

May 11, 2022

zadara

# CREATING HIGHLY AVAILABLE HYBRID CLOUDS

WITH ZADARA, VMWARE & VEEAM

Steve Costigan,  
International Solutions Architect Director

# TABLE OF CONTENTS

INTRODUCTION .....	4
ENVIRONMENT .....	4
ZADARA® .....	4
VEEAM® AVAILABILITY SUITE™ .....	4
VMWARE CLOUD ON AWS .....	5
THE CHALLENGE .....	5
IMPLEMENTING HYBRID CLOUD STORAGE.....	5
DEPLOYMENT OPTIONS .....	6
DEPLOYMENT WALK THROUGH PLANNING.....	9
CREATING VMWARE CLOUD ON AWS SOFTWARE DEFINED DATA CENTER (SDDC).....	10
ZADARA STORAGE SERVICES DEPLOYMENT .....	14
VPSA NAS INTEGRATIONS.....	17
ADDING VEEAM TO OUR HYBRID CLOUD .....	18
VEEAM BACKUP REPOSITORY .....	20
VEEAM BACKUP JOB CREATION .....	23
VMWARE GUEST ISCSI VM SNAPSHOT PROTECTION.....	29
ZADARA VSS HW PROVIDER DEPLOYMENT .....	29
VEEAM / ZADARA VSS HW PROVIDER INTEGRATION.....	37
SUMMARY .....	39

# REVISION HISTORY

VERSION	DATE	DESCRIPTION
1.0	22 Aug 2019	Initial Release

# INTRODUCTION

This guide is intended to provide an insight into the capabilities of utilising Veeam and Zadara to provide a comprehensive suite of data protection capabilities in AWS, VMware Cloud on AWS, VMware ESXi in any deployment, Private Hybrid or Public Cloud.

## ENVIRONMENT

- Single Node VMware Software Defined Data Center in AWS using NSX-T
- Veeam Backup and Replication 9.5 Update 4 on SDDC with Zadara via NSX-T
- VMware SDDC Version 1.5+
- Zadara VPSA All Flash Array, Hybrid Array and Object Storage
- Zadara VSS Hardware Provider

## ZADARA®

Zadara is agile, secure enterprise data storage built for hybrid IT. We help organizations eliminate the cost and complexity traditionally associated with data storage, by providing enterprise data storage solutions as a fully-managed service, with a 100%-uptime guarantee and OpEx consumption-based pricing. Zadara uses industry-standard hardware and patented Zadara software to deliver the power of enterprise-class data storage and management – with the convenience of the cloud. Any data type. Any protocol. Any location. Zadara is available via public clouds, managed service providers, data centers, colocation partners, and on premises in customers' data centers.

Zadara's storage services deliver high performance all-flash, flash cache-accelerated hard disk, and object storage services with unique multi-tenant resource isolation. This ability to deliver multiple storage tiers with consistent and predictable performance, even in shared environments, is the perfect complement to Veeam Availability Suite and VMware Clouds.

## VEEAM® AVAILABILITY SUITE™

Veeam Availability Suite delivers Availability for ALL workloads – virtual, physical and cloud – from a single management console. Veeam helps organizations meet today's service-level and data center Availability objectives for the growing enterprise. Veeam Availability Suite is the premier solution providing the superior data protection capabilities of Veeam Backup & Replication™ paired with the advanced monitoring and reporting of Veeam ONE™ for holistic coverage of all workloads.

# VMWARE CLOUD ON AWS

VMware Cloud on AWS is an integrated cloud offering jointly developed by AWS and VMware delivering a highly scalable, secure and innovative service that allows organizations to seamlessly migrate and extend their on-premises VMware vSphere-based environments to the AWS Cloud running on next-generation Amazon Elastic Compute Cloud (Amazon EC2) bare metal infrastructure. VMware Cloud on AWS is ideal for enterprise IT infrastructure and operations organizations looking to migrate their on-premises vSphere-based workloads to the public cloud, consolidate and extend their data center capacities, and optimize, simplify and modernize their disaster recovery solutions.

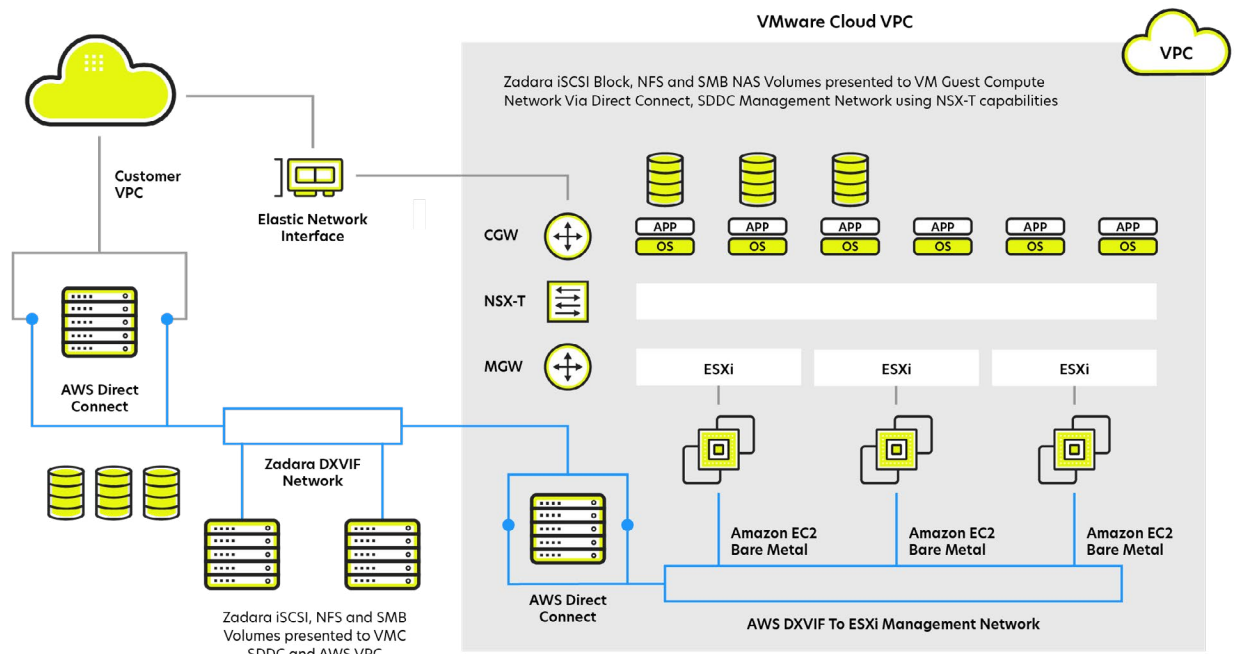
## THE CHALLENGE

The objective is to create a complete, production-ready hybrid cloud environment that offers high-performance, a familiar and proven operating environment, and complete data backup, recovery, and replication capabilities. Achieving these ambitious goals with legacy solutions would generally require extensive custom integration of multiple complex packages. However, by leveraging Zadara Enterprise Storage-As-A-Service, VMware Cloud, and Veeam Availability Suite, our goals can be achieved by a knowledgeable practitioner in just a few hours.

## IMPLEMENTING HYBRID CLOUD STORAGE

For this exercise, we will simultaneously present both iSCSI block and NAS (NFS/SMB) to both the VMware Cloud Virtual Private Cloud (VPC) and an independent VPC comprised of EC2 guests. Backups targeted at object storage can be directed to either Zadara Object Storage or S3.

# DEPLOYMENT OPTIONS



For testing purposes a single node SDDC VMC was defined in AWS London. Setup consists of a single node with 10TB of vSAN SSD storage, split into:

- A vSAN Datastore for vCenter, NSX, and other VMware-provided management VMs
- A Workload Datastore for hosting user VM's

First the environment must be defined in the intended location, in this case AWS London. There are several steps required for this, including defining the management network and the AWS VPC to attach to. You also need access to your AWS environment with permissions to launch a Cloud Formation script, so planning ahead is essential to a smooth deployment.

Once deployment has been completed, which usually takes 120 minutes from submitting the request, you can manage the security setup for the VMC SDDC. This includes allowing access to the vCenter environment via the public IP defined to it, enabling access from the VPC EC2 instances, and enabling access via AWS Direct Connect, if applicable.

Once this is complete you can test accessibility to your VMC vCenter. At this point you may want to enable Direct Connect to your VMware Cloud management network, this is where the up-front planning will come in as it is imperative to avoid accidentally overlapping address ranges. If accessing the VMware Cloud management network via Direct Connect, configure the interfaces to be offered to the VMware-specific account supplied as part of the VMware Cloud install.

This is where it gets interesting. In the current VMware Cloud on AWS, guest VMs are stored on the SSD pool (10 TB raw per ESXi host, approximately 15 TB usable for a 3-node cluster). This storage provides very good performance, but has two key shortcomings:

- Only one vSAN environment is supported. Therefore all backup storage traffic will compete directly with production workloads. This is not good.
- There is no option to provide multiple performance tiers to align application requirements, performance, and cost.

If we don't want to have backup traffic impacting production and/or we want the ability to provide applications with storage capacity that aligns with business requirements, what are the options?

Prior to SDDC v1.5, we were limited by NSX-V functionality to backup applications that could track and backup VMware environments. However, these solutions tend to require complex configuration, particularly when implementing the best practice of automatic offsite backup creation is a design point.

One possible approach would be to deploy a Linux VM, connected to EBS, and running SSH and Perl to expose an NFS export. However this is complex and also creates a single point of failure (the NFS export). Furthermore, going this route will mean that everything is going to be located in the same Availability Zone (AZ) and Data Center, which is a suboptimal backup strategy. Therefore, S3 must also be utilised for protection to meet the requirements, adding yet another layer of complexity. You could address availability concerns by keeping another copy in another AZ, but now you are adding complexity; 2 Linux nodes, 2 sets of EBS. And this is going to need to be SSD as the magnetic option has a 1TB limit. The backup costs are rising precipitously and complexity is increasing geometrically, greatly increasing the risk of failure.

Fortunately, Veeam is a certified VMware Cloud backup solution and listed in the VMware Cloud Marketplace, so we can use it in the VMware Cloud and/or the EC2 VPC, but we still need to connect Veeam to appropriate storage. The simple solution is to present Zadara iSCSI block or SMB NAS volumes to the EC2 instance. This way, even initial backups are automatically stored outside of the AWS data center housing the VMware Cloud and EC2 VPCs.

With NSX-T, we now have an even simpler option; we can now use Zadara's enterprise storage-as-a-service to provide an iSCSI Block volume to a Veeam Backup Server within VMware Cloud environment. This enables a backup repository which has high availability built in, enhancing the availability and integrity of backups.

To further enhance data availability, the backups created on Zadara storage can be replicated to any other Zadara virtual private array or, alternatively via Zadara Backup to Object Storage, to any S3 / Swift compatible Object storage array including Zadara's object storage service or Amazon S3 itself.

