. In NFS exports, the power a root user has over the files and folders can be controlled in Cloud

Volumes Service through export policies and rules and a concept known as root squash.

Root squashing ensures that the root user accessing an NFS mount is squashed to the anonymous numeric user 65534 (see the section "The anonymous user") and is currently only available when using CVS-Performance by selecting Off for root access during export policy rule creation. If the root user is squashed to the anonymous user, it no longer has access to run chown or setuid/setgid commands (the sticky bit) on files or folders in the NFS mount, and files or folders created by the root user show the anon UID as the owner/group. In addition, NFSv4 ACLs cannot be modified by the root user. However, the root user still has access to chmod and deleted files that it does not have explicit permissions for. If you want to limit access to a root user's file and folder permissions, consider using a volume with NTFS ACLs, creating a Windows user named root, and applying the desired permissions to the files or folders.

# The anonymous user

The anonymous (anon) user ID specifies a UNIX user ID or username that is mapped to client requests that arrive without valid NFS credentials. This can include the root user when root squashing is used. The anon user in Cloud Volumes Service is 65534.

This UID is normally associated with the username nobody or nfsnobody in Linux environments. Cloud Volumes Service also uses 65534 as the local UNIX user' pcuser' (see the section "Default local UNIX users and groups"), which is also the default fallback user for Windows to UNIX name mappings when no valid matching UNIX user can be found in LDAP.

Because of the differences in usernames across Linux and Cloud Volumes Service for UID 65534, the name string for users mapped to 65534 might not match when using NFSv4.1. As a result, you might see nobody as the user on some files and folders. See the section "NFSv4.1 and the nobody user/group" for information about this issue and how to resolve it.

# Access control/exports

Initial export/share access for NFS mounts is controlled through host- based export policy rules contained within an export policy. A host IP, host name, subnet, netgroup, or domain is defined to allow access to mount the NFS share and the level of access allowed to the host. Export policy rule configuration options depend on the Cloud Volumes Service level.

For CVS-SW, the following options are available for export-policy configuration:

- Client match. Comma-separated list of IP addresses, comma-separated list of hostnames, subnets, netgroups, domain names.
- **RO/RW access rules.** Select read/write or read only to control level of access to export.CVS-Performance provides the following options:
- Client match. Comma-separated list of IP addresses, comma-separated list of hostnames, subnets, netgroups, domain names.
- RO/RW access rules. Select read/write or read only to control level of access to export.
- Root access (on/off). Configures root squash (see the section "The root user" for details).
- **Protocol type.** This limits access to the NFS mount to a specific protocol version. When specifying both NFSv3 and NFSv4.1 for the volume, either leave both blank or check both boxes.
- Kerberos security level (when Enable Kerberos is selected). Provides the options of krb5, krb5i, and/or krb5p for read-only or read-write access.

# Change ownership (chown) and change group (chgrp)

NFS on Cloud Volumes Service only allows the root user to run chown/chgrp on files and folders. Other users see an Operation not permitted error— even on files they own. If you use root squash (as covered in the section "The root user"), the root is squashed to a nonroot user and is not allowed access to chown and chgrp. There are currently no workarounds in Cloud Volumes Service to allow chown and chgrp for non-root users. If ownership changes are required, consider using dual protocol volumes and set the security style to NTFS to control permissions from the Windows side.

## **Permission management**

Cloud Volumes Service supports both mode bits (such as 644, 777, and so on for rwx) and NFSv4.1 ACLs to control permissions on NFS clients for volumes that use the UNIX security style. Standard permission management is used for these (such as chmod, chown, or nfs4\_setfacl) and work with any Linux client that supports them.

Additionally, when using dual protocol volumes set to NTFS, NFS clients can leverage Cloud Volumes Service name mapping to Windows users, which then are used to resolve the NTFS permissions. This requires an LDAP connection to Cloud Volumes Service to provide numeric-ID-to- username translations because Cloud Volumes Service requires a valid UNIX username to map properly to a Windows username.

# Providing granular ACLs for NFSv3

Mode bit permissions cover only owner, group, and everyone else in the semantics—meaning that there are no granular user access controls in place for basic NFSv3. Cloud Volumes Service does not support POSIX ACLs, nor extended attributes (such as chattr), so granular ACLs are only possible in the following scenarios with NFSv3:

- NTFS security style volumes (CIFS server required) with valid UNIX to Windows user mappings.
- NFSv4.1 ACLs applied using an admin client mounting NFSv4.1 to apply ACLs.

Both methods require an LDAP connection for UNIX identity management and a valid UNIX user and group information populated (see the section "LDAP") and are only available with CVS-Performance instances. To use NTFS security style volumes with NFS, you must use dual-protocol (SMB and NFSv3) or dual-protocol (SMB and NFSv4.1), even if no SMB connections are made. To use NFSv4.1 ACLs with NFSv3 mounts, you must select Both (NFSv3/NFSv4.1) as the protocol type.

Regular UNIX mode bits don't provide the same level of granularity in permissions that NTFS or NFSv4.x ACLs provide. The following table compares the permission granularity between NFSv3 mode bits and NFSv4.1 ACLs. For information about NFSv4.1 ACLs, see nfs4\_acl - NFSv4 Access Control Lists.

NFSv3 mode bits	NFSv4.1 ACLs
<ul> <li>NFSv3 mode bits</li> <li>Set user ID on execution</li> <li>Set group ID on execution</li> <li>Save swapped text (not defined in POSIX)</li> <li>Read permission for owner</li> <li>Write permission for owner on a file; or look up (search) permission for owner in directory</li> <li>Read permission for group</li> <li>Write permission for group</li> <li>Execute permission for group on a file; or look up (search) permission for group</li> <li>Execute permission for group in directory</li> <li>Read permission for others</li> <li>Write permission for others</li> <li>Write permission for others</li> </ul>	NFSv4.1 ACLs Access control entry (ACE) types (Allow/Deny/Audit) * Inheritance flags * directory-inherit * file-inherit * file-inherit * no-propagate-inherit * inherit-only Permissions * read-data (files) / list-directory (directories) * write-data (files) / create-file (directories) * write-data (files) / create-subdirectory (directories) * append-data (files) / create-subdirectory (directories) * delete * delete-child * read-attributes * write-attributes * read-named-attributes
<ul> <li>Execute permission for others on a file; or look up (search) permission for others in directory</li> </ul>	<ul> <li>* write-named-attributes</li> <li>* read-ACL</li> <li>* write-ACL</li> <li>* write-owner</li> <li>* Synchronize</li> </ul>

Finally, NFS group membership (in both NFSv3 and NFSV4.x) is limited to a default maximum of 16 for AUTH\_SYS as per the RPC packet limits. NFS Kerberos provides up to 32 groups and NFSv4 ACLs remove the limitation by way of granular user and group ACLs (up to 1024 entries per ACE).

Additionally, Cloud Volumes Service provides extended group support to extend the maximum supported groups up to 32. This requires an LDAP connection to an LDAP server that contains valid UNIX user and group identities. For more information about configuring this, see Creating and managing NFS volumes in the Google documentation.

# NFSv3 user and group IDs

NFSv3 user and group IDs come across the wire as numeric IDs rather than names. Cloud Volumes Service does no username resolution for these numeric IDs with NFSv3, with UNIX security style volumes using just mode bits. When NFSv4.1 ACLs are present, a numeric ID lookup and/or name string lookup is needed to resolve the ACL properly—even when using NFSv3. With NTFS security style volumes, Cloud Volumes Service must resolve a numeric ID to a valid UNIX user and then map to a valid Windows user to negotiate access rights.

# Security limitations of NFSv3 user and group IDs

With NFSv3, the client and server never have to confirm that the user attempting a read or write with a numeric ID is a valid user; it is just implicitly trusted. This opens the file system up to potential breaches simply by spoofing any numeric ID. To prevent security holes like this, there are a few options available to Cloud Volumes Service.

• Implementing Kerberos for NFS forces users to authenticate with a username and password or keytab file to get a Kerberos ticket to allow access into a mount. Kerberos is available with CVS-Performance instances and only with NFSv4.1.

- Limiting the list of hosts in your export policy rules limits which NFSv3 clients have access to the Cloud Volumes Service volume.
- Using dual-protocol volumes and applying NTFS ACLs to the volume forces NFSv3 clients to resolve numeric IDs to valid UNIX usernames to authenticate properly to access mounts. This requires enabling LDAP and configuring UNIX user and group identities.
- Squashing the root user limits the damage a root user can do to an NFS mount but does not completely remove risk. For more information, see the section "The root user."

Ultimately, NFS security is limited to what the protocol version you are using offers. NFSv3, while more performant in general than NFSv4.1, does not provide the same level of security.

# NFSv4.1

NFSv4.1 provides greater security and reliability as compared to NFSv3, for the following reasons:

- · Integrated locking through a lease-based mechanism
- Stateful sessions
- All NFS functionality over a single port (2049)
- TCP only
- ID domain mapping
- Kerberos integration (NFSv3 can use Kerberos, but only for NFS, not for ancillary protocols such as NLM)

#### NFSv4.1 dependencies

Because of the additionally security features in NFSv4.1, there are some external dependencies involved that were not needed to use NFSv3 (similar to how SMB requires dependencies such as Active Directory).

# NFSv4.1 ACLs

Cloud Volumes Service offers support for NFSv4.x ACLs, which deliver distinct advantages over normal POSIX-style permissions, such as the following:

- · Granular control of user access to files and directories
- Better NFS security
- · Improved interoperability with CIFS/SMB
- Removal of the NFS limitation of 16 groups per user with AUTH\_SYS security
- ACLs bypass the need for group ID (GID) resolution, which effectively removes the GID limitNFSv4.1 ACLs are controlled from NFS clients—not from Cloud Volumes Service. To use NFSv4.1 ACLs, be sure your client's software version supports them and the proper NFS utilities are installed.

# Compatibility between NFSv4.1 ACLs and SMB clients

NFSv4 ACLs are different from Windows file-level ACLs (NTFS ACLs) but carry similar functionality. However, in multiprotocol NAS environments, if NFSv4.1 ACLs are present and you are using dual-protocol access (NFS and SMB on the same datasets), clients using SMB2.0 and later won't be able to view or manage ACLs from Windows security tabs.

## How NFSv4.1 ACLs work

For reference, the following terms are defined:

- · Access control list (ACL). A list of permissions entries.
- Access control entry (ACE). A permission entry in the list.

When a client sets an NFSv4.1 ACL on a file during a SETATTR operation, Cloud Volumes Service sets that ACL on the object, replacing any existing ACL. If there is no ACL on a file, then the mode permissions on the file are calculated from OWNER@, GROUP@, and EVERYONE@. If there are any existing SUID/SGID/STICKY bits on the file, they are not affected.

When a client gets an NFSv4.1 ACL on a file during the course of a GETATTR operation, Cloud Volumes Service reads the NFSv4.1 ACL associated with the object, constructs a list of ACEs, and returns the list to the client. If the file has an NT ACL or mode bits, then an ACL is constructed from mode bits and is returned to the client.

Access is denied if a DENY ACE is present in the ACL; access is granted if an ALLOW ACE exists. However, access is also denied if neither of the ACEs is present in the ACL.

A security descriptor consists of a security ACL (SACL) and a discretionary ACL (DACL). When NFSv4.1 interoperates with CIFS/SMB, the DACL is one-to-one mapped with NFSv4 and CIFS. The DACL consists of the ALLOW and the DENY ACEs.

If a basic chmod is run on a file or folder with NFSv4.1 ACLs set, existing user and group ACLs are preserved, but the default OWNER@, GROUP@, EVERYONE@ ACLs are modified.

A client using NFSv4.1 ACLs can set and view ACLs for files and directories on the system. When a new file or subdirectory is created in a directory that has an ACL, that object inherits all ACEs in the ACL that have been tagged with the appropriate inheritance flags.

If a file or directory has an NFSv4.1 ACL, that ACL is used to control access no matter which protocol is used to access the file or directory.

Files and directories inherit ACEs from NFSv4 ACLs on parent directories (possibly with appropriate modifications) as long as the ACEs have been tagged with the correct inheritance flags.

When a file or directory is created as the result of an NFSv4 request, the ACL on the resulting file or directory depends on whether the file creation request includes an ACL or only standard UNIX file access permissions. The ACL also depends on whether the parent directory has an ACL.

- If the request includes an ACL, that ACL is used.
- If the request includes only standard UNIX file access permissions and the parent directory does not have an ACL, the client file mode is used to set standard UNIX file access permissions.
- If the request includes only standard UNIX file access permissions and the parent directory has a noninheritable ACL, a default ACL based on the mode bits passed into the request is set on the new object.
- If the request includes only standard UNIX file access permissions but the parent directory has an ACL, the ACEs in the parent directory's ACL are inherited by the new file or directory as long as the ACEs have been tagged with the appropriate inheritance flags.

## **ACE** permissions

NFSv4.1 ACLs permissions uses a series of upper- and lower-case letter values (such as rxtncy) to control access. For more information about these letter values, see HOW TO: Use NFSv4 ACL.

## NFSv4.1 ACL behavior with umask and ACL inheritance

NFSv4 ACLs provide the ability to offer ACL inheritance. ACL inheritance means that files or folders created beneath objects with NFSv4.1 ACLs set can inherit the ACLs based on the configuration of the ACL inheritance flag.

Umask is used to control the permission level at which files and folders are created in a directory without administrator interaction. By default, Cloud Volumes Service allows umask to override inherited ACLs, which is expected behavior as per RFC 5661.

## **ACL** formatting

NFSv4.1 ACLs have specific formatting. The following example is an ACE set on a file:

A::ldapuser@domain.netapp.com:rwatTnNcCy

The preceding example follows the ACL format guidelines of:

```
type:flags:principal:permissions
```

A type of A means "allow." The inherit flags are not set in this case, because the principal is not a group and does not include inheritance. Also, because the ACE is not an AUDIT entry, there is no need to set the audit flags. For more information about NFSv4.1 ACLs, see http://linux.die.net/man/5/nfs4\_acl.

If the NFSv4.1 ACL is not set properly (or a name string cannot be resolved by the client and server), the ACL might not behave as expected, or the ACL change might fail to apply and throw an error.

Sample errors include:

```
Failed setxattr operation: Invalid argument
Scanning ACE string 'A:: user@rwaDxtTnNcCy' failed.
```

#### **Explicit DENY**

NFSv4.1 permissions can include explicit DENY attributes for OWNER, GROUP, and EVERYONE. That is because NFSv4.1 ACLs are default-deny, which means that if an ACL is not explicitly granted by an ACE, then it is denied. Explicit DENY attributes override any ACCESS ACEs, explicit or not.

DENY ACEs are set with an attribute tag of D.

In the example below, GROUP@ is allowed all read and execute permissions, but denied all write access.

```
sh-4.1$ nfs4_getfacl /mixed
A::ldapuser@domain.netapp.com:ratTnNcCy
A::OWNER@:rwaDxtTnNcCy
D::OWNER@:
A:g:GROUP@:rxtncy
D:g:GROUP@:waDTC
A::EVERYONE@:rxtncy
D::EVERYONE@:waDTC
```

DENY ACEs should be avoided whenever possible because they can be confusing and complicated; ALLOW ACLs that are not explicitly defined are implicitly denied. When DENY ACEs are set, users might be denied access when they expect to be granted access.

The preceding set of ACEs is equivalent to 755 in mode bits, which means:

- The owner has full rights.
- · Groups have read only.
- Others have read only.

However, even if permissions are adjusted to the 775 equivalent, access can be denied because of the explicit DENY set on EVERYONE.

## NFSv4.1 ID domain mapping dependencies

NFSv4.1 leverages ID domain mapping logic as a security layer to help verify that a user attempting access to an NFSv4.1 mount is indeed who they claim to be. In these cases, the username and group name coming from the NFSv4.1 client appends a name string and sends it to the Cloud Volumes Service instance. If that username/group name and ID string combination does not match, then the user and/or group is squashed to the default nobody user specified in the /etc/idmapd.conf file on the client.

This ID string is a requirement for proper permission adherence, especially when NFSv4.1 ACLs and/or Kerberos are in use. As a result, name service server dependencies such as LDAP servers are necessary to ensure consistency across clients and Cloud Volumes Service for proper user and group name identity resolution.

Cloud Volumes Service uses a static default ID domain name value of defaultv4iddomain.com. NFS clients default to the DNS domain name for its ID domain name settings, but you can manually adjust the ID domain name in /etc/idmapd.conf.

If LDAP is enabled in Cloud Volumes Service, then Cloud Volumes Service automates the NFS ID domain to change to what is configured for the search domain in DNS and clients won't need to be modified unless they use different DNS domain search names.

When Cloud Volumes Service can resolve a username or group name in local files or LDAP, the domain string is used and non-matching domain IDs squash to nobody. If Cloud Volumes Service cannot find a username or group name in local files or LDAP, the numeric ID value is used and the NFS client resolves the name properly (this is similar to NFSv3 behavior).

Without changing the client's NFSv4.1 ID domain to match what the Cloud Volumes Service volume is using, you see the following behavior:

- UNIX users and groups with local entries in Cloud Volumes Service (such as root, as defined in local UNIX users and groups) are squashed to the nobody value.
- UNIX users and groups with entries in LDAP (if Cloud Volumes Service is configured to use LDAP) squashes to nobody if DNS domains are different between NFS clients and Cloud Volumes Service.
- UNIX users and groups with no local entries or LDAP entries use the numeric ID value and resolve to the name specified on the NFS client. If no name exists on the client, only the numeric ID is shown.

The following shows the results of the preceding scenario:

```
# ls -la /mnt/home/prof1/nfs4/
total 8
drwxr-xr-x 2 nobody nobody 4096 Feb 3 12:07 .
drwxrwxrwx 7 root root 4096 Feb 3 12:06 ..
-rw-r--r-- 1 9835 9835 0 Feb 3 12:07 client-user-no-name
-rw-r--r-- 1 nobody nobody 0 Feb 3 12:07 ldap-user-file
-rw-r--r-- 1 nobody nobody 0 Feb 3 12:06 root-user-file
```

When the client and server ID domains match, this is how the same file listing looks:

```
# ls -la
total 8
drwxr-xr-x 2 root root 4096 Feb 3 12:07 .
drwxrwxrwx 7 root root 4096 Feb 3 12:06 ..
-rw-r--r- 1 9835 9835 0 Feb 3 12:07 client-user-no-name
-rw-r--r- 1 apache apache-group 0 Feb 3 12:07 ldap-user-file
-rw-r--r- 1 root root 0 Feb 3 12:06 root-user-file
```

For more information about this issue and how to resolve it, see the section "NFSv4.1 and the nobody user/group."

#### **Kerberos dependencies**

If you plan to use Kerberos with NFS, you must have the following with Cloud Volumes Service:

- Active Directory domain for Kerberos Distribution Center services (KDC)
- Active Directory domain with user and group attributes populated with UNIX information for LDAP functionality (NFS Kerberos in Cloud Volumes Service requires a user SPN to UNIX user mapping for proper functionality.)
- · LDAP enabled on the Cloud Volumes Service instance
- · Active Directory domain for DNS services

#### NFSv4.1 and the nobody user/group

One of the most common issues seen with an NFSv4.1 configuration is when a file or folder is shown in a listing using 1s as being owned by the user:group combination of nobody:nobody.

For example:

```
sh-4.2$ ls -la | grep prof1-file
-rw-r--r- 1 nobody nobody 0 Apr 24 13:25 prof1-file
```

And the numeric ID is 99.

```
sh-4.2$ ls -lan | grep prof1-file
-rw-r--r- 1 99 99 0 Apr 24 13:25 prof1-file
```

In some instances, the file might show the correct owner but nobody as the group.

```
sh-4.2$ ls -la | grep newfile1
-rw-r--r- 1 prof1 nobody 0 Oct 9 2019 newfile1
```

Who is nobody?

The nobody user in NFSv4.1 is different from the nfsnobody user. You can view how an NFS client sees each user by running the id command:

```
# id nobody
uid=99(nobody) gid=99(nobody) groups=99(nobody)
# id nfsnobody
uid=65534(nfsnobody) gid=65534(nfsnobody) groups=65534(nfsnobody)
```

With NFSv4.1, the nobody user is the default user defined by the idmapd.conf file and can be defined as any user you want to use.

```
# cat /etc/idmapd.conf | grep nobody
#Nobody-User = nobody
#Nobody-Group = nobody
```

Why does this happen?

Because security through name string mapping is a key tenet of NFSv4.1 operations, the default behavior when a name string does not match properly is to squash that user to one that won't normally have any access to files and folders owned by users and groups.

When you see nobody for the user and/or group in file listings, this generally means something in NFSv4.1 is misconfigured. Case sensitivity can come into play here.

For example, if user1@CVSDEMO.LOCAL (uid 1234, gid 1234) is accessing an export, then Cloud Volumes Service must be able to find user1@CVSDEMO.LOCAL (uid 1234, gid 1234). If the user in Cloud Volumes Service is USER1@CVSDEMO.LOCAL, then it won't match (uppercase USER1 versus lowercase user1). In

many cases, you can see the following in the messages file on the client:

May 19 13:14:29 centos7 nfsidmap[17481]: nss\_getpwnam: name 'root@defaultv4iddomain.com' does not map into domain 'CVSDEMO.LOCAL' May 19 13:15:05 centos7 nfsidmap[17534]: nss\_getpwnam: name 'nobody' does not map into domain 'CVSDEMO.LOCAL'

The client and server must both agree that a user is indeed who they are claiming to be, so you must check the following to ensure that the user that the client sees has the same information as the user that Cloud Volumes Service sees.

- NFSv4.x ID domain. Client: idmapd.conf file; Cloud Volumes Service uses defaultv4iddomain.com and cannot be changed manually. If using LDAP with NFSv4.1, Cloud Volumes Service changes the ID domain to what the DNS search domain is using, which is the same as the AD domain.
- User name and numeric IDs. This determines where the client is looking for user names and leverages the name service switch configuration—client: nsswitch.conf and/or local passwd and group files; Cloud Volumes Service does not allow modifications to this but automatically adds LDAP to the configuration when it is enabled.
- Group name and numeric IDs. This determines where the client is looking for group names and leverages the name service switch configuration—client: nsswitch.conf and/or local passwd and group files; Cloud Volumes Service does not allow modifications to this but automatically adds LDAP to the configuration when it is enabled.

In almost all cases, if you see nobody in user and group listings from clients, the issue is user or group name domain ID translation between Cloud Volumes Service and the NFS client. To avoid this scenario, use LDAP to resolve user and group information between clients and Cloud Volumes Service.

#### Viewing name ID strings for NFSv4.1 on clients

If you are using NFSv4.1, there is a name-string mapping that takes place during NFS operations, as previously described.

In addition to using /var/log/messages to find an issue with NFSv4 IDs, you can use the nfsidmap -l command on the NFS client to view which usernames have properly mapped to the NFSv4 domain.

For example, this is output of the command after a user that can be found by the client and Cloud Volumes Service accesses an NFSv4.x mount:

```
# nfsidmap -1
4 .id_resolver keys found:
    gid:daemon@CVSDEMO.LOCAL
    uid:nfs4@CVSDEMO.LOCAL
    gid:root@CVSDEMO.LOCAL
    uid:root@CVSDEMO.LOCAL
```

When a user that does not map properly into the NFSv4.1 ID domain (in this case, netapp-user) tries to access the same mount and touches a file, they are assigned nobody:nobody, as expected.

```
# su netapp-user
sh-4.2$ id
uid=482600012 (netapp-user), 2000 (secondary)
sh-4.2$ cd /mnt/nfs4/
sh-4.2$ touch newfile
sh-4.2$ ls -la
total 16
drwxrwxrwx 5 root root
                          4096 Jan 14 17:13 .
drwxr-xr-x. 8 root root
                           81 Jan 14 10:02 ..
-rw-r--r-- 1 nobody nobody
                             0 Jan 14 17:13 newfile
drwxrwxrwx 2 root
                   root 4096 Jan 13 13:20 qtree1
drwxrwxrwx 2 root
                    root 4096 Jan 13 13:13 gtree2
drwxr-xr-x 2 nfs4
                    daemon 4096 Jan 11 14:30 testdir
```

The nfsidmap -l output shows the user pcuser in the display but not netapp-user; this is the anonymous user in our export-policy rule (65534).

```
# nfsidmap -1
6 .id_resolver keys found:
    gid:pcuser@CVSDEMO.LOCAL
    uid:pcuser@CVSDEMO.LOCAL
    gid:daemon@CVSDEMO.LOCAL
    uid:nfs4@CVSDEMO.LOCAL
    gid:root@CVSDEMO.LOCAL
    uid:root@CVSDEMO.LOCAL
```

# SMB

SMB is a network file sharing protocol developed by Microsoft that provides centralized user/group authentication, permissions, locking, and file sharing to multiple SMB clients over an Ethernet network. Files and folders are presented to clients by way of shares, which can be configured with a variety of share properties and offers access control through share-level permissions. SMB can be presented to any client that offers support for the protocol, including Windows, Apple, and Linux clients.

Cloud Volumes Service provides support for the SMB 2.1 and 3.x versions of the protocol.

# Access control/SMB shares

- When a Windows username requests access to the Cloud Volumes Service volume, Cloud Volumes Service looks for a UNIX username using the methods configured by Cloud Volumes Service administrators.
- If an external UNIX identity provider (LDAP) is configured and Windows/UNIX usernames are identical, then Windows usernames will map 1:1 to UNIX usernames without any additional configuration needed. When LDAP is enabled, Active Directory is used to host those UNIX attributes for user and group objects.

- If Windows names and UNIX names do not match identically, then LDAP must be configured to allow Cloud Volumes Service to use the LDAP name mapping configuration (see the section "Using LDAP for asymmetric name mapping").
- If LDAP is not in use, then Windows SMB users map to a default local UNIX user named pcuser in Cloud Volumes Service. This means files written in Windows by users that map to the pcuser show UNIX ownership as pcuser in multiprotocol NAS environments. pcuser here is effectively the nobody user in Linux environments (UID 65534).

In deployments with SMB only, the pcuser mapping still occurs, but it won't matter, because Windows user and group ownership is correctly displayed and NFS access to the SMB-only volume is not allowed. In addition, SMB-only volumes do not support conversion to NFS or dual-protocol volumes after they are created.

Windows leverages Kerberos for username authentication with the Active Directory domain controllers, which requires a username/password exchange with the AD DCs, which is external to the Cloud Volumes Service instance. Kerberos authentication is used when the \\SERVERNAME UNC path is used by the SMB clients and the following is true:

- DNS A/AAAA entry exists for SERVERNAME
- A valid SPN for SMB/CIFS access exists for SERVERNAME

When a Cloud Volumes Service SMB volume is created, the machine account name is created as defined in the section "How Cloud Volumes Service shows up in Active Directory." That machine account name also becomes the SMB share access path because Cloud Volumes Service leverages Dynamic DNS (DDNS) to create the necessary A/AAAA and PTR entries in DNS and the necessary SPN entries on the machine account principal.



For PTR entries to be created, the reverse lookup zone for the Cloud Volumes Service instance IP address must exist on the DNS server.

For example, this Cloud Volumes Service volume uses the following UNC share path: \\cvs-east-433d.cvsdemo.local.

In Active Directory, these are the Cloud Volumes Service-generated SPN entries:

```
PS C:\> setspn /L CVS-EAST-433D
Registered ServicePrincipalNames for CN=CVS-EAST-433D,CN=Computers,DC=cvsdemo,DC=local:
HOST/cvs-east-433d.cvsdemo.local
HOST/CVS-EAST-433D
```

This is the DNS forward/reverse lookup result:

```
PS C:\> nslookup CVS-EAST-433D
Server: activedirectory. region. lab. internal
Address: 10. xx.0. xx
Name: CVS-EAST-433D.cvsdemo.local
Address: 10. xxx.0. x
PS C:\> nslookup 10. xxx.0. x
Server: activedirectory.region.lab.internal
Address: 10.xx.0.xx
Name: CVS-EAST-433D.CVSDEMO.LOCAL
Address: 10. xxx.0. x
```

Optionally, more access control can be applied by enabling/requiring SMB encryption for SMB shares in Cloud Volumes Service. If SMB encryption isn't supported by one of the endpoints, then access is not allowed.

#### Using SMB name aliases

In some cases, it might be a security concern for end users to know the machine account name in use for Cloud Volumes Service. In other cases, you might simply want to provide a simpler access path to your end users. In those cases, you can create SMB aliases.

If you want to create aliases for the SMB share path, you can leverage what is known as a CNAME record in DNS. For example, if you want to use the name \\CIFS to access shares instead of \\cvs-east-433d.cvsdemo.local, but you still want to use Kerberos authentication, a CNAME in DNS that points to the existing A/AAAA record and an additional SPN added to the existing machine account provides Kerberos access.

cifs Properties	?	X
Alias (CNAME) Security		
Alias name (uses parent domain if left blank):		
cifs		
Fully qualified domain name (FQDN):		
cits.cvsdemo.local		
Eully qualified domain name (FQDN) for target host: CVS-EAST-433D.CVSDEMO.LOCAL	Browse	
	<u>D</u> 101130	
OK Cancel	Ap	ply

This is the resulting DNS forward lookup result after adding a CNAME:

```
PS C:\> nslookup cifs
Server: ok-activedirectory.us-east4-a.c.cv-solution-architect-
lab.internal
Address: 10. xx.0. xx
Name: CVS-EAST-433D.cvsdemo.local
Address: 10. xxx.0. x
Aliases: cifs.cvsdemo.local
```

This is the resulting SPN query after adding new SPNs:

#### PS C:\> setspn /L CVS-EAST-433D Registered ServicePrincipalNames for CN=CVS-EAST-433D,CN=Computers,DC=cvsdemo,DC=local: cifs/cifs.cvsdemo.local cifs/cifs HOST/cvs-east-433d.cvsdemo.local HOST/CVS-EAST-433D

In a packet capture, we can see the Session Setup Request using the SPN tied to the CNAME.

431 4.156722	SMB2	308	Negotiate Protocol Response
432 4.156785	SMB2	232	Negotiate Protocol Request
434 4.158108	SMB2	374	Negotiate Protocol Response
435 4.160977	SMB2	1978	Session Setup Request
437 4.166224	SMB2	322	Session Setup Response
438 4.166891	SMB2	152	Tree Connect Request Tree: \\cifs\IPC
439 4.168063	SMB2	138	Tree Connect Response

```
realm: CVSDEMO.LOCAL

sname
name-type: kRB5-NT-SRV-INST (2)
sname-string: 2 items
SNameString: cifs
SNameString: cifs
enc-part
etype: eTYPE-ARCFOUR-HMAC-MD5 (23)
```

## **SMB** authentication dialects

Cloud Volumes Service supports the following dialects for SMB authentication:

- LM
- NTLM
- NTLMv2
- Kerberos

Kerberos authentication for SMB share access is the most secure level of authentication you can use. With AES and SMB encryption enabled, the security level is further increased.

Cloud Volumes Service also supports backward compatibility for LM and NTLM authentication. When Kerberos is misconfigured (such as when creating SMB aliases), share access falls back to weaker authentication methods (such as NTLMv2). Because these mechanisms are less secure, they are disabled in some Active Directory environments. If weaker authentication methods are disabled and Kerberos is not configured properly, share access fails because there is no valid authentication method to fall back to.

For information about configuring/viewing your supported authentication levels in Active Directory, see Network security: LAN Manager authentication level.

#### **Permission models**

#### **NTFS/File permissions**

NTFS permissions are the permissions applied to files and folders in file systems adhering to NTFS logic. You can apply NTFS permissions in Basic or Advanced and can be set to Allow or Deny for access control.

Basic permissions include the following:

Full Control

- Modify
- Read & Execute
- Read
- Write

When you set permissions for a user or group, referred to as an ACE, it resides in an ACL. NTFS permissions use the same read/write/execute basics as UNIX mode bits, but they can also extend to more granular and extended access controls (also known as Special Permissions), such as Take Ownership, Create Folders/Append Data, Write Attributes, and more.

Standard UNIX mode bits do not provide the same level of granularity as NTFS permissions (such as being able to set permissions for individual user and group objects in an ACL or setting extended attributes). However, NFSv4.1 ACLs do provide the same functionality as NTFS ACLs.

NTFS permissions are more specific than share permissions and can be used in conjunction with share permissions. With NTFS permission structures, the most restrictive applies. As such, explicit denials to a user or group overrides even Full Control when defining access rights.

NTFS permissions are controlled from Windows SMB clients.

# Share permissions

Share permissions are more general than NTFS permissions (Read/Change/Full Control only) and control the initial entry into an SMB share—similar to how NFS export policy rules work.

Although NFS export policy rules control access through host-based information such as IP addresses or host names, SMB share permissions can control access by using user and group ACEs in a share ACL. You can set share ACLs either from the Windows client or from the Cloud Volumes Service management UI.

By default, share ACLs and initial volume ACLs include Everyone with Full Control. The file ACLs should be changed but share permissions are overruled by the file permissions on objects in the share.