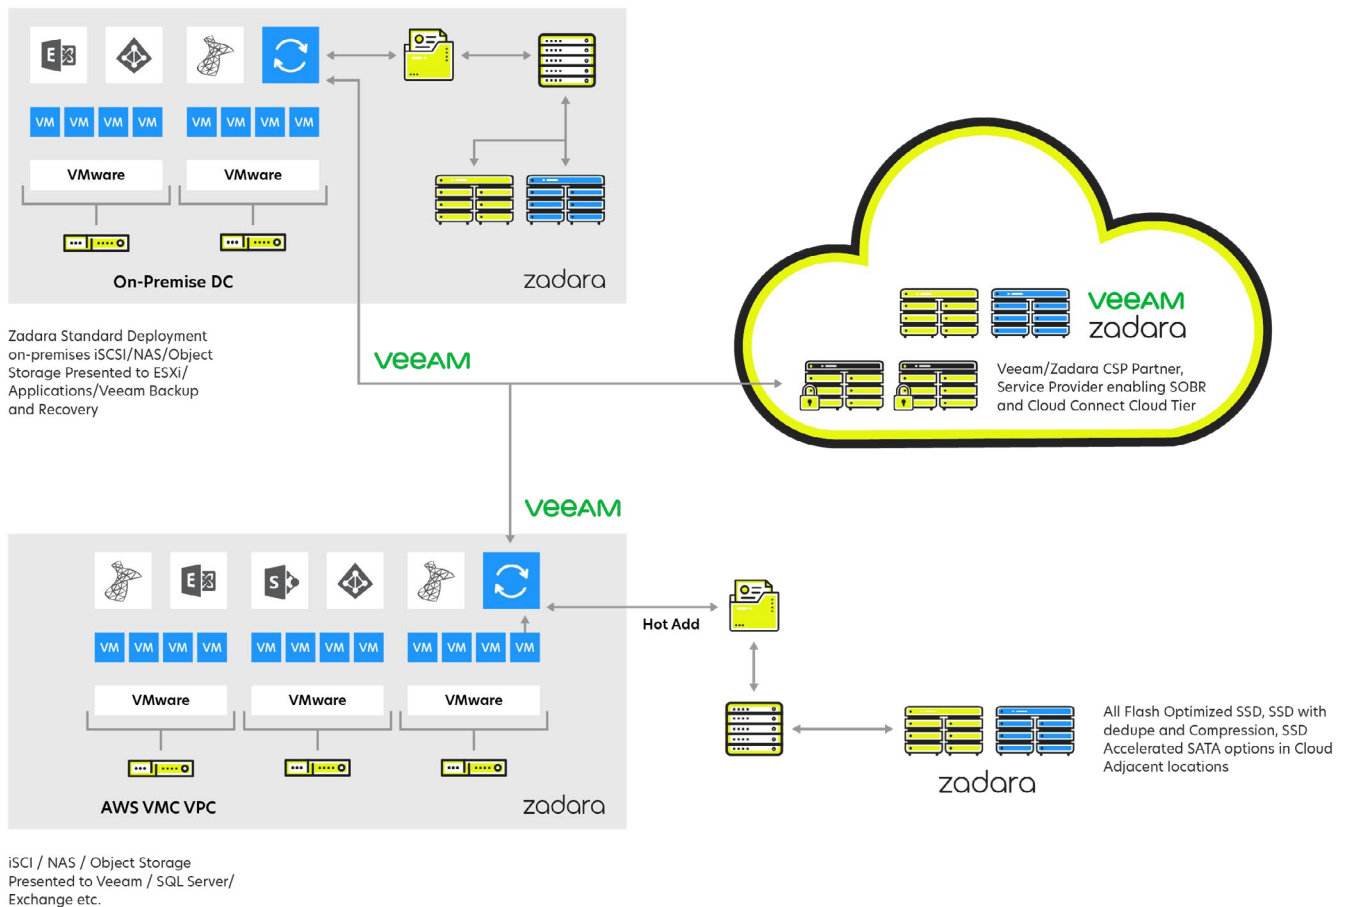
In addition to the object storage service, Zadara storage also offers NAS with NFS and SMB services. Either or both of these services can be used for primary data as well as backup. For this project, we will implement the Active Directory-integrated SMB NAS service. Physics and locking permitting, our service will be capable of serving the same data to applications and users within our VMware Cloud, the EC2 VPC and on-premises.

Consider the possibilities! VDI running in AWS, accessing servers in VMware Cloud, with data persisted on highly available shared storage, offering client-based recovery using the standard "Previous Versions" capability built into every Windows client. And all the shared data can be automatically protected with Zadara's optional integrated McAfee antivirus.

Veeam and Zadara integrated into a hybrid cloud solution, data automatically backed up to S3, set it forget it, but know it will be where you need it when you need it. Need to support Windows and Linux? No problem. Support NFS clients with the same replication capabilities as SMB. Support access to the same files with NFS and SMB. All wrapped up with an integrated backup solution using Windows and Linux Agents from Veeam.

The last piece of our project will be to support Windows Failover Server Clustering (WFSC). For deployments limited to vSAN, there is good news and bad news. The good news is that vSAN 6.7 introduced support for failover clusters via the iSCSI target Service, and the latest VMware Cloud on AWS SDDC enables a shared VMDK capability. The bad news is that the current offerings do not provide snapshots, let alone a Volume Shadow Copy Service (VSS) provider to ensure application consistent snapshots.

Implementing hybrid cloud storage from Zadara, addresses these limitations and provides additional availability options like remote mirroring and multi-zone HA, while enabling all of the benefits of VMware Cloud on AWS and providing a consistent service, independent of how (and where) your Zadara storage is physically deployed.

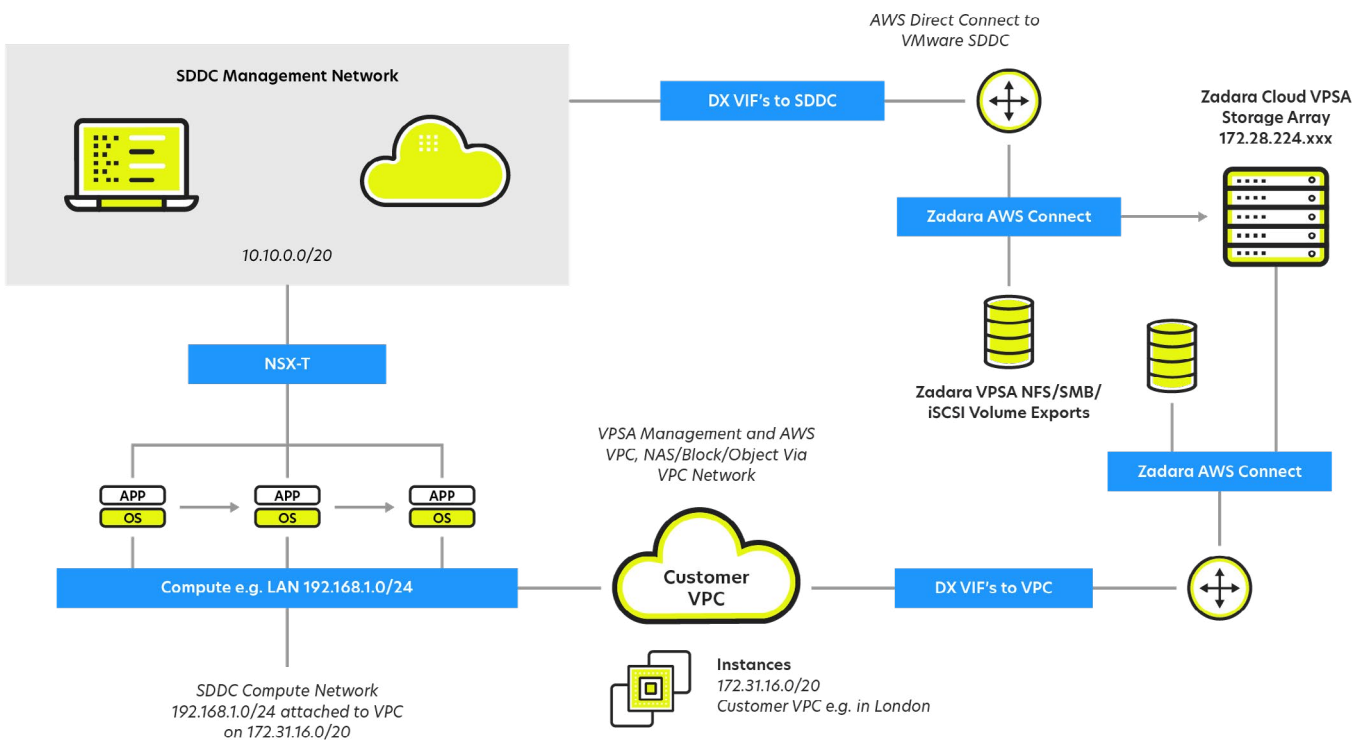




# DEPLOYMENT WALK THROUGH PLANNING

By default the SDDC will utilise a 10.0.0.0/20 network for management functions, ensure that this is not in conflict with your environment or your VPC environment. Check out the Zadara Setup Guide <http://guides.zadara.com/guides.zadara.com/vpsa-guide/1811/getting-started.html>. Zadara virtual private arrays will typically be on a 172.x.x.x network, depending upon the region. However Zadara supports custom networks and this can be set to a customer specific range if required.

Below is a typical top-level Zadara-to-SDDC-to-VPC configuration.



**Note:** SDDC VM's require NSX-T to communicate via DX connections, VM Compute cannot talk to DX routes via VPC, but can talk to EC instances via ENI Interfaces.

# CREATING VMWARE CLOUD ON AWS SOFTWARE DEFINED DATA CENTER (SDDC)

- To create an organization account and request a new SDDC, head over to <https://cloud.vmware.com/vmc-aws>
- Create your SDDC via the Creation Wizard.

The screenshot shows the 'Create Software-Defined Data Center (SDDC)' wizard. The current step is '1. SDDC Properties', which involves giving the SDDC a name, choosing a size, and specifying the AWS region. The configuration is as follows:

Field	Value
AWS Region	EU West (London) - More regions coming later
Deployment	Single Host (Selected), Multi-Host, Stretched Cluster
SDDC Name	Zadara NFS POC
Number of Hosts	1 (Note: 1-host SDDCs expire in 30 days. <a href="#">LEARN MORE</a> )
Host Capacity	2 Sockets, 36 Cores, 512 GB RAM, 10.7 TB Storage
Total Capacity	2 Sockets, 36 Cores, 512 GB RAM, 10.7 TB Storage

Below the configuration fields is a 'NEXT' button. At the bottom of the wizard, there are three additional steps listed:

- 2. Connect to AWS: Specify the AWS account that you want to connect your SDDC with.
- 3. VPC and subnet: Specify the VPC and the subnet to connect to your AWS account.
- 4. Configure Network: Management Subnet (optional).

- Link to your AWS Account. Note that you must have Cloud formation rights in your AWS Account.

The screenshot shows the 'Create Software-Defined Data Center (SDDC)' wizard at step '2. Connect to AWS'. The SDDC name is 'Zadara NFS POC - 1 Hosts - EU West (London)'. The instruction is to 'Specify the AWS account that you want to connect your SDDC with'. A dropdown menu shows 'Connect to a new AWS account'. Below this, there are instructions and a button to open the AWS console.

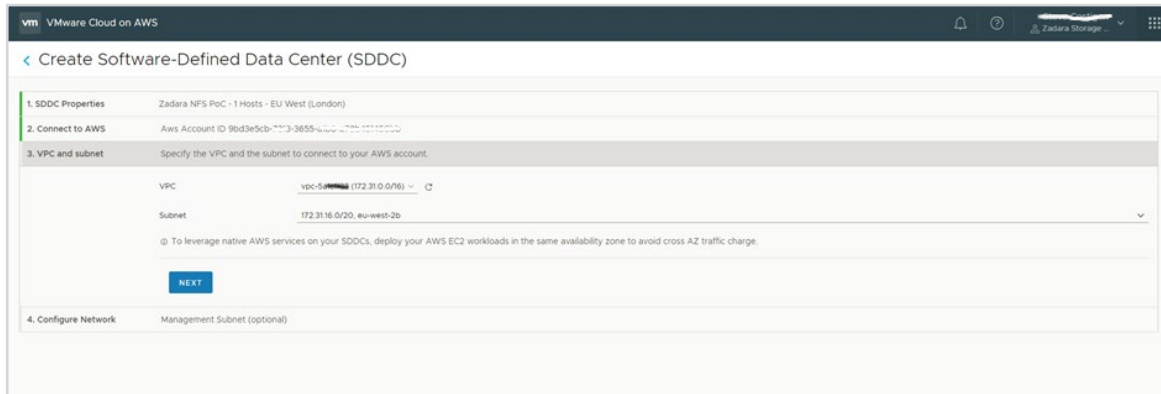
Instructions:

1. Make sure that you have an AWS account. [Click here to create one if needed.](#)
2. Click on the "OPEN AWS CONSOLE WITH CLOUDFORMATION TEMPLATE" button below to launch AWS in a separate browser tab. Each template is time-bounded for security reasons. To restart the task, simply click the button again.  
VMware has defined a CloudFormation template to connect your AWS account. This template creates the IAM roles necessary to allow communication between your SDDC and your AWS account.  
[What is CloudFormation?](#)  
[What are IAM roles?](#)

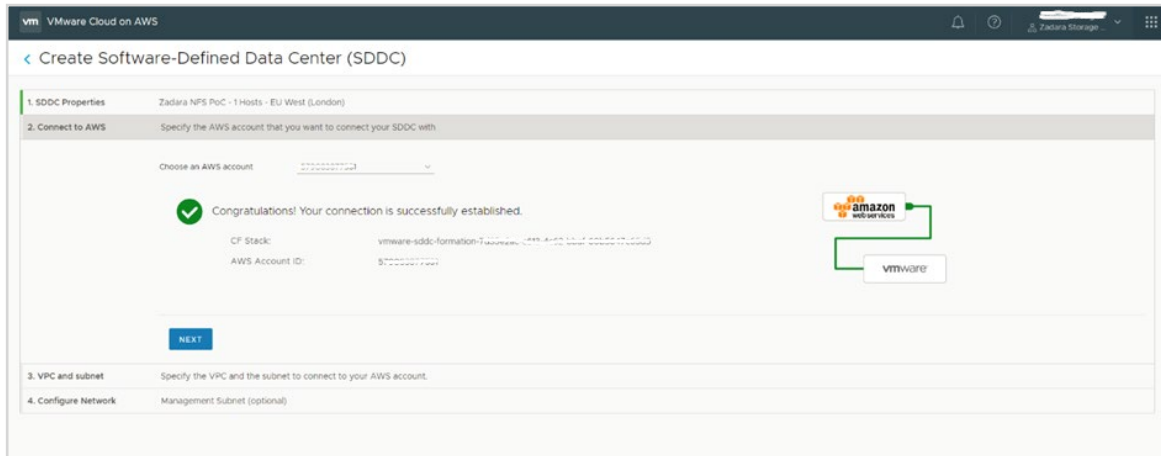
[OPEN AWS CONSOLE WITH CLOUDFORMATION TEMPLATE](#)

3. Log into the AWS account you want to connect with the VMware Cloud on AWS service.
4. On the Create Stack page, check "I acknowledge that AWS CloudFormation might create IAM resources" and click "Create".
5. When the stack creation is completed, close the AWS window and go back to VMC to finish your workflow.