For instance, if a user is only allowed Read access to the Cloud Volumes Service volume file ACL, they are denied access to create files and folders even though the share ACL is set to Everyone with Full Control, as shown in the following figure.

eneral Publish Share Permis	sions Secunty		General Publish Share Permissions Se	ecurity
aroup or user names:			Object name: UCVS-EAST-433D)acta	ticanon-ionamar
Severyone			object name, NOVO-ENSI-455D/ecsta	ucangiy-jonamai
			Group or user names:	
			a cvs-svc	
			Justin Pansi (parisi@cvsdemo.local)	
	A <u>d</u> d	Remove	To change permissions, click Edit.	<u>E</u> dit
Permissions for Everyone	Allow	Deny	Permissions for Justin Parisi	Allow Deny
Full Control			Full control	
Change	Y		Modify	
Read			Read & execute	~
			List folder contents	~
			Read	~
	OK Cancel	Apply	Advanced.	Cancel Appl
→ • ↑ 💄 > Net	OK Cancel	Apply ic-angry-jonarnar	Advanced. OK V U Search	Cancel Appl
→ ✓ ↑ 💄 > Net	OK Cancel twork > cifs > ecstat Name	Apply ic-angry-jonarnar	Аdvanced. ОК V U Search Date modified	Cancel Appl n ecstatic-angry-jonarnar ype S
→ → ↑ 💄 > Net	OK Cancel twork > cifs > ecstat Name	Apply ic-angry-jonarnar	Advanced.	Cancel Appl n ecstatic-angry-jonarnar ype S
→ ✓ ↑ 💄 → Net ✓ Quick access Desktop ✓	OK Cancel twork > cifs > ecstat Name	Apply ic-angry-jonarnar	Advanced.	Cancel Appl
<ul> <li>→ ↑ ↓ &gt; Net</li> <li>✓ Quick access</li> <li>Desktop *</li> <li>✓ Downloads *</li> </ul>	OK Cancel twork > cifs > ecstat Name	ic-angry-jonarnar	Advanced. ОК V U Search Date modified This folder is empty.	Cancel Appl n ecstatic-angry-jonarnar ype S
→ ✓ ↑ 💄 → Net Quick access Desktop * Downloads * Documents * □	OK Cancel twork > cifs > ecstat Name	Apply ic-angry-jonarnar	Advanced.	Cancel Appl
<ul> <li>→ &lt; ↑ ↓ Net</li> <li>♦ Quick access</li> <li>■ Desktop</li> <li>♦ Downloads</li> <li>♥ Documents</li> <li>♥ Pictures</li> </ul>	OK Cancel twork > cifs > ecstat Name Destination Folder Acc	ic-angry-jonarnar	Advanced. ок Search Date modified This folder is empty. Х	Cancel Appl
<ul> <li>→ → ↑ ↓ &gt; Net</li> <li>✓ Quick access</li> <li>Desktop *</li> <li>Downloads *</li> <li>Documents *</li> <li>Pictures *</li> </ul>	OK Cancel twork > cifs > ecstat Name Destination Folder Acc	ic-angry-jonarnar	Advanced. ОК СоК Search Date modified This folder is empty. Х	Cancel Appl
<ul> <li> <p< td=""><td>OK Cancel twork &gt; cifs &gt; ecstat Name Destination Folder Acc</td><td>ic-angry-jonarnar</td><td>Advanced. Advanced.</td><td>Cancel Appl n ecstatic-angry-jonarnar ype S</td></p<></li></ul>	OK Cancel twork > cifs > ecstat Name Destination Folder Acc	ic-angry-jonarnar	Advanced. Advanced.	Cancel Appl n ecstatic-angry-jonarnar ype S
<ul> <li> <p< td=""><td>OK Cancel twork &gt; cifs &gt; ecstat Name Destination Folder Acc</td><td>ic-angry-jonarnar</td><td>Advanced. ОК С С С С С С С С С С С С С</td><td>Cancel Appl n ecstatic-angry-jonarnar ype S</td></p<></li></ul>	OK Cancel twork > cifs > ecstat Name Destination Folder Acc	ic-angry-jonarnar	Advanced. ОК С С С С С С С С С С С С С	Cancel Appl n ecstatic-angry-jonarnar ype S
<ul> <li>A A A A A A A A A A A A A A A A A A A</li></ul>	OK Cancel twork > cifs > ecstat Name Destination Folder Acc	ic-angry-jonarnar	Advanced. ОК ОК Search Date modified This folder is empty. Х m this action -jonarnar	Cancel Appl n ecstatic-angry-jonarnar ype S
<ul> <li>A A A A A A A A A A A A A A A A A A A</li></ul>	OK Cancel twork > cifs > ecstat Name Destination Folder Acc	ic-angry-jonarnar	Advanced.	Cancel Appl n ecstatic-angry-jonarnar ype S
<ul> <li>A A A A A A A A A A A A A A A A A A A</li></ul>	OK Cancel twork > cifs > ecstat Name Destination Folder Acc You need per	ic-angry-jonarnar	Advanced.	Cancel Apple

For best security results, do the following:

- Remove Everyone from the share and file ACLs and instead set share access for users or groups.
- Use groups for access control instead of individual users for ease of management and faster removal/addition of users to share ACLs through group management.
- Allow less restrictive, more general share access to the ACEs on the share permissions and lock down access to users and groups with file permissions for more granular access control.
- Avoid general use of explicit deny ACLs, because they override allow ACLs. Limit use of explicit deny ACLs for users or groups that need to be restricted from access to a file system quickly.
- Make sure that you pay attention to the ACL inheritance settings when modifying permissions; setting the inheritance flag at the top level of a directory or volume with high file counts means that each file below that

directory or volume has inherited permissions added to it, which can create unwanted behavior such as unintended access/denial and long churn of permission modification as each file is adjusted.

## SMB share security features

Snapshot directory visible: Unchecked

When you first create a volume with SMB access in Cloud Volumes Service, you are presented with a series of choices for securing that volume.

Some of these choices depend on the Cloud Volumes Service level (Performance or Software) and choices include:

• Make snapshot directory visible (available for both CVS-Performance and CVS-SW). This option controls whether or not SMB clients can access the Snapshot directory in an SMB share (\\server\share\~snapshot and/or Previous Versions tab). The default setting is Not Checked, which means that the volume defaults to hiding and disallowing access to the ~snapshot directory, and no Snapshot copies appear in the Previous Versions tab for the volume.

Snapshot directory visible: Checked

vs-parisi (\\CVS-EAS	T-C2DB) Properties		🤇 😓 cvs-pa	arisi (\\CVS-EAS	T-C2DB) Properties	
General	Network	Security	Ge	eneral	Network	Security
Previous Versions	Customize	Classification	Previo	Previous Versions Customize Cla		Classification
Previous ver saved autor	rsions come from shade natically to your compu	ow copies, which are ter's hard disk.	Ð	Previous ver saved autor	sions come from shad natically to your compu	ow copies, which are ter's hard disk.
der versions:			Eolder v	ersions:		
Name	Date mod	dified	Nan	ne	Date mo	dified
There are n	o previous versions :	available	∽ To	day (1)		
mere are m	o previous versions e	IVANADIE	<b>•</b>	cvs-parisi	2/1/2022	2 6:56 PM
		and Depterson and			Open	- Pactore
	Open	THESIDIE T				

Hiding Snapshot copies from end users might be desired for security reasons, performance reasons (hiding these folders from AV scans) or preference. Cloud Volumes Service Snapshots are read- only, so even if these Snapshots are visible, end users cannot delete or modify files in the Snapshot directory. File permissions on the files or folders at the time the Snapshot copy was taken apply. If a file or folder's permissions change between Snapshot copies, then the changes also apply to the files or folders in the Snapshot directory. Users and groups can gain access to these files or folders based on permissions. While deletes or modifications of files in the Snapshot directory are not possible, it is possible to copy files or folders out of the Snapshot

directory.

- Enable SMB encryption (available for both CVS-Performance and CVS-SW). SMB encryption is disabled on the SMB share by default (unchecked). Checking the box enables SMB encryption, which means traffic between the SMB client and server is encrypted in-flight with the highest supported encryption levels negotiated. Cloud Volumes Service supports up to AES-256 encryption for SMB. Enabling SMB encryption does carry a performance penalty that might or might not be noticeable to your SMB clients—roughly in the 10-20% range. NetApp strongly encourages testing to see if that performance penalty is acceptable.
- Hide SMB share (available for both CVS-Performance and CVS-SW). Setting this option hides the SMB share path from normal browsing. This means that clients that do not know the share path cannot see the shares when accessing the default UNC path (such as \\CVS-SMB). When the checkbox is selected, only clients that explicitly know the SMB share path or have the share path defined by a Group Policy Object can access it (security through obfuscation).
- Enable access-based enumeration (ABE) (CVS-SW only). This is similar to hiding the SMB share, except the shares or files are only hidden from users or groups that do not have permissions to access the objects. For instance, if Windows user joe is not allowed at least Read access through the permissions, then the Windows user joe cannot see the SMB share or files at all. This is disabled by default, and you can enable it by selecting the checkbox. For more information on ABE, see the NetApp Knowledge Base article How does Access Based Enumeration (ABE) work?
- Enable Continuously Available (CA) share support (CVS-Performance only). Continuously Available SMB shares provide a way to minimize application disruptions during failover events by replicating lock states across nodes in the Cloud Volumes Service backend system. This is not a security feature, but it does offer better overall resiliency. Currently, only SQL Server and FSLogix applications are supported for this functionality.

# Default hidden shares

When an SMB server is created in Cloud Volumes Service, there are hidden administrative shares (using the \$ naming convention) that are created in addition to the data volume SMB share. These include C\$ (namespace access) and IPC\$ (sharing named pipes for communication between programs, such as the remote procedure calls (RPC) used for Microsoft Management Console (MMC) access).

The IPC\$ share contains no share ACLs and cannot be modified—it is strictly used for RPC calls and Windows disallows anonymous access to these shares by default.

The C\$ share allows BUILTIN/Administrators access by default, but Cloud Volumes Service automation removes the share ACL and does not allow access to anyone because access to the C\$ share allows visibility into all mounted volumes in the Cloud Volumes Service file systems. As a result, attempts to navigate to \\SERVER\C\$ fail.

# Accounts with local/BUILTIN administrator/backup rights

Cloud Volumes Service SMB servers maintain similar functionality to regular Windows SMB servers in that there are local groups (such as BUILTIN\Administrators) that apply access rights to select domain users and groups.

When you specify a user to be added to Backup Users, the user is added to the BUILTIN\Backup Operators group in the Cloud Volumes Service instance that uses that Active Directory connection, which then gets the SeBackupPrivilege and SeRestorePrivilege.

When you add a user to Security Privilege Users, the user is given the SeSecurityPrivilege, which is useful in some application use cases, such as SQL Server on SMB shares.

# **Backup Users**

Provide a comma separated list of domain users or a domain group name that require elevated privileges to access volumes created by Cloud Volumes Service.

Accountnames administrator,cvs-svc

# Security Privilege Users

Provide a list of comma separated domain user accounts that require elevated privileges to manage security log for the Active Directory associated with Cloud Volumes Service.

Accountnames administrator,cvs-svc

You can view Cloud Volumes Service local group memberships through the MMC with the proper privileges. The following figure shows users that have been added by using the Cloud Volumes Service console.

Backup Operato	rs Propertie	5		?	X		
General							
Back	Backup Operators						
Description:	Description: Backup Operators group						
Members:							
CVSDEMO	Members:						
A <u>d</u> d	Remove	are not ef user logs	rective until the ne	ext time	the		
	ОК	Cancel	Apply	Н	elp		

The following table shows the list of default BUILTIN groups and what users/groups are added by default.

Local/BUILTIN group	Default members
BUILTIN\Administrators*	DOMAIN\Domain Admins
BUILTIN\Backup Operators*	None
BUILTIN\Guests	DOMAIN\Domain Guests
BUILTIN\Power Users	None
BUILTIN\Domain Users	DOMAIN\Domain Users

\*Group membership controlled in Cloud Volumes Service Active Directory connection configuration.

You can view local users and groups (and group members) in the MMC window, but you cannot add or delete objects or change group memberships from this console. By default, only the Domain Admins group and Administrator are added to the BUILTIN\Administrators group in Cloud Volumes Service. Currently, you cannot modify this.

Computer Management (CVS-EAST-C2DB)	Name Full Name	Description Built-in administrator account	Computer Management (CVS-EAST-C2DB System Tools Event Viewer Shared Folders Shared Folders Shares Stares Stares Local Users and Groups Users Groups	Name Administrators Users Guests Power Users Backup Operators	Description Built-in Administrators group All users Built-in Guets Group Restricted administrative privileges Backup Operators group
Administrators	Properties			? ×	
General	ninistrators				
Description:	Built-in Ad	ministrators grou	0		
Administr	ator 10\Domain Adr	nins Changes to	o a user's group me	embership	
A <u>d</u> d	Remove	are not effe user logs o	ective until the next n.	time the	
	ОК	Cancel	Apply	Help	

## **MMC/Computer Management access**

SMB access in Cloud Volumes Service provides connectivity to the Computer Management MMC, which allows you to view shares, manage share ACLs, ands view/manage SMB sessions and open files.

To use the MMC to view SMB shares and sessions in Cloud Volumes Service, the user logged in currently must be a domain administrator. Other users are allowed access to view or manage the SMB server from MMC and receive a You Do Not Have Permissions dialog box when attempting to view shares or sessions on the Cloud Volumes Service SMB instance.

To connect to the SMB server, open Computer Management, right click Computer Management and then select Connect To Another Computer. This opens the Select Computer dialog box where you can enter the SMB server name (found in the Cloud Volumes Service volume information).

When you view SMB shares with the proper permissions, you see all available shares in the Cloud Volumes Service instance that share the Active Directory connection. To control this behavior, set the Hide SMB Shares option on the Cloud Volumes Service volume instance.

Remember, only one Active Directory connection is allowed per region.

🔚 Computer Management							×
<u>F</u> ile <u>A</u> ction <u>V</u> iew <u>H</u> elp							
Computer Management (CVS-EAST-C2DB)	Share Name	Folder Path	Туре	# Client Connections	Description	Actions	
<ul> <li>W System Tools</li> <li>Task Scheduler</li> </ul>	∰ac\$	C:\	Windows	0		Shares	-
<ul> <li>Event Viewer</li> <li>Shared Folders</li> <li>Sssions</li> <li>Open Files</li> <li>Eccal Users and Groups</li> <li>Performance</li> <li>Device Manager</li> <li>Storage</li> <li>Services and Applications</li> </ul>	<	C:\cvs-pansi C:\dgeyer-smb-test	Windows Windows Windows	1 0 2	2	More Actions	•

Image: System Tools       Image: System Tools       User       Computer Type       # Open Files       Connected Time       Idle Time       Actions         Image: System Tools       Image: System Tools<	imputer Management Action View Help
<ul> <li>Event Viewer</li> <li>gis Shares</li> <li>gis Shares</li> <li>gis Shares</li> <li>gis Shares</li> <li>gis Open Files</li> <li>Services and Applications</li> </ul>	Action View Help Action View Help Particle Analgement (CVS-EAST-C2DB) System Tools Task Scheduler Event Viewer Shared Folders Sessions Device Manager Storage Services and Applications

The following table shows a list of supported/unsupported functionality for the MMC.

Supported functions	Unsupported functions
View shares	<ul> <li>Creating new local users/groups</li> </ul>
<ul> <li>View active SMB sessions</li> </ul>	<ul> <li>Managing/viewing existing local user/groups</li> </ul>
View open files	<ul> <li>View events or performance logs</li> </ul>
<ul> <li>View local users and groups</li> </ul>	Managing storage
<ul> <li>View local group memberships</li> </ul>	<ul> <li>Managing services and applications</li> </ul>
<ul> <li>Enumerate the list of sessions, files, and tree connections in the system</li> </ul>	
<ul> <li>Close open files in the system</li> </ul>	
Close open sessions	
Create/manage shares	

# SMB server security information

The SMB server in Cloud Volumes Service uses a series of options that define security policies for SMB connections, including things such as Kerberos clock skew, ticket age, encryption, and more.

The following table contains a list of those options, what they do, the default configurations, and if they can be modified with Cloud Volumes Service. Some options do not apply to Cloud Volumes Service.
Security option	What it does	Default value	Can change?
Maximum Kerberos Clock Skew (minutes)	Maximum time skew between Cloud Volumes Service and domain controllers. If the time skew exceeds 5 minutes, Kerberos authentication fails. This is set to the Active Directory default value.	5	No
Kerberos Ticket Lifetime (hours)	Maximum time a Kerberos ticket remains valid before requiring a renewal. If no renewal occurs before the 10 hours, you must obtain a new ticket. Cloud Volumes Service performs these renewals automatically. 10 hours is the Active Directory default value.	10	No
Maximum Kerberos Ticket Renewal (days)	Maximum number of days that a Kerberos ticket can be renewed before a new authorization request is needed. Cloud Volumes Service automatically renews tickets for SMB connections. Seven days is the Active Directory default value.	7	No
Kerberos KDC Connection Timeout (secs)	The number of seconds before a KDC connection times out.	3	No
Require Signing for Incoming SMB Traffic	Setting to require signing for SMB traffic. If set to true, clients that do not support signing fail connectivity.	False	
Require Password Complexity for Local User Accounts	Used for passwords on local SMB users. Cloud Volumes Service does not support local user creation, so this option does not apply to Cloud Volumes Service.	True	No

Security option	What it does	Default value	Can change?
Use start_tls for Active Directory LDAP Connections	Used to enable start TLS connections for Active Directory LDAP. Cloud Volumes Service does not currently support enabling this.	False	No
Is AES-128 and AES-256 Encryption for Kerberos Enabled	This controls whether AES encryption is used for Active Directory connections and is controlled with the Enable AES Encryption for Active Directory Authentication option when creating/modifying the Active Directory connection.	False	Yes
LM Compatibility Level	Level of supported authentication dialects for Active Directory connections. See the section "SMB authentication dialects" for more information.	ntlmv2-krb	No
Require SMB Encryption for Incoming CIFS Traffic	Requires SMB encryption for all shares. This is not used by Cloud Volumes Service; instead, set encryption on a per- volume basis (see the section "SMB share security features").	False	No
Client Session Security	Sets signing and/or sealing for LDAP communication. This is not currently set in Cloud Volumes Service but might be needed in future releases to address . Remediation for LDAP authentication issues due to the Windows patch is covered in the section "LDAP channel binding.".	None	No
SMB2 enable for DC connections	Uses SMB2 for DC connections. Enabled by default.	System-default	No

Security option	What it does	Default value	Can change?
LDAP Referral Chasing	When using multiple LDAP servers, referral chasing allows the client to refer to other LDAP servers in the list when an entry is not found in the first server. This is currently not supported by Cloud Volumes Service.	False	No
Use LDAPS for Secure Active Directory Connections	Enables the use of LDAP over SSL. Currently not supported by Cloud Volumes Service.	False	No
Encryption is required for DC Connection	Requires encryption for successful DC connections. Disabled by default in Cloud Volumes Service.	False	No

## Dual-protocol/multiprotocol

Cloud Volumes Service offers the ability to share the same datasets to both SMB and NFS clients while maintaining proper access permissions (dual-protocol). This is done by coordinating identity mapping between protocols and using a centralized backend LDAP server to provide the UNIX identities to Cloud Volumes Service. You can use Windows Active Directory to provide both Windows and UNIX users for ease of use.

## Access control

- Share access controls. Determine which clients and/or user and groups can access a NAS share. For NFS, export policies and rules control client access to exports. NFS exports are managed from the Cloud Volumes Service instance. SMB makes use of CIFS/SMB shares and share ACLs to provide more granular control at the user and group level. You can only configure share-level ACLs from SMB clients by using MMC/Computer Management with an account that has administrator rights on the Cloud Volumes Service instance (see the section "Accounts with local/BUILTIN administrator/backup rights.").
- File access controls. Control permissions at a file or folder level and are always managed from the NAS client. NFS clients can make use of traditional mode bits (rwx) or NFSv4 ACLs. SMB clients leverage NTFS permissions.

The access control for volumes that serve data to both NFS and SMB depends on the protocol in use. For information on permissions with dual protocol, see the section "Permission model."

## **User mapping**

When a client accesses a volume, Cloud Volumes Service attempts to map the incoming user to a valid user in the opposite direction. This is necessary for proper access to be determined across protocols and to ensure that the user requesting access is indeed who they claim to be.

For example, if a Windows user named joe attempts access to a volume with UNIX permissions through SMB,

then Cloud Volumes Service performs a search to find a corresponding UNIX user named joe. If one exists, then files that are written to an SMB share as Windows user joe appears as UNIX user joe from NFS clients.

Alternately, if a UNIX user named joe attempts access to a Cloud Volumes Service volume with Windows permissions, then the UNIX user must be able to map to a valid Windows user. Otherwise, access to the volume is denied.

Currently, only Active Directory is supported for external UNIX identity management with LDAP. For more information about configuring access to this service, see Creating an AD connection.

#### **Permission model**

When using dual-protocol setups, Cloud Volumes Service makes use of security styles for volumes to determine the type of ACL. These security styles are set based on which NAS protocol is specified, or in the case of dual protocol, is a choice made at the time of Cloud Volumes Service volume creation.

- If you are only using NFS, Cloud Volumes Service volumes use UNIX permissions.
- If you are only using SMB, Cloud Volumes Service volumes use NTFS permissions.

If you are creating a dual-protocol volume, you can choose the ACL style at volume creation. This decision should be made based on the desired permissions management. If your users manage permissions from Windows/SMB clients, select NTFS. If your users prefer using NFS clients and chmod/chown, use UNIX security styles.

#### **Considerations for creating Active Directory connections**

Cloud Volumes Service provides the ability to connect your Cloud Volumes Service instance to an external Active Directory server for identity management for both SMB and UNIX users. Creating an Active Directory connection is required to use SMB in Cloud Volumes Service.

The configuration for this provides several options that require some consideration for security. The external Active Directory server can be an on-premises instance or cloud native. If you are using an on-premises Active Directory server, don't expose the domain to the external network (such as with a DMZ or an external IP address). Instead, use secure private tunnels or VPNs, one-way forest trusts, or dedicated network connections to the on-premises networks with Private Google Access. See the Google Cloud documentation for more information about best practices using Active Directory in Google Cloud.



CVS-SW requires Active Directory servers to be located in the same region. If a DC connection is attempted in CVS-SW to another region, the attempt fails. When using CVS-SW, be sure to create Active Directory sites that include the Active Directory DCs and then specify sites in Cloud Volumes Service to avoid cross-region DC connection attempts.

#### **Active Directory credentials**

When SMB or LDAP for NFS is enabled, Cloud Volumes Service interacts with the Active Directory controllers to create a machine account object to use for authentication. This is no different from how a Windows SMB client joins a domain and requires the same access rights to Organizational Units (OUs) in Active Directory.

In many cases, security groups do not allow the use of a Windows administrator account on external servers such as Cloud Volumes Service. In some cases, the Windows Administrator user is disabled entirely as a security best practice.

#### Permissions needed to create SMB machine accounts

To add Cloud Volumes Service machine objects to an Active Directory, an account that either has administrative rights to the domain or has delegated permissions to create and modify machine account objects to a specified OU is required. You can do this with the Delegation of Control Wizard in Active Directory by creating a custom task that provides a user access to creation/deletion of computer objects with the following access permissions provided:

- Read/Write
- Create/Delete All Child Objects
- Read/Write All Properties
- Change/Reset Password

Doing this automatically adds a security ACL for the defined user to the OU in Active Directory and minimizes the access to the Active Directory environment. After a user has been delegated, that username and password can be provided as Active Directory Credentials in this window.



The username and password that is passed to the Active Directory domain leverages Kerberos encryption during the machine account object query and creation for added security.

#### Active Directory connection details

The Active Directory Connection Details provide fields for administrators to give specific Active Directory schema information for machine account placement, such as the following:

- Active Directory Connection Type. Used to specify whether the Active Directory connection in a region is used for volumes of either Cloud Volumes Service or CVS-Performance service type. If this is set incorrectly on an existing connection, it might not work properly when used or edited.
- Domain. The Active Directory domain name.
- Site. Limits Active Directory servers to a specific site for security and performance considerations. This is necessary when multiple Active Directory servers span regions because Cloud Volumes Service does not currently support allowing Active Directory authentication requests to Active Directory servers in a different region than the Cloud Volumes Service instance. (For instance, the Active Directory domain controller is in a region that only CVS-Performance supports but you want an SMB share in a CVS-SW instance.)
- DNS servers. DNS servers to use in name lookups.
- NetBIOS name (optional). If desired, the NetBIOS name for the server. This what is used when new machine accounts are created using the Active Directory connection. For instance, if the NetBIOS name is set to CVS-EAST then the machine account names will be CVS-EAST-{1234}. See the section "How Cloud Volumes Service shows up in Active Directory" for more information.
- **Organizational Unit (OU).** The specific OU to create the computer account. This is useful if you're delegating control to a user for machine accounts to a specific OU.
- **AES Encryption.** You can also check or uncheck the Enable AES Encryption for AD Authentication checkbox. Enabling AES encryption for Active Directory authentication provides extra security for Cloud Volumes Service to Active Directory communication during user and group lookups. Before enabling this option, check with your domain administrator to confirm that the Active Directory domain controllers support AES authentication.

()

By default, most Windows servers do not disable weaker ciphers (such as DES or RC4-HMAC), but if you choose to disable weaker ciphers, confirm Cloud Volumes Service Active Directory connection has been configured to enable AES. Otherwise, authentication failures occur. Enabling AES encryption doesn't disable weaker ciphers but instead adds support for AES ciphers to the Cloud Volumes Service SMB machine account.

## Kerberos realm details

This option does not apply to SMB servers. Rather, it is used when configuring NFS Kerberos for the Cloud Volumes Service system. When these details are populated, the NFS Kerberos realm is configured (similar to a krb5.conf file on Linux) and is used when NFS Kerberos is specified on the Cloud Volumes Service volume creation, as the Active Directory connection acts as the NFS Kerberos Distribution Center (KDC).



Non-Windows KDCs are currently unsupported for use with Cloud Volumes Service.

## Region

A region enables you to specify the location where the Active Directory connection resides. This region must be the same region as the Cloud Volumes Service volume.

 Local NFS Users with LDAP. In this section, there is also an option to Allow Local NFS Users with LDAP. This option must be left unselected if you want to extend your UNIX user group membership support beyond the 16-group limitation of NFS (extended groups). However, using extended groups requires a configured LDAP server for UNIX identities. If you don't have an LDAP server, leave this option unselected. If you have an LDAP server and want to also use local UNIX users (such as root), select this option.

#### Backup users

This option enables you to specify Windows users that have backup permissions to the Cloud Volumes Service volume. Backup privileges (SeBackupPrivilege) are necessary for some applications to properly backup and restore data in NAS volumes. This user has a high level of access to data in the volume, so you should consider enabling auditing of that user access. After it is enabled, audit events display in Event Viewer > Windows Logs > Security.

🐻 Event Properties - Event 4674, Sec	urity-Auditing		×
General Details			
Friendly View <u>X</u> ML View			
SubjectUserName	parisi		
SubjectDomainNa	ame CVSDEMO		
SubjectLogonId (	0x31de4904		
ObjectServer	Security		
ObjectType	-		
ObjectName	e		4
Handleld	0x1174		
AccessMask	1048577		
PrivilegeList	SeBackupPrivilege		
ProcessId	0x498		
ProcessName	C:\Windows\System32\wbem\WmiPrvSE.exe	~	
Сору		Cl	ose

## Security privilege users

This option enables you to specify Windows users that have security modification permissions to the Cloud Volumes Service volume. Security privileges (SeSecurityPrivilege) are necessary for some applications (such as SQL Server) to properly set permissions during installation. This privilege is needed to manage the security log. Although this privilege is not as powerful as SeBackupPrivilege, NetApp recommends auditing user access of users with this privilege level if needed.

For more information, see Special privileges assigned to new logon.

## How Cloud Volumes Service shows up in Active Directory

Cloud Volumes Service shows up in Active Directory as a normal machine account object. The naming conventions are as follows.

- CIFS/SMB and NFS Kerberos create separate machine account objects.
- NFS with LDAP enabled creates a machine account in Active Directory for Kerberos LDAP binds.
- Dual protocol volumes with LDAP share the CIFS/SMB machine account for LDAP and SMB.
- CIFS/SMB machine accounts use a naming convention of NAME-1234 (random four digit ID with hyphen appended to <10 character name) for the machine account. You can define NAME by the NetBIOS name setting on the Active Directory connection (see the section "Active Directory connection details").
- NFS Kerberos uses NFS-NAME-1234 as the naming convention (up to 15 characters). If more than 15 characters are used, the name is NFS-TRUNCATED-NAME-1234.

- NFS-only CVS-Performance instances with LDAP enabled create an SMB machine account for binding to the LDAP server with the same naming convention as CIFS/SMB instances.
- When an SMB machine account is created, default hidden admin shares (see the section "Default hidden shares") are also created (c\$, admin\$, ipc\$), but those shares have no ACLs assigned and are inaccessible.
- The machine account objects are placed in CN=Computers by default, but a you can specify a different OU when necessary. See the section "Permissions needed to create SMB machine accounts" for information about what access rights are needed to add/remove machine account objects for Cloud Volumes Service.

When Cloud Volumes Service adds the SMB machine account to Active Directory, the following fields are populated:

- cn (with the specified SMB server name)
- dNSHostName (with SMBserver.domain.com)
- msDS-SupportedEncryptionTypes (Allows DES\_CBC\_MD5, RC4\_HMAC\_MD5 if AES encryption is not enabled; if AES encryption is enabled, DES\_CBC\_MD5, RC4\_HMAC\_MD5, AES128\_CTS\_HMAC\_SHA1\_96, AES256\_CTS\_HMAC\_SHA1\_96 are allowed for Kerberos ticket exchange with the machine account for SMB)
- name (with the SMB server name)
- sAMAccountName (with SMBserver\$)
- servicePrincipalName (with host/smbserver.domain.com and host/smbserver SPNs for Kerberos)

If you want to disable weaker Kerberos encryption types (enctype) on the machine account, you can change the msDS-SupportedEncryptionTypes value on the machine account to one of the values in the following table to allow AES only.

msDS-SupportedEncryptionTypes value	Enctype enabled
2	DES_CBC_MD5
4	RC4_HMAC
8	AES128_CTS_HMAC_SHA1_96 only
16	AES256_CTS_HMAC_SHA1_96 only
24	AES128_CTS_HMAC_SHA1_96 and AES256_CTS_HMAC_SHA1_96
30	DES_CBC_MD5, RC4_HMAC, AES128_CTS_HMAC_SHA1_96 and AES256_CTS_HMAC_SHA1_96

To enable AES encryption for SMB machine accounts, click Enable AES Encryption for AD Authentication when creating the Active Directory connection.

To enable AES encryption for NFS Kerberos, see the Cloud Volumes Service documentation.

## Other NAS Infrastructure service dependencies (KDC, LDAP, and DNS)

When using Cloud Volumes Service for NAS shares, there might be external dependencies required for proper functionality. These dependencies are in play under specific circumstances. The following table shows various configuration options and what,

if any, dependencies are required.

Configuration	Dependencies required
NFSv3 only	None
NFSv3 Kerberos only	Windows Active Directory: * KDC * DNS * LDAP
NFSv4.1 only	Client ID mapping configuration (/etc/idmap.conf)
NFSv4.1 Kerberos only	<ul> <li>Client ID mapping configuration (/etc/idmap.conf)</li> <li>Windows Active Directory: KDC DNS LDAP</li> </ul>
SMB only	Active Directory: * KDC * DNS
Multiprotocol NAS (NFS and SMB)	<ul> <li>Client ID mapping configuration (NFSv4.1 only; /etc/idmap.conf)</li> <li>Windows Active Directory: KDC DNS LDAP</li> </ul>

## Kerberos keytab rotation/password resets for machine account objects

With SMB machine accounts, Cloud Volumes Service schedules periodic password resets for the SMB machine account. These password resets occur using Kerberos encryption and operate on a schedule of every fourth Sunday at a random time between 11PM and 1AM. These password resets change the Kerberos key versions, rotate the keytabs stored on the Cloud Volumes Service system, and help maintain a greater level of security for SMB servers running in Cloud Volumes Service. Machine account passwords are randomized and are not known to administrators.

For NFS Kerberos machine accounts, password resets take place only when a new keytab is created/exchanged with the KDC. Currently, this is not possible to do in Cloud Volumes Service.

## Network ports for use with LDAP and Kerberos

When using LDAP and Kerberos, you should determine the network ports in use by these services. You can find a complete list of ports in use by Cloud Volumes Service in the Cloud Volumes Service documentation on security considerations.

## LDAP

Cloud Volumes Service acts as an LDAP client and uses standard LDAP search queries for user and group lookups for UNIX identities. LDAP is necessary if you intend to use users and groups outside the standard default users provided by Cloud Volumes Service. LDAP is also necessary if you plan on using NFS Kerberos

with user principals (such as user1@domain.com). Currently, only LDAP using Microsoft Active Directory is supported.

To use Active Directory as a UNIX LDAP server, you must populate the necessary UNIX attributes on users and groups you intend to use for UNIX identities. Cloud Volumes Service uses a default LDAP schema template that queries attributes based on RFC-2307-bis. As a result, the following table shows the bare minimum necessary Active Directory attributes to populate for users and groups and what each attribute is used for.

For more information about setting LDAP attributes in Active Directory, see Managing dual-protocol access.

Attribute	What it does
uid*	Specifies the UNIX user name
uidNumber*	Specifies the UNIX user's numeric ID
gidNumber*	Specifies the UNIX user's primary group numeric ID
objectClass*	Specifies what type of object is being used; Cloud Volumes Service requires "user" to be included in the list of object classes (is included in most Active Directory deployments by default).
name	General information about the account (real name, phone number, and so on—also known as gecos)
unixUserPassword	No need to set this; not used in UNIX identity lookups for NAS authentication. Setting this puts the configured unixUserPassword value in plaintext.
unixHomeDirectory	Defines path to UNIX home directories when a user authenticates against LDAP from a Linux client. Set this if you want to use LDAP for UNIX home directory functionality.
loginShell	Defines path to the bash/profile shell for Linux clients when a user authenticates against LDAP.

\*Denotes attribute is required for proper functionality with Cloud Volumes Service. Remaining attributes are for client-side use only.

Attribute	What it does
cn*	Specifies the UNIX group name. When using Active Directory for LDAP, this is set when the object is first created, but it can be changed later. This name cannot be the same as other objects. For instance, if your UNIX user named user1 belongs to a group named user1 on your Linux client, Windows doesn't allow two objects with the same cn attribute. To work around this, rename the Windows user to a unique name (such as user1-UNIX); LDAP in Cloud Volumes Service uses the uid attribute for UNIX user names.
gidNumber*	Specifies the UNIX group numeric ID.

Attribute	What it does
objectClass*	Specifies what type of object is being used; Cloud Volumes Service requires group to be included in the list of object classes (this attribute is included in most Active Directory deployments by default).
memberUid	Specifies which UNIX users are members of the UNIX group. With Active Directory LDAP in Cloud Volumes Service, this field is not necessary. The Cloud Volumes Service LDAP schema uses the Member field for group memberships.
Member*	Required for group memberships/secondary UNIX groups. This field is populated by adding Windows users to Windows groups. However, if the Windows groups don't have UNIX attributes populated, they are not included in the UNIX user's group membership lists. Any groups that need to be available in NFS must populate the required UNIX group attributes listed in this table.