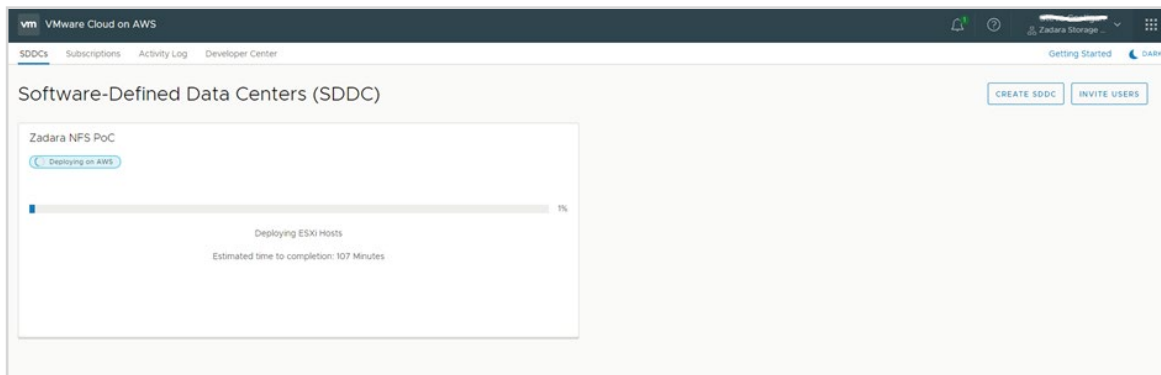
- Identify the VPC and subnet for your environment that you want to connect to your VMC SDDC.

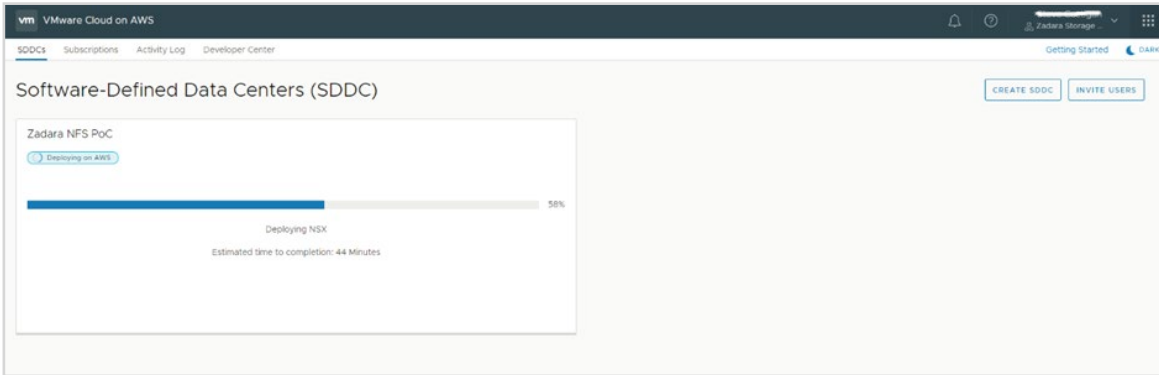
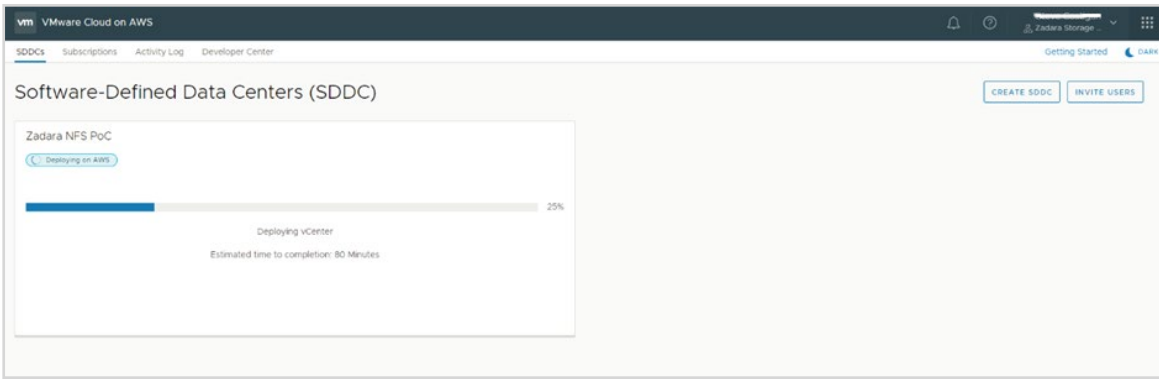


- Once Connection has been successfully established, you will see the following success message:

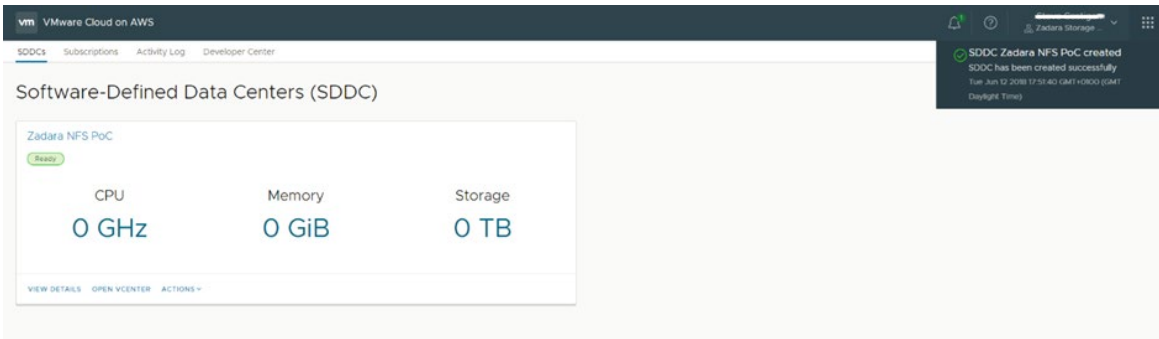


- Once All Networking information has been submitted and the request to create has been completed you will see the following screens:





- Upon completion you will see the following:



- When logging into the VMware in AWS console, you will see the following for each SDDC that you have defined:

**Software-Defined Data Centers (SDDC)**

Zadara-NSX-T-PoC EU (London)

Ready Expires in 16 days

CPU	Memory	Storage
83 GHz	512 GiB	10 TB

VIEW DETAILS OPEN VCENTER ACTIONS

- View Details displays a summary and allows you to set the networking firewall rules.

BACK TO LIST Zadara-NSX-T-PoC EU (London) ACTIONS OPEN VCENTER

Summary Networking & Security Add Ons Troubleshooting Settings Support

Capacity and Usage | eu-west-2b

Hosts	CPU	Memory	Storage
1	83 GHz	512 GiB	10 TB

- Networking and Security displays an overview of connections and how they are related to various network environments.

Zadara-NSX-T-PoC EU (London)

Summary Networking & Security Add Ons Troubleshooting Settings Support

**Overview**

Network Segments > VPN NAT Security Edge Firewall Distributed Firewall Inventory Groups Services Tools IPFIX Port Mirroring System DNS Public IPs Direct Connect Connected VPC

Overview

SDDC Public IP: 35.178.218.205

**Management Gateway** vCenter NSX

Appliance Subnet: 10.10.12.0/22  
Infrastructure Subnet: 10.10.0.0/20

11 Edge Firewall Rules, 9 Groups

**Compute Gateway** Workloads

2 Segments, 6 Edge Firewall Rules & 0 Distributed Firewall Rules, 2 Groups, 0 Public IPs

Internet

On Prem

Amazon VPC wpc-5atcfc3s

No IPsec VPN Configured


Direct Connect

Note the Public IP is the Public IP of the vCenter mapped by VMware Cloud on AWS for management. The management network is defined previously, during SDDC setup. In the Security Tab you need to define access rules for management and access to vCenter. By default, everything is denied. Create new objects for specific IP's to enable access and then grant access to vCenter, for instance:

⋮	Matt_VC	Matt_VC	vCenter	ICMP (ALL ICMP), SSO (TCP 7444), HTTPS (TCP 443) <a href="#">show less</a>	Allow	Disabled
---	---------	---------	---------	--	-------	----------

At this point you can begin setting up the SDDC environment and begin the process of requesting the Zadara virtual private array that we are going to connect.

Before requesting the Zadara storage service, we need to get the AWS Account ID for the SDDC. Click on Direct Connect in the SDDC Networking and Security tab.



## ZADARA STORAGE SERVICES DEPLOYMENT

For the next step, head over to <https://manage.zadarastorage.com> and request a new AWS VPSA to be created in the AWS region containing the just-deployed SDDC. Alternatively, you can sign in to your AWS Account and make the request via marketplace.

[https://aws.amazon.com/marketplace/pp/B07K4WPL9X?qid=1564583550014&sr=0-1 &ref\\_=srh\\_res\\_product\\_title](https://aws.amazon.com/marketplace/pp/B07K4WPL9X?qid=1564583550014&sr=0-1 &ref_=srh_res_product_title).

Complete the details in the signup form or once complete send an email to [support@zadarastorage.com](mailto:support@zadarastorage.com). Document the virtual array name, your account name, and the AWS Account ID that you want the Direct Connect Virtual Interfaces (DXVI) offered to, as well as the BGP ASN assigned to your environment.

**Note:** The AWS VPC Account ID and the AWS VMware Cloud Account ID's are different. You need to supply the VMware-supplied account details.

Once they have been offered, head back into the SDDC Direct Connect pane and accept the 2 virtual interfaces.

After a few minutes you should see something similar to this, showing the learned and advertised routes from the SDDC and the Zadara sides:

Virtual Interface Name	Virtual Interface ID	Direct Connect ID	State	BGP Status	
Zadara_AWS_SDDC	dxvif-fg...	dxcon-fg...	Attached	Up	DELETE
Zadara_AWS_SDDC	dxvif-fg...	dxcon-fg...	Attached	Up	DELETE

REFRESH

Advertised BGP Routes

192.168.1.0/24 192.168.100.0/24 10.10.12.0/22

Learned BGP Routes

172.28.224.0/22 10.51.0.0/16 10.52.0.0/16 10.53.0.0/16

You will also need to configure rules to allow access between the VM Networks and the VPC's, the VM Network and the Direct Connect allowing specific services as required to the VPSA Arrays.