\*Denotes attribute is required for proper functionality with Cloud Volumes Service. Remaining attributes are for client-side use only.

## LDAP bind information

To query users in LDAP, Cloud Volumes Service must bind (login) to the LDAP service. This login has readonly permissions and is used to query LDAP UNIX attributes for directory lookups. Currently, LDAP binds are possible only by using an SMB machine account.

You can only enable LDAP for CVS-Performance instances and use it for NFSv3, NFSv4.1, or dual-protocol volumes. An Active Directory connection must be established in the same region as the Cloud Volumes Service volume for successful deployment of the LDAP-enabled volume.

When LDAP is enabled, the following occurs in specific scenarios.

- If only NFSv3 or NFSv4.1 is used for the Cloud Volumes Service project, then a new machine account is created in the Active Directory domain controller, and the LDAP client in Cloud Volumes Service binds to Active Directory by using the machine account credentials. No SMB shares are created for the NFS volume and default hidden administrative shares (see the section "Default hidden shares") have share ACLs removed.
- If dual-protocol volumes are used for the Cloud Volumes Service project, then only the single machine account created for SMB access is used to bind the LDAP client in Cloud Volumes Service to Active Directory. No additional machine accounts are created.
- If dedicated SMB volumes are created separately (either before or after NFS volumes with LDAP are enabled), then the machine account for LDAP binds is shared with the SMB machine account.
- If NFS Kerberos is also enabled, two machine accounts are created—one for SMB shares and/or LDAP binds and one for NFS Kerberos authentication.

## **LDAP** queries

Although LDAP binds are encrypted, LDAP queries are passed over the wire in plaintext by using the common LDAP port 389. This well-known port cannot currently be changed in Cloud Volumes Service. As a result,

someone with access to packet sniffing in the network can see user and group names, numeric IDs, and group memberships.

However, Google Cloud VMs cannot sniff other VM's unicast traffic. Only VMs actively participating in LDAP traffic (that is, being able to bind) can see traffic from the LDAP server. For more information about packet sniffing in Cloud Volumes Service, see the section "Packet sniffing/trace considerations."

## LDAP client configuration defaults

When LDAP is enabled in a Cloud Volumes Service instance, an LDAP client configuration is created with specific configuration details by default. In some cases, options either do not apply to Cloud Volumes Service (not supported) or are not configurable.

LDAP client option	What it does	Default value	Can change?
LDAP Server List	Sets LDAP server names or IP addresses to use for queries. This is not used for Cloud Volumes Service. Instead, Active Directory Domain is used to define LDAP servers.	Not set	No
Active Directory Domain	Sets the Active Directory Domain to use for LDAP queries. Cloud Volumes Service leverages SRV records for LDAP in DNS to find LDAP servers in the domain.	Set to the Active Directory domain specified in the Active Directory connection.	No
Preferred Active Directory Servers	Sets the preferred Active Directory servers to use for LDAP. Not supported by Cloud Volumes Service. Instead, use Active Directory sites to control LDAP server selection.	Not set.	No
Bind using SMB Server Credentials	Binds to LDAP by using the SMB machine account. Currently, the only supported LDAP bind method in Cloud Volumes Service.	True	No
Schema Template	The schema template used for LDAP queries.	MS-AD-BIS	No
LDAP Server Port	The port number used for LDAP queries. Cloud Volumes Service currently uses only the standard LDAP port 389. LDAPS/port 636 is not currently supported.	389	No

LDAP client option	What it does	Default value	Can change?
Is LDAPS Enabled	Controls whether LDAP over Secure Sockets Layer (SSL) is used for queries and binds. Currently not supported by Cloud Volumes Service.	False	No
Query Timeout (sec)	Timeout for queries. If queries take longer than the specified value, queries fail.	3	No
Minimum Bind Authentication Level	The minimum supported bind level. Because Cloud Volumes Service uses machine accounts for LDAP binds and Active Directory does not support anonymous binds by default, this option does not come into play for security.	Anonymous	No
Bind DN	The user/distinguished name (DN) used for binds when simple bind is used. Cloud Volumes Service uses machine accounts for LDAP binds and does not currently support simple bind authentication.	Not set	No
Base DN	The base DN used for LDAP searches.	The Windows domain use for the Active Directory connection, in DN format (that is, DC=domain, DC=local).	No
Base search scope	The search scope for base DN searches. Values can include base, onelevel, or subtree. Cloud Volumes Service only supports subtree searches.	Subtree	No
User DN	Defines the DN where user searches start for LDAP queries. Currently not supported for Cloud Volumes Service, so all user searches start at the base DN.	Not set	No

LDAP client option	What it does	Default value	Can change?
User search scope	The search scope for user DN searches. Values can include base, onelevel, or subtree. Cloud Volumes Service does not support setting the user search scope.	Subtree	No
Group DN	Defines the DN where group searches start for LDAP queries. Currently not supported for Cloud Volumes Service, so all group searches start at the base DN.	Not set	No
Group search scope	The search scope for group DN searches. Values can include base, onelevel, or subtree. Cloud Volumes Service does not support setting the group search scope.	Subtree	No
Netgroup DN	Defines the DN where netgroup searches start for LDAP queries. Currently not supported for Cloud Volumes Service, so all netgroup searches start at the base DN.	Not set	No
Netgroup search scope	The search scope for netgroup DN searches. Values can include base, onelevel, or subtree. Cloud Volumes Service does not support setting the netgroup search scope.	Subtree	No
Use start_tls over LDAP	Leverages Start TLS for certificate based LDAP connections over port 389. Currently not supported by Cloud Volumes Service.	False	No
Enable netgroup-by-host lookup	Enables netgroup lookups by hostname rather than expanding netgroups to list all members. Currently not supported by Cloud Volumes Service.	False	No

LDAP client option	What it does	Default value	Can change?
Netgroup-by-host DN	Defines the DN where netgroup-by-host searches start for LDAP queries. Netgroup-by-host is currently not supported for Cloud Volumes Service.	Not set	No
Netgroup-by-host search scope	The search scope for netgroup-by-host DN searches. Values can include base, onelevel or subtree. Netgroup-by-host is currently not supported for Cloud Volumes Service.	Subtree	No
Client session security	Defines what level of session security is used by LDAP (sign, seal, or none). LDAP signing is supported by CVS- Performance, if requested by Active Directory. CVS- SW does not support LDAP signing. For both service types, sealing is currently not supported.	None	No
LDAP referral chasing	When using multiple LDAP servers, referral chasing allows the client to refer to other LDAP servers in the list when an entry is not found in the first server. This is currently not supported by Cloud Volumes Service.	False	No
Group membership filter	Provides a custom LDAP search filter to be used when looking up group membership from an LDAP server. Not currently supported with Cloud Volumes Service.	Not set	No

## Using LDAP for asymmetric name mapping

Cloud Volumes Service, by default, maps Windows users and UNIX users with identical usernames bidirectionally without special configuration. As long as Cloud Volumes Service can find a valid UNIX user (with LDAP), then 1:1 name mapping occurs. For instance, if Windows user johnsmith is used, then, if Cloud Volumes Service can find a UNIX user named johnsmith in LDAP, name mapping succeeds for that user, all files/folders created by johnsmith show the correct user ownership, and all ACLs affecting johnsmith are

honored regardless of the NAS protocol in use. This is known as symmetric name mapping.

Asymmetric name mapping is when the Windows user and UNIX user identity don't match. For instance, if Windows user johnsmith has a UNIX identity of jsmith, Cloud Volumes Service needs a way to be told about the variation. Because Cloud Volumes Service currently doesn't support creation of static name mapping rules, LDAP must be used to look up the identity of the users for both Windows and UNIX identities to ensure proper ownership of files and folders and expected permissions.

By default, Cloud Volumes Service includes LDAP in the ns-switch of the instance for the name map database, so that to provide name mapping functionality by using LDAP for asymmetric names, you only need to modify some of the user/group attributes to reflect what Cloud Volumes Service looks for.

The following table shows what attributes must be populated in LDAP for asymmetric name mapping functionality. In most cases, Active Directory is already configured to do this.

Cloud Volumes Service attribute	What it does	Value used by Cloud Volumes Service for name mapping
Windows to UNIX objectClass	Specifies the type of object being used. (That is, user, group, posixAccount, and so on)	Must include user (can contain multiple other values, if desired.)
Windows to UNIX attribute	that defines the Windows username at creation. Cloud Volumes Service uses this for Windows to UNIX lookups.	No change needed here; sAMAccountName is the same as the Windows login name.
UID	Defines the UNIX username.	Desired UNIX username.

Cloud Volumes Service currently does not use domain prefixes in LDAP lookups, so multiple domain LDAP environments do not function properly with LDAP namemap lookups.

The following example shows a user with the Windows name asymmetric, the UNIX name unix-user, and the behavior it follows when writing files from both SMB and NFS.

The following figure shows how LDAP attributes look from the Windows server.

asymmetric Properties

Published (	Certificates	Member Of	Passwor	d Replication	Dial-in	Object
Security	Er	nvironment	Sess	sions	Remote o	ontrol
General	Address	Account	Profile	Telephones	Orga	nization
Remote	Desktop Se	ervices Profile	С	OM+	Attribute I	Editor

## Attributes:

Attribute	Value	^
name	asymmetric	
objectCategory	CN=Person,CN=Schema,CN=Configuration,	
objectClass	top; person; organizationalPerson; user	
objectGUID	de489556-dd7b-43a3-98fa-2722f79d67ed	
objectSid	S-1-5-21-3552729481-4032800560-2279794	
primaryGroupID	513 = ( GROUP_RID_USERS )	
pwdLastSet	1/19/2017 1:56:34 PM Eastern Standard Tim	
replPropertyMetaData	AttID Ver Loc.USN Org.DSA	
sAMAccountName	asymmetric	
sAMAccountType	805306368 = ( NORMAL_USER_ACCOUNT	
uid	unix-user	
uidNumber	1207	

From an NFS client, you can query the UNIX name but not the Windows name:

```
# id unix-user
uid=1207(unix-user) gid=1220(sharedgroup) groups=1220(sharedgroup)
# id asymmetric
id: asymmetric: no such user
```

When a file is written from NFS as unix-user, the following is the result from the NFS client:

```
sh-4.2$ pwd
/mnt/home/ntfssh-4.2$ touch unix-user-file
sh-4.2$ ls -la | grep unix-user
-rwx----- 1 unix-user sharedgroup 0 Feb 28 12:37 unix-user-nfs
sh-4.2$ id
uid=1207(unix-user) gid=1220(sharedgroup) groups=1220(sharedgroup)
```

From a Windows client, you can see that the owner of the file is set to the proper Windows user:

```
PS C:\ > Get-Acl \\demo\home\ntfs\unix-user-nfs | select Owner
Owner
-----
NTAP\asymmetric
```

Conversely, files created by the Windows user asymmetric from an SMB client show the proper UNIX owner, as shown in the following text.

SMB:

PS Z:\ntfs> echo TEXT > asymmetric-user-smb.txt

NFS:

```
sh-4.2$ ls -la | grep asymmetric-user-smb.txt
-rwx----- 1 unix-user sharedgroup 14 Feb 28 12:43 asymmetric-
user-smb.txt
sh-4.2$ cat asymmetric-user-smb.txt
TEXT
```

#### LDAP channel binding

Because of a vulnerability with Windows Active Directory domain controllers, Microsoft Security Advisory ADV190023 changes how DCs allow LDAP binds.

The impact for Cloud Volumes Service is the same as for any LDAP client. Cloud Volumes Service does not currently support channel binding. Because Cloud Volumes Service supports LDAP signing by default through negotiation, LDAP channel binding should not be an issue. If you do have issues binding to LDAP with channel binding enabled, follow the remediation steps in ADV190023 to allow LDAP binds from Cloud Volumes Service to succeed.

#### DNS

Active Directory and Kerberos both have dependencies on DNS for host name to IP/IP to host name resolution. DNS requires port 53 to be open. Cloud Volumes Service does not make any modifications to DNS records, nor does it currently support the use of dynamic DNS on network interfaces.

You can configure Active Directory DNS to restrict which servers can update DNS records. For more information, see Secure Windows DNS.

Note that resources within a Google project default to using Google Cloud DNS, which isn't connected with Active Directory DNS. Clients using Cloud DNS cannot resolve UNC paths returned by Cloud Volumes Service. Windows clients joined to the Active Directory domain are configured to use Active Directory DNS and can resolve such UNC paths.

To join a client to Active Directory, you must configure its DNS configuration to use Active Directory DNS.

Optionally, you can configure Cloud DNS to forward requests to Active Directory DNS. See Why can't my client resolve the SMB NetBIOS name? for more information.



Cloud Volumes Service does not currently support DNSSEC and DNS queries are performed in plaintext.

#### File access auditing

Currently not supported for Cloud Volumes Service.

#### **Antivirus protection**

You must perform antivirus scanning in Cloud Volumes Service at the client to a NAS share. There is currently no native antivirus integration with Cloud Volumes Service.

#### Service operation

The Cloud Volumes Service team manages the backend services in Google Cloud and uses multiple strategies to secure the platform and prevent unwanted access.

Each customer gets their own unique subnet that has access fenced off from other customers by default, and every tenant in Cloud Volumes Service gets their own namespace and VLAN for total data isolation. After a user is authenticated, the Service Delivery Engine (SDE) can only read configuration data specific to that tenant.

#### **Physical security**

With proper preapproval, only onsite engineers and NetApp-badged Field Support Engineers (FSEs) have access to the cage and racks for physical work. Storage and network management is not permitted. Only these onsite resources are able to perform hardware maintenance tasks.

For onsite engineers, a ticket is raised for the statement of work (SOW) that includes the rack ID and device location (RU) and all other details are included in the ticket. For NetApp FSEs, a site visitation ticket must be raised with the COLO and the ticket includes the visitor's details, date, and time for auditing purposes. The SOW for the FSE is communicated internally to NetApp.

## **Operations team**

The operations team for Cloud Volumes Service consists of Production Engineering and a Site Reliability Engineer (SRE) for Cloud Volume Services and NetApp Field Support Engineers and Partners for hardware. All operations team members are accredited for work in Google Cloud and detailed records of work are maintained for every ticket raised. In addition, there is a stringent change control and approval process in place to ensure each decision is appropriately scrutinized.

The SRE team manages the control plane and how the data is routed from UI requests to backend hardware and software in Cloud Volumes Service. The SRE team also manages system resources, such as volume and inode maximums. SREs are not allowed to interact with or have access to customer data. SREs also provide coordination with Return Material Authorizations (RMAs), such as new disk or memory replacement requests for the backend hardware.

#### **Customer responsibilities**

Customers of Cloud Volumes Service manage their organization's Active Directory and user role management as well as the volume and data operations. Customers can have administrative roles and can delegate permissions to other end users within the same Google Cloud project using the two predefined roles that NetApp and Google Cloud provide (Administrator and Viewer).

The administrator can peer any VPC within the customer project to Cloud Volumes Service that the customer determines to be appropriate. It is the responsibility of the customer to manage access to their Google Cloud marketplace subscription and to manage the VPCs that have access to the data plane.

## Malicious SRE protection

One concern that could arise is how does Cloud Volumes Service protect against scenarios in which there is a malicious SRE or when SRE credentials have been compromised?

Access to the production environment is with a limited number of SRE individuals only. Administrative privileges are further restricted to a handful of experienced administrators. All actions performed by anyone in the Cloud Volumes Service production environment are logged and any anomalies to the baseline or suspicious activities are detected by our security information and event management (SIEM) threat intelligence platform. As a result, malicious actions can be tracked and mitigated before too much damage is done to the Cloud Volumes Service backend.

#### Volume life cycle

Cloud Volumes Service manages only the objects within the service—not the data within the volumes. Only clients accessing the volumes can manage the data, the ACLs, file owners, and so on. The data in these volumes is encrypted at rest and access is limited to tenants of the Cloud Volumes Service instance.

The volume lifecycle for Cloud Volumes Service is create-update-delete. Volumes retain Snapshot copies of volumes until the volumes are deleted, and only validated Cloud Volumes Service administrators can delete volumes in Cloud Volumes Service. When a volume deletion is requested by an administrator, an additional step of entering the volume name is required to verify the deletion. After a volume is deleted, the volume is gone and cannot be recovered.

In cases where a Cloud Volumes Service contract is terminated, NetApp marks volumes for deletion after a specific time period. Before that time period expires, you can recover volumes at the customer's request.

## Certifications

Cloud Volumes Services for Google Cloud is currently certified to ISO/IEC 27001:2013 and ISO/IEC 27018:2019 standards. The service also recently received its SOC2 Type I attestation report. For information about the NetApp commitment to data security and privacy, see Compliance: Data security and data privacy.

#### GDPR

Our commitments to privacy and compliance with GDPR are available in a number of our customer contracts, such as our Customer Data Processing Addendum, which includes the Standard Contractual Clauses provided by the European Commission. We also make these commitments in our Privacy Policy, backed by the core values set out in our corporate Code of Conduct.

#### Additional information and contact information

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

· Google Cloud documentation for Cloud Volumes Service

https://cloud.google.com/architecture/partners/netapp-cloud-volumes/

· Google private service access

https://cloud.google.com/vpc/docs/private-services-access?hl=en\_US

NetApp product documentation

https://www.netapp.com/support-and-training/documentation/

Cryptographic Validation Module Program—NetApp CryptoMod

https://csrc.nist.gov/projects/cryptographic-module-validation-program/certificate/4144

The NetApp Solution for Ransomware

https://www.netapp.com/pdf.html?item=/media/16716-sb-3938pdf.pdf&v=202093745

• TR-4616: NFS Kerberos in ONTAP

https://www.netapp.com/pdf.html?item=/media/19384-tr-4616.pdf

#### Contact us

Let us know how we can improve this technical report.

Contact us at doccomments@netapp.com. Include TECHNICAL REPORT 4918 in the subject line.

## **BlueXP Backup and Recovery**

## BlueXP backup and recovery for VMs

#### 3-2-1 Data Protection for VMware with SnapCenter Plug-in and BlueXP backup and recovery for VMs

The 3-2-1 backup strategy is an industry accepted data protection method, providing a comprehensive approach to safeguarding valuable data. This strategy is reliable and ensures that even if some unexpected disaster strikes, there will still be a copy of the data available.

Author: Josh Powell - NetApp Solutions Engineering

## Overview

The strategy is comprised of three fundamental rules:

- 1. Keep at least three copies of your data. This ensures that even if one copy is lost or corrupted, you still have at least two remaining copies to fall back on.
- Store two backup copies on different storage media or devices. Diversifying storage media helps protect against device-specific or media-specific failures. If one device gets damaged or one type of media fails, the other backup copy remains unaffected.
- 3. Finally, ensure that at least one backup copy is offsite. Offsite storage serves as a fail-safe against localized disasters like fires or floods that could render onsite copies unusable.

This solution document covers a 3-2-1 backups solution using SnapCenter Plug-in for VMware vSphere (SCV) to create primary and secondary backups of our on-premises virtual machines and BlueXP backup and

recovery for virtual machines to backup a copy of our data to cloud storage or StorageGRID.

## **Use Cases**

This solution addresses the following use cases:

- Backup and restore of on-premises virtual machines and datastores using using SnapCenter Plug-in for VMware vSphere.
- Backup and restore of on-premises virtual machines and datastores, hosted on ONTAP clusters, and backed up to object storage using BlueXP backup and recovery for virtual machines.

## NetApp ONTAP Data Storage

ONTAP is NetApp's industry leading storage solution that offers unified storage whether you access over SAN or NAS protocols. The 3-2-1 backup strategy ensures on-premises data is protected on more than one media type and NetApp offers platforms ranging from high-speed flash to lower-cost media.



For more information on all of NetApp's hardware platform's check out NetApp Data Storage.

## SnapCenter Plug-in for VMware vSphere

The SnapCenter Plugin for VMware vSphere is a data protection offering which is tightly integrated with VMware vSphere and allows easy management of backup and restores for virtual machines. As part of that solution, SnapMirror provides a fast and reliable method to create a second immutable backup copy of virtual machine data on a secondary ONTAP storage cluster. With this architecture in place, virtual machine restore operations can easily be initiated from either the primary or secondary backup locations.

SCV is deployed as a linux virtual appliance using an OVA file. The plug-in now uses a remote plug-in architecture. The remote plug-in runs outside of the vCenter server and is hosted on the SCV virtual appliance.

For detailed information on SCV refer to SnapCenter Plug-in for VMware vSphere documentation.

## BlueXP backup and recovery for virtual machines

BlueXP backup and recovery is a cloud based tool for data management that provides a single control plane for a wide range of backup and recovery operations across both on-premises and cloud environments. Part of the NetApp BlueXP backup and recovery suite is a feature that integrates with the SnapCenter Plugin for VMware vSphere (on-premises) to extend a copy of the data to object storage in the cloud. This establishes a third copy of the data offsite that is sourced from the primary or secondary storage backups. BlueXP backup and recovery makes it easy to set up storage policies that transfer copies of your data from either of these two on-prem locations. Choosing between the primary and secondary backups as the source in BlueXP Backup and Recovery will result in one of two topologies being implemented:

**Fan-out Topology** – When a backup is initiated by the SnapCenter Plug-in for VMware vSphere, a local snapshot is immediately taken. SCV then initiates a SnapMirror operation that replicates the most recent snapshot to the Secondary ONTAP cluster. In BlueXP Backup and Recovery, a policy specifies the primary ONTAP cluster as the source for a snapshot copy of the data to be transferred to object storage in your cloud provider of choice.



**Cascading Topology** – Creating the primary and secondary data copies using SCV is identical to the fan-out topology mentioned above. However, this time a policy is created in BlueXP Backup and Recovery specifying that the backup to object storage will originate from the secondary ONTAP cluster.



BlueXP backup and recovery can create backup copies of on-premises ONTAP snapshots to AWS Glacier, Azure Blob, and GCP Archive storage.







# AWS Glacier Azure GCP and Deep Glacier Blob Archive Archive Storage

In addition, you can use NetApp StorageGRID as the object storage backup target. For more on StorageGRID refer to the StorageGRID landing page.

#### **Solution Deployment Overview**

This list provides the high level steps necessary to configure this solution and execute backup and restore operations from SCV and BlueXP backup and recovery:

- 1. Configure SnapMirror relationship between the ONTAP clusters to be used for primary and secondary data copies.
- 2. Configure SnapCenter Plug-In for VMware vSphere.
  - a. Add Storage Systems
  - b. Create backup policies
  - c. Create resource groups
  - d. Run backup first backup jobs
- 3. Configure BlueXP backup and recovery for virtual machines
  - a. Add working environment
  - b. Discover SCV and vCenter appliances
  - c. Create backup policies
  - d. Activate backups
- 4. Restore virtual machines from primary and secondary storage using SCV.
- 5. Restore virtual machines from object storage using BlueXP backup and restore.

#### Prerequisites

The purpose of this solution is to demonstrate data protection of virtual machines running in VMware vSphere and located on NFS Datastores hosted by NetApp ONTAP. This solution assumes the following components are configured and ready for use:

- 1. ONTAP storage cluster with NFS or VMFS datastores connected to VMware vSphere. Both NFS and VMFS datastores are supported. NFS datastores were utilized for this solution.
- 2. Secondary ONTAP storage cluster with SnapMirror relationships established for volumes used for NFS datastores.
- 3. BlueXP connector installed for cloud provider used for object storage backups.
- 4. Virtual machines to be backed are on NFS datastores residing on the primary ONTAP storage cluster.
- Network connectivity between the BlueXP connector and on-premises ONTAP storage cluster management interfaces.
- 6. Network connectivity between the BlueXP connector and on-premises SCV appliance VM and between the BlueXP connecter and vCenter.
- 7. Network connectivity between the on-premises ONTAP intercluster LIFs and the object storage service.
- 8. DNS configured for management SVM on primary and secondary ONTAP storage clusters. For more information refer to Configure DNS for host-name resolution.

#### **High Level Architecture**

The testing / validation of this solution was performed in a lab that may or may not match the final deployment environment.



Cloud Provider

## **Solution Deployment**

In this solution, we provide detailed instructions for deploying and validating a solution that utilizes SnapCenter Plug-in for VMware vSphere, along with BlueXP backup and recovery, to perform the backup and recovery of Windows and Linux virtual machines within a VMware vSphere cluster located in an on-premises data center. The virtual machines in this setup are stored on NFS datastores hosted by an ONTAP A300 storage cluster. Additionally, a separate ONTAP A300 storage cluster serves as a secondary destination for volumes replicated using SnapMirror. Furthermore, object storage hosted on Amazon Web Services and Azure Blob were employed as targets for a third copy of the data.

We will go over creating SnapMirror relationships for secondary copies of our backups managed by SCV and configuration of backup jobs in both SCV and BlueXP backup and recovery.

For detailed information on SnapCenter Plug-in for VMware vSphere refer to the SnapCenter Plug-in for VMware vSphere documentation.

For detailed information on BlueXP backup and recovery refer to the BlueXP backup and recovery documentation.

#### Establish SnapMirror relationships between ONTAP Clusters

SnapCenter Plug-in for VMware vSphere uses ONTAP SnapMirror technology to manage the transport of secondary SnapMirror and/or SnapVault copies to a secondary ONTAP Cluster.

SCV backup policies have the option of using SnapMirror or SnapVault relationships. The primary difference is that when using the SnapMirror option, the retention schedule configured for backups in the policy will be the same at the primary and secondary locations. SnapVault is designed for archiving and when using this option a separate retention schedule can be established with the SnapMirror relationship for the snapshot copies on the secondary ONTAP storage cluster.

Setting up SnapMirror relationships can be done in BlueXP where many of the steps are automated, or it can be done using System Manager and the ONTAP CLI. All of these methods are discussed below.

## Establish SnapMirror relationships with BlueXP

The following steps must be completed from the BlueXP web console:

Begin by logging into the BlueXP web console and navigating to the Canvas.

1. Drag and drop the source (primary) ONTAP storage system onto the destination (secondary) ONTAP storage system.

My working environments	My estate		
ronment			
	NTAPSelect On-Premises ONTAP 1.3 ITIB	On-Premises ONTAP 173.74TiB Capacity	
	Capacity		ots-demo On-Premises ONTAP 3TIB Capacity
E13A300 On-Premises ONTAP 75.21T/B			
Capacity			ANF

2. From the menu that appears select **Replication**.



3. On the **Destination Peering Setup** page select the destination Intercluster LIFs to be used for the connection between storage systems.

Replication Setup	Destination Peering Setup					
	Replication require	Select the destination LIFs you wo s an initial connection between the two v For more information about LIF select	uld like to use for cluster peering setup. vorking environments which is called a cli tions, see Cloud Manager documentation	uster peer relationship.		
CVO_InterCluster_B	CVO_InterCluster_A	zoneb-n1	zoneb-n2	☑ intercluster_node_1	✓ intercluster_node_2	
ntaphci-a300-02 : a0a-3510 172.21.254.212/24 up	ntaphci-a300-01 : a0a-3510 172.21.254.211/24 up	<pre>     traphci-a300-01 :         a0a-3484     172.21.228.21/24 up </pre>	<pre>     traphci-a300-02 :         a0a-3484     172.21.228.22/24 up </pre>	<b>P</b> ntaphci-a300-01 : a0a-181 10.61.181.193/24 up	<b>P</b> ntaphci-a300-01 : a0a-181 10.61.181.194/24   up	

4. On the **Destination Volume Name** page, first select the source volume and then fill out the destination volume name and select the destination SVM and aggregate. Click on **Next** to continue.

		Select the volume that you want to r	eplicate
E13A300			
CDM01	ONLINE	Data	CONLINE
FO torage VM Name F502 learng Policy None olume Type RW	CAPACITY 206 GB Allocated Disk Used	NFO CAPY Storage VM Name F502 Thering Policy None Volume Type RW	S12 GB Allocated
Demo	= ONLINE	Demo02_01	CONLINE
FO torage VM Name zonea ering Policy None eling Policy PM	250 GB Allocated	INFO CAPY Storage VM Name Demo Thering Policy None Webma Tuna 055	S00 GB Disk Used

## Destination Volume Name

Destination Volume Name

Demo\_copy

Destination Storage VM

EHC\_NFS

**Destination Aggregate** 

EHCAggr01

5. Choose the max transfer rate for replication to occur at.

You should limi	it the transfe	r rate. An unlimited rate mi	zht
rou snoulu inni	it the dansie	r deer var anniheed vate mig	5
negatively impa might impact ye	act the perfo our Internet	rmance of other application performance.	is and it
negatively impa might impact ye	act the perfor	rmance of other application performance.	is and it

6. Choose the policy that will determine the retention schedule for secondary backups. This policy can be created beforehand (see the manual process below in the **Create a snapshot retention policy** step) or can be changed after the fact if desired.

Replication Setup	Replication Policy	
↑ Previous Step	Default Policies Additional Policies	
CloudBackupService-1674046623282 Original Policy Name: CloudBackupService-1674046623282 Creates a SnapiVault relationship which replicates Snapshot copies with the following labels to the destination volume: hourly (1), adia (15), weekly (4) (# of retained Snapshot copies in parenthesis)	CloudBackupService-1674047424679 Custom Policy - No Comment More info	CloudBackupService-1674047718637 Custom Policy - No Comment More info

7. Finally, review all information and click on the **Go** button to start the replication setup process.

↑ Previous Step			Review your selection and star	t the replication process		
	Source	Destination	Source Volume Allocated Size:	250 GB	Destination Aggregate:	EHCAggr01
			Source Volume Used Size:	1.79 GB	Destination Storage VM:	EHC_NFS
	E13A300	ntaphci-a300e9u25	Source Thin Provisioning:	Yes	Max Transfer Rate:	100 MB/s
	=		Destination Volume Allocated Si	e: 250 GB	SnapMirror Policy:	Mirror
			Destination Thin Provisioning:	No	Replication Schedule:	One-time copy
	Demo	Demo_copy				

## Establish SnapMirror relationships with System Manager and ONTAP CLI

All required steps for establishing SnapMirror relationships can be accomplished with System Manager or the ONTAP CLI. The following section provides detailed information for both methods:

#### Record the source and destination Intercluster logical interfaces

For the source and destination ONTAP clusters, you can retrieve the inter-cluster LIF information from System Manager or from the CLI.

1. In ONTAP System Manager, navigate to the Network Overview page and retrieve the IP addresses of Type: Intercluster that are configured to communicate with the AWS VPC where FSx is installed.

Buckets												
Qtrees		Natural Interfacer	Ocutante									
Quotas		Network Interfaces	Portues									
Storage VHs		+ Add								Q Search 👲 De	ownload ♥ Filter	ide 🗸
Tiers												
NETWORK	A	Name	Status	Storage VM	IPspace	Address 0	Current Node	Current Port	Portset	Protocols	Туре	Thre
Overview		vesam_repo	0	Backup	Default	10.61.181.179	E13A300_1	a0a-181		SMB/CIPS, NFS, S3	Data	
Ethernet Ports		CM01	0		Default	10.61.181.180	E13A300_1	181-666			Cluster/Node Mgmt	.0
FC Ports												1
EVENTS & JOBS		HC_N3	0		Default	10.61.181.183	E13A300_1	a0a-181			Intercluster, Cluster/Node Mgmt	0
PROTECTION	~	HC_92	0		Default	10.61.181.184	E13A300_2	181-60tt			Intercluster, Cluster/Node Mgmt	ಂ
	10.22	lif_ora_sym_614	0	ora_tvm	Default	10.61.181.185	E13A300_1	a0a-181		SMB/CIFS, NFS, FL	Data	0

2. To retrieve the Intercluster IP addresses using the CLI run the following command:

ONTAP-Dest::> network interface show -role intercluster

To establish cluster peering between ONTAP clusters, a unique passphrase entered at the initiating ONTAP cluster must be confirmed in the other peer cluster.

1. Set up peering on the destination ONTAP cluster using the cluster peer create command. When prompted, enter a unique passphrase that is used later on the source cluster to finalize the creation process.

```
ONTAP-Dest::> cluster peer create -address-family ipv4 -peer-addrs
source_intercluster_1, source_intercluster_2
Enter the passphrase:
Confirm the passphrase:
```

2. At the source cluster, you can establish the cluster peer relationship using either ONTAP System Manager or the CLI. From ONTAP System Manager, navigate to Protection > Overview and select Peer Cluster.



- 3. In the Peer Cluster dialog box, fill out the required information:
  - a. Enter the passphrase that was used to establish the peer cluster relationship on the destination ONTAP cluster.

- b. Select Yes to establish an encrypted relationship.
- c. Enter the intercluster LIF IP address(es) of the destination ONTAP cluster.
- d. Click Initiate Cluster Peering to finalize the process.

		•	Rem
STORAGE VM PERMISSIONS		PASSPHRASE 1	
All storage VMs (incl ×			•
Storage VMs created in the future also will be giv permissions.	ven	It cannot be determined from the this relationship was encrypted. encrypted? 2 Yes No	e passphrase whethe Is the relationship
		To generate passphrase, La	unch Remote Cluster
		Intercluster Network Interfac	es IP Addresses
		172.30.15.42	
		172.30.14.28	
			Cancel
		+ Add	
4			

4. Verify the status of the cluster peer relationship from the destination ONTAP cluster with the following command:



The next step is to set up an SVM relationship between the destination and source storage virtual machines that contain the volumes that will be in SnapMirror relationships.

1. From the destination ONTAP cluster, use the following command from the CLI to create the SVM peer relationship:

```
ONTAP-Dest::> vserver peer create -vserver DestSVM -peer-vserver
Backup -peer-cluster OnPremSourceSVM -applications snapmirror
```

- 2. From the source ONTAP cluster, accept the peering relationship with either ONTAP System Manager or the CLI.
- 3. From ONTAP System Manager, go to Protection > Overview and select Peer Storage VMs under Storage VM Peers.



- 4. In the Peer Storage VM's dialog box, fill out the required fields:
  - The source storage VM
  - The destination cluster
  - The destination storage VM



5. Click Peer Storage VMs to complete the SVM peering process.

SnapCenter manages retention schedules for backups that exist as snapshot copies on the primary storage system. This is established when creating a policy in SnapCenter. SnapCenter does not manage retention policies for backups that are retained on secondary storage systems. These policies are managed separately through a SnapMirror policy created on the secondary FSx cluster and associated with the destination volumes that are in a SnapMirror relationship with the source volume.