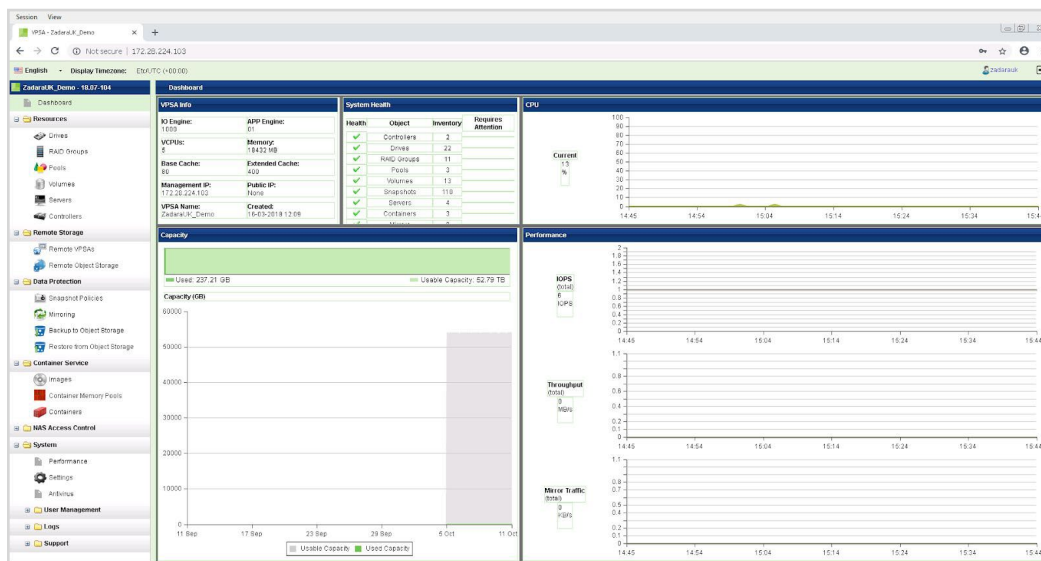
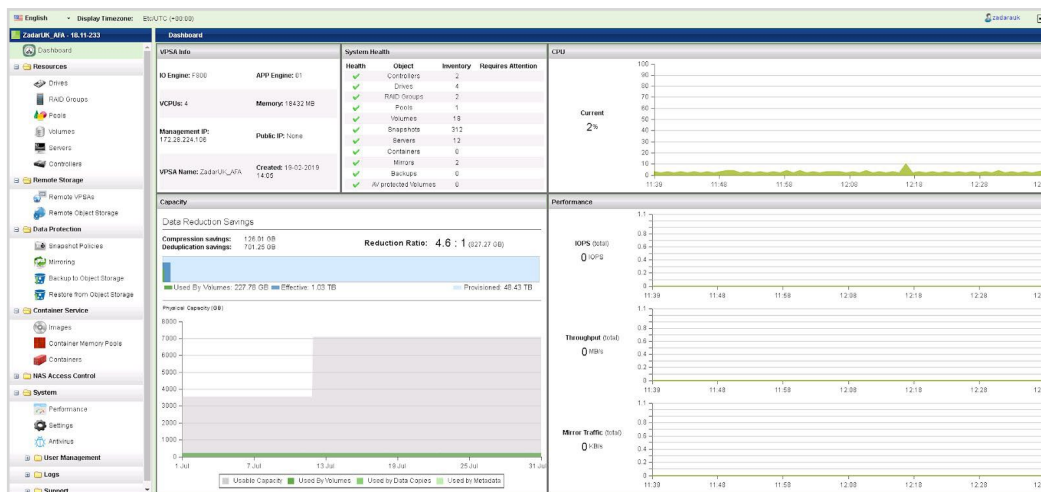
VM_VPSA	Any	Connected VPC Prefixes, DirectConnect Prefixes, S3 Prefixes <a href="#">show less</a>	Any	Allow	VPC interface, Direct Connect interface <a href="#">show less</a>	Disabled
---------	-----	--	-----	-------	--	----------

**Note:** At this point you may need to ensure that traffic can flow via the VPC ENI to the SDDC. Ensure that the rules applied allow two way communication between the EC2 instances and any configured VM's in the compute gateway.

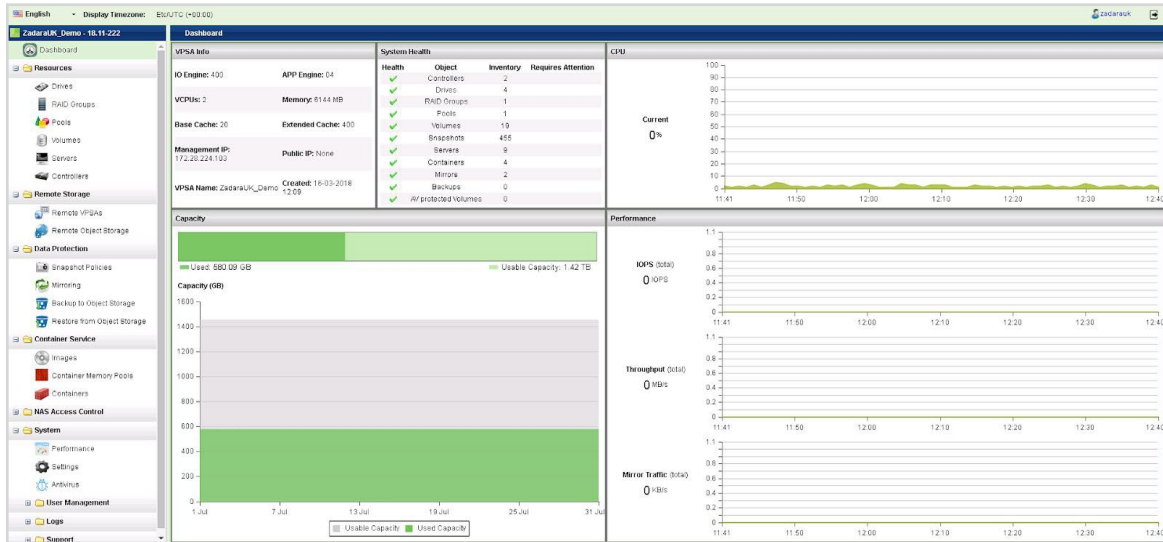
If your VPC also needs iSCSI, NFS, SMB storage from a Zadara VPSA Array, then additional DXVIs will need to be provisioned – speak to your Zadara account manager about this.

At this point you can either continue managing and provisioning the Zadara storage, or you can schedule an “onboarding” session with Zadara support to help get you up on running.

If we now log in to our VPSA you should see a Dashboard similar to this (this example is using a Zadara all-flash array).



Or this if using a Zadara hybrid storage array:



The Servers Tab allows us to register new Clients and setup access to volumes. Here you can see that we are presenting NFS volumes to ESX Servers as a datastore. We also have some Windows clients defined in both the AWS VPC and the ESX Compute Network via NSX-T.

Name	iSCSI / FC Connectivity	IP or CIDR Block	iSCSI IQN	IPsec iSCSI	IPsec NFS	Registered	OS
AWS_Demo_Server	Active	172.31.31.38	iqn.1991-05.com.microsoftwin-ub...	Disabled	Disabled	no	Windows
ESX_NFS_Servers		10.10.0.0/20		Disabled	Disabled	no	ESXi
ZADARA-WIN-DC1	Active	192.168.100.10	iqn.1991-05.com.microsofzadara-...	Disabled	Disabled	yes	Microsoft Windows Server 2016 ...
WIN-VBEAM-01	Active	192.168.100.11	iqn.1991-05.com.microsoftwin-vee...	Disabled	Disabled	yes	Windows

Servers can be registered manually or automatically via a setup script from the virtual private array.

Below, we can see the range of volumes presented from our virtual private array to both EC2 instances, ESXi Hosts, and VM Guests, all from a single Zadara VPSA. Actual deployment and capabilities may vary by region. Check with your Zadara team for capabilities in each region.



Name ↑	Capacity	Status	Protection	Data Type	Pool	Server(s)
AD_Block_Test	1 TB	In-use		BLOCK	R10_SATA_Pool	ZADARA-WIN-DC1
EC_Test	1 TB	In-use		BLOCK	R10_SATA_Pool	AWS_Demo_Server
ESX_Guest_Test	10 GB	In-use		BLOCK	R10_SATA_Pool	ZADARA-WIN-DC1
ESX_Guest_Vol2	25 GB	Available		BLOCK	R10_SATA_Pool	
imeter_test	10 GB	In-use		BLOCK	R10_SATA_Pool	AWS_Demo_Server
NextCloud	100 GB	In-use	🔒	File-System	R10_SATA_Pool	nextcloud_zcs
SMBTest	10 GB	In-use	🔒	File-System	SSD_Pool1	Multiple(3)
Veeam_Backups	1 TB	In-use		BLOCK	R10_SATA_Pool	WIN-VEEAM-01
Veeam_SMB	10 TB	In-use	🔒	File-System	R10_SATA_Pool	WIN-VEEAM-01
VMware_Templates	200 GB	In-use	🔒	BLOCK	R10_SATA_Pool	AWS_Demo_Server
VMwareOS1	1 TB	In-use		File-System	SSD_Pool1	Multiple(2)
VMwareOS2	10 TB	In-use		File-System	SSD_Pool1	Multiple(2)

Here you can see the two Zadara-presented NFS Datastores attached to the SDDC running in VMware for AWS. As of this writing, mounting NFS volumes along with NSX-T is in development / preview mode. Check out the VMware Cloud on AWS roadmap for more details <https://cloud.vmware.com/vmc-aws/roadmap>. There are a number of ways that virtual machines can be created or imported. VMs can be imported from NFS mounts. They can be imported from Veeam backups that have been replicated to and mounted by an EC2 instance, and subsequently copied to the vSAN SDDC datastore. For environments that do not implement Veeam on the primary site, VMs can also be imported from Zadara Backup to Object Storage backups stored on Zadara Object Storage, Amazon S3, or Google Cloud. This option requires an ESXi environment, either customer-owned or via a Zadara service partner, to mount native ESX Datastores prior to restoring contained virtual machines.

## VPSA NAS INTEGRATIONS

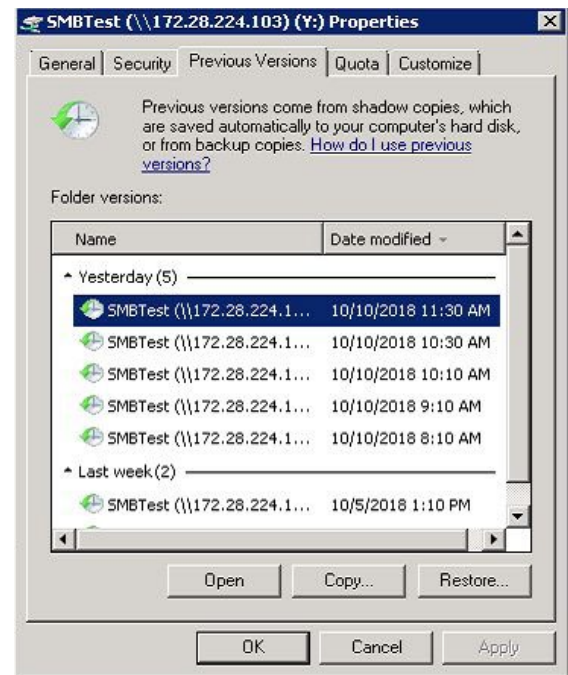
If you look closely at the VPSA volumes image above, you will notice that this example is running an open source NextCloud Sync and Share Docker container service to access a Zadara volume running within this virtual array. This configuration provides sync and share capability to any EC2 and SDDC guests. The Sync and Share can also be exposed via the public internet using a reverse proxy or via additional Zadara interconnects, or DXVI's to the customer site. Providing AWS-external connectivity allows these shares to be used as an offsite file backup repository, complete with snapshot capabilities, for unstructured data.

This example virtual array is also serving Active Directory-integrated SMB Volumes to clients with from AD controllers in EC2 and VM Guest instances. This capability is extremely useful when deploying VMware Horizon 7 View to the SDDC environment. Deploying these AD-integrated block devices enables roaming profiles and home directories to be centrally stored and backed up. Add Zadara's support for user and group quotas, optional integrated antivirus protection, and full suite of replication capabilities (snapshots, local clones, remote clones, and remote mirroring) and we have a hybrid cloud environment ideally suited for VDI use cases.

This hybrid cloud storage configuration also enables Windows Clients in either the SDDC or the VMware Cloud to make use of the Windows Previous Versions restores, natively, from within the standard windows explorer tools for self service data recovery.

Combining the snapshots with replication enables a number of advanced options features such as:

- Data protection via geographic distribution
- On-premises-to-cloud disaster recovery
- Cloud-to-on-premises disaster recovery
- Populating test environments with production data
- Replicating between different types of storage pools, e.g. from accelerated hard disk to all-flash pools
- Replicating production to enable rapid testing with independent copies of production data



## ADDING VEEAM TO OUR HYBRID CLOUD

Deploying Zadara's Backup to Object Store capabilities can be a great approach for organizations with modest backup and recovery requirements. But, for enterprise-wide data availability services when "Availability For Any App, Any Data, Across Any Cloud" is the requirement, Veeam Availability Suite the solution of choice.

In particular, Veeam's capabilities for backing up VMware Cloud on AWS, combined with straightforward integration for Zadara's enterprise storage-as-a-service, provides an excellent path for delivering hybrid cloud-aware enterprise grade data availability services.