When creating a SnapCenter policy, you have the option to specify a secondary policy label that is added to the SnapMirror label of each snapshot generated when a SnapCenter backup is taken.



On the secondary storage, these labels are matched to policy rules associated with the destination volume for the purpose of enforcing retention of snapshots.

The following example shows a SnapMirror label that is present on all snapshots generated as part of a policy used for daily backups of our SQL Server database and log volumes.

Select secondary replication options	6
Select secondary replication options	

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label	Custom Label 🔹	0
	sql-daily	
Error retry count	3 🗘 🚯	

For more information on creating SnapCenter policies for a SQL Server database, see the SnapCenter documentation.

You must first create a SnapMirror policy with rules that dictate the number of snapshot copies to retain.

1. Create the SnapMirror Policy on the FSx cluster.

ONTAP-Dest::> snapmirror policy create -vserver DestSVM -policy PolicyName -type mirror-vault -restart always

2. Add rules to the policy with SnapMirror labels that match the secondary policy labels specified in the SnapCenter policies.

```
ONTAP-Dest::> snapmirror policy add-rule -vserver DestSVM -policy
PolicyName -snapmirror-label SnapMirrorLabelName -keep
#ofSnapshotsToRetain
```

The following script provides an example of a rule that could be added to a policy:
```
ONTAP-Dest::> snapmirror policy add-rule -vserver sql_svm_dest
-policy Async SnapCenter SQL -snapmirror-label sql-ondemand -keep 15
```



Create additional rules for each SnapMirror label and the number of snapshots to be retained (retention period).

#### **Create destination volumes**

To create a destination volume on ONTAP that will be the recipient of snapshot copies from our source volumes, run the following command on the destination ONTAP cluster:

```
ONTAP-Dest::> volume create -vserver DestSVM -volume DestVolName
-aggregate DestAggrName -size VolSize -type DP
```

#### Create the SnapMirror relationships between source and destination volumes

To create a SnapMirror relationship between a source and destination volume, run the following command on the destination ONTAP cluster:

```
ONTAP-Dest::> snapmirror create -source-path
OnPremSourceSVM:OnPremSourceVol -destination-path DestSVM:DestVol -type
XDP -policy PolicyName
```

#### Initialize the SnapMirror relationships

Initialize the SnapMirror relationship. This process initiates a new snapshot generated from the source volume and copies it to the destination volume.

To create a volume, run the following command on the destination ONTAP cluster:

ONTAP-Dest::> snapmirror initialize -destination-path DestSVM:DestVol

#### Configure the SnapCenter Plug-in for VMware vSphere

Once installed, the SnapCenter Plug-in for VMware vSphere can be accessed from the vCenter Server Appliance Management interface. SCV will manage backups for the NFS datastores mounted to the ESXi hosts and that contain the Windows and Linux VMs.

Review the Data protection workflow section of the SCV documentation for more information on the steps involved in configuring backups.

To configure backups of your virtual machines and datastores the following steps will need to be completed from the plug-in interface.

Discover the ONTAP storage clusters to be used for both primary and secondary backups.

1. In the SnapCenter Plug-in for VMware vSphere navigate to **Storage Systems** in the left-hand menu and click on the **Add** button.

SnapCenter Plug-in for VMware vSphere INSTANCE 10.61.181.201:8080 v

🔄 Dashboard	Storage Systems		
🍺 Settings	🛖 Add 🥖 🗄	dit 🗙 Delete 🕞 Export	
BResource Groups	Name	Display Name	
Policies	E 10.61.181.180	E13A300	
Storage Systems	Anthos	Anthos	
	Backup	Backup	
Guest File Restore	Demo	Demo	
e.	172.21.146.13	FS02	
	170 0414046	6 oro cum	

2. Fill out the credentials and platform type for the primary ONTAP storage system and click on Add.

Add	Storage	System
-----	---------	--------

Platform	All Flash FAS	
Authentication Method	Credentials	O Certificate
Username	admin	
Password	•••••	
Protocol	HTTPS	
Port	443	
Timeout	60	Seconds
Preferred IP	Preferred IP	
Event Management Syste	m(EMS) & AutoSupport Setting	g
Log Snapcenter server	events to syslog	

Policies specify the retention period, frequency and replication options for the backups managed by SCV.

Review the Create backup policies for VMs and datastores section of the documentation for more information.

To create backup policies complete the following steps:

1. In the SnapCenter Plug-in for VMware vSphere navigate to **Policies** in the left-hand menu and click on the **Create** button.

SnapCenter Plug-in for VMware vSphere INSTANCE 10.61.181.201:8080 v

🟠 Dashboard	Policies	
อ Settings	🕂 Create 🥒 Edit 🗙 Rem	ove 🕞 Export
🔃 Resource Groups	<i>▲Name</i>	VM Co
Policies	Daily	No
Storage Systems	FCD	No
itorage of stems	Hourly	No
👩 Guest File Restore	Monthly	No
	On Demand	h1-

2. Specify a name for the policy, retention period, frequency and replication options, and snapshot label.

### New Backup Policy

Name	Daily
Description	description
Retention	Days to keep 🔹 30 🖕 🚺
Frequency	Daily
Replication	🗌 Update SnapMirror after backup 🕧
	🗹 Update SnapVault after backup 🕧
	Snapshot label Daily
Advanced $\vee$	VM consistency ()
	<ul> <li>Include datastores with independent disks</li> </ul>
	Scripts 👔
	Enter script path

When creating a policy in the SnapCenter Plug-in you will see options for SnapMirror and SnapVault. If you choose SnapMirror, the retention schedule specified in the policy will be the same for both the primary and secondary snapshots. If you choose SnapVault, the retention schedule for the secondary snapshot will be based on a separate schedule implemented with the SnapMirror relationship. This is useful when you wish longer retention periods for secondary backups.

(j)

(;)

Snapshot labels are useful in that they can be used to enact policies with a specific retention period for the SnapVault copies replicated to the secondary ONTAP cluster. When SCV is used with BlueXP Backup and Restore, the Snapshot label field must either be blank or <u>match</u> the label specified in the BlueXP backup policy.

3. Repeat the procedure for each policy required. For example, separate policies for daily, weekly, and monthly backups.

Resource groups contain the datastores and virtual machines to be included in a backup job, along with the associated policy and backup schedule.

Review the Create resource groups section of the documentation for more information.

To create resource groups complete the following steps.

1. In the SnapCenter Plug-in for VMware vSphere navigate to **Resource Groups** in the left-hand menu and click on the **Create** button.

🔄 Dashboard	Resource Groups	
👩 Settings	Create / Edit Y Delete	Run Now
Resource Groups	Name	Des
Policies	SMBC	
Storage Systems	Oracle_Servers	
Cuart File Partora	Demo	
Guest File Restore		

SpanCenter Plug-in for VMware vSphere INSTANCE 10, 61 191 201:8080 -

- 2. In the Create Resource Group wizard, enter a name and description for the group, as well as information required to receive notifications. Click on **Next**
- 3. On the next page select the datastores and virtual machines that wish to be included in the backup job and then click on **Next**.

### Create Resource Group

1. General into & nouncation	Scope:	Datastores	· .		
2. Resource 3. Spanning disks	Datacenter:	Datastores Virtual Machines			
4. Policies		Tags	ntity name		
5. Schedules	Available entit	es			Selected entities
6. Summary	🗐 Demo				INFS_SCV
	🗐 DemoD	S			SVERT
	📒 destinat	ion		*	
	🗐 esxi7-ho	-01 Local		>	
	🗐 esxi7-ho	-02 Local		<	
	🗐 esxi7-ho	-03 Local		"	
	🖾 aavi7 ha	041.0001		1	



You have the option to select specific VMs or entire datastores. Regardless of which you choose, the entire volume (and datastore) is backed up since the backup is the result of taking a snapshot of the underlying volume. In most cases, it is easiest to choose the entire datastore. However, if you wish to limit the list of available VMs when restoring, you can choose only a subset of VMs for backup.

4. Choose options for spanning datastores for VMs with VMDKs that reside on multiple datastores and then click on **Next**.

#### Create Resource Group

<ul> <li>1. General info &amp; notification</li> </ul>	<ul> <li>Always exclude all spanning datastores</li> </ul>
<ul> <li>2. Resource</li> </ul>	This means that only the datastores directly added to the resource group and the primary datastore of VMs directly added to the resource group will be backed up.
3. Spanning disks	unectly added to the resource group will be backed up
4. Policies	<ul> <li>Always include all spanning datastores</li> </ul>
5. Schedules	All datastores spanned by all included VMs are included in this backup
6. Summary	Manually select the spanning datastores to be included You will need to modify the list every time new VMs are added There are no spanned entities in the selected virtual entities list.

 $(\mathbf{i})$ 

i

BlueXP backup and recovery does not currently support backing up VMs with VMDKs that span multiple datastores.

5. On the next page select the policies that will be associated with the resource group and click on **Next**.

#### Create Resource Group

<ul> <li>1. General info &amp; notification</li> </ul>	+ Cre	eate			
2. Resource		Name	VM Consistent	Include independent di	Schedule
3. Spanning disks		Daily	No	No	Daily
4. Policies		FCD	No	Yes	On Demand Only
5. Schedules		Monthly	No	No	Monthly
5. Summary		On Demand	No	No	On Demand Only
		Weekly	No	No	Weekly

When backing up SCV managed snapshots to object storage using BlueXP backup and recovery, each resource group can only be associated with a single policy.

6. Select a schedule that will determine at what times the backups will run. Click on Next.

Create Resource Gro	up			
<ul> <li>4. General info &amp; notification</li> </ul>				
<ul> <li>2. Resource</li> </ul>	Daily	•	Туре	Daily
<ul> <li>3. Spanning disks</li> </ul>			Every	1 Day(s)
<ul> <li>4. Policies</li> </ul>			Starting	06/23/2023
5. Schedules				
6. Summary			At	07 🗘 00 🗘 PM

7. Finally, review the summary page and then on **Finish** to complete the resource group creation.

#### Run a backup job

In this final step, run a backup job and monitor its progress. At least one backup job must be successfully completed in SCV before resources can be discovered from BlueXP backup and recovery.

- 1. In the SnapCenter Plug-in for VMware vSphere navigate to **Resource Groups** in the left-hand menu.
- 2. To initiate a backup job, select the desired resource group and click the **Run Now** button.

#### SnapCenter Plug-in for VMware vSphere INSTANCE 10.61.181.201:8080 ~

Dashboard	Resource Groups			
😰 Settings	📥 Create 🥖 Edit 🛛 💥 Delete	Run Now	C Suspend	
Resource Groups	Name	De	scription	
🝓 Policies	Win01			
Storage Systems	SMBC			
Guest File Restore	Oracle_Servers			
ouestime restore	Demo			
>	SQL_Servers_Dally			
	SQL_Servers_Weekly			

3. To monitor the backup job, navigate to **Dashboard** on the left hand menu. Under **Recent Job Activities** click on the Job ID number to monitor the job progress.

Job Details : 2614	Ċ X
Validate Retention Settings	
Quiescing Applications	
🤣 Retrieving Metadata	
Creating Snapshot copy	
Onquiescing Applications	
Registering Backup	
Backup Retention	
🤣 Clean Backup Cache	
🤣 Send EMS Messages	
(Job 2616)SnapVault Update	
Q Running, Start Time: 07/31/2023 07:24:40 PM.	~
	CLOSE DOWNLOAD JOB LOGS

#### Configure Backups to Object Storage in BlueXP backup and recovery

For BlueXP to manage the data infrastructure effectively, it requires the prior installation of a Connector. The Connector executes the actions involved in discovering resources and managing data operations.

For more information on the BlueXP Connector refer to Learn about Connectors in the BlueXP documentation.

Once the connector is installed for the cloud provider being utilized, a graphic representation of the object storage will be viewable from the Canvas.

To configure BlueXP backup and recovery to backup data managed by SCV on-premises, complete the following steps:

The first step is to add the on-premises ONTAP storage systems to BlueXP

1. From the Canvas select **Add Working Environment** to begin.

Canvas My working environments
Add Working Environment

2. Select **On-Premises** from the choice of locations and then click on the **Discover** button.

	Choose a Location				
	Microsoft Azure	aws Amazon Web Services Select	Google Cloud Platform	On-Premises	
On-Premis	ses ONTAP			Discover 🗸 🗸	

3. Fill out the credentials for the ONTAP storage system and click the **Discover** button to add the working environment.

10.61.181.180		
·		
User Name		
admin		
Password		
•••••	0	

To discover the on-premises datastore and virtual machine resources, add info for the SCV data broker and credentials for the vCenter management appliance.

1. From the BlueXP left-hand menu selection **Protection > Backup and recovery > Virtual Machines** 

1	NetApp BlueXP		
8	Storage	•	ironment
9	Health	•	
Ŧ	Protection	¥	
	Backup and recovery	습	Volumes
	Disaster recovery (Beta)	습	Applications
	Replication	☆	Virtual Machines
ବ	Governance	•	Kubernetes
۲	Mobility	•	Job Monitoring
••	Extensions	•	Reports

2. From the Virtual Machines main screen access the **Settings** drop down menu and select **SnapCenter Plug-in for VMware vSphere**.

SnapCente	r Plug-i	n for VMware vSphere
Policies		
	SnapCente Policies	SnapCenter Plug-i Policies

3. Click on the **Register** button and then enter the IP address and port number for the SnapCenter Plugin appliance and the username and password for the vCenter management appliance. Click on the **Register** button to begin the discovery process.

SnapCenter Plug-in for VMware vSphere	Username
10.61.181.201	administrator@vsphere.local

4. The progress of jobs can be monitored from the Job Monitoring tab.

	Job Name: Discover Virtual Resources from SnapCenter Plugin for VMWare vSphere Job ld: 559167ba-8876-45db-9131-b918a165d0a1					
Ot Jot	ner Type	Jul 31 2023, 9:18:22 pm Start Time	Jul 31 2023, 9:18:2 End Time	6 pm 🥏 Su Job Sta	uccess atus	
ıb-Jobs(2)						Collapse All
Job Name	\$1	Job ID 🛟	Start Time	End Time	Duration	• 🗘
Discover Virtual R	esources from SnapCenter Plu	559167ba-8876-45db	Jul 31 2023, 9:18:22 pm	Jul 31 2023, 9:18:26 pm	4 Seconds	
Discoveri	ng Virtual Resources	99446761-f997-4c80-8	Jul 31 2023, 9:18:22 pm	Jul 31 2023, 9:18:24 pm	2 Seconds	
Registerir	g Datastores	b7ab4195-1ee5-40ff-9a	Jul 31 2023, 9:18:24 pm	Jul 31 2023, 9:18:26 pm	2 Seconds	

5. Once discovery is complete you will be able to view the datastores and virtual machines across all discovered SCV appliances.

image::bxp-scv-hybrid-23.png[View available resources]

In BlueXP backup and recovery for virtual machines, create policies to specify the retention period, backup source and the archival policy.

For more information on creating policies refer to Create a policy to back up datastores.

1. From the BlueXP backup and recovery for virtual machines main page, access the **Settings** drop down menu and select **Policies**.



- 2. Click on Create Policy to access the Create Policy for Hybrid Backup window.
  - a. Add a name for the policy
  - b. Select the desired retention period
  - c. Select if backups will be sourced from the primary or secondary on-premises ONTAP storage system
  - d. Optionally, specify after what period of time backups will be tiered to archival storage for additional cost savings.

Policy Details	Policy Name	
Policy Detailo	12 week - daily backups	
Retention ()		
	Daily	
	Backups to retain SnapMir 84 Daily	rror Label
	Weekly	Setup Retention Weekly
	D Monthly	Setup Retention Monthly
	O Secondary	
Archival Policy	Backups reside in standard storage for frequently you can tier backups to archival storage for furthe	accessed data. Optionally, er cost optimization.
	Tier Backups to Archival	
	Archival After (Days)	
	Cancel	ate
(i) The Snapl policy too.	Mirror Label entered here is used to The label name must match the lat	o identify which backups to apply bel name in the corresponding on

The final step is to activate data protection for the individual datastores and virtual machines. The following steps outline how to activate backups to AWS.

For more information refer to Back up datastores to Amazon Web Services.

1. From the BlueXP backup and recovery for virtual machines main page, access the settings drop down for the datastore to be backed up and select **Activate Backup**.

	6 Datastores						
	Filter By 🕂					Q VM View	Settings   ▼
	Datastore	Datastore	īype ≎   vCenter	0   Poli	icy Name 🗘 🍦	Protection Status	• I
	NFS_SCV	NFS	vcsa7-hc.so	ldc.netapp.com		Unprotected	••••
	OTS_DS01	NFS	172.21.254	160 1 Ye	ear Daily LTR	Protected	View Details
	SCV_WKLD	NFS	vcsa7-hc.sd	ldc.netapp.com 1 Ye	ear Daily LTR	Protected	•••
2. A	Assign the policy to be used for the data protection operation and click on <b>Next</b> .						
1	•	Assign Policy	(2) Add Working Enviro	nments (3) Sele	ct Provider (4) Col	nfigure Provider	(5) Review
				Assign Polic	су		
	21 Policies						
	Poli	icy Name	SnapMirror Label	Retention Count	Backup Source	Archiv	al Policy
	5 Ye	ear Daily LTR	daily	daily : 1830	Primary	Not Ac	tive
	🕑 5 Ye	ear Daily LTR	daily	daily : 1830	Primary	Not Act	ive
	<b>7 Ye</b>	ear Weekly LTR	weekly	weekly : 370	Primary	Not Ac	tive

3. At the **Add Working Environments** page the datastore and working environment with a check mark should appear if the working environment has been previously discovered. If the working environment has not been previously discovered you can add it here. Click on **Next** to continue.

Assign	Policy 2 Add Working Environ	ments (3) Select Provider	4 Configure Provider	5 Review
Provide ONTAP cluster (w	Add V	Vorking Environment	S ng environment details will appe	ar for all volumes that reside
SVM	Volume	Unitple working environments when	ironment	sters.
EHC_NFS	NFS_SCV	OnPremV	VorkingEnvironment-6MzE27u1	Edit

4. At the Select Provider page click on AWS and then click on the Next button to continue.

Assign Policy	Add Working Environments	3 Select Provider (4) Configure P	Provider (5) Review
	Selec	t Provider	
aws		0	StorageGRID
Amazon Web Services	Microsoft Azure	Google Cloud Platform	StorageGRID

 Fill out the provider specific credential information for AWS including the AWS access key and secret key, region, and archival tier to be used. Also, select the ONTAP IP space for the on-premises ONTAP storage system. Click on Next.

Assign Policy Add Working Environments	Select Provider 4 Configure Provider 5 Revie
Con Cloud Manager needs the fol	figure Provider owing details to connect with the cloud provider.
Provider Information	Location and Connectivity
AWS Account	Region
	US East (N. Virginia)
AWS Access Key	IP space for Environment OnPremWorkingEnvironment-6MzE27u1
Enter AWS Access Key Required	Default
AWS Secret Key	
Enter AWS Secret Key	Archival Tier
Required	Glacier

6. Finally, review the backup job details and click on the **Activate Backup** button to initiate data protection of the datastore.

#### Restoring Virtual Machines in the case of data loss

Ensuring the safeguarding of your data is only one aspect of comprehensive data protection. Equally crucial is the ability to promptly restore data from any location in the event of data loss or a ransomware attack. This capability is vital for maintaining seamless business operations and meeting recovery point objectives.

NetApp offers a highly adaptable 3-2-1 strategy, providing customized control over retention schedules at the

 $(\checkmark)$ 

Add Working Environments

	Rev	view
	Policy	5 Year Daily LTR
	SVM	EHC_NFS
	Volumes	NFS_SCV
	Working Environment	OnPremWorkingEnvironment-6MzE27u1
	Backup Source	Primary
	Cloud Service Provider	AWS
	AWS Account	
	AWS Access Key	
	Region	US East (N. Virginia)
	IP space	Default
	Tier Backups to Archival	No
	Previous	Activate Backup
D	At this point data transfer may not imm scans for any outstanding snapshots er storage.	ediately begin. BlueXP backup and recovery very hour and then transfers them to object

Select Provider

(~)

 $(\checkmark)$ 

Configure Provider



primary, secondary, and object storage locations. This strategy provides the flexibility to tailor data protection approaches to specific needs.

This section provides an overview of the data restoration process from both the SnapCenter Plug-in for VMware vSphere and BlueXP backup and recovery for virtual machines.

#### Restoring Virtual Machines from SnapCenter Plug-in for VMware vSphere

For this solution virtual machines were restored to original and alternate locations. Not all aspects of SCV's data restoration capabilities will be covered in this solution. For in depth information on all that SCV has to offer refer to the Restore VMs from backups in the product documentation.

Complete the following steps to restore a virtual machine restore from primary or secondary storage.

- 1. From the vCenter client navigate to **Inventory > Storage** and click on the datastore that contains the virtual machines you wish to restore.
- 2. From the **Configure** tab click on **Backups** to access the list of available backups.

vSphere Client Q, Search in at environments					C &	Administrator	EVSPHERELOCAL ~	•	0
	Summary Monitor	Permission Files Hosts VMs							
<ul> <li>South State State (State State Stat</li></ul>	Alarm Defections Tacheduled Tasks Genetal	Backups	(+ lopint				(m)		
DemoDS	Device Backing	Name Same	Locatoria	Stati Tate	End Trees	Mounted	Policy	SMapp S	lapahó,
destruction	Connectivity with Hosts Inactivant Acceleration Copatility with ShapCenter Physics for VMwa., v Inaccent Science	soc.mm, may 67.112023, 112430.0751. Compared	Pitmary & Secondary	1012020 7 24 35 PM	1/19/2003 7 24 10 PM	No	Dwly	No	
exi3-8c-01 cocal		routine pay, 67-11 2021, 09 200000. Completed	Premary & Secondary	TOD 2023 IS 25:00 AM	T19/2023 9 28:02 AM	No	Detty	No	
exi7+0-421.008     exi7+0-421.008     exi7+0-421.008     exi7+0-441.008		10-, mm, mm, 01 10 2023, 04 30:00 ft. Compared	Promary & Security	700/2022 9:29:00 AM	1130/2023 939402 AM	No	Dely	160	
		ico men.min.07/29-2023.09/380005. Completed	Premary & Secondary	P29/2023 8 39:00 AM	109/2023 9:39 02 AM	hip	Daty	No	
		104, mere, may 07 38 2071, 24 76 00 8. Compared	Primary & Secondary	7/20/2023 (9.28 00 AM	7/20/2023 9:39:02 AM	340.1	Dively	No.	
exo7-bc-05 Local	The second s	soc.mmil.mety.0123.2023.3025.5020. Completed	Primary & Successivy	1/2//2023 10 29:10 AM	7/27/2023 10:25:52 XM	No	Deely	100	
eso7-ac-06 Local	Electrical -	scy, mine, may \$177,2021, 39,5728.06. Completed	Prenary & Secondary	7/21/2023 US7.28 AM	7272023 85730 AM	No	Dety	No	
🗐 iso		10x, mms, may, 07.27.2022, 00.39.00.05. Completed	Primary & Secondary	7/27/2023 8:38:00 AM	WA 65464 (1927202) k3822 AM	tio	Dwy	.No.	
目 NPS_SCV		and strength of 22 2022, 00 24 h Dark Completed	Primary & Secondary	7/21/2023 8 34 Yi AM	7/27/2023 9.3412 AM	No	Dety	1911	
I NPS_SOL									
ER SOV DEMO									

3. Click on a backup to access the list of VMs and then select a VM to restore. Click on **Restore**.

SCV_DEMO	NS						
Summary Monitor Configur	e Permissions File	es Hosts	VMs				
Alarm Definitions Scheduled Tasks General Device Backing Connectivity with Hosts Hardware Acceleration	Name: scv_dem Time Stamp: Mon Jul 3 Mounted: No Policy: Daily VMware snapshot: No Entities	o_daily_07-31-202 11 2023 19:24:36 G	3_19.24.36.0755 MT-0400 (Eastern Daylight Time)				
Capability sets SnapCenter Plug-In for VMwa ✓ Resource Groups Backups	The following entities are included in the backup:scv_demo_daily_07-31-2023_19.24.36.0755 Select an entity and click Restore to restore it.						
	Entity Name	Quiesced	UUD	Location			
	SQLSRV-07	No	5032d1f2-2591-7f7b-46e3-8dbd4a6b2fb4	[SCV_DEMO] SQLSRV-07/SQLSRV-07.vmx			
	scv_restore_test	Yes	50323c8e-04a3-5acf-a2df-a6bc0ced0419	[SCV_DEMO] scv_restore_test/scv_restore_test.vmx			
	SQLSRV-06	No	50327515-8cce-5942-0f85-350ad39bce42	[SCV_DEMO] SQLSRV-06/SQLSRV-06.vmx			
	SQLSRV-08	No	5032b2a9-e1af-c56a-6923-6dbd0eeb6327	[SCV_DEMO] SQLSRV-08/SQLSRV-08.vmx			
	SQLSRV-05	No	50326625-dd29-af23-2fd5-fe04e0a57a69	[SCV_DEMO] SQLSRV-05/SQLSRV-05.vmx			
	SCV_DEMO	No	netfs://172.21.118.112///SCV_DEMO	SCV:/vol/SCV_DEMO			

 From the Restore wizard select to restore the entire virtual machine or a specific VMDK. Select to install to the original location or alternate location, provide VM name after restore, and destination datastore. Click Next.

<ul> <li>✓ 1. Select scope</li> </ul>	Restore scope	Entire virtual machine	
2. Select location	Restart VM		
3. Summary	Restore Location	Original Location	
		(This will restore the entire VM to the original Hypervisor with the origi	inal
		settings. Existing VM will be unregistered and replaced with this VM.)	
		Alternate Location	
		(This will create a new VM on selected vCenter and Hypervisor with the	he
		customized settings.)	
	Destination vCenter Server	10.61.181.210 -	
	Destination ESXi host	esxi7-hc-04.sddc.netapp.com	
	Network	Management 181	
	vivi fidine alter restore	Sul_SRV_08_restored	
Choose to backup fr	rom the primary or second	BACK NEXT FINISH ary storage location.	CANCE
Choose to backup fr Restore	rom the primary or second	BACK NEXT FINISH	CANCE
Choose to backup fr Restore < 1. Select scope	rom the primary or second	BACK NEXT FINISH ary storage location.	CANCE
Choose to backup fr Restore < 1. Select scope 2. Select location	rom the primary or second	BACK NEXT FINISH ary storage location.	CANCE
Choose to backup fr Restore < 1. Select scope 2. Select location 3. Summary	rom the primary or second Destination datastore Lo SCV_DEMO ((	BACK NEXT FINISH ary storage location.	CANCE
Choose to backup fr Restore < 1. Select scope 2. Select location 3. Summary	rom the primary or second Destination datastore Lo SCV_DEMO ((	BACK NEXT FINISH ary storage location. ations rimary) SCV:SCV_DEMO mary SCV:SCV_DEMO econdary) EHC_NFS:SCV_DEMO_dest	CANCE
Choose to backup fr Restore < 1. Select scope 2. Select location 3. Summary	rom the primary or second Destination datastore Lo SCV_DEMO (f	BACK NEXT FINISH ary storage location. ations rimary) SCV:SCV_DEMO econdary) EHC_NFS:SCV_DEMO_dest	CANCE
Choose to backup fr Restore < 1. Select scope 2. Select location 3. Summary	rom the primary or second Destination datastore Lo SCV_DEMO ((	BACK NEXT FINISH ary storage location. ations rimary) SCV:SCV_DEMO mary SCV:SCV_DEMO econdary) EHC_NFS:SCV_DEMO_dest	CANCE

#### Restoring Virtual Machines from BlueXP backup and recovery for virtual machines

BlueXP backup and recovery for virtual machines allows restores of virtual machines to their original location. Restore functions are accessed through the BlueXP web console.

For more information refer to Restore virtual machines data from the cloud.

To restore a virtual machine from BlueXP backup and recovery, complete the following steps.

1. Navigate to **Protection > Backup and recovery > Virtual Machines** and click on Virtual Machines to view the list of virtual machines available to be restored.

Backup and recovery	Volumes	Restore	Applications	Virtual M	lachines	Kubernet	es Job Monitoring	Reports
	$\bigcirc$	4 Working En	vironments		<b>6</b> Datastores		J14 Virtual Machines	

2. Access the settings drop down menu for the VM to be restored and select

4 Virtual Machines					
ilter By +				Q 🛑 VM Vie	w Settings   V
Virtual Machine	O Datastore Type	≎ vCenter ∽	Policy Name	≎   Protection Status ≎   Last Bacl	kup ⊜
SQLSRV-08	NFS	vcsa7-hc.sddc.netap		Unprotected	•••
SQLSRV-04	NFS	vcsa7-hc.sddc.netap	1 Year Daily LTR	Protected Jul 31, 20	23, 7:2
OracleSrv_03	NFS	vcsa7-hc.sddc.netap		Unprotected	Restore

3. Select the backup to restore from and click on Next.

	Backup Name	\$	Backup Time	٥
0	SQL_Servers_Daily_07-31-2023_19.23.39.0938		Jul 31, 2023, 7:23:42 PM	
	SQL_Servers_Daily_07-31-2023_16.40.00.0661		Jul 31, 2023, 4:40:03 PM	
	SQL_Servers_Daily_07-30-2023_16.40.00.0690		Jul 30, 2023, 4:40:03 PM	

- 4. Review a summary of the backup job and click on **Restore** to start the restore process.
- 5. Monitor the progress of the restore job from the **Job Monitoring** tab.

		<b>/ dol</b> dol	lame: Restore 17 files fro	e6620fdbf		
	Restore Files Job Type	NFS_SQL Restore Content	17 Files Content Files	NFS_SQL 3 I Restore to Job	In Progress Status	
						Expan
B	Restore Content					^
aws	ots-demo Working Environment Name	NAS_VOLS SVM Name	NFS_SQL Volume Name	SQL_Servers_Daily_07-31-202 Backup Name	3 Jul 31 2023, 7:24:0 Backup Time	3 pm
5	Restore from					~
aws	AWS	us-east-1	982589175402	netapp-backup-d56250b0-24a	id	
	Provider	Region	Account ID	Bucket/Container Name		

#### Conclusion

The 3-2-1 backup strategy, when implemented with SnapCenter Plug-in for VMware vSphere and BlueXP backup and recovery for virtual machines, offers a robust, reliable, and cost-effective solution for data protection. This strategy not only ensures data redundancy and accessibility but also provides the flexibility of restoring data from any location and from both on-premises ONTAP storage systems and cloud based object storage.

The use case presented in this documentation focuses on proven data protection technologies that highlight the integration between NetApp, VMware, and the leading cloud providers. The SnapCenter Plug-in for VMware vSphere provides seamless integration with VMware vSphere, allowing for efficient and centralized management of data protection operations. This integration streamlines the backup and recovery processes for virtual machines, enabling easy scheduling, monitoring, and flexible restore operations within the VMware ecosystem. BlueXP backup and recovery for virtual machines provides the one (1) in 3-2-1 by providing secure, air-gapped backups of virtual machine data to cloud based object storage. The intuitive interface and logical workflow provide a secure platform for long-term archival of critical data.

#### **Additional Information**

To learn more about the technologies presented in this solution refer to the following additional information.

- SnapCenter Plug-in for VMware vSphere documentation
- BlueXP documentation

### VMware Sovereign Cloud

#### VMware Resources for Sovereign Cloud

#### NetApp and VMware Sovereign Cloud

#### **Overview of VMware Sovereign Cloud**

The concept of sovereignty is emerging as a necessary component of cloud computing for many entities that process and maintain highly sensitive data, such as national and state governments, and highly regulated industries, such finance and healthcare. National governments are also looking to expand digital economic capability and reduce reliance on multi-national firms for their cloud services.

#### VMware Sovereign Cloud Initiative

VMware defines a sovereign cloud as one that:

- Protects and unlocks the value of critical data (e.g., national data, corporate data, and personal data) for both private and public sector organizations
- · Delivers a national capability for the digital economy
- · Secures data with audited security controls
- · Ensures compliance with data privacy laws
- Improves control of data by providing both data residency and data sovereignty with full jurisdictional control

#### Partnering with a Trusted VMware Sovereign Cloud Service Provider

To ensure success, organizations must work with partners they trust and that are capable of hosting authentic and autonomous sovereign cloud platforms. VMware Cloud Providers recognized within the VMware Sovereign Cloud initiative commit to designing and operating cloud solutions based on modern, software-defined architectures that embody key principles and best practices outlined in the VMware Sovereign Cloud framework.

- Data Sovereignty and Jurisdictional Control All data is resident and subject to the exclusive control and authority of the nation state where that data was collected. Operations are fully managed within the jurisdiction
- **Data Access and Integrity** Cloud infrastructure is resilient and available in at least two data center locations within the jurisdiction with secure and private connectivity options available.
- **Data Security and Compliance** Information security management system controls are certified against an industry recognized global (or regional) standard and audited regularly.
- **Data Independence and Mobility** Support for modern application architectures to prevent vendor cloud lock-in and enable application portability and independence

For more information from VMware, please visit:

- VMware Sovereign Cloud Overview
- What is VMware Sovereign Cloud?
- Introducing the New VMware Sovereign Cloud Initiative
- VMware Sovereign Cloud Technical White Paper

#### Netpp with VMware Sovereign Cloud: Use Cases

NetApp provides support for VMware Sovereign Cloud concepts through the integration of several NetApp technologies.

Use the following link(s) to discover more about the NetApp technology integrations with VMware Sovereign Cloud:

NetApp StorageGRID as an Object Store Extension

#### NetApp StorageGRID as an Object Store Extension

NetApp has collaborated with VMware to integrate NetApp StorageGRID into VMware Cloud Director in support of the VMware Sovereign Cloud. This plug-in to VMware Cloud Director enables service providers to use StorageGRID as their object storage offering (regardless of use case) and allows StorageGRID management through the same VMware multi-tenant solution (VMware Cloud Director) used by service providers to manage other parts of their offering catalog.

Partners that deliver VMware Sovereign Clouds can choose NetApp StorageGRID to help them managed and maintain cloud environments with unstructured data. Its universal compatibility in its native support for industry-standard APIs, like Amazon S3 API, helps ensure smooth interoperability across diverse cloud environments, and unique innovations such as automated lifecycle management helps ensure more cost-effective safeguarding, storage, and long-term preservation of customers' unstructured data.