More specifically, creating a Veeam backup infrastructure running in the SDDC, with storage presented from Zadara to the backup servers as iSCSI Block or SMB Shares, enables a backup schedule to be defined that places both second and third data copies outside of the all-flash vSAN environment. Leveraging storage that is independent of the vSAN production datastores offers three key advantages:

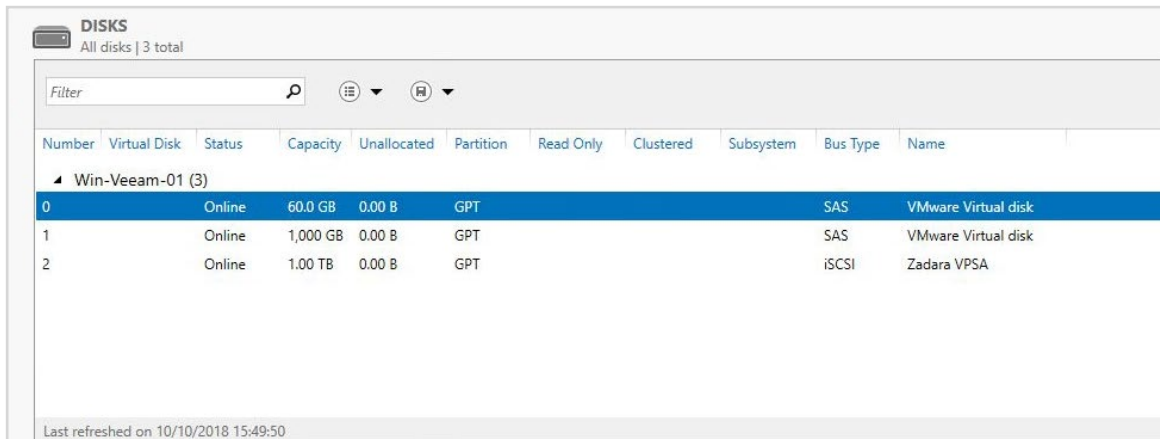
- Backup storage I/O will not compete with production workloads
- Higher capacity storage pools enable longer retention time
- Leveraging hard disk capacity reduces backup storage costs

Let's get started...

First we create the Veeam Windows Server as VM Guest, install the required updates, start the iSCSI service and register the hosts with the Zadara VPSA.

Then we map an iSCSI volume that will be used to store the VSAN backups, format it and prepare to install the Veeam backup software. Download the Veeam 9.5 Update 4 image, mount it and begin the installation.

**Note:** Update 3a+ is required to support SDDC in VMware



Number	Virtual Disk	Status	Capacity	Unallocated	Partition	Read Only	Clustered	Subsystem	Bus Type	Name
Win-Veeam-01 (3)										
0		Online	60.0 GB	0.00 B	GPT				SAS	VMware Virtual disk
1		Online	1,000 GB	0.00 B	GPT				SAS	VMware Virtual disk
2		Online	1.00 TB	0.00 B	GPT				iSCSI	Zadara VPSA

Last refreshed on 10/10/2018 15:49:50

Outside of the VMware Cloud on AWS SDDC we could create 2 types of repository disks, 1 running from VMFS via an NFS, or Block mounted Datastore and a direct iSCSI Volume to a Zadara Block Volume.

# VEEAM BACKUP REPOSITORY

Here we look at the details for the iSCSI Repository Created "Zadara Block Repository" which is mounted on the Veeam Backup Server as drive F:\

**Edit Backup Repository**

**Name**  
Type in a name and description for this backup repository.

**Name:**  
Zadara Block Repository

**Description:**  
Created by ZADARA\Administrator at 10/10/2018 16:03.

< Previous   **Next >**   Finish   Cancel

**Edit Backup Repository**

**Type**  
Choose type of backup repository you want to create.

**Microsoft Windows server**  
Microsoft Windows server with internal or directly attached storage. Data mover process running directly on the server allows for improved backup efficiency, especially over slow links.


**Linux server**  
Linux server with internal, directly attached, or mounted NFS storage. Data mover process running directly on the server allows for more efficient backups, especially over slow links.

**Shared folder**  
CIFS (SMB) share. When backing up over slow links, we recommend that you specify a gateway server located in the same site with the shared folder.

**Deduplicating storage appliance**  
Advanced integration with Dell EMC Data Domain, ExaGrid and HPE StoreOnce. For basic integration, use the Shared folder option above.


< Previous   **Next >**   Finish   Cancel



**Edit Backup Repository** ✕

 **Server**  
Choose server backing your repository. You can select server from the list of managed servers added to the console.

Name	Repository server:			Add New...																																							
Type	Win-Veeam-01.zadara.local (Backup server)			Populate																																							
Server	<table border="1"> <thead> <tr> <th>Path</th> <th>Capacity</th> <th>Free</th> </tr> </thead> <tbody> <tr> <td>C:\</td> <td>59.4 GB</td> <td>29.8 GB</td> </tr> <tr> <td>E:\</td> <td>999.8 GB</td> <td>984.5 GB</td> </tr> <tr style="background-color: #e0f0ff;"> <td>F:\</td> <td>1023.8 GB</td> <td>983.5 GB</td> </tr> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> <tr><td> </td><td> </td><td> </td></tr> </tbody> </table>				Path	Capacity	Free	C:\	59.4 GB	29.8 GB	E:\	999.8 GB	984.5 GB	F:\	1023.8 GB	983.5 GB																											
Path	Capacity	Free																																									
C:\	59.4 GB	29.8 GB																																									
E:\	999.8 GB	984.5 GB																																									
F:\	1023.8 GB	983.5 GB																																									
Repository	<p>&lt; Previous   <b>Next &gt;</b>   Finish   Cancel</p>																																										

**Edit Backup Repository** ✕

 **Repository**  
Type in path to the folder where backup files should be stored, and set repository load control options.

Name	Location		
Type	Path to folder:		
Server	F:\	Browse...	
Repository	 Capacity: <b>1023.8 GB</b>  Free space: <b>983.5 GB</b>	Populate	
Mount Server	<b>Load control</b> Running too many concurrent tasks against the same repository may reduce overall performance, and cause I/O operations to timeout. Control storage device saturation with the following settings:		
Review	<input checked="" type="checkbox"/> Limit maximum concurrent tasks to: <input type="text" value="4"/>		
Apply	<input type="checkbox"/> Limit read and write data rates to: <input type="text"/> MB/s		
Click Advanced to customize repository settings			Advanced...
<p>&lt; Previous   Next &gt;   Finish   Cancel</p>			

Enable and setup vPower NFS so that change blocks on instant recovered VM's can be stored and tracked, a separate volume can be used but a new folder was used for this environment.

**Edit Backup Repository** [Close]

**Mount Server**  
Specify a server to mount backups to for file-level restores. vPower NFS service allows for running virtual machines directly from backup files, enabling advanced functionality such as Instant VM Recovery, SureBackup and On-Demand Sandbox.

**Name** Mount server:  
Win-Veeam-01.zadara.local (Backup server)

**Type**

**Server**  Enable vPower NFS service on the mount server (recommended)  
Specify vPower NFS write cache location on the mount server. Make sure the selected volume has enough free disk space available to store changed disk blocks of instantly recovered VMs.

**Repository** Folder: F:\PowerNFS

**Mount Server**

**Review**

**Apply**

Click Ports to change NFS server and backup mount listener ports

< Previous    Next >    Finish    Cancel

**Edit Backup Repository** [Close]

**Review**  
Please review the settings, and click Apply to continue.

**Name** Repository type: **Windows**

**Type** Mount host: **Win-Veeam-01.zadara.local**

**Server** Account: **Not set**

**Repository** Backup folder: **F:\**

**Mount Server** Write throughput: **Not limited**

**Review** Max parallel tasks: **4**

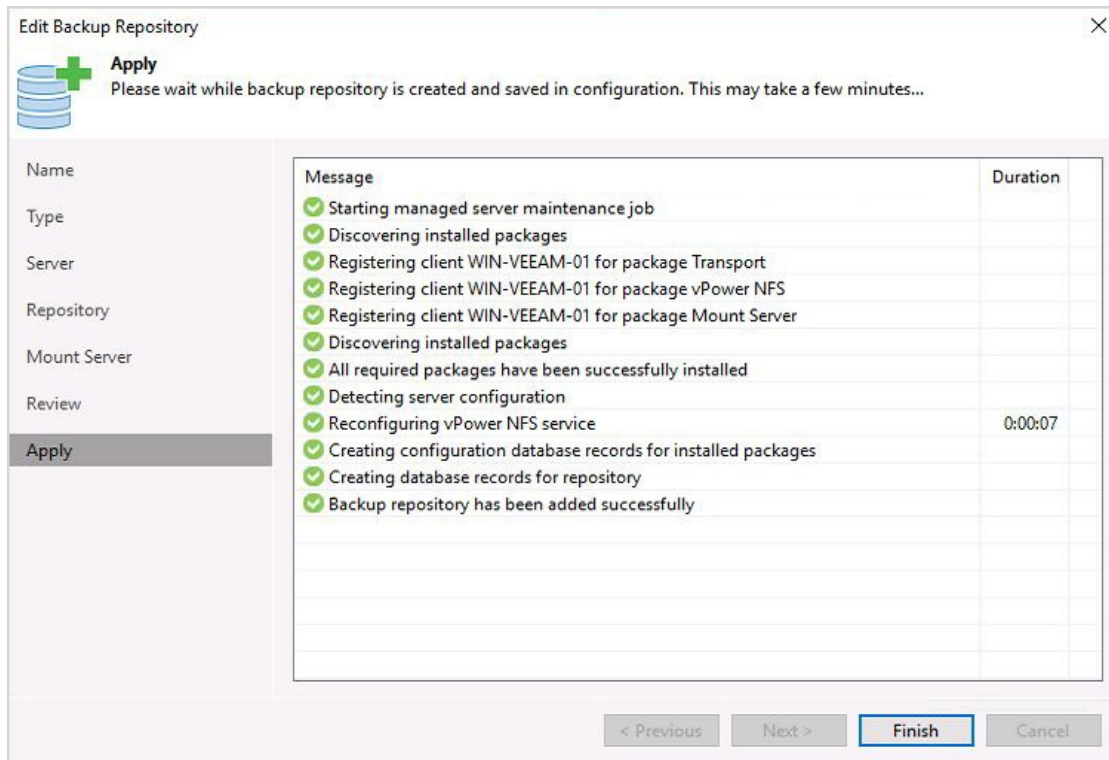
**Apply**

The following components will be processed on server Win-Veeam-01.zadara.local:

Component name	Status
Transport	already exists
vPower NFS	already exists
Mount Server	already exists


Import existing backups automatically  
 Import guest file system index

< Previous    **Apply**    Finish    Cancel




## VEEAM BACKUP JOB CREATION


Create a Backup Job. Below we will see the details of a created Job to backup an AD Server.



**Edit Backup Job [Backup Job AD Server]** X

 **Virtual Machines**  
Select virtual machines to process via container, or granularly. Container provides dynamic selection that automatically changes as you add new VM into container.

**Name** | Virtual machines to backup:

Name	Type	Size	
 Windows 2016 ServerEval1	Virtual M...	58.7 GB	


**Virtual Machines** | **Storage** | **Guest Processing** | **Schedule** | **Summary**

**Add...** **Remove** **Exclusions...** **Up** **Down** **Recalculate**

Total size: **58.7 GB**

**< Previous** **Next >** **Finish** **Cancel**

**Edit Backup Job [Backup Job AD Server]** X

 **Storage**  
Specify processing proxy server to be used for source data retrieval, backup repository to store the backup files produced by this job and customize advanced job settings if required.

**Name** | **Virtual Machines** | **Storage** | **Guest Processing** | **Schedule** | **Summary**

**Backup proxy:**  
 **Choose...**

**Backup repository:**  
 **Map backup**

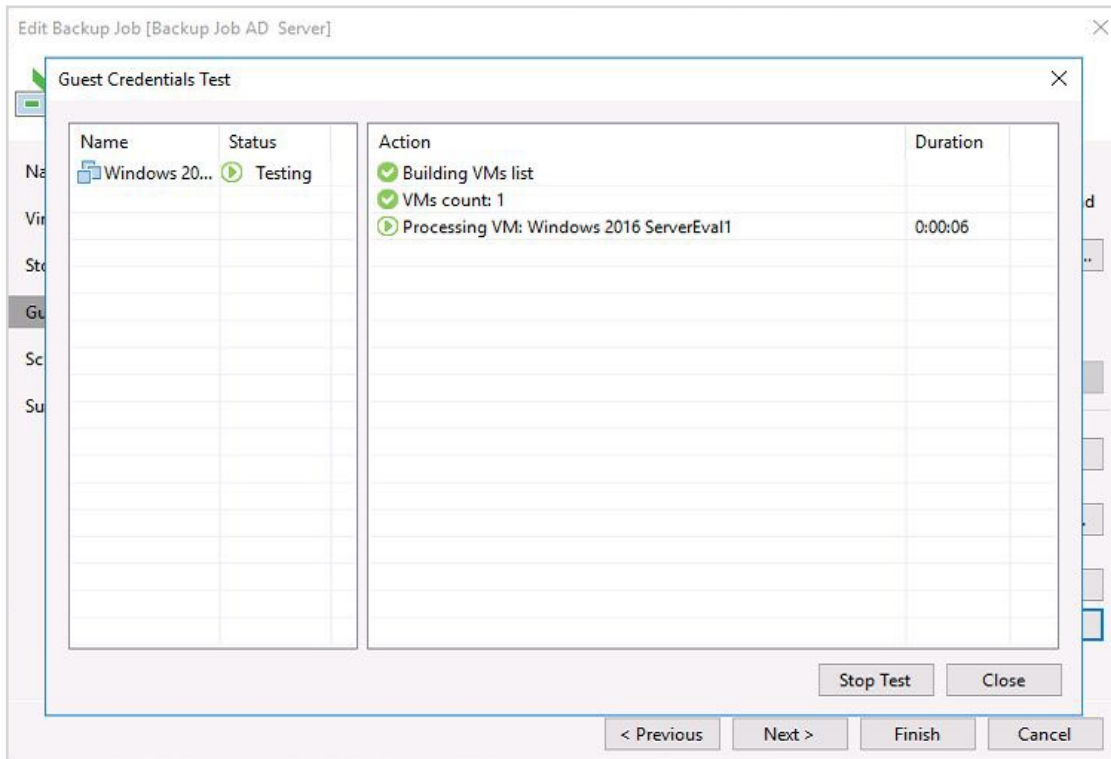
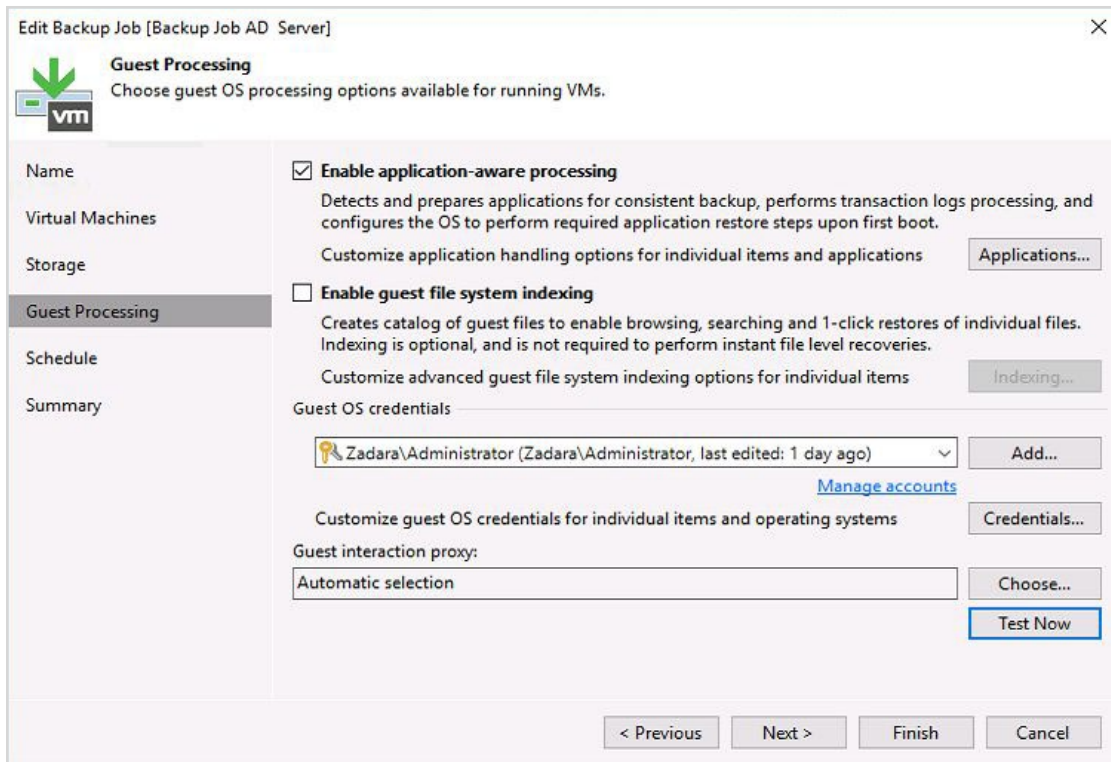
**Retention policy**  
Restore points to keep on disk:  ⓘ

**Configure secondary destinations for this job**  
Copy backups produced by this job to another backup repository, or to tape. Best practices recommend maintaining at least 2 backups of production data, with one of them being off-site.

Advanced job settings include backup mode, compression and deduplication, block size, notification settings, automated post-job activity and other settings. **Advanced**

**< Previous** **Next >** **Finish** **Cancel**





**Edit Backup Job [Backup Job AD Server]**

**Schedule**  
Specify the job scheduling options. If you do not set the schedule, the job will need to be controlled manually.

**Name**  Run the job automatically

**Virtual Machines**  Daily at this time: 22:00 Everyday Days...

**Storage**  Monthly at this time: 22:00 Fourth Saturday Months...

**Guest Processing**  Periodically every: 1 Hours Schedule...

**Schedule**  After this job: Backup Job 1 (Created by ZADARA\Administrator at 10/10/2018 15:46.)

**Summary**

**Automatic retry**

Retry failed items processing: 3 times

Wait before each retry attempt for: 10 minutes

**Backup window**

Terminate job if it exceeds allowed backup window Window...

If the job does not complete within allocated backup window, it will be terminated to prevent snapshot commit during production hours.

< Previous Apply Finish Cancel

**Edit Backup Job [Backup Job AD Server]**

**Summary**  
The job's settings have been saved successfully. Click Finish to exit the wizard.

**Name** Summary:

**Virtual Machines** Name: Backup Job AD Server

**Storage** Target Path: F:\

**Guest Processing** Type: VMware Backup

**Schedule** Enable application-aware processing

**Summary** Source items:  
Windows 2016 ServerEval1 (vcenter.sddc-35-178-92-134.vmwarevmc.com)

Command line to start the job on backup server:  
"C:\Program Files\Veeam\Backup and Replication\Backup\Veeam.Backup.Manager.exe" backup f6665bd8-a651-4690-97c8-acfea72aeb5d

Run the job when I click Finish

< Previous Next > Finish Cancel

Job statistics from Backup Job Incremental:

Backup Job AD Server (Incremental) 1 of 1 VMs

Job progress: 100%

SUMMARY		DATA		STATUS	
Duration:	02:19	Processed:	59.7 GB (100%)	Success:	1 ✓
Processing rate:	37 MB/s	Read:	3.4 GB	Warnings:	0
Bottleneck:	Proxy	Transferred:	1.4 GB (2.4x)	Errors:	0

THROUGHPUT (ALL TIME)

					Speed: 171.9 MB/s
--	--	--	--	--	-------------------

NAME	STATUS	ACTION	DURATI...
Windows 2016 Server...	✓ Success	✓ Job started at 12/10/2018 08:55:06 ✓ Building VMs list ✓ VM size: 60.0 GB (59.7 GB used) ✓ Changed block tracking is enabled ✓ Processing Windows 2016 ServerEval1 ✓ All VMs have been queued for processing ✓ Load: Source 58% > Proxy 80% > Network 70% > Target 1% ✓ Primary bottleneck: Proxy ✓ Job finished at 12/10/2018 08:57:26	00:02 01:50

Hide Details OK

## VEEAM BACKUP FILES SHOWING FULL AND INCREMENTAL FILES

Backup Job AD Server

This PC > Zadara\_Volume (F:) > Backup Job AD Server

Name	Date modified	Type	Size
Backup Job AD Server	12/10/2018 08:57	Veeam Backup & Replication backup chain metadata file	22 KB
Backup Job AD ServerD2018-10-10T161048	10/10/2018 16:14	Veeam Backup & Replication full backup file	18,407,936 ...
Backup Job AD ServerD2018-10-10T162524	10/10/2018 16:26	Veeam Backup & Replication incremental backup file	113,984 KB
Backup Job AD ServerD2018-10-12T085528	12/10/2018 08:57	Veeam Backup & Replication incremental backup file	1,594,496 KB

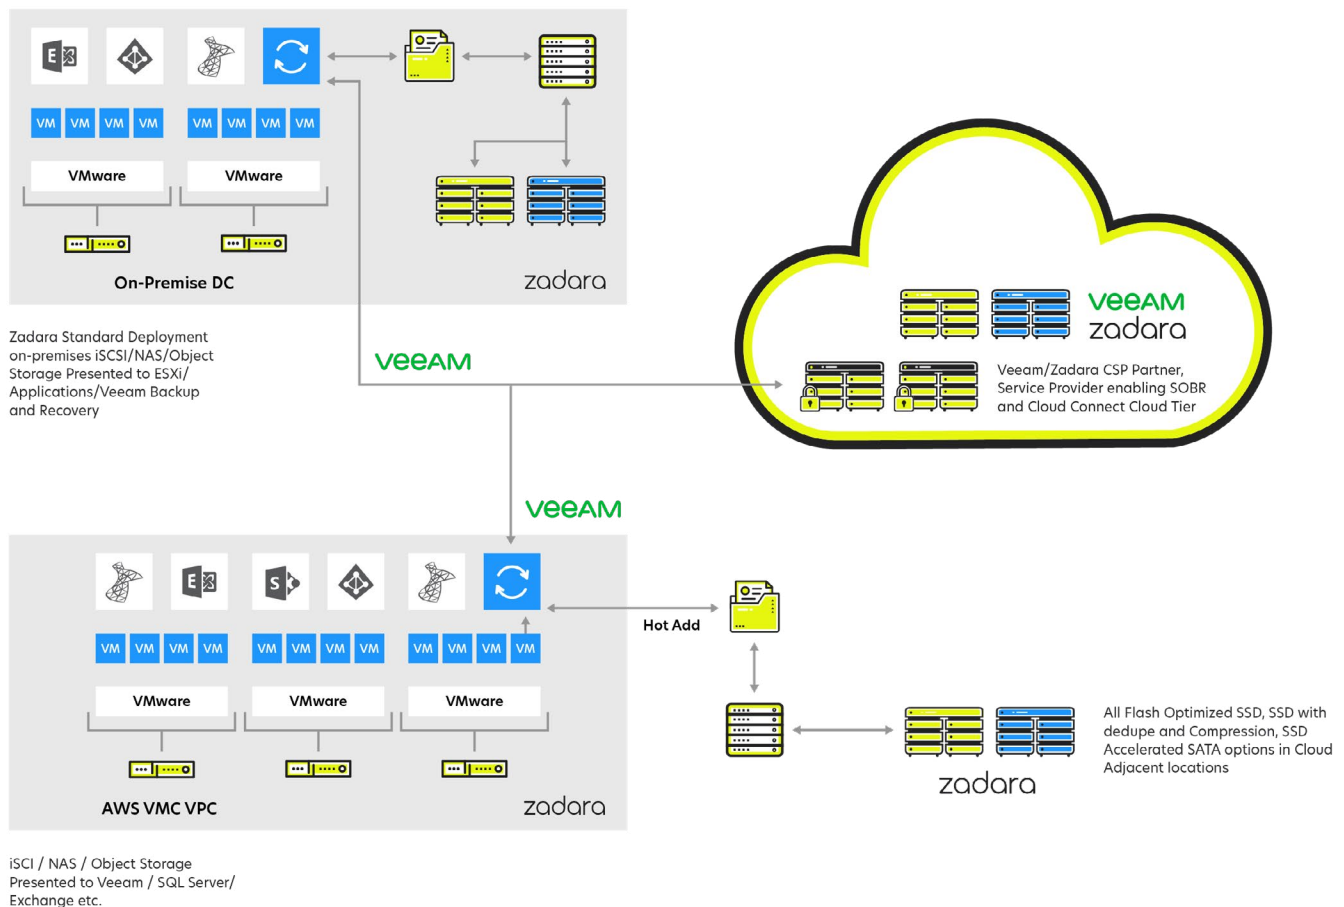
## VEEAM CAPACITY TIER SCALE-OUT BACKUP REPOSITORY

Veeam 9.5 update 4 provides the ability to add a Scale-Out Backup Repository (SOBR) utilizing AWS S3 and qualified S3 compatible Object Storage. This offers the option of combining Zadara as an offsite, off-cloud, primary backup location and utilizing AWS S3 for a multiple Availability Zone(local or geographically remote) target. This provides added capabilities and options for both On-Premises and Public Cloud solutions for long term data retention and disaster recovery options.

Users can also make use of their complete Zadara platform to provide the Object Scale Out Backup Repository as well as the Primary Data Copy see “[Veeam - Configure VPSA Object Storage as an Object Storage Repository](https://support.zadara.com/hc/en-us/articles/360027853491-Veeam-Configure-VPSA-Object-Storage-as-an-Object-Storage-Repository)” (<https://support.zadara.com/hc/en-us/articles/360027853491-Veeam-Configure-VPSA-Object-Storage-as-an-Object-Storage-Repository>) for best practices.