NetApp's Sovereign Cloud integration with Cloud Director providers customers with:

- Assurance that sensitive data, including metadata, remains under sovereign control while preventing access by foreign authorities that could violate data privacy laws.
- Increased security and compliance that protects applications and data from rapidly evolving attack vectors while maintaining continuous compliance with a trusted local. infrastructure, built-in frameworks, and local experts.
- Future-proofed infrastructure to react quickly to changing data privacy regulations, security threats, and geopolitics.
- The ability to unlock the value of data with secure data sharing and analysis to drive innovation without violating privacy laws. Data integrity is protected to ensure accurate insights.

For more information on the StorageGRID integration, check out the following:

NetApp Announcement

# NetApp Hybrid Multicloud with Red Hat OpenShift Container workloads

#### NetApp Hybrid Multicloud solutions for Red Hat OpenShift Container workloads

NetApp is seeing a significant increase in customers modernizing their legacy enterprise applications and building new applications using containers and orchestration platforms built around Kubernetes. Red Hat OpenShift Container Platform is one example that we see adopted by many of our customers.

#### Overview

As more and more customers begin adopting containers within their enterprises, NetApp is perfectly positioned to help serve the persistent storage needs of their stateful applications and classic data management needs such as data protection, data security, and data migration. However, these needs are met using different strategies, tools, and methods.

**NetApp ONTAP** based storage options listed below, deliver security, data protection, reliability, and flexibility for containers and Kubernetes deployments.

- Self-managed storage in on-premises:
  - NetApp Fabric Attached Storage (FAS), NetApp All Flash FAS Arrays (AFF), NetApp All SAN Array (ASA) and ONTAP Select
- Provider-managed storage in on-premises:
  - NetApp Keystone provides Storage as a Service (STaaS)
- Self-managed storage in the cloud:
  - NetApp Cloud Volumes ONTAP(CVO) provide self managed storage in the hyperscalers
- Provider-managed storage in the cloud:
  - Cloud Volumes Service for Google Cloud (CVS), Azure NetApp Files (ANF), Amazon FSx for NetApp ONTAP offer fully managed storage in the hyperscalers

ONTAP feature highlights

Storage Administration	Performance & Scalability			
<ul> <li>Multi-tenancy</li> <li>FlexVol &amp; FlexGroup</li> <li>LUN</li> <li>Quotas</li> <li>ONTAP CLI &amp; API</li> <li>System Manager &amp; BlueXP</li> </ul>	FlexCache     Inconnect, session trunking, multipathing     FlexClone     Scale-out clusters			
Availability & Resilience	Access Protocols			
<ul> <li>Multi-AZ HA deployment (MetroCluster)</li> <li>SnapShot &amp; SnapRestore</li> <li>SnapMirror Cloud</li> <li>SnapMirror</li> </ul>	<ul> <li>NFS –v3, v4, v4.1, v4.2</li> <li>iSCSI</li> <li>SMB – v2, v3</li> <li>Multi-protocol access</li> </ul>			
Storage Efficiency	Security & Compliance			
<ul> <li>Deduplication &amp; Compression</li> <li>Thin provisioning</li> <li>Compaction</li> <li>Data Tiering (Fabric Pool)</li> </ul>	<ul> <li>Fpolicy &amp; Vscan</li> <li>LDAP &amp; Kerberos</li> <li>Active Directory integration</li> <li>Certificate based authentication</li> </ul>			

**NetApp BlueXP** enables you to manage all of your storage and data assets from a single control plane/interface.

You can use BlueXP to create and administer cloud storage (for example, Cloud Volumes ONTAP and Azure NetApp Files), to move, protect, and analyze data, and to control many on-prem and edge storage devices.

**NetApp Astra Trident** is a CSI Compliant Storage Orchestrator that enable quick and easy consumption of persistent storage backed by a variety of the above-mentioned NetApp storage options. It is an open-source software maintained and supported by NetApp.

#### Astra Trident CSI feature highlights



Business critical container workloads need more than just persistent volumes. Their data management requirements require protection and migration of the application kubernetes objects as well.

Application data includes kubernetes objects in addition to the user data: Some examples are as follows:

- kubernetes objects such as pods specs, PVCs, deployments, services
- custom config objects such as config maps and secrets
- persistent data such as Snapshot copies, backups, clones
- custom resources such as CRs and CRDs

÷.

**NetApp Astra Control**, available as both fully-managed and self-managed software, provides orchestration for robust application data management. Refer to the Astra documentation for additional details on the Astra family of products.

This reference documentation provides validation of migration and protection of container-based applications, deployed on RedHat OpenShift container platform, using NetApp Astra Control Center. In addition, the solution provides high-level details for the deployment and the use of Red Hat Advanced Cluster Management (ACM) for managing the container platforms. The document also highlights the details for the integration of NetApp storage with Red Hat OpenShift container platforms using Astra Trident CSI provisioner. Astra Control Center is deployed on the hub cluster and is used to manage the container applications and their persistent storage lifecycle. Finally, it provides a solution for replication and failover and fail-back for container workloads on managed Red Hat OpenShift clusters in AWS (ROSA) using Amazon FSx for NetApp ONTAP (FSxN) as persistent storage.

#### Value propositions of NetApp Hybrid Multicloud solutions for Red Hat OpenShift Container workloads

Most customers do not just start out building Kubernetes based environments without any existing infrastructure. Perhaps they are a traditional IT shop running most of their enterprise applications on virtual machines (in large VMware environments for example). Then they start building small container-based environments to satisfy the needs of their modern application development teams. These initiatives usually start small and begin to

become more pervasive as the teams learn these new technologies and skills, and begin to recognize the many benefits of adopting them.

The good news for customers is that NetApp can serve the needs of both environments. This set of solutions for hybrid multicloud with Red Hat OpenShift will empower NetApp customers to adopt modern cloud technologies and services without having to overhaul their entire infrastructure and organization. Whether customer applications and data are hosted on-premises, in cloud, run on virtual machines, or on containers, NetApp can provide consistent data management, protection, security, and portability. With these new solutions, the same value NetApp has delivered in on-premises data center environments for decades will be available across the enterprise entire data horizon, without requiring significant investment to retool, acquire new skills, or build new teams. NetApp is positioned well to help customers solve these business challenges regardless of what phase of their cloud journey they are in.

NetApp Hybrid Multi-Cloud with Red Hat Openshift:

- Gives customers validated designs and practices which demonstrate the best ways for customers to manage, protect, secure, and migrate their data and applications when using Red Hat OpenShift with NetApp based storage solutions.
- Present best practices for customers running Red Hat OpenShift with NetApp storage in VMware environments, bare metal infrastructure, or a combination of both.
- Demonstrate strategies and options for both on-prem and cloud environments, as well as hybrid environments where both are used.

#### Supported Solutions of NetApp Hybrid Multicloud for Red Hat OpenShift Container workloads

The solution tests and validates Migration & Centralized Data Protection with OpenShift container platform (OCP), OpenShift Advanced Cluster Manager (ACM), NetApp ONTAP, NetApp BlueXP and NetApp Astra Control Center (ACC).

For this solution, the following scenarios are tested and validated by NetApp. The solution is separated into multiple scenarios based on the following characteristics:

- on-premises
- cloud
  - self-managed OpenShift clusters and self-managed NetApp storage
  - $\circ\,$  provider-managed OpenShift clusters and provider-managed NetApp storage

#### We will be building out additional solutions and use cases in the future.

#### Scenario 1: Data protection and migration within the on-premises environment using ACC

#### On-premises: self-managed OpenShift clusters and self-managed NetApp storage

- Using ACC, create Snapshot copies, backups and restores for data protection.
- Using ACC, perform a SnapMirror replication of container applications.



Scenario 2: Data protection and migration from the on-premises environment to AWS environment using ACC

## On-premises: Self-managed OpenShift cluster and self-managed storage AWS Cloud: Self-managed OpenShift cluster and self-managed storage

- Using ACC, perform backups and restores for data protection.
- Using ACC, perform a SnapMirror replication of container applications.

#### Scenario 2



Scenario 3: Data protection and migration from the on-premises environment to AWS environment

#### On-premises: Self-managed OpenShift cluster and self-managed storage AWS Cloud: Provider-managed OpenShift cluster (ROSA) and provider-managed storage (FSxN)

- Using BlueXP, perform replication of persistent volumes (FSxN).
- Using OpenShift GitOps, recreate application metadata.

Scenario 3



## Scenario 4: Data protection and migration from the on-premises environment to GCP environment using ACC

#### On-premises: Self-managed OpenShift cluster and self-managed storage Google Cloud: Self-managed OpenShift cluster and self-managed storage

- Using ACC, perform backups and restores for data protection.
- Using ACC, perform a SnapMirror replication of container applications.



For considerations when using ONTAP in a MetroCluster configuration, refer here.

### Scenario 5: Data protection and migration from the on-premises environment to Azure environment using ACC

On-premises: Self-managed OpenShift cluster and self-managed storage Azure Cloud: Self-managed OpenShift cluster and self-managed storage

- Using ACC, perform backups and restores for data protection.
- Using ACC, perform a SnapMirror replication of container applications.

			Virtual Network
vSphere Cluster Zone A	vSphere Cluster Zone B	vSphere Cluster Zone C	Master nodes subnet OpenShift cluster Master nodes
OpenShift cluster 1	OpenShift cluster 1	OpenShift cluster Advanced Cluster Manager	Worker nodes subnet
Trident CSI	Trident CSI	Astra Control Center	Trident CSI
MetroCl	uster	SnapMin	Cloud Volumes ONTAP

For considerations when using ONTAP in a MetroCluster configuration, refer here.

#### Versions of various components used in the solution validation

The solution tests and validates Migration & Centralized Data Protection with OpenShift container platform, OpenShift Advanced Cluster Manager, NetApp ONTAP, and NetApp Astra Control Center.

Scenarios 1, 2 and 3 of the solution were validated using the versions as shown in the table below:

Component	Version	
VMware	vSphere Client version 8.0.0.10200 VMware ESXi, 8.0.0, 20842819	
Hub Cluster	OpenShift 4.11.34	
Source and Destination Clusters	OpenShift 4.12.9 on-premises and in AWS	
NetApp Astra Trident	Trident Server and Client 23.04.0	
NetApp Astra Control Center	ACC 22.11.0-82	
NetApp ONTAP	ONTAP 9.12.1	
AWS FSx for NetApp ONTAP	Single AZ	

Scenario 4 of the solution was validated using the versions as shown in the table below:

Component	Version	
VMware	vSphere Client version 8.0.2.00000 VMware ESXi, 8.0.2, 22380479	
Hub Cluster	OpenShift 4.13.13	
Source and Destination Clusters	OpenShift 4.13.12 on-premises and in Google Cloud	
NetApp Astra Trident	Trident Server and Client 23.07.0	
NetApp Astra Control Center	ACC 23.07.0-25	
NetApp ONTAP	ONTAP 9.12.1	
Cloud Volumes ONTAP	Single AZ, Single node,9.14.0	

Scenario 5 of the solution was validated using the versions as shown in the table below:

Component	Version
VMware	vSphere Client version 8.0.2.00000 VMware ESXi, 8.0.2, 22380479
Source and Destination Clusters	OpenShift 4.13.25 on-premises and in Azure
NetApp Astra Trident	Trident Server and Client and Astra Control Provisioner 23.10.0
NetApp Astra Control Center	ACC 23.10
NetApp ONTAP	ONTAP 9.12.1
Cloud Volumes ONTAP	Single AZ, Single node,9.14.0

#### Supported NetApp Storage integrations with Red Hat Open Shift Containers

Whether the Red Hat Open Shift containers are running on VMware or in the hyperscalers, NetApp Astra Trident can be used as the CSI provisioner for the various types of backend NetApp storage that it supports.

The following diagram depicts the various backend NetApp storage that can be integrated with OpenShift clusters using NetApp Astra Trident.


ONTAP Storage Virtual Machine (SVM) provides secure multi-tenancy. A Single OpenShift cluster can connect to single SVM or multiple SVMs or even to multiple ONTAP clusters. Storage class filters the backend storage based on parameters or by labels. Storage administrators define the parameters to connect to storage system using trident backend configuration. On successful connection establishment, it creates the trident backend and populates the information which the storage class can filter.

The relationship between the storageclass and backend is shown below.



Application owner requests persistent volume using storage class. The storage class filters the backend storage.

The relationship between the pod and backend storage is shown below.



### **Container Storage Interface (CSI) Options**

On vSphere environments, customers can pick VMware CSI driver and/or Astra Trident CSI to integrate with ONTAP. With VMware CSI, the persistent volumes are consumed as local SCSI disks, whereas with Trident, it is consumed with network.

As VMware CSI does not support RWX access modes with ONTAP, applications need to use Trident CSI if

RWX mode is required. With FC based deployments, VMware CSI is preferred and SnapMirror Business Continuity (SMBC) provides zone level high availability.

## VMware CSI supports

- Core Block based datastores (FC, FCoE, iSCSI, NVMeoF)
- Core File based datastores (NFS v3, v4)
- vVol datastores (block and file)

## Trident has following drivers to support ONTAP

- ontap-san (dedicated volume)
- ontap-san-economy (shared volume)
- ontap-nas (dedicated volume)
- ontap-nas-economy (shared volume)
- ontap-nas-flexgroup (dedicated large scale volume)

For both VMware CSI and Astra Trident CSI, ONTAP supports nconnect, session trunking, kerberos, etc. for NFS and multipathing, chap authentication, etc. for block protocols.

In AWS, FSx for NetApp ONTAP (FSxN) can be deployed in single Availability Zone (AZ) or in Multi AZ. For production workloads that requires high availability, multi-AZ provides zonal level fault tolerance and has better NVMe read cache compared to single AZ. For more info, check AWS performance guidelines. To save cost on disaster recovery site, single AZ FSx ONTAP can be utilized. image::rhhc\_storage\_options\_fsxn\_options.png["Replication between Multi-AZ and Single-AZ"]

For number of SVMs that are supported by FSx ONTAP, refer managing FSx ONTAP storage virtual machine

# NetApp Hybrid Multicloud solutions for Red Hat OpenShift Container workloads

NetApp is seeing a significant increase in customers modernizing their legacy enterprise applications and building new applications using containers and orchestration platforms built around Kubernetes. Red Hat OpenShift Container Platform is one example that we see adopted by many of our customers.

# Overview

As more and more customers begin adopting containers within their enterprises, NetApp is perfectly positioned to help serve the persistent storage needs of their stateful applications and classic data management needs such as data protection, data security, and data migration. However, these needs are met using different strategies, tools, and methods.

**NetApp ONTAP** based storage options listed below, deliver security, data protection, reliability, and flexibility for containers and Kubernetes deployments.

- Self-managed storage in on-premises:
  - NetApp Fabric Attached Storage (FAS), NetApp All Flash FAS Arrays (AFF), NetApp All SAN Array (ASA) and ONTAP Select
- Provider-managed storage in on-premises:

- NetApp Keystone provides Storage as a Service (STaaS)
- Self-managed storage in the cloud:
  - NetApp Cloud Volumes ONTAP(CVO) provide self managed storage in the hyperscalers
- Provider-managed storage in the cloud:
  - Cloud Volumes Service for Google Cloud (CVS), Azure NetApp Files (ANF), Amazon FSx for NetApp ONTAP offer fully managed storage in the hyperscalers

9

Storage Administration	Performance & Scalability
<ul> <li>Multi-tenancy</li> <li>FlexVol &amp; FlexGroup</li> <li>LUN</li> <li>Quotas</li> <li>ONTAP CLI &amp; API</li> <li>System Manager &amp; BlueXP</li> </ul>	FlexCache     Inconnect, session trunking, multipathing     FlexClone     Scale-out clusters
Availability & Resilience	Access Protocols
Multi-AZ HA deployment     (MetroCluster)     SnapShot & SnapRestore     SnapMirror Cloud     SnapMirror	NFSv3, v4, v4.1, v4.2     iSCSI     SMB - v2, v3     Multi-protocol access
Storage Efficiency	Security & Compliance
<ul> <li>Deduplication &amp; Compression</li> <li>Compaction</li> <li>Data Tiering (Fabric Pool)</li> </ul>	Fpolicy & Vscan     LDAP & Kerberos     Active Directory integration     Certificate based authentication

**NetApp BlueXP** enables you to manage all of your storage and data assets from a single control plane/interface.

You can use BlueXP to create and administer cloud storage (for example, Cloud Volumes ONTAP and Azure NetApp Files), to move, protect, and analyze data, and to control many on-prem and edge storage devices.

**NetApp Astra Trident** is a CSI Compliant Storage Orchestrator that enable quick and easy consumption of persistent storage backed by a variety of the above-mentioned NetApp storage options. It is an open-source software maintained and supported by NetApp.

# **ONTAP** feature highlights

# Astra Trident CSI feature highlights



Business critical container workloads need more than just persistent volumes. Their data management requirements require protection and migration of the application kubernetes objects as well.

Application data includes kubernetes objects in addition to the user data: Some examples are as follows:

- kubernetes objects such as pods specs, PVCs, deployments, services
- custom config objects such as config maps and secrets
- persistent data such as Snapshot copies, backups, clones
- custom resources such as CRs and CRDs

÷

**NetApp Astra Control**, available as both fully-managed and self-managed software, provides orchestration for robust application data management. Refer to the Astra documentation for additional details on the Astra family of products.

This reference documentation provides validation of migration and protection of container-based applications, deployed on RedHat OpenShift container platform, using NetApp Astra Control Center. In addition, the solution provides high-level details for the deployment and the use of Red Hat Advanced Cluster Management (ACM) for managing the container platforms. The document also highlights the details for the integration of NetApp storage with Red Hat OpenShift container platforms using Astra Trident CSI provisioner. Astra Control Center is deployed on the hub cluster and is used to manage the container applications and their persistent storage lifecycle. Finally, it provides a solution for replication and failover and fail-back for container workloads on managed Red Hat OpenShift clusters in AWS (ROSA) using Amazon FSx for NetApp ONTAP (FSxN) as persistent storage.

### NetApp Solution with Red Hat OpenShift Container platform workloads on VMware

If customers have a need to run their modern containerized applications on infrastructure in their private data centers, they can do so. They should plan and deploy the Red Hat OpenShift container platform (OCP) for a successful production-ready environment for deploying their container workloads. Their OCP clusters can be deployed on VMware or bare metal.

NetApp ONTAP storage delivers data protection, reliability, and flexibility for container deployments. Astra Trident serves as the dynamic storage provisioner to consume persistent ONTAP storage for customers' stateful applications. Astra Control Center can be used to orchestrate the many data management requirements of stateful applications such as data protection, migration, and business continuity.

With VMware vSphere, NetApp ONTAP tools provides a vCenter Plugin which can be utilized to provision datastores. Apply tags and use it with OpenShift for storing the node configuration and data. NVMe based storage provides lower latency and high performance.

This solution provides details for data protection and migration of container workloads using Astra Control Center. For this solution, the container workloads are deployed on Red Hat OpenShift clusters on vSphere within the on-premises environment.

NOTE: We will provide a solution for container workloads on OpenShift clusters on bare metal in the future.

## Data protection and migration solution for OpenShift Container workloads using Astra Control Center



## Deploy and configure the Red Hat OpenShift Container platform on VMware

This section describes a high-level workflow of how to set up and manage OpenShift clusters and manage stateful applications on them. It shows the use of NetApp ONTAP storage arrays with the help of Astra Trident to provide persistent volumes. Details are provided about the use of Astra Control Center to perform data protection and migration activities for the stateful applications.



There are several ways of deploying Red Hat OpenShift Container platform clusters. This highlevel description of the setup provides documentation links for the specific method that was used. You can refer to the other methods in the relevant links provided in the resources section.

Here is a diagram that depicts the clusters deployed on VMware in a data center.

	Region: Datacenter 1		Region: Datacenter 2
vSphere Cluster Zone A	vSphere Cluster Zone B	vSphere Cluster Zone C	vSphere Cluster Zone A
OpenShift cluster 1	OpenShift cluster 1	Advanced Cluster Manager	OpenShift cluster 2
Trident CSI	Trident CSI	Astra Control Center	Trident CSI
		Sn Sn	apMirror

The setup process can be broken down into the following steps:

## Deploy and configure a CentOS VM

( 🖓 )

- It is deployed in the VMware vSphere environment.
- This VM is used for deploying some components such as NetApp Astra Trident and NetApp Astra Control Center for the solution.
- A root user is configured on this VM during installation.

## Deploy and configure an OpenShift Container Platform cluster on VMware vSphere (Hub Cluster)

Refer to the instructions for the Assisted deployment method to deploy an OCP cluster.

Remember the following:

- Create ssh public and private key to provide to the installer. These keys will be used to login to the master and worker nodes if needed.

- Download the installer program from the assisted installer. This program is used to boot the VMs that you create in the VMware vSphere environment for the master and worker nodes.

- VMs should have the minimum CPU, memory, and hard disk requirement. (Refer to the vm create commands on this page for the master and the worker nodes which provide this information)

- The diskUUID should be enabled on all VMs.
- Create a minimum of 3 nodes for master and 3 nodes for worker.

- Once they are discovered by the installer, turn on the VMware vSphere integration toggle button.

This is installed using the Advanced Cluster Management Operator on the Hub Cluster. Refer to the instructions here.

### Install an internal Red Hat Quay registry on the Hub Cluster.

- An internal registry is required to push the Astra image. A Quay internal registry is installed using the Operator in the Hub cluster.
- Refer to the instructions here

### Install two additional OCP clusters (Source and Destination)

- The additional clusters can be deployed using the ACM on the Hub Cluster.
- Refer to the instructions here.

### Configure NetApp ONTAP storage

- Install an ONTAP cluster with connectivity to the OCP VMs in VMWare environment.
- · Create an SVM.
- Configure NAS data lif to access the storage in SVM.

### Install NetApp Trident on the OCP clusters

- Install NetApp Trident on all three clusters: Hub, source, and destination clusters
- Refer to the instructions here.
- Create a storage backend for ontap-nas .
- Create a storage class for ontap-nas.
- Refer to instructions here.

### Install NetApp Astra Control Center

- NetApp Astra Control Center is installed using the Astra Operator on the Hub Cluster.
- Refer to the instructions here.

#### Points to remember:

- \* Download NetApp Astra Control Center image from the support site.
- \* Push the image to an internal registry.
- \* Refer to instructions here.

### **Deploy an Application on Source Cluster**

Use OpenShift GitOps to deploy an application. (eg. Postgres, Ghost)

After you add a cluster to Astra Control management, you can install apps on the cluster (outside of Astra Control) and then go to the Applications page in Astra Control to define the apps and their resources. Refer to Start managing apps section of Astra Control Center.

The next step is to use the Astra Control Center for Data protection and Data migration from the source to the destination cluster.

# Data protection using Astra

This page shows the data protection options for Red Hat OpenShift Container based applications running on VMware vSphere using Astra Control Center (ACC).

As users take their journey of modernizing their applications with Red Hat OpenShift, a data protection strategy should be in place to protect them from accidental deletion or any other human errors. Often a protection strategy is also required for regulatory or compliance purposes to protect their data from a diaster.

The requirements of data protection varies from reverting back to a point in time copy to automatically failing over to a different fault domain without any human intervention. Many customers pick ONTAP as their preferred storage platform for their Kubernetes applications because of its rich features like multitenancy, multi-protocol, high performance and capacity offerings, replication and caching for multi-site locations, security and flexibility.

Data protection in ONTAP can be achieved using ad-hoc or policy controlled

- Snapshot
- backup and restore

Both Snapshot copies and backups protect the following types of data:

- The application metadata that represents the state of the application
- Any persistent data volumes associated with the application
- Any resource artifacts belonging to the application

## Snapshot with ACC

A point in time copy of data can be captured using Snapshot with ACC. Protection policy defines the number of Snapshot copies to keep. Minimum schedule option available is hourly. Manual, on-demand Snapshot copies can be taken at any time and at shorter intervals than scheduled Snapshot copies. Snapshot copies are stored on the same provisioned volume as the app.

## **Configuring Snapshot with ACC**

Ams						
( furthease )	abon -				e	Adres
National Contract	de Arrestation starsk © memy		e shumot	9 Anti-Labor Participan	<ul> <li>System print printment</li> </ul>	<u>.</u>
- Over Instance	And the second s	0				
Automatic State	Cara primetras Incorpo Permitina - Tamotina I	unit Arrige Sam				
	Arrest 4 Antipart 1			W. mark	O manual B tarts	an 2 Autoine
					141	Address of the local
	30 MH	1000	Terral state	(In tribulate) Co-Daringed	Classed #	Reference.
	optime child with the	122	3 million	(D) the Schemater	metabacia tati org	4
	gine sugerier michaeler	( Chinana	10 million	() in terms	100 a (mg (4) 10 an 171)	10
	C and manual strandomstate	2	12-1471	0	2012/06/10 1414 171	÷.
	glass insights the particular	-12	0.000	O do Liman	1013000 Pt 1021-172	10

#### **Backup and Restore with ACC**

A backup is based on a Snapshot. ACC can take Snapshot copies using CSI and perform backup using the point in time Snapshot copy. The backup is stored in an external object store (any s3 compatible including ONTAP S3 at a different location). Protection policy can be configured for scheduled backups and the number of backup versions to keep. The minimum RPO is one hour.

### Restoring an application from a backup using ACC

ACC restores application from the S3 bucket where the backups are store.

🗳 4me								
() hardined	e) gbort -						e ann	*
O hannatara O thema	-dj. 200	n in actione scarses (in monthly		B) this second	9 errication restitution G toosing price uniquest	C behavior print a	() ()	
Out at an a	helinen Autoren		0 mmmt					
C Later	Data posteriore desage d	formation formation formation (A)/10	ey tain					
3 met	Alter * Contare *				T and .	O Institut	B tongs of Autom	1
i sessi	- 100	Nie (	1000 April 1	(in Scholars / On Jaman)	Rodar	count + 1	Automatica	in a
30. page 7.	and sale of a	8 mm	1 T	(i) co-canacan	***	appear to or	Andread applications Andread Sandrage	ì

#### Application specific execution hooks

In addition, execution hooks can be configured to run in conjunction with a data protection operation of a managed app. Even though storage array level data protection features are available, often additional steps are needed to make backups and restores, application consistent. The app-specific additional steps could be: - before or after a Snapshot copy is created.

- before or after a backup is created.
- after restoring from a Snapshot copy or backup.

Astra Control can execute these app-specific steps coded as custom scripts called execution hooks.

NetApp Verda GitHub project provides execution hooks for popular cloud-native applications to make protecting applications straightforward, robust, and easy to orchestrate. Feel free to contribute to that project if you have enough information for an application that is not in the repository.

Sample execution hook for pre-Snapshot of a redis application.

OOK DETAILS ?				EXECUTION HOOKS
Operation Pre-snapshot		Hook arguments (optional) 1 pre × Enter hook arguments	2	Execution hooks allow Astra Control to execute your own custom scripts before or after snapshot.
look name edis-pre-snapshot				Read more in Manage application execution hooks
NTAINER IMAGES				
Apply to all container images				
Use a regular expression to target container images t	or the hook.			
Container image names to match edis				
RIPT 7				
+ Add			\Xi Search	
Name 4				
mariadb_mysqLsh				
postgresqLsh				
redis_hook.sh				

#### **Replication with ACC**

For regional protection or for a low RPO and RTO solution, an application can be replicated to another Kubernetes instance running at a different site, preferably in another region. ACC utilizes ONTAP async SnapMirror with RPO as low as 5 minutes. Replication is done by replicating to ONTAP and then a fail over creates the Kubernetes resources in the destination cluster.



Note that replication is different from the backup and restore where the backup goes to S3 and restore is performed from S3. Refer xref:./rhhc/ here to get additional details about the differences between the two types of data protection.

Refer here for SnapMirror setup instructions.

#### **SnapMirror with ACC**

	⊘ ghost ∧			c	Adum
Applications Clusters Cloud Instances	-& APPERCATION STATUS @ History	<ul> <li>Sity protected</li> </ul>	APPLICATION PROTECTION     Protection polycy configured	Bygication policy configured	Snipshot Back op Cone Nexton
tackends	ghad some	O up-later)			fall-over Reverse replice
lucken	Data perturbition Storage Resources Execution books Activity	fasies			transpr
Account.	Configure 1			C Stagehold D Sach	ign 2 August
elwiy uport	Source Shorts	Contraction S S S S S S S S S S S S S		Replication relationship	her to

san-economy and nas-economy storage drivers do not support replication feature. Refer here for additional details.

### Demo video:

( i

a ....

Demonstration video of disaster recovery with Astra Control Center

#### Data protection with Astra Control Center

#### Business Continuity with MetroCluster

Most of our hardware platform for ONTAP has high availability features to protect from device failures avoiding the need to perform diaster recovery. But to protect from fire or any other disaster and to continue the business with zero RPO and low RTO, often a MetroCluster solution is used.

Customers who currently have an ONTAP system can extend to MetroCluster by adding supported ONTAP systems within the distance limitations for providing zone level disaster recovery. Astra Trident, the CSI (Container Storage Interface) supports NetApp ONTAP including MetroCluster configuration as well as other options like Cloud Volumes ONTAP, Azure NetApp Files, AWS FSx for NetApp ONTAP, etc. Astra Trident provides five storage driver options for ONTAP and all are supported for MetroCluster configuration. Refer here for additional details about ONTAP storage drivers supported by Astra Trident.

The MetroCluster solution requires layer 2 network extension or capability to access the same network address from both fault domains. Once MetroCluster configuration is in place, the solution is transparent to application owners as all the volumes in the MetroCluster svm are protected and get the benefits of SyncMirror (zero RPO).



For Trident Backend Configuration (TBC), do not specify the dataLIF and SVM when using MetroCluster configuration. Specify SVM management IP for managementLIF and use vsadmin role credentials.

Details on Astra Control Center Data Protection features are available here

## Data migration using Astra Control Center

This page shows the data migration options for container workloads on Red Hat OpenShift clusters with Astra Control Center (ACC).

Kubernetes Applications are often required to be moved from one environment to another. To migrate an application along with its persistent data, NetApp ACC can be utilized.

## Data Migration between different Kubernetes environment

ACC supports various Kubernetes flavors including Google Anthos, Red Hat OpenShift, Tanzu Kubernetes Grid, Rancher Kubernetes Engine, Upstream Kubernetes, etc. For additional details, refer here.

To migrate application from one cluster to another, you can use one of the following features of ACC:

- replication
- backup and restore
- clone

**()** 

Refer to the data protection section for the replication and backup and restore options.

Refer here for additional details about cloning.



Astra Replication feature is only supported with Trident Container Storage Interface (CSI). However, replication is not supported by nas-economy & san-economy drivers.

## Performing data replication using ACC

	🔞 ghost 🔿		C. Anto	ni i
i Applications Clusters	-APPLICATION STATUS	S APPLIC Distrymolected Distance (or	Atton PROTECTION Regulation poles configured Com	to to
	Definition <b>b</b> ghost zonez	Conv Organization	falle Konst	ner ner
) hicken	Data penterbien Storage Resources Execution hooks Archity	Tanks	Loss.	raje
	Configure +		🖸 Snapshots 😫 Backupa	d Austral
, recommend	Source ghot © Headity Part server	Centruiden Ghatt Sientty Gigenstate? Time Integrated	Replication relationship STATUS (1) prainting   totalished SCREDUX Replicate screptulat every 5 descuber to (2) order-charter? EAST SYNC Synt disarbox 30 secures	

# NetApp Hybrid Multicloud solutions for Red Hat OpenShift Container workloads

NetApp is seeing a significant increase in customers modernizing their legacy enterprise applications and building new applications using containers and orchestration platforms built around Kubernetes. Red Hat OpenShift Container Platform is one example that we see adopted by many of our customers.

## Overview

As more and more customers begin adopting containers within their enterprises, NetApp is perfectly positioned to help serve the persistent storage needs of their stateful applications and classic data management needs such as data protection, data security, and data migration. However, these needs are met using different strategies, tools, and methods.

**NetApp ONTAP** based storage options listed below, deliver security, data protection, reliability, and flexibility for containers and Kubernetes deployments.

- Self-managed storage in on-premises:
  - NetApp Fabric Attached Storage (FAS), NetApp All Flash FAS Arrays (AFF), NetApp All SAN Array (ASA) and ONTAP Select
- Provider-managed storage in on-premises:

- NetApp Keystone provides Storage as a Service (STaaS)
- Self-managed storage in the cloud:
  - NetApp Cloud Volumes ONTAP(CVO) provide self managed storage in the hyperscalers
- Provider-managed storage in the cloud:
  - Cloud Volumes Service for Google Cloud (CVS), Azure NetApp Files (ANF), Amazon FSx for NetApp ONTAP offer fully managed storage in the hyperscalers

Storage Administration	Performance & Scalability
<ul> <li>Multi-tenancy</li> <li>FlexVol &amp; FlexGroup</li> <li>LUN</li> <li>Quotas</li> <li>ONTAP CLI &amp; API</li> <li>System Manager &amp; BlueXP</li> </ul>	FlexCache     rconnect, session trunking, multipathing     FlexClone     Scale-out clusters
Availability & Resilience	Access Protocols
<ul> <li>Multi-AZ HA deployment (MetroCluster)</li> <li>SnapShot &amp; SnapRestore</li> <li>SnapMirror Cloud</li> <li>SnapMirror</li> </ul>	NFS -v3, v4, v4.1, v4.2     iSCSI     SMB - v2, v3     Multi-protocol access
Storage Efficiency	Security & Compliance
<ul> <li>Deduplication &amp; Compression</li> <li>Compaction</li> <li>Compaction</li> <li>Data Tiering (Fabric Pool)</li> </ul>	Fpolicy & Vscan     LDAP & Kerberos     Active Directory integration     Certificate based authentication

**NetApp BlueXP** enables you to manage all of your storage and data assets from a single control plane/interface.

You can use BlueXP to create and administer cloud storage (for example, Cloud Volumes ONTAP and Azure NetApp Files), to move, protect, and analyze data, and to control many on-prem and edge storage devices.