Zadara and Veeam provide joint support for using Zadara’s object storage service via Zadara’s S3 compatibility. Please see the interim compatibility list at <https://forums.veeam.com/object-storage-f52/unofficial-compatibility-list-for-veeam-cloud-tier-t56956.html>.

The following diagram illustrates the value that integrating Veeam and Zadara provides; a multiple copy, multi-site, all tier storage and backup solution.



# VMWARE GUEST ISCSI VM SNAPSHOT PROTECTION

Zadara Provides a Volume Shadow Service (VSS) provider to enable application consistent snapshots. This allows snapshots and clones from one VM Guest to be instantly presented to another VM running in the same SDDC, a different SDDC within AWS, or to an AWS EC2 instance running in a connected VPC. The value of this capability extends beyond simply providing a backup solution. VSS-integration makes it possible to immediately present application consistent, zero capacity clones. Use cases for these clones include:

- Ransomware protection
- Instant population of test environments with complete sets of application data
- Verification of backed up data prior to offsite transfer

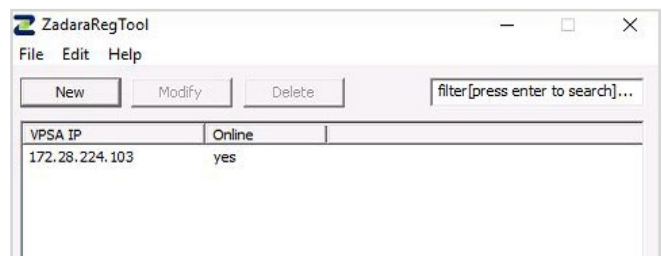
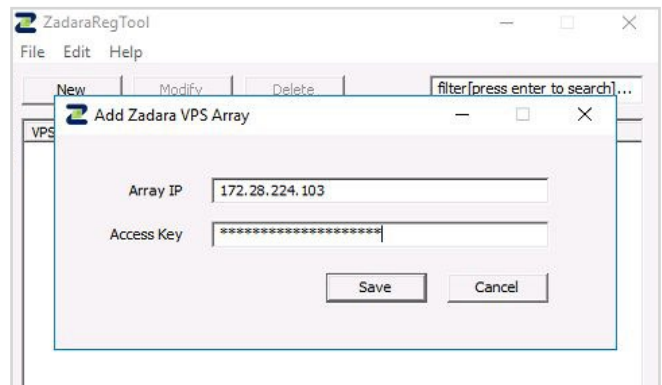
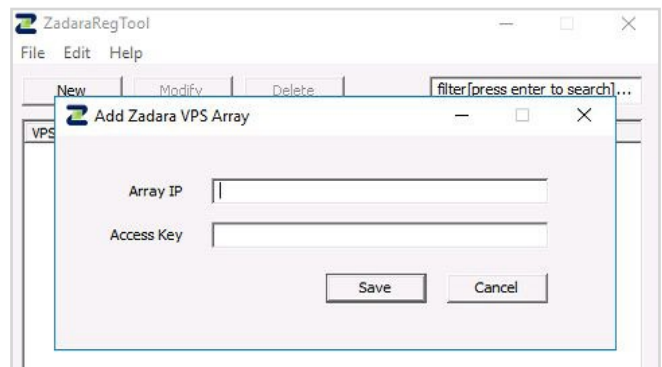
## ZADARA VSS HW PROVIDER DEPLOYMENT

The Zadara VSS Hardware Provider can be downloaded from the following URL [https://zadara-storage-software.s3.amazonaws.com/ZadaraHardwareProvider\\_x64.msi](https://zadara-storage-software.s3.amazonaws.com/ZadaraHardwareProvider_x64.msi) for 64bit Windows Clients and [https://zadara-storage-software.s3.amazonaws.com/ZadaraHardwareProvider\\_x86.msi](https://zadara-storage-software.s3.amazonaws.com/ZadaraHardwareProvider_x86.msi) for 32bit Windows Clients.

Once downloaded and installed you will need to configure access to your VPSA Array API.

- Create a specific user within the VPSA Array (refer to user documentation).
- Login as the user and present the IP address of the VPSA and the access key for your user with permissions to create snapshots and clones.

Microsoft has released a number of CLI VSS tools over the years. Here we will use vssadmin and DiskShadow to demonstrate some of the capabilities. Open a command prompt as an Administrator and execute "vssadmin list providers." You should see something like the below image showing the ZadaraVssProvider.



```

Administrator: Command Prompt
Provider Id: {89300202-3cec-4981-9171-19f59559e0f2}
Version: 1.0.0.1

Provider name: 'Microsoft Software Shadow Copy provider 1.0'
Provider type: System
Provider Id: {b5946137-7b9f-4925-af80-51abd60b20d5}
Version: 1.0.0.7

C:\Users\Administrator>vssadmin list providers
vssadmin 1.1 - Volume Shadow Copy Service administrative command-line tool
(C) Copyright 2001-2013 Microsoft Corp.

Provider name: 'ZadanaVssProvider'
Provider type: Hardware
Provider Id: {6a111feb-1ea9-496e-ab0e-b85ff8a42e46}
Version: 1.0.12.29

Provider name: 'Microsoft File Share Shadow Copy provider'
Provider type: Fileshare
Provider Id: {89300202-3cec-4981-9171-19f59559e0f2}
Version: 1.0.0.1

Provider name: 'Microsoft Software Shadow Copy provider 1.0'
Provider type: System
Provider Id: {b5946137-7b9f-4925-af80-51abd60b20d5}
Version: 1.0.0.7

C:\Users\Administrator>

```

Check the Volumes with VSSAdmin, Volumes E: and F: are iSCSI presented Zadara Volumes of 10Gb and 1TB respectively.

```

Administrator: Command Prompt

C:\Users\Administrator>vssadmin List Volumes
vssadmin 1.1 - Volume Shadow Copy Service administrative command-line tool
(C) Copyright 2001-2013 Microsoft Corp.

Volume path: C:\
Volume name: \\?\Volume{8054f192-5511-4065-877a-2e5244847a40}\
Volume path: E:\
Volume name: \\?\Volume{f7265667-489c-4b27-83b5-b692b14dfcac}\
Volume path: F:\
Volume name: \\?\Volume{04392ffd-3c9d-40db-8690-55001b45f209}\
Volume path: \\?\Volume{82e38205-c71f-4528-93e3-bcb2a37e06db}\
Volume name: \\?\Volume{82e38205-c71f-4528-93e3-bcb2a37e06db}\


C:\Users\Administrator>

```

Zadara 4K 100% Read 8 x SSD R10 Pool.

Name	Total Capacity	LUN	Access Type	Status
ESX_QuestL_Test	10 GB	0	iSCSI	In-use
SMBTest	10 GB		SMB	In-use
AD_Block_Test	1 TB	1	iSCSI	In-use

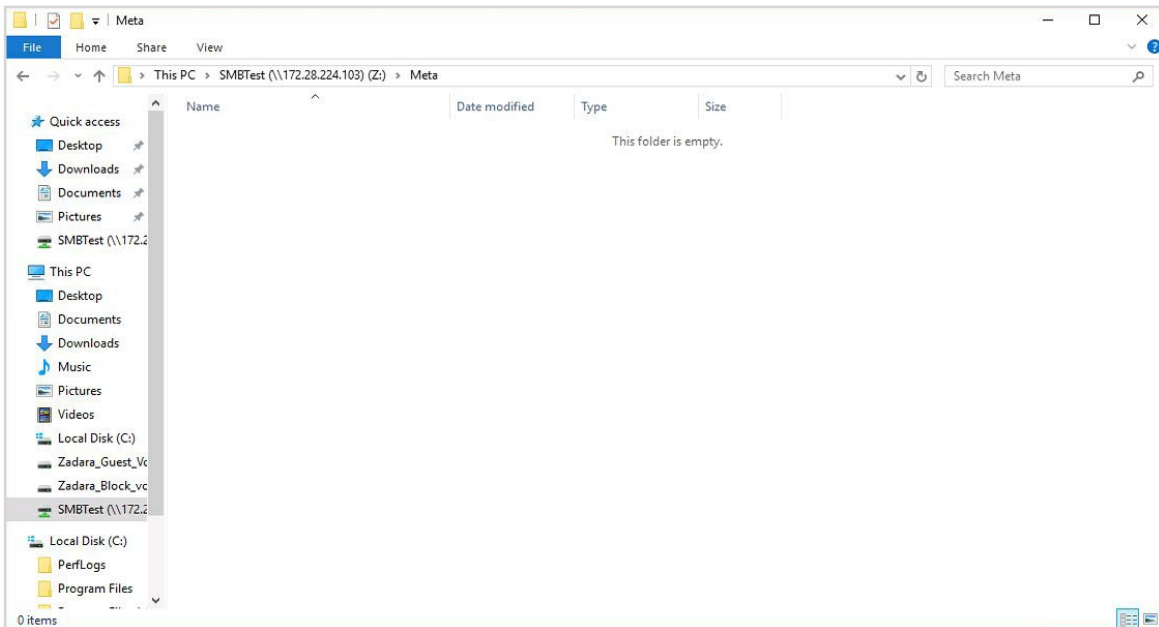
Before Creating the VSS Snapshot there are no Snapshots or Clones of this volume.



## LET'S USE DISKSHADOW TO CREATE A VSS AWARE SNAPSHOT

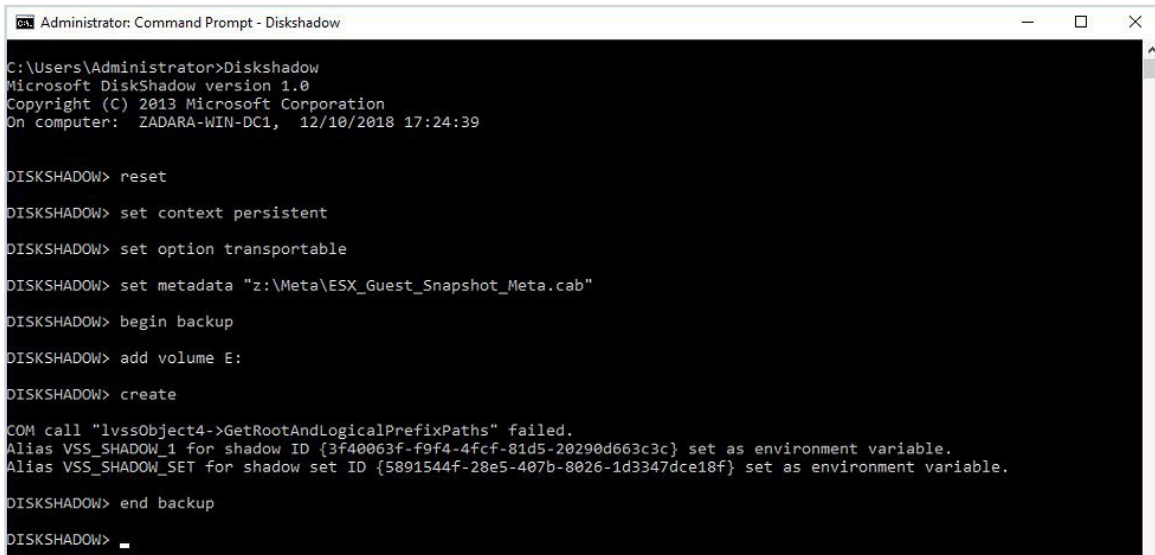
First prepare the environment and use a Zadara NAS volume to hold the transportable snapshot metadata cab file.

Here we create a Meta Folder on our SMBTest NAS Share.



You may see an error on VSS Object Creation in some Windows Installations. This is a known Windows error and does not interfere with the VSS snapshot creation. It is a Microsoft BCD error (see this technote <https://support.microsoft.com/en-gb/help/3025158/diskshadow-error-when-you-try-to-create-a-vss-snapshot-in-windows-serv>).