**NetApp Astra Trident** is a CSI Compliant Storage Orchestrator that enable quick and easy consumption of persistent storage backed by a variety of the above-mentioned NetApp storage options. It is an open-source software maintained and supported by NetApp.

# **ONTAP** feature highlights

9

# Astra Trident CSI feature highlights



Business critical container workloads need more than just persistent volumes. Their data management requirements require protection and migration of the application kubernetes objects as well.

Application data includes kubernetes objects in addition to the user data: Some examples are as follows:

- kubernetes objects such as pods specs, PVCs, deployments, services
- custom config objects such as config maps and secrets
- persistent data such as Snapshot copies, backups, clones
- custom resources such as CRs and CRDs

**NetApp Astra Control**, available as both fully-managed and self-managed software, provides orchestration for robust application data management. Refer to the Astra documentation for additional details on the Astra family of products.

This reference documentation provides validation of migration and protection of container-based applications, deployed on RedHat OpenShift container platform, using NetApp Astra Control Center. In addition, the solution provides high-level details for the deployment and the use of Red Hat Advanced Cluster Management (ACM) for managing the container platforms. The document also highlights the details for the integration of NetApp storage with Red Hat OpenShift container platforms using Astra Trident CSI provisioner. Astra Control Center is deployed on the hub cluster and is used to manage the container applications and their persistent storage lifecycle. Finally, it provides a solution for replication and failover and fail-back for container workloads on managed Red Hat OpenShift clusters in AWS (ROSA) using Amazon FSx for NetApp ONTAP (FSxN) as persistent storage.

## NetApp Solution with Red Hat OpenShift Container platform workloads in Hybrid Cloud

Customers may be at a point in their modernization journey when they are ready to move some select workloads or all workloads from their data centers to the cloud. They may choose to use self-managed OpenShift containers and self-managed NetApp storage in the cloud for various reasons. They should plan and deploy the Red Hat OpenShift container platform (OCP) in the cloud for a successful production-ready environment for

÷

migrating their container workloads from their data centers. Their OCP clusters can be deployed on VMware or Bare Metal in their data centers and on AWS, Azure or Google Cloud in the cloud environment.

NetApp Cloud Volumes ONTAP storage delivers data protection, reliability, and flexibility for container deployments in AWS, Azure and in Google Cloud. Astra Trident serves as the dynamic storage provisioner to consume the persistent Cloud Volumes ONTAP storage for customers' stateful applications. Astra Control Center can be used to orchestrate the many data management requirements of stateful applications such as data protection, migration, and business continuity.

## Data protection and migration solution for OpenShift Container workloads in a hybrid cloud using Astra Control Center

On-premises and AWS image::rhhc-self-managed-aws.png[]

On-premises and Google Cloud image::rhhc-self-managed-gcp.png[]

On-premises and Azure Cloud image::rhhc-self-managed-azure.png[]

# Deploy and configure the Red Hat OpenShift Container platform on AWS

This section describes a high-level workflow of how to set up and manage OpenShift Clusters in AWS and deploy stateful applications on them. It shows the use of NetApp Cloud Volumes ONTAP storage with the help of Astra Trident to provide persistent volumes. Details are provided about the use of Astra Control Center to perform data protection and migration activities for the stateful applications.



There are several ways of deploying Red Hat OpenShift Container platform clusters on AWS. This high-level description of the setup provides documentation links for the specific method that was used. You can refer to the other methods in the relevant links provided in the resources section.

Here is a diagram that depicts the clusters deployed on AWS and connected to the data center using a VPN.



The setup process can be broken down into the following steps:

## Install an OCP cluster on AWS from the Advanced Cluster Management.

- Create a VPC with a site-to-site VPN connection (using pfsense) to connect to the on-premises network.
- · On-premises network has internet connectivity.
- Create 3 private subnets in 3 different AZs.
- Create a Route 53 private hosted zone and a DNS resolver for the VPC.

Create OpenShift Cluster on AWS from the Advanced Cluster Management (ACM) Wizard. Refer to instructions here.



You can also create the cluster in AWS from the OpenShift Hybrid Cloud console. Refer here for instructions.



When creating the cluster using the ACM, you have the ability to customize the installation by editing the yaml file after filling in the details in the form view. After the cluster is created, you can ssh login to the nodes of the cluster for troubleshooting or additional manual configuration. Use the ssh key you provided during installation and the username core to login.

- Install the connector in on-premises VMware environment. Refer to instructions here.
- Deploy a CVO instance in AWS using the connector. Refer to instructions here.



The connector can also be installed in the cloud environment. Refer here for additional information.

### Install Astra Trident in the OCP Cluster

- Deploy Trident Operator using Helm. Refer to instructions here
- Create a backend and a storage class. Refer to instructions here.

### Add the OCP cluster on AWS to the Astra Control Center.

Add the OCP cluster in AWS to Astra Control Center.

#### Using CSI Topology feature of Trident for multi-zone architectures

Cloud providers, today, enable Kubernetes/OpenShift cluster administrators to spawn nodes of the clusters that are zone based. Nodes can be located in different availability zones within a region, or across various regions. To facilitate the provisioning of volumes for workloads in a multi-zone architecture, Astra Trident uses CSI Topology. Using the CSI Topology feature, access to volumes can be limited to a subset of nodes, based on regions and availability zones. Refer here for additional details.

Kubernetes supports two volume binding modes:

- When **VolumeBindingMode is set to Immediate** (default), Astra Trident creates the volume without any topology awareness. Persistent Volumes are created without having any dependency on the requesting pod's scheduling requirements.

- When **VolumeBindingMode set to WaitForFirstConsumer**, the creation and binding of a Persistent Volume for a PVC is delayed until a pod that uses the PVC is scheduled and created. This way, volumes are created to meet the scheduling constraints that are enforced by topology requirements.

Astra Trident storage backends can be designed to selectively provision volumes based on availability zones (Topology-aware backend). For StorageClasses that make use of such a backend, a volume would only be created if requested by an application that is scheduled in a supported region/zone. (Topology-aware StorageClass) Refer here for additional details.

#### Deploy and configure the Red Hat OpenShift Container platform on GCP

This section describes a high-level workflow of how to set up and manage OpenShift Clusters in GCP and deploy stateful applications on them. It shows the use of NetApp Cloud Volumes ONTAP storage with the help of Astra Trident to provide persistent volumes. Details are provided about the use of Astra Control Center to perform data protection and migration activities for the stateful applications.

Here is a diagram that shows the clusters deployed on GCP and connected to the data center using a VPN.



There are several ways of deploying Red Hat OpenShift Container platform clusters in GCP. This high-level description of the setup provides documentation links for the specific method that was used. You can refer to the other methods in the relevant links provided in the resources section.

The setup process can be broken down into the following steps:

(i)

- Ensure that you have met all the prerequisites stated here.
- For the VPN connectivity between on-premises and GCP, a pfsense VM was created and configured. For instructions, see here.
  - The remote gateway address in pfsense can be configured only after you have created a VPN gateway in Google Cloud Platform.
  - The remote network IP addresses for the Phase 2 can be configured only after the OpenShift cluster installation program runs and creates the infrastructure components for the cluster.
  - The VPN in Google Cloud can only be configured after the infrastructure components for the cluster are created by the installation program.
- Now install the OpenShift cluster on GCP.
  - Obtain the installation program and the pull secret and deploy the cluster following the steps provided in the documentation here.
  - The installation creates a VPC network in Google Cloud Platform. It also creates a private zone in Cloud DNS and adds A records.
    - Use the CIDR block address of the VPC network to configure the pfsense and establish the VPN connection. Ensure firewalls are setup correctly.
    - Add A records in the DNS of the on-premises environment using the IP address in the A records of the Google Cloud DNS.
  - The installation of the cluster completes and will provide a kubeconfig file and username and password to login to the console of the cluster.

### Deploy Cloud Volumes ONTAP in GCP using BlueXP.

- Install a connector in Google Cloud. Refer to instructions here.
- Deploy a CVO instance in Google Cloud using the connector. Refer to instructions here. https://docs.netapp.com/us-en/bluexp-cloud-volumes-ontap/task-getting-started-gcp.html

### Install Astra Trident in the OCP Cluster in GCP

- There are many methods to deploy Astra Trident as shown here.
- For this project, Astra Trident was installed by deploying Astra Trident Operator manually using the instructions here.
- Create backend and a storage classes. Refer to instructions here.

#### Add the OCP cluster on GCP to the Astra Control Center.

- Create a separate KubeConfig file with a cluster role that contains the minimum permissions necessary for a cluster to be managed by Astra Control. The instructions can be found here.
- Add the cluster to Astra Control Center following the instructions here

## Using CSI Topology feature of Trident for multi-zone architectures

Cloud providers, today, enable Kubernetes/OpenShift cluster administrators to spawn nodes of the clusters that are zone based. Nodes can be located in different availability zones within a region, or across various regions. To facilitate the provisioning of volumes for workloads in a multi-zone architecture, Astra Trident uses CSI Topology. Using the CSI Topology feature, access to volumes can be limited to a subset of nodes, based on regions and availability zones. Refer here for additional details.

Kubernetes supports two volume binding modes:

- When **VolumeBindingMode is set to Immediate** (default), Astra Trident creates the volume without any topology awareness. Persistent Volumes are created without having any dependency on the requesting pod's scheduling requirements.

- When **VolumeBindingMode set to WaitForFirstConsumer**, the creation and binding of a Persistent Volume for a PVC is delayed until a pod that uses the PVC is scheduled and created. This way, volumes are created to meet the scheduling constraints that are enforced by topology requirements.

Astra Trident storage backends can be designed to selectively provision volumes based on availability zones (Topology-aware backend). For StorageClasses that make use of such a backend, a volume would only be created if requested by an application that is scheduled in a supported region/zone. (Topology-aware StorageClass) Refer here for additional details.

## **Demonstration Video**

OpenShift Cluster installation on Google Cloud Platform

## Importing OpenShift clusters into Astra Control Center

### Deploy and configure the Red Hat OpenShift Container platform on Azure

This section describes a high-level workflow of how to set up and manage OpenShift Clusters in Azure and deploy stateful applications on them. It shows the use of NetApp Cloud Volumes ONTAP storage with the help of Astra Trident/Astra Control Provisioner to provide persistent volumes. Details are provided about the use of Astra Control Center to perform data protection and migration activities for the stateful applications.

Here is a diagram that shows the clusters deployed on Azure and connected to the data center using a VPN.

l	kegion: Datacenter 1		Virtual Network
vSphere Cluster Zone A	vSphere Cluster Zone B	vSphere Cluster Zone C	Master nodes subnet OpenShift cluster Master nodes
OpenShift cluster 1	OpenShift cluster 1	OpenShift cluster Advanced Cluster Manager	Worker nodes subnet
Trident CSI	Trident CSI	Astra Control Center	Trident CSI 🔞
		SnapMirro	Cloud Volumes ONTAP

There are several ways of deploying Red Hat OpenShift Container platform clusters in Azure. This high-level description of the setup provides documentation links for the specific method that was used. You can refer to the other methods in the relevant links provided in the resources section.

The setup process can be broken down into the following steps:

(i)

- Ensure that you have met all the prerequisites stated here.
- Create a VPN, subnets and network security groups and a private DNS zone. Create VPN gateway and site-to-site VPN Connection.
- For the VPN connectivity between on-premises and Azure, a pfsense VM was created and configured. For instructions, see here.
- Obtain the installation program and the pull secret and deploy the cluster following the steps provided in the documentation here.
- The installation of the cluster completes and will provide a kubeconfig file and username and password to login to the console of the cluster.

A sample install-config.yaml file is given below.

```
apiVersion: v1
baseDomain: sddc.netapp.com
compute:
- architecture: amd64
  hyperthreading: Enabled
  name: worker
  platform:
    azure:
      encryptionAtHost: false
      osDisk:
        diskSizeGB: 512
        diskType: "StandardSSD LRS"
      type: Standard D2s v3
      ultraSSDCapability: Disabled
      #zones:
      #- "1"
      #- "2"
      #- "3"
  replicas: 3
controlPlane:
  architecture: amd64
  hyperthreading: Enabled
  name: master
  platform:
    azure:
      encryptionAtHost: false
      osDisk:
        diskSizeGB: 1024
        diskType: Premium LRS
      type: Standard D8s v3
      ultraSSDCapability: Disabled
  replicas: 3
```

```
metadata:
  creationTimestamp: null
  name: azure-cluster
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0/16
  networkType: OVNKubernetes
  serviceNetwork:
  - 172.30.0.0/16
platform:
  azure:
    baseDomainResourceGroupName: ocp-base-domain-rg
    cloudName: AzurePublicCloud
    computeSubnet: ocp-subnet2
    controlPlaneSubnet: ocp-subnet1
    defaultMachinePlatform:
      osDisk:
        diskSizeGB: 1024
        diskType: "StandardSSD LRS"
      ultraSSDCapability: Disabled
    networkResourceGroupName: ocp-nc-us-rg
    #outboundType: UserDefinedRouting
    region: northcentralus
    resourceGroupName: ocp-cluster-ncusrg
    virtualNetwork: ocp vnet ncus
publish: Internal
pullSecret:
```

### Deploy Cloud Volumes ONTAP in Azure using BlueXP.

- Install a connector in in Azure. Refer to instructions here.
- Deploy a CVO instance in Azure using the connector. Refer to instructions link:https://docs.netapp.com/us-en/bluexp-cloud-volumes-ontap/task-getting-started-azure.html [here.]

### Install Astra Control Provisioner in the OCP Cluster in Azure

- For this project, Astra Control Provisioner (ACP) was installed on all the clusters (on-prem cluster, onprem cluster where Astra Control Center is deployed and the cluster in Azure). Learn more about the Astra Control Provisioner here.
- Create backend and a storage classes. Refer to instructions here.

### Add the OCP cluster on Azure to the Astra Control Center.

- Create a separate KubeConfig file with a cluster role that contains the minimum permissions necessary for a cluster to be managed by Astra Control. The instructions can be found here.
- Add the cluster to Astra Control Center following the instructions here

## Using CSI Topology feature of Trident for multi-zone architectures

Cloud providers, today, enable Kubernetes/OpenShift cluster administrators to spawn nodes of the clusters that are zone based. Nodes can be located in different availability zones within a region, or across various regions. To facilitate the provisioning of volumes for workloads in a multi-zone architecture, Astra Trident uses CSI Topology. Using the CSI Topology feature, access to volumes can be limited to a subset of nodes, based on regions and availability zones. Refer here for additional details.

Kubernetes supports two volume binding modes:

- When **VolumeBindingMode is set to Immediate** (default), Astra Trident creates the volume without any topology awareness. Persistent Volumes are created without having any dependency on the requesting pod's scheduling requirements.

- When **VolumeBindingMode set to WaitForFirstConsumer**, the creation and binding of a Persistent Volume for a PVC is delayed until a pod that uses the PVC is scheduled and created. This way, volumes are created to meet the scheduling constraints that are enforced by topology requirements.

Astra Trident storage backends can be designed to selectively provision volumes based on availability zones (Topology-aware backend). For StorageClasses that make use of such a backend, a volume would only be created if requested by an application that is scheduled in a supported region/zone. (Topology-aware StorageClass) Refer here for additional details.

## **Demonstration Video**

Using Astra Control for Failover and Failback of applications

## Data protection using Astra Control Center

This page shows the data protection options for Red Hat OpenShift Container based applications running on VMware vSphere or in the cloud using Astra Control Center (ACC).

As users take their journey of modernizing their applications with Red Hat OpenShift, a data protection strategy should be in place to protect them from accidental deletion or any other human errors. Often a protection strategy is also required for regulatory or compliance purposes to protect their data from a diaster.

The requirements of data protection varies from reverting back to a point in time copy to automatically failing over to a different fault domain without any human intervention. Many customers pick ONTAP as their preferred storage platform for their Kubernetes applications because of its rich features like multitenancy, multi-protocol, high performance and capacity offerings, replication and caching for multi-site locations, security and flexibility.

Customers may have a cloud environment setup as their data center extension, so that they can leverage the benefits of the cloud as well as be well positioned to move their workloads at a future time. For such customers, backing up of their OpenShift applications and their data to the cloud environment becomes an

inevitable choice. They can then restore the applications and the associated data either to an OpenShift cluster in the cloud or in their data center.

### Backup and Restore with ACC

Application owners can review and update the applications discovered by ACC. ACC can take Snapshot copies using CSI and perform backup using the point in time Snapshot copy. Backup destination can be an object store in the cloud environment. Protection policy can be configured for scheduled backups and the number of backup versions to keep. The minimum RPO is one hour.

## Restoring an application from a backup using ACC



### Application specific execution hooks

Even though storage array level data protection features are available, often additional steps are needed to make backups and restores application consistent. The app-specific additional steps could be:

- before or after a Snapshot copy is created.
- before or after a backup is created.
- after restoring from a Snapshot copy or backup.

Astra Control can execute these app-specific steps coded as custom scripts called execution hooks.

NetApp's open source project Verda provides execution hooks for popular cloud-native applications to make protecting applications straightforward, robust, and easy to orchestrate. Feel free to contribute to that project if you have enough information for an application that is not in the repository.

Sample execution hook for pre-Snapshot of a redis application.

OK DETAILS 2			EXECUTION HOOKS
re-snapshot	Hook arguments (optional) 1 pre × Enter hook arguments	1	Execution hooks allow Astra Control to execute your own custom scripts before or after snapshot.
ook name edis-pre-snapshot			Read more in Manage application executio hooks
NTAINER IMAGES ?			
Apply to all container images			
lse a regular expression to target container images for the hook.			
ontainer image names to match : •dis			
LIPT P			
+ Add		₹ Search	
Name 4			
mariadb_mysql.sh			
o postgresql.sh			
redis_hook.sh			

## **Replication with ACC**

For regional protection or for a low RPO and RTO solution, an application can be replicated to another Kubernetes instance running at a different site, preferably in another region. ACC utilizes ONTAP async SnapMirror with RPO as low as 5 minutes.

Refer here for SnapMirror setup instructions.

### SnapMirror with ACC

530

	() ghost o			c	Advert
Applications Clusters Cloud Instances	Ar APPLICATION STATUS	<ul> <li>Fully protocold</li> <li>Output</li> </ul>	APPLICATION PROTECTION     Protection polycy to influend	Seglication policy configured	Snigsbet Back op Clone Auston
tackends.	<ul> <li>ghod zonez</li> </ul>	O organization			fallover Noverse replicat
luckm	Duta protection Storage Resources Execution hooks Activity	Tanks			University
	Configure +			û Stanber 🔒 Back	ge 2 Gebo
utony apport	Source ghost @ republic phose public phose public phose public phose public phose public phose public Attended	I		Replication relationship STATUS Status Agencial subplicit form Agencial subplicit form Status Agencial Status S	net Ro

 $(\mathbf{i})$ 

san-economy and nas-economy storage drivers do not support replication feature. Refer here for additional details.

### Demo video:

Demonstration video of disaster recovery with Astra Control Center

Data protection with Astra Control Center

Details on Astra Control Center Data Protection features are available here

Disaster recovery (Failover and Failback using replication) with ACC

Using Astra Control for Failover and Failback of applications

### Data migration using Astra Control Center

This page shows the data migration options for container workloads on Red Hat OpenShift clusters with Astra Control Center (ACC). Specifically, customers can use ACC to

- move some selected workloads or all workloads from their on-premises data centers to the cloud

- clone their apps to the cloud either for testing purposes or move from the data center to the cloud

### **Data Migration**

To migrate application from one environment to another, you can use one of the following features of ACC:

replication

- backup and restore
- clone

Refer to the data protection section for the replication and backup and restore options.

Refer here for additional details about cloning.



Astra Replication feature is only supported with Trident Container Storage Interface (CSI). However, replication is not supported by nas-economy & san-economy drivers.

## Performing data replication using ACC

D. CHARGE	⊚ ghost ∧			C.	Advert
<ul> <li>Applications</li> <li>Clusters</li> <li>Cloud instances</li> </ul>	Ar Application status	<ul> <li>Fully protocold</li> </ul>	APPLICATION PROTECTION	Reglication policy configured	Snapsher Back op Come Nastore
à tuckente	Definition Carlor generation (Carlor Carlor	-tand			fall-per Reverse replicatio
] Buckett	Data pentaction Storage Resources Execution hooks Activity fasks				benange
, Account	Configure +			O Staphets D fach	ge Z Septor
. Activity Engineers	Source ghot @ isource in galanter in galan	E		Replication relationship STATUS D thatty: Latationed Stellbus Replicate suppliers every 5 mend C conclusions LAST SYNC Specification 20 selects	wi #0

# NetApp Hybrid Multicloud solutions for Red Hat OpenShift Container workloads

NetApp is seeing a significant increase in customers modernizing their legacy enterprise applications and building new applications using containers and orchestration platforms built around Kubernetes. Red Hat OpenShift Container Platform is one example that we see adopted by many of our customers.

## Overview

As more and more customers begin adopting containers within their enterprises, NetApp is perfectly positioned to help serve the persistent storage needs of their stateful applications and classic data management needs such as data protection, data security, and data migration. However, these needs are met using different strategies, tools, and methods.

**NetApp ONTAP** based storage options listed below, deliver security, data protection, reliability, and flexibility for containers and Kubernetes deployments.

• Self-managed storage in on-premises:

- NetApp Fabric Attached Storage (FAS), NetApp All Flash FAS Arrays (AFF), NetApp All SAN Array (ASA) and ONTAP Select
- Provider-managed storage in on-premises:
  - NetApp Keystone provides Storage as a Service (STaaS)
- Self-managed storage in the cloud:
  - · NetApp Cloud Volumes ONTAP(CVO) provide self managed storage in the hyperscalers
- Provider-managed storage in the cloud:
  - Cloud Volumes Service for Google Cloud (CVS), Azure NetApp Files (ANF), Amazon FSx for NetApp ONTAP offer fully managed storage in the hyperscalers

Storage Administration	Performance & Scalability				
<ul> <li>Multi-tenancy</li> <li>ONTAP CLI &amp; API</li> <li>FlexVol &amp; FlexGroup</li> <li>System Manager &amp; BlueXP</li> <li>LUN</li> <li>Quotas</li> </ul>	FlexCache     flexCache     roconnect, session trunking, multipathing     FlexClone     Scale-out clusters				
Availability & Resilience	Access Protocols				
Multi-AZ HA deployment (MetroCluster)     SnapShot & SnapRestore     SnapMirror Cloud     SnapMirror	NFS -v3, v4, v4.1, v4.2     iSCSI     SMB - v2, v3     Multi-protocol access				
Storage Efficiency	Security & Compliance				
<ul> <li>Deduplication &amp; Compression</li> <li>Compaction</li> <li>Data Tiering (Fabric Pool)</li> </ul>	Fpolicy & Vscan     LDAP & Kerberos     Active Directory integration     Certificate based authentication				

**NetApp BlueXP** enables you to manage all of your storage and data assets from a single control plane/interface.

You can use BlueXP to create and administer cloud storage (for example, Cloud Volumes ONTAP and Azure NetApp Files), to move, protect, and analyze data, and to control many on-prem and edge storage devices.

**NetApp Astra Trident** is a CSI Compliant Storage Orchestrator that enable quick and easy consumption of persistent storage backed by a variety of the above-mentioned NetApp storage options. It is an open-source software maintained and supported by NetApp.

# **ONTAP** feature highlights

# Astra Trident CSI feature highlights

<ul> <li>CSI specific</li> <li>CSI NetApp<sup>®</sup> Snapshot<sup>™</sup> copies and volume creation from CSI Snapshot copies</li> <li>CSI topology</li> <li>Volume expansion</li> </ul>	Security <ul> <li>Dynamic-export policy management</li> <li>iSCSI initiator-groups dynamic management</li> <li>iSCSI bidirectional CHAP</li> </ul>
Control     Storage and performance     consumption     Cross Namespace Volume     Monitoring     Access	Installation methods <ul> <li>Binary</li> <li>Operator</li> <li>Helm chart</li> <li>GitOps</li> </ul>
Choose your access mode         • RWO (ReadWriteOnce, i.e 1⇔1)       • RWOP (ReadWriteOnce POD)         • RWX (ReadWriteMany, i.e 1⇔n)         • ROX (ReadOnlyMany)	Choose your protocol <ul> <li>NFS</li> <li>SMB</li> <li>iSCSI</li> </ul>

Business critical container workloads need more than just persistent volumes. Their data management requirements require protection and migration of the application kubernetes objects as well.

Application data includes kubernetes objects in addition to the user data: Some examples are as follows:

- kubernetes objects such as pods specs, PVCs, deployments, services
- custom config objects such as config maps and secrets
- persistent data such as Snapshot copies, backups, clones
- custom resources such as CRs and CRDs

**NetApp Astra Control**, available as both fully-managed and self-managed software, provides orchestration for robust application data management. Refer to the Astra documentation for additional details on the Astra family of products.

This reference documentation provides validation of migration and protection of container-based applications, deployed on RedHat OpenShift container platform, using NetApp Astra Control Center. In addition, the solution provides high-level details for the deployment and the use of Red Hat Advanced Cluster Management (ACM) for managing the container platforms. The document also highlights the details for the integration of NetApp storage with Red Hat OpenShift container platforms using Astra Trident CSI provisioner. Astra Control Center is deployed on the hub cluster and is used to manage the container applications and their persistent storage lifecycle. Finally, it provides a solution for replication and failover and fail-back for container workloads on managed Red Hat OpenShift clusters in AWS (ROSA) using Amazon FSx for NetApp ONTAP (FSxN) as persistent storage.

### NetApp Solution with Managed Red Hat OpenShift Container platform workloads on AWS

Customers may be "born in the cloud" or may be at a point in their modernization journey when they are ready to move some select workloads or all workloads from their data centers to the cloud. They may choose to use provider-managed OpenShift containers and provider-managed NetApp storage in the cloud for running their workloads. They should plan and deploy the Managed Red Hat OpenShift container clusters (ROSA) in the cloud for a successful production-ready environment for their container workloads. When they are in AWS cloud, they could also deploy FSx for NetApp ONTAP for the storage needs.

FSx for NetApp ONTAP delivers data protection, reliability, and flexibility for container deployments in AWS. Astra Trident serves as the dynamic storage provisioner to consume the persistent FSxN storage for customers' stateful applications.

As ROSA can be deployed in HA mode with control plane nodes spread across multiple availability zones, FSx ONTAP can also be provisioned with Multi-AZ option which provides high availability and protect against AZ failures.



There are no data transfer charges when accessing an Amazon FSx file system from the file system's preferred Availability Zone (AZ). For more info on pricing, refer here.

## Data protection and migration solution for OpenShift Container workloads



## Deploy and configure the Managed Red Hat OpenShift Container platform on AWS

This section describes a high-level workflow of setting up the Managed Red Hat OpenShift clusters on AWS(ROSA). It shows the use of Managed FSx for NetApp ONTAP (FSxN) as the storage backend by Astra Trident to provide persistent volumes. Details are provided about the deployment of FSxN on AWS using BlueXP. Also, details are provided about the use of BlueXP and OpenShift GitOps (Argo CD) to perform data protection and migration activities for the stateful applications on ROSA clusters.

Here is a diagram that depicts the ROSA clusters deployed on AWS and using FSxN as the backend storage.



This solution was verified by using two ROSA clusters in two VPCs in AWS. Each ROSA cluster was integrated with FSxN using Astra Trident. There are several ways of deploying ROSA clusters and FSxN in AWS. This high-level description of the setup provides documentation links for the specific method that was used. You can refer to the other methods in the relevant links provided in the resources section.

The setup process can be broken down into the following steps:

### Install ROSA clusters

- Create two VPCs and set up VPC peering connectivity between the VPCs.
- Refer here for instructions to install ROSA clusters.

#### Install FSxN

(i)

- Install FSxN on the VPCs from BlueXP.
   Refer here for BlueXP account creation and to get started.
   Refer here for installing FSxN.
   Refer here for creating a connector in AWS to manage the FSxN.
- Deploy FSxN using AWS. Refer here for deployment using AWS console.

• Use Helm chart to install Trident on ROSA clusters. url for the Helm chart: https://netapp.github.io/trident-helm-chart

## Integration of FSxN with Astra Trident for ROSA clusters



OpenShift GitOps can be utilized to deploy Astra Trident CSI to all managed clusters as they get registered to ArgoCD using ApplicationSet.





- Refer here for details about creating backend and storage class.
- Make the storage class created for FsxN with Trident CSI as default from OpenShift Console. See screenshot below:

Create StorageGlas	Create					StorageClasses		C Administrator
					Name  Search by name /	>	Home	
olicy 1	Reclaim				Provisioner 1	Name 1	>	Operators
1	Delete				csitrident.netappio	SO foxn-nas - Default	>	Workloads
1	Delete				kubernetes.io/aws-ebs	<b>GD</b> 992	>	Networking
1	Delete				ebs.csi.aws.com	SD gp2-csi	<u></u>	<b>.</b>
1	Delete				ebs.csi.avs.com	<b>60</b> gp 3		Storage
1	Delete				ebs.csi.aws.com	🚱 gp3-csi	ms.	PersistentVolumes PersistentVolumeClair
	and a						115	PersistentVolumeClair StorageClasses
								StorageClasses VolumeSnapshots

### Deploy an application using OpenShift GitOps (Argo CD)

- Install OpenShift GitOps operator on the cluster. Refer to instructions here.
- SetUp a new Argo CD instance for the cluster. Refer to instructions here.

Open the console of Argo CD and deploy an app. As an example, you can deploy a Jenkins App using Argo CD with a Helm Chart. When creating the application, the following details were provided: Project: default cluster: https://kubernetes.default.svc Namespace: Jenkins The url for the Helm Chart: https://charts.bitnami.com/bitnami

Helm Parameters: global.storageClass: fsxn-nas

### **Data protection**

This page shows the data protection options for Managed Red Hat OpenShift on AWS (ROSA) clusters using Astra Control Service. Astra Control Service (ACS) provides an easy-to-use graphical user-interface with which you can add clusters, define applications running on them, and perform application aware data management activities. ACS functions can also be accessed using an API that allows for automation of workflows.

Powering Astra Control (ACS or ACC) is NetApp Astra Trident. Astra Trident integrates several types of Kubernetes clusters such as Red Hat OpenShift, EKS, AKS, SUSE Rancher, Anthos etc., with various flavors of NetApp ONTAP storage such as FAS/AFF, ONTAP Select, CVO, Google Cloud Volumes Service, Azure
NetApp Files and Amazon FSx for NetApp ONTAP.

This section provides details for the following data protection options using ACS:

- A video showing Backup and Restore of a ROSA application running in one region and restoring to another region.
- A video showing Snapshot and Restore of a ROSA application.
- Step-by-step details of installing a ROSA cluster, Amazon FSx for NetApp ONTAP, using NetApp Astra Trident to integrate with storage backend, installing a postgresql application on ROSA cluster, using ACS to create a snapshot of the application and restoring the application from it.
- A blog showing step-by-step details of creating and restoring from a snapshot for a mysql application on a ROSA cluster with FSx for ONTAP using ACS.

#### Backup/Restore from Backup

The following video shows the backup of a ROSA application running in one region and restoring to another region.

#### FSx NetApp ONTAP for Red Hat OpenShift Service on AWS

#### Snapshot/Restore from snapshot

The following video shows taking a snapshot of a ROSA application and restoring from the snapshot after.

Snapshot/Restore for Applications on Red Hat OpenShift Service on AWS (ROSA)clusters with Amazon FSx for NetApp ONTAP storage

#### Blog

• Using Astra Control Service for data management of apps on ROSA clusters with Amazon FSx storage

#### Step-by-Step Details to create snapshot and restore from it

#### **Prerequisite setup**

- AWS account
- Red Hat OpenShift account
- · IAM user with appropriate permissions to create and access ROSA cluster
- AWS CLI
- ROSA CLI
- OpenShift CLI(oc)
- · VPC with subnets and appropriate gateways and routes
- ROSA Cluster installed into the VPC
- · Amazon FSx for NetApp ONTAP created in the same VPC
- Access to the ROSA cluster from OpenShift Hybrid Cloud Console

#### **Next Steps**

1. Create an admin user and login to the cluster.

- 2. Create a kubeconfig file for the cluster.
- 3. Install Astra Trident on the cluster.
- 4. Create a backend, storage class and snapshot class configuration using the Trident CSI provisioner.
- 5. Deploy a postgresql application on the cluster.
- 6. Create a database and add a record.
- 7. Add the cluster into ACS.
- 8. Define the application in ACS.
- 9. Create a snapshot using ACS.
- 10. Delete the database in the postgresql application.
- 11. Restore from a snapshot using ACS.
- 12. Verify your app has been restored form the snapshot.

# 1. Create an admin user and login to the cluster

Access the ROSA cluster by creating an admin user with the following command : (You need to create an admin user only if you did not create one at the time of installation)

rosa create admin --cluster=<cluster-name>

The command will provide an output that will look like the following. Login to the cluster using the oc login command provided in the output.

W: It is recommended to add an identity provider to login to this cluster. See 'rosa create idp --help' for more information. I: Admin account has been added to cluster 'my-rosa-cluster'. It may take up to a minute for the account to become active. I: To login, run the following command: oc login https://api.my-rosa-cluster.abcd.p1.openshiftapps.com:6443 \ --username cluster-admin \ --password FWGYL-2mkJI-00000-00000



You can also login to the cluster using a token. If you already created an admin-user at the time of cluster creation, you can login to the cluster from the Red Hat OpenShift Hybrid Cloud console with the admin-user credentials. Then by clicking on the top right corner where it displays the name of the logged in user, you can obtain the oc login command (token login) for the command line.

# 2. Create a kubeconfig file for the cluster

Follow the procedures here to create a kubeconfig file for the ROSA cluster. This kubeconfig file will be used later when you add the cluster into ACS.

#### 3. Install Astra Trident on the cluster

Install Astra Trident (latest version) on the ROSA cluster. To do this, you can follow any one of the procedures

given here. To install Trident using helm from the console of the cluster, first create a project called Trident.

E SpenSi	<b>t</b> hift Service on AWS				=	<b>\$</b> 2	Ð	0	cluster-ac	imin <del>-</del>
Projects									Create	Project
▼ Filter ▼ Nam	e 👻 trident									
Name trident X C	lear all filters									
Name 1	Display name	Status	1	Requester	I		Created	1.1		
PR trident	trident	Active		rosaadmin			🛛 Feb 1	2, 2024, 9	9:54 PM	I

Then from the Developer view, create a Helm chart repository. For the URL field use

'https://netapp.github.io/trident-helm-chart'. Then create a helm release for Trident operator.

Project trident 👻
Create Helm Chart Repository Add helm chart repository.
Configure via:   Form view O YAML view
Scope type
<ul> <li>Namespaced scoped (ProjectHelmChartRepository)</li> </ul>
Add Helm Chart Repository in the selected namespace.
Cluster scoped (HelmChartRepository)
Add Helm Chart Repository at the cluster level and in all namespaces.
Name *
trident
A unique name for the Helm Chart repository.
Display name
Astra Trident
A display name for the Helm Chart repository.
Description
NetApp Astra Trident
A description for the Helm Chart repository.
Disable usage of the repo in the developer catalog.
URL .
https://netapp.github.io/trident-helm-chart

Project: trident 🔹		
Developer Catalog > Helm Cha	rts	
Helm Charts		
Browse for charts that help man catalog. Alternatively, develope	nage complex installations and upgrades. Clus rs can try to configure their own custom Helm	ter administrators can custo Chart repository.
All items	All items	
CI/CD	0	
Languages	Q Filter by keyword	A-Z -
Other		
Chart Repositories	Helm Charts	
Astra Trident (1)	TRIDENT	
<ul> <li>OpenShift Helm</li> <li>Charts (87)</li> </ul>	Trident Operator	
	A Helm chart for deploying	
Source	NetApp's Trident CSI storage	
Community (33)	provisioner using the Trident	
Partner (42)		
Red Hat (12)		

Verify all trident pods are running by going back to the Administrator view on the console and selecting pods in the trident project.

<b>Red Hat</b> OpenShift Service	on AWS					
🏟 Administrator 👻	Project: trident 🔻					
Home >	Pods					
Operators >	▼ Filter ▼ Name	Search by name	1			
Workloads 🗸	Name †	Status 1	Ready 1	Restarts 👔	Owner 1	Mem
Pods	P trident-controller- 69cff44ddf-4dqnj	C Running	6/6	0	RS trident-controller- 69cff44ddf	3
Deployments DeploymentConfigs	P trident-node-linux- 4b6fm	C Running	2/2	0	DS trident-node-linux	ω.
StatefulSets	P trident-node-linux- 4sckw	C Running	2/2	0	<b>DS</b> trident-node-linux	-
ConfigMaps	P trident-node-linux- 7I42w	C Running	2/2	0	DS trident-node-linux	-
CronJobs	P trident-node-linux- dbhp4	C Running	2/2	0	DS trident-node-linux	-
Jobs DaemonSets	P trident-node-linux- gj5km	C Running	2/2	0	<b>OS</b> trident-node-linux	-
ReplicaSets	P trident-node-linux- r79c8	C Running	2/2	0	DS trident-node-linux	
HorizontalPodAutoscalers	P trident-node-linux- tzwdp	C Running	2/2	0	DS trident-node-linux	~
PodDisruptionBudgets	P trident-node-linux- vdvxt	C Running	2/2	0	DS trident-node-linux	•
Networking >	P trident-operator- 7f7fd45c68-6crcb	C Running	1/1	0	RS trident-operator- 7f7fd45c68	2

# 4. Create a backend, storage class and snapshot class configuration using the Trident CSI provisioner

Use the yaml files shown below to create a trident backend object, storage class object and the Volumesnapshot object. Be sure to provide the credentials to your Amazon FSx for NetApp ONTAP file system you created, the management LIF and the vserver name of your file system in the configuration yaml for the backend. To get those details, go to the AWS console for Amazon FSx and select the file system, navigate to the Administration tab. Also, click on update to set the password for the fsxadmin user.



You can use the command line to create the objects or create them with the yaml files from the hybrid cloud console.

ESx > File systems > fs-049f9a23aac951429		
fsx-for-rosa (fs-049f9a23aac951429)		
▼ Summary		
File system ID fs-049f9a23aac951429 Lifecycle state Available File system type ONTAP Deployment type Single-AZ	SSD storage capacity Update 1	Availability Zones us-west-2b 🗇 Creation time 2024-02-12T20:15:23-05:00
Network & security Monitoring & performance Administration Stora	ge virtual machines Volumes Backups Updates Tags	
ONTAP administration		
Management endpoint - DNS name management.fs-049f9a23aac951429.fsx.us-west-2.amazonaws.com 🗗 Inter-cluster endpoint - DNS name intercluster.fs-049f9a23aac951429.fsx.us-west-2.amazonaws.com 🗗	Management endpoint - IP address 10.49.9.135 🗗 Inter-cluster endpoint - IP address 10.49.9.49 🗗 10.49.9.251 🗗	ONTAP administrator username fsxadmin 🗗 ONTAP administrator password Update

#### **Trident Backend Configuration**

```
apiVersion: v1
kind: Secret
metadata:
 name: backend-tbc-ontap-nas-secret
type: Opaque
stringData:
 username: fsxadmin
 password: <password>
___
apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
 name: ontap-nas
spec:
 version: 1
  storageDriverName: ontap-nas
  managementLIF: <management lif>
  backendName: ontap-nas
  svm: fsx
  credentials:
    name: backend-tbc-ontap-nas-secret
```

# Storage Class

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
   name: ontap-nas
provisioner: csi.trident.netapp.io
parameters:
   backendType: "ontap-nas"
   media: "ssd"
   provisioningType: "thin"
   snapshots: "true"
allowVolumeExpansion: true
```

#### snapshot class

```
apiVersion: snapshot.storage.k8s.io/v1
kind: VolumeSnapshotClass
metadata:
    name: trident-snapshotclass
driver: csi.trident.netapp.io
deletionPolicy: Delete
```

Verify that the backend, storage class and the trident-snapshotclass objects are created by issuing the commands shown below.

[ec2-user@ip	-10-49-11-13	2 storage]\$ kul	pectl get	tbc -n	trident				
NAME	BACKEND NAME	BACKEND UUI	)			PHASE	STAT	JS	
ontap-nas	ontap-nas	8a5e4583-2da	ic-46bb-b	01e-fa70	3816f121	Bound	Succ	ess	
[ec2-user@ip	-10-49-11-13	2 storage]\$ kut	ectl get	sc					
NAME	PROVISIO	NER	RECLAIM	POLICY	VOLUMEBI	NDINGMODE		ALLOWVOLUMEEXPANSION	AGE
gp2	kubernet	es.io/aws-ebs	Delete		WaitForF	irstConsu	ımer	true	3h23m
gp2-csi	ebs.csi.	aws.com	Delete		WaitForF	irstConsu	mer	true	3h19m
gp3 (default	:) ebs.csi.	aws.com	Delete		WaitForF	irstConsu	mer	true	3h23m
gp3-csi	ebs.csi.	aws.com	Delete		WaitForF	irstConsu	mer	true	3h19m
ontap-nas	csi.trid	lent.netapp.io	Delete		Immediat	e		true	141m
[ec2-user@ip	-10-49-11-13	2 storage]\$ kut	ectl get	Volumes	mapshotcl	ass			
NAME		DRIVER		DELETIC	ONPOLICY	AGE			
csi-aws-vsc		ebs.csi.aws.com	n	Delete		3h19m			
trident-snap	oshotclass	csi.trident.net	tapp.io	Delete		6m56s			
[ec2-user@ip	-10-49-11-13	2 storage]\$							

At this time, an important modification you need to make is to set ontap-nas as the default storage class instead of gp3 so that the postgresql app you deploy later can use the default storage class. In the Openshift console of your cluster, under Storage select StorageClasses. Edit the annotation of the current default class to be false and add the annotation storageclass.kubernetes.io/is-default-class set to true for the ontap-nas storage class.

■ <sup>Sed</sup> Hat OpenShift Servi				0	cluster-admin <del>v</del>
StorageClasses	Edit annotations Key storageclass.kubernetes.io/is	Value false	•		Create StorageClass
Name 1	Add more	_		Reclaim po	licy 1
SC gp2-csi			Cancel Save	Delete	1
SC gp3 – Default	ebs.csi.av	vs.com		Delete	1
SC gp3-csi	ebs.csi.au	vs.com		Delete	1
SC ontap-nas	csitriden	t.netapp.io		Delete	1

StorageClasses		Create StorageClass
Name   Search by name  I		
Name 1	Provisioner 1	Reclaim policy 1
SC gp2	kubernetes.io/aws-ebs	Delete
SC gp2-csi	ebs.csi.aws.com	Delete #
SC gp3	ebs.csi.aws.com	Delete #
SC gp3-csi	ebs.csi.aws.com	Delete
SC ontap-nas - Default	csitrident.netapp.io	Delete :

# 5. Deploy a postgresql application on the cluster

You can deploy the application from the command line as follows:

helm install postgresql bitnami/postgresql -n postgresql --create-namespace

[ec2-user@ip-10-	49-11-132 astra]\$ helm install postgresql bitnami/postgresql -n postgresqlcreate-namespace
NAME: postgresql	(un Fab 13 14/46-16 2024
NAMESPACE: posta	ue res la tradit dece
STATUS: deployed	
REVISION: 1	
TEST SUITE: None	
NOTES:	
CHART MARE: post	gresq1
APP VERSION: 16	44.049 2.0
** Please be nat	ient while the chart is being deployed **
Statement and the second	
PostgreSQL can b	e accessed via port \$432 on the following DMS names from within your cluster:
postgresql.p	ostgresql.svc.cluster.local - Read/Write connection
To get the passw	word for "postgres" run:
export POSTG	RES_PASSWORD=\$(kubect1 get secretnamespace postgresq1 postgresq1 -o jsonpath="{.data.postgres-password}"   base64 -d)
To connect to yo	our database run the following command:
kubectl run	postgresql-clientrmtty -irestart='Never'namespace postgresqlimage docker.io/bitnami/postgresql:16.2.0-debian-11-r1env="PG
command	psqlhost postgresql -U postgres -d postgres -p 5432
> NOTE: If y the error "psql:	ou access the container using bash, make sure that you execute "/opt/bitnami/scripts/postgresql/entrypoint.sh /bin/bash" in order to avoid local user with ID 1001} does not exist"
To connect to yo	our database from outside the cluster execute the following commands:
kubectl port PGPASSWORD="	-forwardnamespace postgresql svc/postgresql 5432:5432 & \$POSTGRES_PASSWORD" psqlhost 127.0.0.1 -U postgres -d postgres -p 5432
WARNING: The con case, old PVC w [ec2-usen@ip-10-	ofigured password will be ignored on new installation in case when previous PostgreSQL release was deleted through the helm command. In that ill have an old password, and setting it through helm won't take effect. Deleting persistent volumes (PVs) will solve the issue. 49-11-132 astra]5 _
	If you do not see the application pods running, then there might be an error caused due to
	in you do not do ano application pous running, then there might be an error baused due to
	security context constraints.

image::rhhc-scc-error.png[]

(i)

Fix the error by editing the runAsUser and fsGroup fields in statefuleset.apps/postgresql object with the uid that is in the output of the oc get project command as shown below. image::rhhc-scc-fix.png[]

postgresql app should be running and using persistent volumes backed by Amazon FSx for NetApp ONTAP storage.

postgresql-0 1/1 Running 0 2m46s	NAME
1 2 01 40 40 44 432 14	postgresql-0
[ec2-user@ip-10-49-11-132 astra]\$	[ec2-user@ip-1

[ec2-user@ip-10-49-	11-132 st	orage]\$ kubectl get pvc -n postgresql				
NAME	STATUS	VOLUME	CAPACITY	ACCESS MODES	STORAGECLASS	AGE
data-postgresql-0	Bound	pvc-dd09524a-de75-4825-9424-03a9b91195ca	8Gi	RWO	ontap-nas	4m2s
[ec2-user@ip-10-49-	11-13 <mark>2</mark> st	orage]\$				

6. Create a database and add a record



#### 7. Add the cluster into ACS

Log in to ACS. Select cluster and click on Add. Select other and upload or paste the kubeconfig file.

6			
ROVIDER			
Microsoft Azure	Soogle Cloud Platform	aws Amazon Web Services	Other
BECONFIG			
Please ensure that the	cubeconfig used for this cluster has a	long-lived token associated with it.	
Provide Astra Control access to	your Kubernetes clusters by entering a	kubaconfin cradential Follow these instr	uctions (2 on how to create
Provide Astra Control access to a dedicated admin-role kubeco	your Kubernetes clusters by entering a nfig.	kubeconfig credential. Follow these instr	uctions 🖸 on how to create
Provide Astra Control access to a dedicated admin-role kubect Upload file Paste or type	your Kubernetes clusters by entering a onfig.	kubeconfig credential. Follow these instr	uctions 🖸 on how to create
Provide Astra Control access to a dedicated admin-role kubect Upload file Paste or type XJuZXR1cy5pby9z2XJ2aWN1Y	your Kubernetes clusters by entering a infig. MNjb3VudC9zZXJ2aWN1LWFjY291bnQul	kubeconfig credential. Follow these instr bmFt2SI6ImFzdHJhY29udHJvbC1z2XJ2a	uctions C on how to create
Provide Astra Control access to a dedicated admin-role kubeco Upload file Paste or type XJuZXR1cy5pby9zZXJ2aWN1Y 1cm51dGVzLm1vL3N1cnZpY2V ToiLCJzdWI101JzeXN02W06c	your Kubernetes clusters by entering a infig. WNjb3VudC9zZXJ2aWN1LWFjY291bnQu hY2NvdW50L3N1cnZpY2UtYWNjb3VudC 2VydmljZWFjY291bnQ62GVmYXVadDph	kubeconfig credential. Follow these instr bmFt2SI6ImFzdHJhY29udHJvbC1z2XJ2a 51aWQ101I4NzFh0TI4MC0wMTEyLTRmYzA c3RyYWNvbnRyb2wtc2Vydmlj2S1hY2Nvd	Uctions C on how to create
Provide Astra Control access to a dedicated admin-role kubeco Upload file Paste or type XJuZXR1cySpby9zZXJ2aWN1Y 1cm51dGVzLm1vL3N1cnZpY2V To1LCJzdW1101JzeXN0ZW06c LkW-8ZDY0ShQ5Uo1aSbJ- 0SIdSr0EbvfcQ3tSf40VC72n 3XWHFZ2cTXXpdKqtzWfmBLXY	your Kubernetes clusters by entering a unfig. WNjb3VudC9zZXJ2aWN1LWFjY291bnQu hY2NvdW50L3N1cn2pY2UtYWNjb3VudC1 2Vydmlj2WFjY291bnQ62GVmYXVadDph M4BqYbN8cm0y0V81pF30G7tYA9XAIdw1 huN1CzBMY7S5MVhB2WD_eixptN02a1	kubeconfig credential. Follow these instr bmFt2SI6ImFzdHJhY29udHJvbClz2XJ2a SlaWQ101I4NzFhOTI4McOwMTEyLTRmYzA c3RyYWNvbnRyb2wtc2Vydmlj2SlhY2Nvd X98xAXJ00T2U0G2xbyLWfOqLCFDx3_uS9 vaWmI2jrUQL0_q8Uj2Exe9vVH1KPxfb0C	Uctions C on how to create WN1LWFjY291bnQilCJrdWJ tCWFkNS0z2DISNzA2N2NiN W50In0.M7-IRxcaKOe7S- uqU63t8LLmeenCBiOm9PaD xU4TvHncbathvL6m21N7Om

Click **Next** and select ontap-nas as the default storage class for ACS. Click **Next**, review the details and **Add** the cluster.

🛱 Add clu	ister	STEP 2/3: STORAG	E		×
STORAGE					
✓ Assign a	new default storage class				
The following	storage classes are available on	the cluster.			
Set default	Storage class	Storage provisioner	Reclaim policy	Binding mode	Eligibility
	gp2	kubernetesio/aws-ebs	Delete	waitForFirstConsumer	
0	gp2-csi	ebs.csi.aws.com	Delete	WaitForFirstConsumer	⊘ Eligible
0	gp3	ebs.csi.aws.com	Delete	WaitForFirstConsumer	⊘ Eligible
0	gp3-csi	ebs.csi.aws.com	Delete	WaitForFirstConsumer	⊘ Eligible
•	ontap-nas Default	csi.trident.netapp.io	Delete	Immediate	⊘ Eligible
		← Back Ne	xt →		

#### 8. Define the application in ACS

Define the postgresql application in ACS. From the landing page, select **Applications**, **Define** and fill in the appropriate details. Click **Next** a couple of times, Review the details and click **Define**. The application gets added to ACS.

🛱 Add clu	ister	STEP 2/3: STORA	GE		×					
STORAGE										
✓ Assign a	✓ Assign a new default storage class									
The following	storage classes are available on	the cluster.								
Set default	Storage class	Storage provisioner	Reclaim policy	Binding mode	Eligibility					
<u> </u>	gpz	kubernetes.io/aws-ebs	Delete	waitForFirstConsumer						
	gp2-csi	ebs.csi.aws.com	Delete	WaitForFirstConsumer	⊘ Eligible					
	gp3	ebs.csi.aws.com	Delete	WaitForFirstConsumer	⊘ Eligible					
	gp3-csi	ebs.csi.aws.com	Delete	WaitForFirstConsumer	Sligible					
. •	ontap-nas Default	csi.trident.netapp.io	Delete	Immediate	⊘ Eligible					
		← Back	Next →							

#### 9. Create a snapshot using ACS

There are many ways to create a snapshot in ACS. You can select the application and create a snapshot from the page that shows the details of the application. You can click on Create snapshot to create an on-demand snapshot or configure a protection policy.

Create an on-demand snapshot by simply clicking on **Create snapshot**, providing a name, reviewing the details, and clicking on **Snapshot**. The snapshot state changes to Healthy after the operation is completed.

[뉴] Dashboard	Data protection Storage Resource	es Execution hooks	Activity Tasks	Snapshots
Applications	Actions  Configure protection	n policy	- Search	
Clusters			0–0 of 0 entries	
Cloud instances	Name State On-	Schedule / On-Demand	Created †	Actions
Buckets				
& Account				
Activity				
Support	After you	You don't have any snap	oshots will be listed here	
NetApp		Create snapshot	Will be listed here	
				enerio
[2] Dashboard				
Applications		<b>S</b> AI	PPLICATION PROTECTION	1
Clusters	Available	(i) Partially protecte	d \Lambda No scheduled protectio	in policy
Cloud instances	Definition postgresql	Cluster S api-r	rosa-cluster1-nn5w-p1	
Buckets	Data protection Storage Resource	es Execution hooks	Activity Tasks	
Account	Actions   Configure protection	n policy	- Search	
E Activity			1-1 of 1 entries	
Support	Name	State On-Schedu	le / On-Demand Created 🕇	Actions
🔇 🗖 NetApp	postgresql-snapshot-20240213154610	⊘ Healthy  ◎ On-Dem	nand 2024/02/13 15:4	48 UTC 🔋

#### 10. Delete the database in the postgresql application

Log back into postgresql, list the available databases, delete the one you created previously and list again to ensure that the database has been deleted.

postgres=# \1 List of databases | ICU Locale | ICU Rules | Access priv | Owner | Encoding | Locale Provider | Collate | Name Ctype UTF8 libc en US.UTF-8 | en US.UTF-8 postgres erp postgres postgres UTF8 libc en US.UTF-8 | en US.UTF-8 1 template0 | postgres | UTF8 libc en\_US.UTF-8 | en\_US.UTF-8 | =c/postgres postgres=CTc UTF8 en\_US.UTF-8 | en\_US.UTF-8 template1 | postgres libc -c/postgres | postgres=CTc/ (4 rows) postgres=# DROP DATABASE erp; DROP DATABASE postgres=# \1 List of databases | Encoding | Locale Provider | | ICU Locale | ICU Rules | Name Owner Collate Ctype Access priv postgres | UTF8 libc en\_US.UTF-8 | en\_US.UTF-8 postgres template0 | postgres | UTF8 libc en\_US.UTF-8 | en\_US.UTF-8 | =c/postgres postgres=CTc template1 | postgres UTF8 libc en\_US.UTF-8 | en\_US.UTF-8 =c/postgres I postgres=CTc (3 rows)

# 11. Restore from a snapshot using ACS

To restore the application from a snapshot, go to ACS UI landing page, select the application and select Restore. You need to pick a snapshot or a backup from which to restore. (Typically, you would have multiple created based on a policy that you have configured). Make appropriate choices in the next couple of screens and then click on **Restore**. The application status moves from Restoring to Available after it has been restored from the snapshot.

[뉴 Dashboard	© postgresql ^			с (	Actions 🗸
<ul> <li>Applications</li> <li>Clusters</li> <li>Cloud instances</li> </ul>	APPLICATION STATUS	() Part	S APPLICATION PROT	ECTION scheduled protect	Snapshot Back up Clone Restore
Buckets	Definition postgresql		Cluster 3 api-rosa-cluster1-nn5w-p1-op	<u>p</u>	Unmanage
오 Account I Activity 당 Support	Data protection         Storage         Resources           Actions <ul> <li>Configure protection p</li> <li>Configure protection p</li> </ul>	Execution hoo	ks Activity Tasks	1–1 of 1 entr	
	Name	State	On-Schedule / On-Demand	Created †	Actions
<ul> <li>In NetApp</li> </ul>	postgresql-snapshot-20240213164912	⊘ Healthy	On-Demand	2024/02/13 16	50 UTC []

RESTORE TYPE					
Restore the application to new namespaces on any available	ble cluster or to original na	mespaces on the original cluster.			
Restore to new namespaces     Restore to original namespaces					
RESTORE SOURCE					
Select a snapshot or backup to restore the application to	a previous state.				
	Time range 🗸	Filter	Snapshots 🔒 Backups		
Application snapshot	Snapshot state	On-Schedule / On-Demand	Created ↑		
postgresql-snapshot-20240213164912	⊘ Healthy	On-Demand	2024/02/13 16:50 UTC		
	Ca	ncel Next →			

Ƙ Dashboard	🍥 postgresql \land		C Actions	~	
<ul> <li>Applications</li> <li>Clusters</li> <li>Cloud instances</li> </ul>	- →- APPLICATION STATUS	(i) <u>P</u>	S APPLICATION PROT artially protected A No	ECTION scheduled protection policy	(1)
Buckets	Definition postgresql		Cluster api-rosa-cluster1-nn5w-p1-o	<b>p</b>	
& Account	Data protection Storage Resource	s Execution ho	ooks Activity Tasks		
E Activity	Actions   Configure protection	policy	- Search		
Support				1–1 of 1 entries	
	Name	State	On-Schedule / On-Demand	Created †	Actions
In NetApp	postgresql-snapshot-20240213164912	⊘ Healthy	On-Demand	2024/02/13 16:50 UTC	(1)

# 12. Verify your app has been restored from the snapshot

Login to the postgresql client and you should now see the table and the record in the table that you previously had. That's it. Just by clicking a button, your application has been restored to a previous state. That is how easy we make it for our customers with Astra Control.

[ec2-user@i 2.0-debian- Warning: wo legeEscalat pod or cont text.seccom If you don'	p-10-49-11- 11-r1env uld violate ion=false), ainer "post pProfile.ty t see a com	132 ~]\$ kub ="PGPASSWORJ PodSecurity unrestrictor gresql-clier pe to "Runt mand prompt	ectl run postgresq) D=\$POSTGRES_PASSWOR y "restricted:v1.24 ed capabilities (co nt" must set securi imeDefault" or "Loo , try pressing ento	L-clientrm RD"command I": allowPrivi Dontainer "post ityContext.run calhost") er.	tty -irest psqlhost legeEscalation gresql-client" AsNonRoot-true)	<pre>cart='Never' - postgresql -L postgresql -L i= false (con must set secu , seccompProf </pre>	-namespace p I postgres -d itainer "post irityContext. file (pod or	ostgresqlimage docker.io/bitnami/postgresql:16.   postgres -p 5432  gresql-client" must set securityContext.allowPrivi capabilities.drop=["ALL"]), runAsNonRoot != true ( container "postgresql-client" must set securityCon
Sources and	Mis			List of da	tabases			
Name	Owner:	Encoding	Locale Provider	Collate	Ctype	ICU Locale	ICU Rules	Access privileges
erp postgres template8	postgres postgres postgres	1 UTF8 1 UTF8 1 UTF8 1 UTF8	11bc 11bc 11bc	en_US.UTF-8 en_US.UTF-8 en_US.UTF-8	en_US.UTF-8 en_US.UTF-8 en_US.UTF-8 en_US.UTF-8			-c/postgres + postgres-CTc/postgres rc/nostgres +
Comparents.	hoselbes			en_ostant-o	en_051017-0			postgres=CTc/postgres
(4 rows) postgres=# You are now erp=# \dt L Schema     public   p (1 row) erp=# SELEC id   first	\c erp connected ist of rela Name   Ty ersons   ta T * from PE name   last	to database tions pe   Owner ble   postgr name 	"erp" as user "pos n res	stgres".				
(1 row)	1 Doe							Activate Windows

#### **Data migration**

This page shows the data migration options for container workloads on Managed Red Hat OpenShift clusters using FSx for NetApp ONTAP for persistent storage.

#### Data Migration

Red Hat OpenShift service on AWS as well as FSx for NetApp ONTAP (FSxN) are part of their service portfolio by AWS. FSxN is available on Single AZ or Multi-AZ options.

Multi-Az option provides data protection from availability zone failure.

FSxN can be integrated with Astra Trident to provide persistent storage for applications on ROSA clusters.

#### Integration of FSxN with Trident using Helm chart

#### ROSA Cluster Integration with Amazon FSx for ONTAP

The migration of container applications involves:

- Persistent volumes: this can be accomplished using BlueXP. Another option is to use Astra Control Center to handle container application migrations from on-premises to the cloud environment. Automation can be used for the same purpose.
- Application metadata: this can be accomplished using OpenShift GitOps (Argo CD).

#### Failover and Fail-back of applications on ROSA cluster using FSxN for persistent storage

The following video is a demonstration of application failover and fail-back scenarios using BlueXP and Argo CD.

#### Failover and Fail-back of applications on ROSA cluster



# Data protection for Container Apps in OpenShift Container Platform using OpenShift API for Data Protection (OADP)

Author: Banu Sundhar, NetApp

This section of the reference document provides details for creating backups of Container Apps using the OpenShift API for Data Protection (OADP) with Velero on NetApp ONTAP S3 or NetApp StorageGRID S3. The backups of namespace scoped resources including Persistent Volumes(PVs) of the app are created using CSI Astra Trident Snapshots.

The persistent storage for container apps can be backed by ONTAP storage integrated to the OpenShift Cluster using Astra Trident CSI. In this section we use OpenShift API for Data Protection (OADP) to perform backup of apps including its data volumes to

- ONTAP Object Storage
- StorageGrid

We then restore from the backup when needed. Please note that the app can be restored only to the cluster from where the backup was created.

OADP enables backup, restore, and disaster recovery of applications on an OpenShift cluster. Data that can be protected with OADP include Kubernetes resource objects, persistent volumes, and internal images.



Red Hat OpenShift has leveraged the solutions developed by the OpenSource communities for data protection. Velero is an open-source tool to safely backup and restore, perform disaster recovery, and migrate Kubernetes cluster resources and persistent volumes. To use Velero easily, OpenShift has developed the OADP operator and the Velero plugin to integrate with the CSI storage drivers. The core of the OADP APIs that are exposed are based on the Velero APIs. After installing the OADP operator and configuring it, the backup/restore operations that can be performed are based on the operations exposed by the Velero API.



OADP 1.3 is available from the operator hub of OpenShift cluster 4.12 and later. It has a built-in Data Mover that can move CSI volume snapshots to a remote object store. This provides portability and durability by

moving snapshots to an object storage location during backup. The snapshots are then available for restoration after disasters.

# The following are the versions of the various components used for the examples in this section

- OpenShift Cluster 4.14
- OADP Operator 1.13 provided by Red Hat
- Velero CLI 1.13 for Linux
- Astra Trident 24.02
- ONTAP 9.12
- postgresql installed using helm.

#### Astra Trident CSI OpenShift API for Data Protection (OADP) Velero

# Data protection for Container Apps in OpenShift Container Platform using OpenShift API for Data Protection (OADP)

Author: Banu Sundhar, NetApp

This section of the reference document provides details for creating backups of Container Apps using the OpenShift API for Data Protection (OADP) with Velero on NetApp ONTAP S3 or NetApp StorageGRID S3. The backups of namespace scoped resources including Persistent Volumes(PVs) of the app are created using CSI Astra Trident Snapshots.

The persistent storage for container apps can be backed by ONTAP storage integrated to the OpenShift Cluster using Astra Trident CSI. In this section we use OpenShift API for Data Protection (OADP) to perform backup of apps including its data volumes to

- ONTAP Object Storage
- StorageGrid

We then restore from the backup when needed. Please note that the app can be restored only to the cluster from where the backup was created.

OADP enables backup, restore, and disaster recovery of applications on an OpenShift cluster. Data that can be protected with OADP include Kubernetes resource objects, persistent volumes, and internal images.



Red Hat OpenShift has leveraged the solutions developed by the OpenSource communities for data protection. Velero is an open-source tool to safely backup and restore, perform disaster recovery, and migrate Kubernetes cluster resources and persistent volumes. To use Velero easily, OpenShift has developed the OADP operator and the Velero plugin to integrate with the CSI storage drivers. The core of the OADP APIs that are exposed are based on the Velero APIs. After installing the OADP operator and configuring it, the backup/restore operations that can be performed are based on the operations exposed by the Velero API.



OADP 1.3 is available from the operator hub of OpenShift cluster 4.12 and later. It has a built-in Data Mover that can move CSI volume snapshots to a remote object store. This provides portability and durability by

moving snapshots to an object storage location during backup. The snapshots are then available for restoration after disasters.

# The following are the versions of the various components used for the examples in this section

- OpenShift Cluster 4.14
- OADP Operator 1.13 provided by Red Hat
- Velero CLI 1.13 for Linux
- Astra Trident 24.02
- ONTAP 9.12
- postgresql installed using helm.

#### Astra Trident CSI OpenShift API for Data Protection (OADP) Velero

# Installation of OpenShift API for Data Protection (OADP) Operator

This section outlines the installation of OpenShift API for Data Protection (OADP) Operator.

#### Prerequisites

- A Red Hat OpenShift cluster (later than version 4.12) installed on bare-metal infrastructure with RHCOS worker nodes
- A NetApp ONTAP cluster integrated with the cluster using Astra Trident
- A Trident backend configured with an SVM on ONTAP cluster
- · A StorageClass configured on the OpenShift cluster with Astra Trident as the provisioner
- · Trident Snapshot class created on the cluster
- · Cluster-admin access to Red Hat OpenShift cluster
- · Admin access to NetApp ONTAP cluster
- An application eg. postgresql deployed on the cluster
- An admin workstation with tridentctl and oc tools installed and added to \$PATH

#### Steps to install OADP Operator

1. Go to the Operator Hub of the cluster and select Red Hat OADP operator. In the Install page, use all the default selections and click install. On the next page, again use all the defaults and click Install. The OADP operator will be installed in the namespace openshift-adp.

Home	OperatorHub			
Operators 🗸	Discover Operators from the Kub	bernetes community and Red Hat partners, curated b	by Red Hat. You can purchase commercial so	oftware through Red Hat Ma
OperatorHub	optional add-ons and shared ser	vices to your developers. After installation, the Oper-	ator capabilities will appear in the Developer	r Catalog providing a sen-se
Installed Operators	All Items Al/Machine Learning	All Items		
Workloads >	Application Runtime	Q OADP X		
Virtualization >	Cloud Provider	Red Hat	Community	
Networking >	Database Developer Tools		OADP Operator	
Storage >	Development Tools Drivers and plugins	provided by Red Hat	provided by Red Hat	
Builds >	Integration & Delivery Logging & Tracing	Protection) operator sets up and installs Data Protection	Protection) operator sets up and installs Velero on the OpenShift	
Observe >	Modernization & Migration			



Project: All Projects 🔻							
Installed Operators Installed Operators are represented by ClusterServiceVersions within this Namespace. For more information, see the Understanding Operators documentation Operator and ClusterServiceVersion using the Operator SDK C.							
Name	Search by name						
Name	OpenShift Virtualization 4.14.4 provided by Red Hat	Namespace I openshift-cnv	Managed Namespaces	Status Succeeded Up to date			
۲	OADP Operator 1.3.0 provided by Red Hat	NS openshift-adp	NS openshift-adp	Succeeded Up to date			
4	Package Server 0.0.1-snapshot provided by	NS openshift-operator-lifecycle- manager	NS openshift-operator-lifecycle- manager	Succeeded			

# Prerequisites for Velero configuration with Ontap S3 details

After the installation of the operator succeeds, configure the instance of Velero. Velero can be configured to use S3 compatible Object Storage. Configure ONTAP S3 using the procedures shown in the Object Storage Management section of ONTAP documentation. You will need the following information from your ONTAP S3 configuration to integrate with Velero.

- A Logical Interface (LIF) that can be used to access S3
- · User credentials to access S3 that includes the access key and the secret access key
- · A bucket name in S3 for backups with access permissions for the user
- For secure access to the Object storage, TLS certificate should be installed on the Object Storage server.

#### Prerequisites for Velero configuration with StorageGrid S3 details

Velero can be configured to use S3 compatible Object Storage. You can configure StorageGrid S3 using the procedures shown in the StorageGrid documentation. You will need the following information from your StorageGrid S3 configuration to integrate with Velero.

- The endpoint that can be used to access S3
- User credentials to access S3 that includes the access key and the secret access key
- · A bucket name in S3 for backups with access permissions for the user
- For secure access to the Object storage, TLS certificate should be installed on the Object Storage server.

#### Steps to configure Velero

• First, create a secret for an ONTAP S3 user credential or StorageGrid Tenant user credentials. This will be used to configure Velero later. You can create a secret from the CLI or from the web console. To create a secret from the web console, select Secrets, then click on Key/Value Secret. Provide the values for the credential name, key and the value as shown. Be sure to use the Access Key Id and Secret Access Key of your S3 user. Name the secret appropriately. In the sample below, a secret with ONTAP S3 user credentials named ontap-s3-credentials is created.

Installed Operators	Project: openshift-adp 🔻			
Workloads 🗸 🗸	Secrets			Create 👻
Pods				Key/value secret
Deployments	▼ Filter ▼ Name ▼ Search by name	Size		Image pull secret
DeploymentConfigs	Name 1 Type 1	S 1	Created 1	Source secret
StatefulSets	S builder-dockercfg-7g8ww kubernetes.io/do	ockercfg 1	Apr 11, 2024, 10:52 AN	Webhook secret
Secrets	S builder-token-rm4s kubernetes.io/se	rvice-account-token 4	Apr 11, 2024, 10:52 AN	From YAML

Edit key/value sec	ret
Key/value secrets let you inject variables.	sensitive data into your application as files or environment
Secret name *	
ontap-s3-credentials	
Unique name of the new secret	
Key *	
cloud	
Value	
	Browse
Drag and drop file with your val	e here or browse to upload it.
[default] aws_access_key_id= aws_secret_access_key=	
<ul> <li>Add key/value</li> </ul>	

To create a secret named sg-s3-credentials from the CLI you can use the following command.

# oc create secret generic sg-s3-credentials --namespace openshift-adp --from-file cloud=cloud-credentials.txt Where credentials.txt file contains the Access Key Id and the Secret Access Key of the S3 user in the following format: [default] aws\_access\_key\_id=< Access Key ID of S3 user> aws\_secret\_access\_key=<Secret Access key of S3 user>

• Next, to configure Velero, select Installed Operators from the menu item under Operators, click on OADP operator, and then select the **DataProtectionApplication** tab.

Home	>	Installed Operators				
Operators	•	Installed Operators are represented by C	IusterServiceVersions within this Namesp	ace. For more information, see the Und	erstanding Operators documentation 🗗. Or cr	reate an Operator and ClusterServiceVersion using th
OperatorHub		Operator SDK 🗹				
Installed Operators		Name   Search by name	7			
Workloads	>	Name 1	Managed Namespaces 1	Status	Last updated	Provided APIs
Virtualization	>	OADP Operator 1.3.0 provided by Red Hat	NS openshift-adp	Succeeded Up to date	Apr 11, 2024, 10:53 AM	BackupRepository : Backup
Networking	>					BackupStorageLocation DeleteBackupRequest View 11 more

Click on Create DataProtectionApplication. In the form view, provide a name for the DataProtection Application or use the default name.

Project: openshift-adp	•				
Installed Operators > Operators > OADP Operator 13.0 provided by R	ator details led Hat				Actions 👻
ServerStatusRequest	VolumeSnapshotLocation	DataDownload	DataUpload	CloudStorage	DataProtectionApplication
DataProtection	Applications				Create DataProtectionApplication

Now go to the YAML view and replace the spec information as shown in the yaml file examples below.

Sample yaml file for configuring Velero with ONTAP S3 as the backupLocation

```
spec:
 backupLocations:
    - velero:
        config:
          insecureSkipTLSVerify: 'false' ->use this for https
communication with ONTAP S3
          profile: default
          region: us-east-1
          s3ForcePathStyle: 'true' ->This allows use of IP in s3URL
          s3Url: 'https://10.61.181.161' ->Ensure TLS certificate for S3
is configured
        credential:
          key: cloud
          name: ontap-s3-credentials -> previously created secret
        default: true
        objectStorage:
          bucket: velero -> Your bucket name previously created in S3 for
backups
          prefix: container-demo-backup ->The folder that will be created
in the bucket
          caCert: <base64 encoded CA Certificate installed on ONTAP
Cluster with the SVM Scope where the bucker exists>
        provider: aws
  configuration:
    nodeAgent:
      enable: true
      uploaderType: kopia
      #default Data Mover uses Kopia to move snapshots to Object Storage
    velero:
      defaultPlugins:
        - csi ->This plugin to use CSI snapshots
        - openshift
        - aws
        - kubevirt -> This plugin to use Velero with OIpenShift
Virtualization
```

Sample yaml file for configuring Velero with StorageGrid S3 as the backupLocation

```
spec:
 backupLocations:
    - velero:
        config:
          insecureSkipTLSVerify: 'true'
          profile: default
          region: us-east-1 ->region of your StorageGrid system
          s3ForcePathStyle: 'True'
          s3Url: 'https://172.21.254.25:10443' ->the IP used to access S3
        credential:
          key: cloud
          name: sg-s3-credentials ->secret created earlier
        default: true
        objectStorage:
          bucket: velero
          prefix: demobackup
        provider: aws
 configuration:
    nodeAgent:
      enable: true
      uploaderType: kopia
    velero:
      defaultPlugins:
        - csi
        - openshift
        - aws
        - kubevirt
```

The spec section in the yaml file should be configured appropriately for the following parameters similar to the example above

#### backupLocations

ONTAP S3 or StorageGrid S3 (with its credentials and other information as shown in the yaml) is configured as the default BackupLocation for velero.

#### snapshotLocations

If you use Container Storage Interface (CSI) snapshots, you do not need to specify a snapshot location because you will create a VolumeSnapshotClass CR to register the CSI driver. In our example, you use Astra Trident CSI and you have previously created VolumeSnapShotClass CR using the Trident CSI driver.

#### **Enable CSI plugin**

Add csi to the defaultPlugins for Velero to back up persistent volumes with CSI snapshots. The Velero CSI plugins, to backup CSI backed PVCs, will choose the VolumeSnapshotClass in the cluster that has **velero.io/csi-volumesnapshot-class** label set on it. For this

- You must have the trident VolumeSnapshotClass created.
- · Edit the label of the trident-snapshotclass and set it to

velero.io/csi-volumesnapshot-class=true as shown below.



Ensure that the snapshots can persist even if the VolumeSnapshot objects are deleted. This can be done by setting the **deletionPolicy** to Retain. If not, deleting a namespace will completely lose all PVCs ever backed up in it.

```
apiVersion: snapshot.storage.k8s.io/v1
kind: VolumeSnapshotClass
metadata:
   name: trident-snapshotclass
driver: csi.trident.netapp.io
deletionPolicy: Retain
```

VolumeSnapshotClasses > VolumeSnapshotClass details	
vsc trident-snapshotclass	
Details YAML Events	
VolumeSnapshotClass details	
Name	
trident-snapshotclass	
Labels Ec	dit 🖋
velero.io/csi-volumesnapshot-class=true	
Annotations	
1 annotation 🖋	
Driver	
csi.trident.netapp.io	
Deletion policy	
Retain	

Ensure that the DataProtectionApplication is created and is in condition:Reconciled.

Project: o	penshift-adp 🝷							
Installed Op	DADP Operator detail DADP Operator 3.2 provided by Red Hat	ls						Actions 👻
Schedule	ServerStatusReq	uest VolumeSnapsho	otLocation	DataDownload	DataUpload	CloudStorage	DataProtection	Application
DataP Name	rotectionAppli	cations				I	Create DataProtecti	onApplication
Name	I	Kind I	Status	ī	Labels I	Last up	odated 1	
OPA v	velero-container- backup-ontap	DataProtectionApplication	Condition	Reconciled	No labels	🕲 Juli	15, 2024, 2:31 PM	I

The OADP operator will create a corresponding BackupStorageLocation. This will be used when creating a backup.

Project: openshift-adp 🔹					
Installed Operators > Operator d OADP Operator 1.3.2 provided by Red Ha	etails st				Actions 👻
kupRepository Backup	BackupStorageLocation	DeleteBackupRequest	DownloadRequest F	PodVolumeBackup	PodVolumeRestore
Name  Search by name Name	Kind I	Status 1	Labels 1	Create I	3ackupStorageLocation
ESI velero-container- backup-ontap-1	BackupStorageLocation	Phase: Available	app.kubernetes.io/componen app.kubernet=velero-contai app.kubernetes.io/m=oadp- app.kubernetes=oadp-oper	t=bsi I Jul 15, 2024, 2:3	IPM E
			openshift.io/oadp=True openshift.io/oadp-registry=Tr	rue	

# Creating on-demand backup for Apps in OpenShift Container Platform

This section outlines how to create on-demand backup for VMs in OpenShift Virtualization.

#### Steps to create a backup of an App

To create an on-demand backup of an app (app metadata and persistent volumes of the app), click on the **Backup** tab to create a Backup Custom Resource (CR). A sample yaml is provided to create the Backup CR. Using this yaml, the app and its persistent storage in the specified namespace will be backed up. Additional parameters can be set as shown in the documentation.

A snapshot of the persistent volumes and the app resources in the namespace specified will be created by the CSI. This snapshot will be stored in the backup location specified in the yaml. The backup will remain in the system for 30 days as specified in the ttl.

```
spec:
    csiSnapshotTimeout: 10m0s
    defaultVolumesToFsBackup: false
    includedNamespaces:
        - postgresql ->namespace of the app
    itemOperationTimeout: 4h0m0s
    snapshotMoveData: false
    storageLocation: velero-container-backup-ontap-1 -->this is the
    backupStorageLocation previously created when Velero is configured.
    ttl: 720h0m0s
```

Once the backup completes, its Phase will show as completed.

Installed C	Operators OADP 1.3.2 pro	<ul> <li>Operato</li> <li>ovided by</li> </ul>	rator details or Red Hat							Actions 👻
Details	YA	ML	Subscr <mark>i</mark> ption	Events	All instances	BackupReposito	ory Backup	BackupStorageLocation	DeleteBa	ckupReques
Backu	ups								Cr	eate Backup
Name	• s	earch by	name	1						
Name	e 1		Kind	1	Status	I	Labels 1	Last updated	1	
B ba	ackup-p	ostgresql	-ontaps3 Backu	p	Phase	Completed	velero.io/sto =ve	elero-container 😗 Jul 16, 2024	4, 10:01 AM	I

You can inspect the backup in the Object storage with the help of an S3 browser application. The path of the backup shows up in the configured bucket with the prefix name (velero/container-demo-backup). You can see the contents of the backup includes the volume snapshots, logs, and other metadata of the application.



In StorageGrid, you can also use the S3 console that is available from the Tenant Manager to view the backup objects.

lame	Size	Туре	Last Modified	Storage Class
<b>`</b> .				
backup-postgresql-ontaps3.tar.gz	384.66 KB	GZ File	7/16/2024 10:01:20 AM	STANDARD
velero-backup.json	3.30 KB	JSON File	7/16/2024 10:01:20 AM	STANDARD
backup-postgresql-ontaps3-csi-volumesnap	731 bytes	GZ File	7/16/2024 10:01:19 AM	STANDARD
backup-postgresql-ontaps3-csi-volumesnap	760 bytes	GZ File	7/16/2024 10:01:19 AM	STANDARD
backup-postgresql-ontaps3-resource-list.jso	823 bytes	GZ File	7/16/2024 10:01:19 AM	STANDARD
backup-postgresql-ontaps3-itemoperations.j	378 bytes	GZ File	7/16/2024 10:01:19 AM	STANDARD
backup-postgresql-ontaps3-volumesnapshot	29 bytes	GZ File	7/16/2024 10:01:19 AM	STANDARD
backup-postgresql-ontaps3-podvolumeback	29 bytes	GZ File	7/16/2024 10:01:19 AM	STANDARD
backup-postgresql-ontaps3-results.gz	49 bytes	GZ File	7/16/2024 10:01:19 AM	STANDARD
backup-postgresql-ontaps3-csi-volumesnap	429 bytes	GZ File	7/16/2024 10:01:19 AM	STANDARD
hackup-postaresal-ontans3-logs az	12 01 KB	G7 File	7/16/2024 10:01:19 AM	STANDARD

#### Creating scheduled backups for Apps

To create backups on a schedule, you need to create a Schedule CR.

The schedule is simply a Cron expression allowing you to specify the time at which you want to create the backup. A sample yaml to create a Schedule CR is shown below.

```
apiVersion: velero.io/v1
kind: Schedule
metadata:
    name: schedule1
    namespace: openshift-adp
spec:
    schedule: 0 7 * * *
    template:
        includedNamespaces:
            - postgresql
        storageLocation: velero-container-backup-ontap-1
```

The Cron expression 0 7 \* \* \* means a backup will be created at 7:00 every day. The namespaces to be included in the backup and the storage location for the backup are also specified. So instead of a Backup CR, Schedule CR is used to create a backup at the specified time and frequency.

Once the schedule is created, it will be Enabled.

Project: openshift-adp	•				
Installed Operators > Operator OADP Operator 1.3.2 provided by Rev	or details d Hat				Actions 👻
PodVolumeRestore Re	store Schedule	ServerStatusRequest	VolumeSnapshotLocation	DataDownload DataUpload	CloudStorage
4					•
Schedules				a	eate Schedule
Name 👻 Search by na	me				
Name 1	Kind 1	Status 1	Labels 1	Last updated	
S schedule1	Schedule	Phase: 🔮 En	abled No labels	🕲 Jul 16, 2024, 10:32 AM	ŧ

Backups will be created according to this schedule, and can be viewed from the Backup tab.

- 10	roject: ope <mark>nsh</mark> ift-a	adp 💌					
In	Stalied Operators > OADP Oper 132 provided	Operator details rator í by Red Hat					Actions 👻
ż	All instances	BackupRepository	Backup	BackupStorageLocation	DeleteBackupRequest	DownloadRequest	PodVolumeBackup
4							
F	Backups						Create Backup
E	Name	n by name	<i>I</i>	Status	Labels 1	Last updated	Create Backup
E	Name  Search Name  Search Name  Search Name  Search	n by name Kind I esql-ontaps3 Backup	2	Status I Phase: 📀 Completed	Labels 1 velero.io/sto =velero-contail	Last updated I ner	Create Backup

# Migrate an App from one cluster to another

Velero's backup and restore capabilities make it a valuable tool for migrating your data between clusters. This section describes how to migrate apps(s) from one cluster to another by creating a backup of the app in Object storage from one cluster and then restoring the app from the same object storage to another cluster.

#### **Prerequisites on Cluster 1**

- Astra Trident must be installed on the cluster.
- A trident backend and Storage class must be created.
- OADP operator must be installed on the cluster.
- The DataProtectionApplication should be configured.

Use the following spec to configure the DataProtectionApplication object.

```
spec:
  backupLocations:
    - velero:
        config:
          insecureSkipTLSVerify: 'false'
          profile: default
          region: us-east-1
          s3ForcePathStyle: 'true'
          s3Url: 'https://10.61.181.161'
        credential:
          key: cloud
          name: ontap-s3-credentials
        default: true
        objectStorage:
          bucket: velero
          caCert: <base-64 encoded tls certificate>
          prefix: container-backup
        provider: aws
  configuration:
    nodeAgent:
      enable: true
      uploaderType: kopia
    velero:
      defaultPlugins:
        - csi
        - openshift
        - aws
        - kubevirt
```

• Create an application on the cluster and take a backup of this application. As an example, install a postgres application.

[root@localhos	t ~]# oc	get nodes	5							
NAME	STATUS	ROLES		AGE	VERSION					
ocp6-master1	Ready	control-	plane, maste	r 3d13h	v1.27.15+6147	7456				
ocp6-master2	Ready	worker		3d12h	v1.27.15+6147	7456				
ocp6-master3	Ready	control-	plane,maste	r 3d13h	v1.27.15+6147	7456				
ocp6-worker1	Ready	worker		3d12h	v1.27.15+6147	7456				
ocp6-worker2	Ready	worker		3d12h	v1.27.15+6147	7456				
ocp6-worker3	Ready	control-	plane, maste	r 3d12h	v1.27.15+6147	7456				
[root@localhos	t ~]# he	1m install	postgresq1	bitnami/p	ostgresql -n po	ostgresqlcre	ate namespac	e^C		
[root@localhos	t ~]# oc	get pods	-n postgres	q1						
NAME	READY	STATUS	RESTARTS	AGE						
postgresql-0	1/1	Running	0	4h53m						ÿ
[root@localhos	t ~]# oc	get pvc -	n postgresq	1						1
NAME	ST	ATUS VOL	.UME			CAPACITY	ACCESS MODE	S STORAGECLASS	AGE	
data-postgresq	1-0 Bo	und pvc	-f7a3c772-0	e61-49cb-a	3d0-7c7b2ec87dd	:6 8Gi	RWO	ontap-nas	4h53m	li li
[root@localhos	t ~]# oc	get pv -n	postgresql							
NAME				CAPACITY	ACCESS MODES	RECLAIM POLIC	Y STATUS	CLAIM		STORAGECLASS
REASON AGE										
pvc-2e9e982f-5 4h55m	4a4-4e7b	-8eae-a589	e0d9d819	1Gi	RWO	Delete	Bound	trident/basic		ontap-nas
pvc-f7a3c772-0 4h53m	e61-49cb	-a3d0-7c7b	2ec87dc6	BGi	RWO	Delete	Bound	postgresq1Ødätätj Go to Set	collings to activate	ontap-nas Windows.
[root@localhos	t ~]# _									1

• Use the following spec for the backup CR:

```
spec:
    csiSnapshotTimeout: 10m0s
    defaultVolumesToFsBackup: false
    includedNamespaces:
        - postgresql
    itemOperationTimeout: 4h0m0s
    snapshotMoveData: true
    storageLocation: velero-sample-1
    ttl: 720h0m0s
```

Project: ope	nshift-adp	•				
Installed Opera OAI 1.4.0	ators > Operator DP Operator provided by Red	or details d Hat				Actions 👻
Repository	Backup	BackupStorageLocation	n DeleteBackupReques	t DownloadRequest	PodVolumeBackup	PodVolumeRes
Backups	6					Create Backup
Name 👻	Search by nar	me /				
Name 1			Kind 1		Status	
B backup			Backup		Activate Wind Go to Settings to a	OWS pleted ctivate Windows.

You can click on the **All instances** tab to see the different objects being created and moving through different phases to finally come to the backup **completed** phase.

A backup of the resources in the namespace postgresql will be stored in the Object Storage location (ONTAP S3) specified in the backupLocation in the OADP spec.

#### **Prerequisites on Cluster 2**

- Astra Trident must be installed on cluster 2.
- The postgresql app must NOT be already installed in the postgresql namespace.
- OADP operator must be installed on cluster 2, and the BackupStorage Location must be pointing to the same object storage location where the backup was stored from the first cluster.
- The Backup CR must be visible from the second cluster.

NAME	READY	STATUS	RESTARTS	AGE
trident-controller-6799cfb77f-8rzvk	6/6	Running	6	2d7h
trident-node-linux-7wvjz	2/2	Running	2	2d7h
trident-node-linux-8vvm2	2/2	Running	0	2d7h
trident-node-linux-bgs6f	2/2	Running	2	2d7h
trident-node-linux-njwb8	2/2	Running	0	2d7h
trident-node-linux-scqjl	2/2	Running	0	2d7h
trident-node-linux-swr69	2/2	Running	2	2d7h
trident-operator-b88b86fc8-7fk68 [root@localbost ~]#	1/1	Running	1	2d7h




1.4.0 provided	<b>ator</b> by Red Hat							Actions 🝷
Details YAML	Subscription	Events	All instances	BackupRepository	Backup	BackupStorageLocation	DeleteBackupRequest	DownloadRequest
3ackups								Create Backup
Name 🔻 Search	by name	1						
Name 1		Kind I		Status 1		Labels 1	Last updated 1	
					a.	(		

Restore the app on this cluster from the backup. Use the following yaml to create the Restore CR.

```
apiVersion: velero.io/v1
kind: Restore
apiVersion: velero.io/v1
metadata:
   name: restore
   namespace: openshift-adp
spec:
   backupName: backup
   restorePVs: true
```

When the restore is completed, you will see that the postgresql app is running on this cluster and is associated with the pvc and a corresponding pv. The state of the app is the same as when the backup was taken.

Project: ope	nshift-adp 💌					
Installed Opera OA 1.4.0	ators > Operator details DP Operator 0 provided by Red Hat					Actions 👻
eLocation	DeleteBackupRequest	DownloadRequ	est PodVolumeBackup	PodVolumeRestore	Restore Sc	hedule Server
Restore	S					Create Restore
Name 🔻	Search by name	Z				
Name 1			Kind		Status 1	
R restore			Restore		Activate Windo Phase: Comp Go to Settings to act	WS leted wate Windows.

1.4.0 pr	ovided by Red Hat					Actions 🝷
Location I	DeleteBackupRequest	DownloadRequest	PodVolumeBackup	PodVolumeRestore	Restore Sc	hedule Serve
Name 👻 S	Search by name	7			,	

# Restore an App from a backup

This section describes how to restore apps(s) from a backup.

# Prerequisites

To restore from a backup, let us assume that the namespace where the app existed got accidentally deleted.



To restore from the backup that we just created, we need to create a Restore Custom Resource (CR). We need to provide it a name, provide the name of the backup that we want to restore from and set the restorePVs to true. Additional parameters can be set as shown in the documentation. Click on Create button.

Pro	ect: openshift-adp 🔹						
Insta	OADP Operator 1.3.0 provided by Red Ha	etails					Actions 💌
est	DownloadRequest	PodVolumeBackup	PodVolumeRestore	Restore	Schedule	ServerStatusRequest	VolumeSnap •
Re	stores					C	reate Restore

```
apiVersion: velero.io/v1
kind: Restore
apiVersion: velero.io/v1
metadata:
   name: restore
   namespace: openshift-adp
spec:
   backupName: backup-postgresql-ontaps3
   restorePVs: true
```

When the phase shows completed, you can see that the app has been restored to the state when the snapshot was taken. The app is restored to the same namespace.

Proje	ect: openshift-adp 🔹						
Instal	Ied Operators         > Operator def           OADP Operator         1.3.0 provided by Red Hat	tails					Actions 👻
est	DownloadRequest	PodVolumeBackup	PodVolumeRestore	Restore	Schedule	ServerStatusRequest	VolumeSr
Re	stores						Create Restore
Nan	e Search by name	1					
Ν	lame 🌐	Kind 1		Status 🚦	Labe	is 1	
(	restore1	Restore		Phase: 🛇 Comp	oleted No la	bels	:

[root@localhos	t ~]#				
[root@localhos	t ~]# oc	get pods	-n postgre	sql	
No resources fo	ound in p	postgresql	namespace		
[root@localhos	t ~]# oc	get pods	-n postgre	sql	
NAME	READY	STATUS		RESTARTS	AGE
postgresql-0	0/1	Container	Creating	0	16s
[root@localhos	t ~]# oc	get pods	-n postgre	sql	
NAME	READY	STATUS	RESTARTS	AGE	
postgresql-0	0/1	Running	0	22s	
[root@localhos	t ~]# oc	get pods	-n postgre	sql	
NAME	READY	STATUS	RESTARTS	AGE	
postgresql-0	0/1	Running	0	29s	
[root@localhos	t ~]# oc	get pods	-n postgre	sql	
NAME	READY	STATUS	RESTARTS	AGE	
postgresql-0	1/1	Running	0	37s	
[root@localhos	t~]#				

To restore the App to a different namespace, you can provide a namespaceMapping in the yaml definition of the Restore CR.

The following sample yaml file creates a Restore CR to restore an App and its persistent storage from the postgresql namespace, to the new namespace postgresql-restored.

```
apiVersion: velero.io/v1
kind: Restore
metadata:
   name: restore-to-different-ns
   namespace: openshift-adp
spec:
   backupName: backup-postgresql-ontaps3
   restorePVs: true
   includedNamespaces:
        postgresql
   namespaceMapping:
        postgresql: postgresql-restored
```

When the phase shows completed, you can see that the app has been restored to the state when the snapshot was taken. The App is restored to a different namespace as specified in the yaml.

[root@localhos	t ~]# oc	get pods	-n postgres	ql
No resources f	ound in p	postgresql	namespace.	
[root@localhos	t ~]# oc	get pods	-n postgres	ql-restored
NAME	READY	STATUS	RESTARTS	AGE
postgresql-0	0/1	Running	0	19s
[root@localhos	t ~]# oc	get pods	-n postgres	ql-restored
NAME	READY	STATUS	RESTARTS	AGE
postgresql-0	0/1	Running	0	22s
[root@localhos	t ~]# oc	get pods	-n postgres	ql-restored
NAME	READY	STATUS	RESTARTS	AGE
postgresql-0	1/1	Running	0	36s
[root@localhos	t ~]#			

Velero provides a generic ability to modify the resources during restore by specifying json patches. The json patches are applied to the resources before they are restored. The json patches are specified in a configmap and the configmap is referenced in the restore command. This feature enables you to restore using different storage class.

In the example below, the app, during deployment uses ontap-nas as the storage class for its persistent volumes. A backup of the app named backup-postgresql-ontaps3 is created.

Project: postgresql 🔹		
tata-postgresql-0 ⊗ ∞unit		
Details YAML Events VolumeSnapshots		
PersistentVolumeClaim details		
8 GIB		
$\bigcirc$		
lame		Status © Bound
na-posigiesq=0		Requested capacity
iametopoce IS postgresol		8 GiB
nbels	Edit 🖊	Capacity 8 GiB
appluberretexiq/componentsprimary) (appluberretexiq/instancespostgress)) (appluberretexiq/inamespostgress)) (velexiq/backup-namesbackup-postgress)-postgress) (velexiq/instancespostgress)		Used
(veleoiq/volume-snapshot-namesvelen-date-postgesof-0-74ph)		13 MB
unotations		Access modes ReadWiteOnce
and an all and a second		Volume mode
la selector		Storana Clarear
Created at		Storagecesses ontap-nas-on-#300e9u25
9 Dia 16, 2024, 209 PM		PersistentVolumes
Vaner No owner		
oject openshift-adp 🔹		
Inited Operators > Operator details		
1.32 provided by Red Hat		

Details YAML Subscription	Events All instances	BackupRepository	Backup	BackupStorageLocation	DeleteBackupRequest	DownloadRequest	PodVolumeBackup	PodVolumeRestore	Restore
Backups									C
Nome • Search by name	X								
Nome	Kind 1		5	tetus	Labels 1		Lest updated	1	

Simulate a loss of the app by uninstalling the app.

To restore the VM using a different storage class, for example, ontap-nas-eco storage class, you need to do the following two steps:

#### Step 1

Create a config map (console) in the openshift-adp namespace as follows: Fill in the details as shown in the screenshot: Select namespace : openshift-adp Name: change-ontap-sc (can be any name) Key: change-ontap-sc-config.yaml: Value:

```
version: v1
resourceModifierRules:
- conditions:
    groupResource: persistentvolumeclaims
    resourceNameRegex: "data-postgresql*"
    namespaces:
    - postgresql
    patches:
    - operation: replace
    path: "/spec/storageClassName"
    value: "ontap-nas-eco"
```

onfig maps hold key-value pairs that can be used in pods to read application configuration.	
Configure via:   Form view O YAML view	
lame *	
change-ontap-sc	
unique name for the ConfigMap within the project	
) Immutable	
nmutable, if set to true, ensures that data stored in the ConfigMap cannot be updated	
lata	
lata contains the configuration data that is in UTF-8 range	
	C Remove key/value
ley *	
change-ontap-sc.yaml	
/a)110	
aue	Provide
	browse
rag and drop file with your value here or browse to upload it.	
version: vi	
- conditions:	
groupResource: persistentvolumeclaims	
resourceNameRegex: "data-postgresql*"	
namespaces:	
- postgresql	
patches:	
- operation: replace	
path: "/spec/storageClassName"	
value: "ontap-nas-eco"	

The resulting config map object should look like this (CLI):

[root@localhos	st ~]# kubectl describe cm/change-ontap-sc -n openshift-adp
Name:	change-ontap-sc
Namespace:	openshift-adp
Labels:	<none></none>
Annotations:	<none></none>
Data	
====	
change-ontap-s	sc.yaml:
version: v1	
resourceModifi	ierRules:
- conditions:	
groupResc	ource: persistentvolumeclaims
resource	NameRegex: "data-postgresql*"
namespace	25:
- postgre	esql
patches:	
- operation:	: replace
path: "/sp	pec/storageClassName"
value: "or	ntap-nas-eco"
BinaryData	
====	
Events: <none< td=""><td>2&gt;</td></none<>	2>
[root@localhos	st ~]#

This config map will apply the resource modifier rule when the restore is created. A patch will be applied to replace the storage class name to ontap-nas-eco for all persistent volume claims starting with rhel.

# Step 2

To restore the VM use the following command from the Velero CLI:

```
#velero restore create restore1 --from-backup backup1 --resource
-modifier-configmap change-storage-class-config -n openshift-adp
```

The app is restored in the same namespace with the persistent volume claims created using the storage class ontap-nas-eco.

[root@localhos	t~]#	oc get po	ods -n postgre	sql					
NAME	READY	STATUS	6 RESTARTS	AGE					
postgresql-0	1/1	Runnin	ng 0	11m					
[root@localhos	t~]#	oc get pv	/c -n postgres	q1					
NAME		STATUS	VOLUME			CAPACITY	ACCESS MODES	STORAGECLASS	AGE
data-postgresq	1-0	Bound	pvc-33526ea4-	37c2-4180-a9f6-	fb47aea3b4e2	8Gi	RWO	ontap-nas-eco	11m
[root@localhos	t ~]#								

### Deleting backups and restores in using Velero

This section outlines how to delete backups and restores of Apps in OpenShift container platform using Velero.

#### List all backups

You can list all Backup CRs by using the OC CLI tool or the Velero CLI tool. Download the Velero CLI as given in the instructions in the Velero documentation.

[root@localhost ~]# oc get	backups -n c	penshift;	-adp				
NAME	AGE						
backup-postgresql-ontaps3	23h						
backup2	26s						
schedule1-20240717070005	6h42m						
[root@localhost ~]# velero	get backups	-n opensi	hift-adp				i i
NAME	STATUS	ERRORS	WARNINGS	CREATED	EXPIRES	STORAGE LOCATION	SELECTOR
backup-postgresql-ontaps3	Completed	0	0	2024-07-16 10:01:08 -0400 F	DT 29d	velero-container-backup-ontap-1	<none></none>
backup2	Completed	0	0	2024-07-17 09:42:32 -0400 E	DT 29d	velero-container-backup-ontap-1	<none></none>
schedule1-20240717070005	Completed	0	0	2024-07-17 03:00:05 -0400 F	DT 29d	velero-container-backup-ontap-1	<none></none>
[root@localhost ~]# _							

#### Deleting a backup

You can delete a Backup CR without deleting the Object Storage data by using the OC CLI tool. The backup will be removed from the CLI/Console output. However, since the corresponding backup is not removed from the object storage, it will re-appear in the CLI/console output.

[root@localhost ~]# oc del@	ete backup backup2 -n openshift-adp
backup.velero.io "backup2"	deleted
[root@localhost ~]# oc get NAME	backups -n openshift-adp AGE
backup-postgresql-ontaps3	23h
schedule1-20240717070005	6h49m
[root@localhost ~]# oc get NAME	backups -n openshift-adp AGE
backup-postgresql-ontaps3	23h
backup2	24s
schedule1-20240717070005	6h50m
[root@localhost ~]# _	

If you want to delete the Backup CR AND the associated object storage data, you can do so by using the Velero CLI tool.

[root@localhost ~]# velero	get backups	-n opensk	nift-adp				1
NAME	STATUS	ERRORS	WARNINGS	CREATED	EXPIRES	STORAGE LOCATION	SELECTOR
backup-postgresql-ontaps3	Completed	0	0	2024-07-16 10:01:08 -0400 EDT	29d	velero-container-backup-ontap-1	<none></none>
backup2	Completed	0	0	2024-07-17 09:42:32 -0400 EDT	29d	velero-container-backup-ontap-1	<none></none>
schedule1-20240717070005	Completed	0	0	2024-07-17 03:00:05 -0400 EDT	29d	velero-container-backup-ontap-1	<none></none>
[root@localhost ~]# velero	delete backu	p backup2	2 -n openshi	ft-adp			
Are you sure you want to co	ontinue (Y/N)	? Y					
Request to delete backup "b	backup2" subm	itted suc	cessfully.				
The backup will be fully de	eleted after	all assoc	iated data	(disk snapshots, backup files,	restores) a	are removed.	
[root@localhost ~]# velero	get backups	-n opensh	nift-adp				
NAME	STATUS	ERRORS	WARNINGS	CREATED	EXPIRES	STORAGE LOCATION	SELECTOR
backup-postgresql-ontaps3	Completed	0	0	2024-07-16 10:01:08 -0400 EDT	29d	velero-container-backup-ontap-1	<none></none>
schedule1-20240717070005	Completed	0	0	2024-07-17 03:00:05 -0400 EDT	29d	velero-container-backup-ontap-1	<none></none>
[root@localhost ~]#							

# **Deleting the Restore**

You can delete the Restore CR Object by using either the OC CLI or the Velero CLI

[root@loca	lhost ~]# velero get restore	e -n openshif	t-adp								
NAME	BACKUP	STATUS	STARTED	COMPLETED	ERRORS	WARNINGS	CREATED		SELECTOR		
restore	backup-postgresql-ontaps3	Completed	2024-07-16 14:59:22 -0400 EDT	2024-07-16 14:59:45 -0400 EDT		10	2024-07-16 14:59:22 -	0400 EDT	<none></none>		
restore1	backup-postgresql-ontaps3	Completed	2024-07-16 16:36:37 -0400 EDT	2024-07-16 16:36:59 -0400 EDT			2024-07-16 16:36:37 -	0400 EDT	<none></none>		
[root@loca	lhost ~]# velero restore del	lete restore1	-n openshift-adp								
Are you sure you want to continue (Y/N)? Y											
Request to delete restore "restore1" submitted successfully.											
The restore will be fully deleted after all associated data (restore files in object storage) are removed.											
(root@localhost ~]# velero get restore -n openshift-adp											
NAME	BACKUP	STATUS	STARTED	COMPLETED	ERRORS	WARNINGS	CREATED		SELECTOR		
restore	backup-postgresql-ontaps3	Completed	2024-07-16 14:59:22 -0400 EDT	2024-07-16 14:59:45 -0400 EDT	0	10	2024-07-16 14:59:22 -0	400 EDT	<none></none>		
[root@local	lhost ~]#										
[root@local	<pre>lhost ~]# oc delete restore</pre>	restore -n o	penshift-adp								
restore.ve	lero.io "restore" deleted										
(root@localhost ~]# oc get restore -n openshift-adp											
No resources found in openshift-adp namespace.											
root@localhost ~]# velero get restore -n openshift-adp											
[root@local	lhost ~]#							Activate V	Vindows		
								10100 1			

# **Copyright information**

Copyright © 2024 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, 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) (FEB 2014).

# **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.