Here are the set of commands used to create the VSS Aware Snapshot as a Transportable Snapshot:



```
C:\Users\Administrator>Diskshadow
Microsoft DiskShadow version 1.0
Copyright (C) 2013 Microsoft Corporation
On computer: ZADARA-WIN-DC1, 12/10/2018 17:24:39

DISKSHADOW> reset

DISKSHADOW> set context persistent

DISKSHADOW> set option transportable

DISKSHADOW> set metadata "z:\Meta\ESX_Guest_Snapshot_Meta.cab"

DISKSHADOW> begin backup

DISKSHADOW> add volume E:

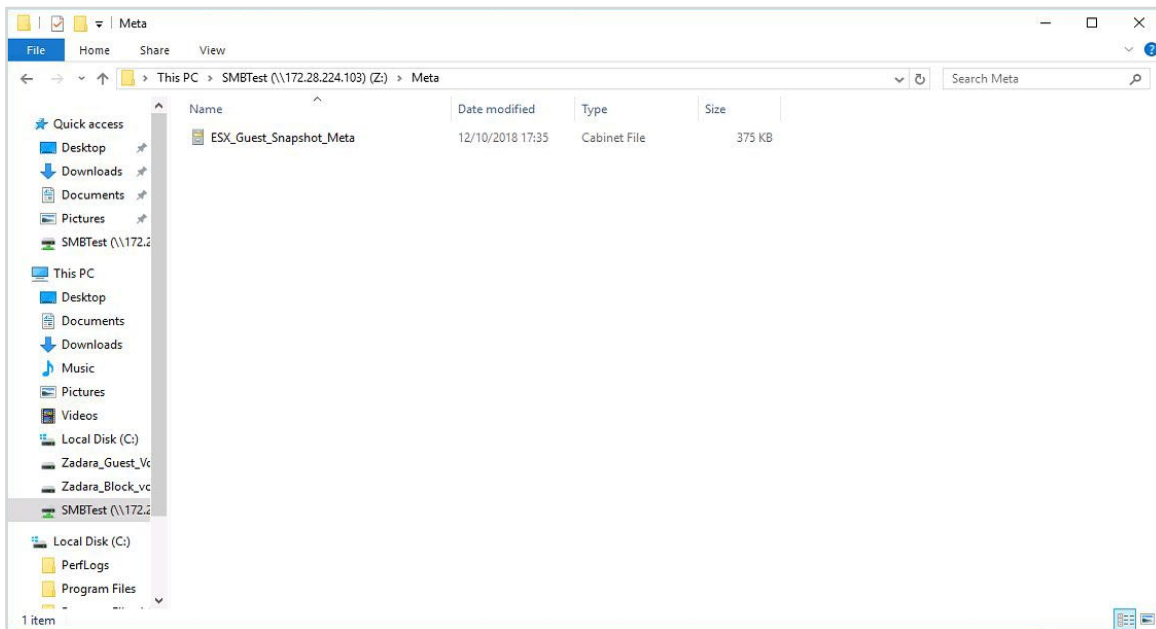
DISKSHADOW> create

COM call "lvssObject4->GetRootAndLogicalPrefixPaths" failed.
Alias VSS_SHADOW_1 for shadow ID {3f40063f-f9f4-4fcf-81d5-20290d663c3c} set as environment variable.
Alias VSS_SHADOW_SET for shadow set ID {5891544f-28e5-407b-8026-1d3347dce18f} set as environment variable.

DISKSHADOW> end backup

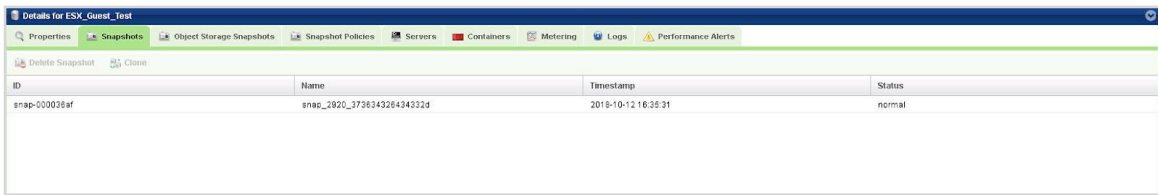
DISKSHADOW>
```

Here is the Metadata created during snapshot creation on the SMBTest NAS Share:





Now in our VPSA Array we can see the creation of the snapshot, but there is no clone at this point. It is just an application consistent point in time.

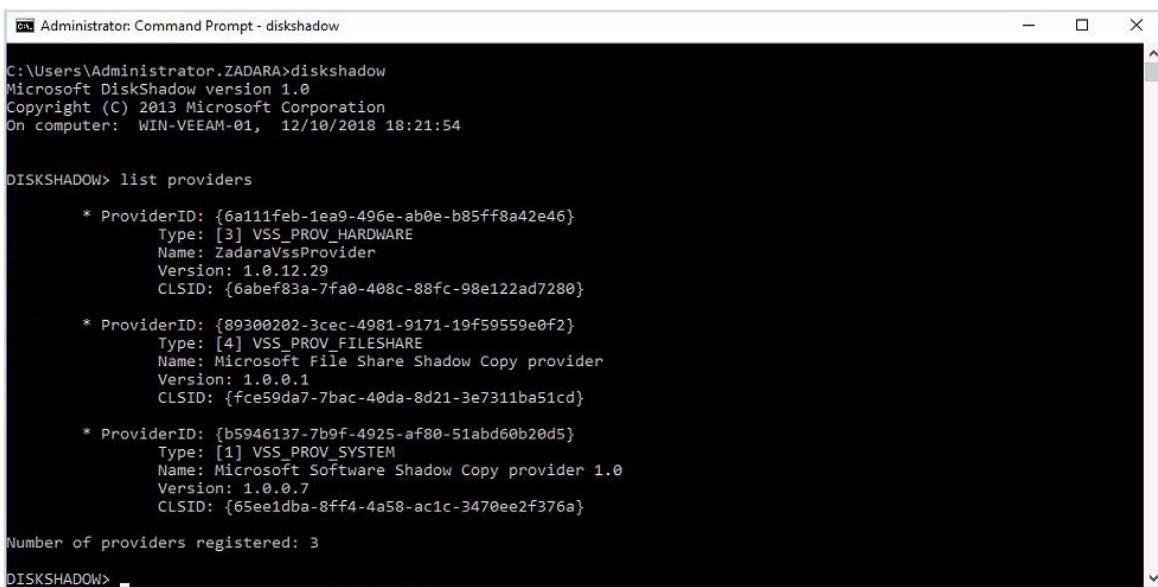


The screenshot shows the 'Details for ESX\_Guest\_Test' window in the Veeam Backup & Replication console. The 'Snapshots' tab is active, displaying a table with one snapshot entry.

ID	Name	Timestamp	Status
snap-000038f	snac_2920_373834328434322d	2018-10-12 16:35:31	normal

On our Test / Dev, Backup Server we can now create a Clone and access the data.

## LET'S USE DISKSHADOW THIS TIME TO LIST THE PROVIDERS INSTEAD OF VSSADMIN



```
Administrator: Command Prompt - diskshadow
C:\Users\Administrator.ZADARA>diskshadow
Microsoft DiskShadow version 1.0
Copyright (C) 2013 Microsoft Corporation
On computer: WIN-VEEAM-01, 12/10/2018 18:21:54

DISKSHADOW> list providers

* ProviderID: {6a111feb-1ea9-496e-ab0e-b85ff8a42e46}
  Type: [3] VSS_PROV_HARDWARE
  Name: ZadaraVssProvider
  Version: 1.0.12.29
  CLSID: {6abef83a-7fa0-408c-88fc-98e122ad7280}

* ProviderID: {89300202-3cec-4981-9171-19f59559e0f2}
  Type: [4] VSS_PROV_FILESHARE
  Name: Microsoft File Share Shadow Copy provider
  Version: 1.0.0.1
  CLSID: {fce59da7-7bac-40da-8d21-3e7311ba51cd}

* ProviderID: {b5946137-7b9f-4925-af80-51abd60b20d5}
  Type: [1] VSS_PROV_SYSTEM
  Name: Microsoft Software Shadow Copy provider 1.0
  Version: 1.0.0.7
  CLSID: {65ee1dba-8ff4-4a58-ac1c-3470ee2f376a}

Number of providers registered: 3

DISKSHADOW>
```

Now import the Metadata from the Cab file into the VM with Diskshadow.

```
Administrator: Command Prompt - diskshadow
On computer: WIN-VEEAM-01, 12/10/2018 18:21:54

DISKSHADOW> list providers

* ProviderID: {6a111feb-1ea9-496e-ab0e-b85ff8a42e46}
  Type: [3] VSS_PROV_HARDWARE
  Name: ZadaraVssProvider
  Version: 1.0.12.29
  CLSID: {6abef83a-7fa0-408c-88fc-98e122ad7280}

* ProviderID: {89300202-3cec-4981-9171-19f59559e0f2}
  Type: [4] VSS_PROV_FILESHARE
  Name: Microsoft File Share Shadow Copy provider
  Version: 1.0.0.1
  CLSID: {fce59da7-7bac-40da-8d21-3e7311ba51cd}

* ProviderID: {b5946137-7b9f-4925-af80-51abd60b20d5}
  Type: [1] VSS_PROV_SYSTEM
  Name: Microsoft Software Shadow Copy provider 1.0
  Version: 1.0.0.7
  CLSID: {65ee1dba-8ff4-4a58-ac1c-3470ee2f376a}

Number of providers registered: 3

DISKSHADOW> load metadata "z:\Meta\ESX_Guest_Snapshot_Meta.cab"
Alias VSS_SHADOW_1 for value {3f40063f-f9f4-4fcf-81d5-20290d663c3c} set as an environment variable.
Alias VSS_SHADOW_SET for value {5891544f-28e5-407b-8026-1d3347dce18f} set as an environment variable.

DISKSHADOW>
```

```
Administrator: Command Prompt - diskshadow

DISKSHADOW> list providers

* ProviderID: {6a111feb-1ea9-496e-ab0e-b85ff8a42e46}
  Type: [3] VSS_PROV_HARDWARE
  Name: ZadaraVssProvider
  Version: 1.0.12.29
  CLSID: {6abef83a-7fa0-408c-88fc-98e122ad7280}

* ProviderID: {89300202-3cec-4981-9171-19f59559e0f2}
  Type: [4] VSS_PROV_FILESHARE
  Name: Microsoft File Share Shadow Copy provider
  Version: 1.0.0.1
  CLSID: {fce59da7-7bac-40da-8d21-3e7311ba51cd}

* ProviderID: {b5946137-7b9f-4925-af80-51abd60b20d5}
  Type: [1] VSS_PROV_SYSTEM
  Name: Microsoft Software Shadow Copy provider 1.0
  Version: 1.0.0.7
  CLSID: {65ee1dba-8ff4-4a58-ac1c-3470ee2f376a}

Number of providers registered: 3

DISKSHADOW> load metadata "z:\Meta\ESX_Guest_Snapshot_Meta.cab"
Alias VSS_SHADOW_1 for value {3f40063f-f9f4-4fcf-81d5-20290d663c3c} set as an environment variable.
Alias VSS_SHADOW_SET for value {5891544f-28e5-407b-8026-1d3347dce18f} set as an environment variable.

DISKSHADOW> import

DISKSHADOW>
```

Check the Shadow Data has imported that this has come from our other Server.

```
Administrator: Command Prompt - diskshadow
Version: 1.0.0.7
CLSID: {65ee1dba-8ff4-4a58-ac1c-3470ee2f376a}

Number of providers registered: 3

DISKSHADOW> load metadata "z:\Meta\ESX_Guest_Snapshot_Meta.cab"
Alias VSS_SHADOW_1 for value {3f40063f-f9f4-4fcf-81d5-20290d663c3c} set as an environment variable.
Alias VSS_SHADOW_SET for value {5891544f-28e5-407b-8026-1d3347dce18f} set as an environment variable.

DISKSHADOW> import

DISKSHADOW> list shadows all

Querying all shadow copies on the computer ...

* Shadow copy ID = {3f40063f-f9f4-4fcf-81d5-20290d663c3c}           %VSS_SHADOW_1%
- Shadow copy set: {5891544f-28e5-407b-8026-1d3347dce18f}       %VSS_SHADOW_SET%
- Original count of shadow copies = 1
- Original volume name: \\?\Volume{f7265667-489c-4b27-83b5-b692b14dfcac}\ [volume not on this machine]
- Creation time: 12/10/2018 17:35:31
- Shadow copy device name: \\?\Volume{35fa8b4e-cd81-11e8-9c66-005056875b46}
- Originating machine: Zadara-WIN-DC1.zadara.local
- Service machine: Zadara-WIN-DC1.zadara.local
- Not exposed
- Provider ID: {6a111feb-1ea9-496e-ab0e-b85ff8a42e46}
- Attributes: Transportable No_Auto_Release Persistent Hardware Imported

Number of shadow copies listed: 1

DISKSHADOW>
```

Now we need to expose this ShadowCopy Volume to the OS.

```
Administrator: Command Prompt - diskshadow
Number of providers registered: 3

DISKSHADOW> load metadata "z:\Meta\ESX_Guest_Snapshot_Meta.cab"
Alias VSS_SHADOW_1 for value {3f40063f-f9f4-4fcf-81d5-20290d663c3c} set as an environment variable.
Alias VSS_SHADOW_SET for value {5891544f-28e5-407b-8026-1d3347dce18f} set as an environment variable.

DISKSHADOW> import

DISKSHADOW> list shadows all

Querying all shadow copies on the computer ...

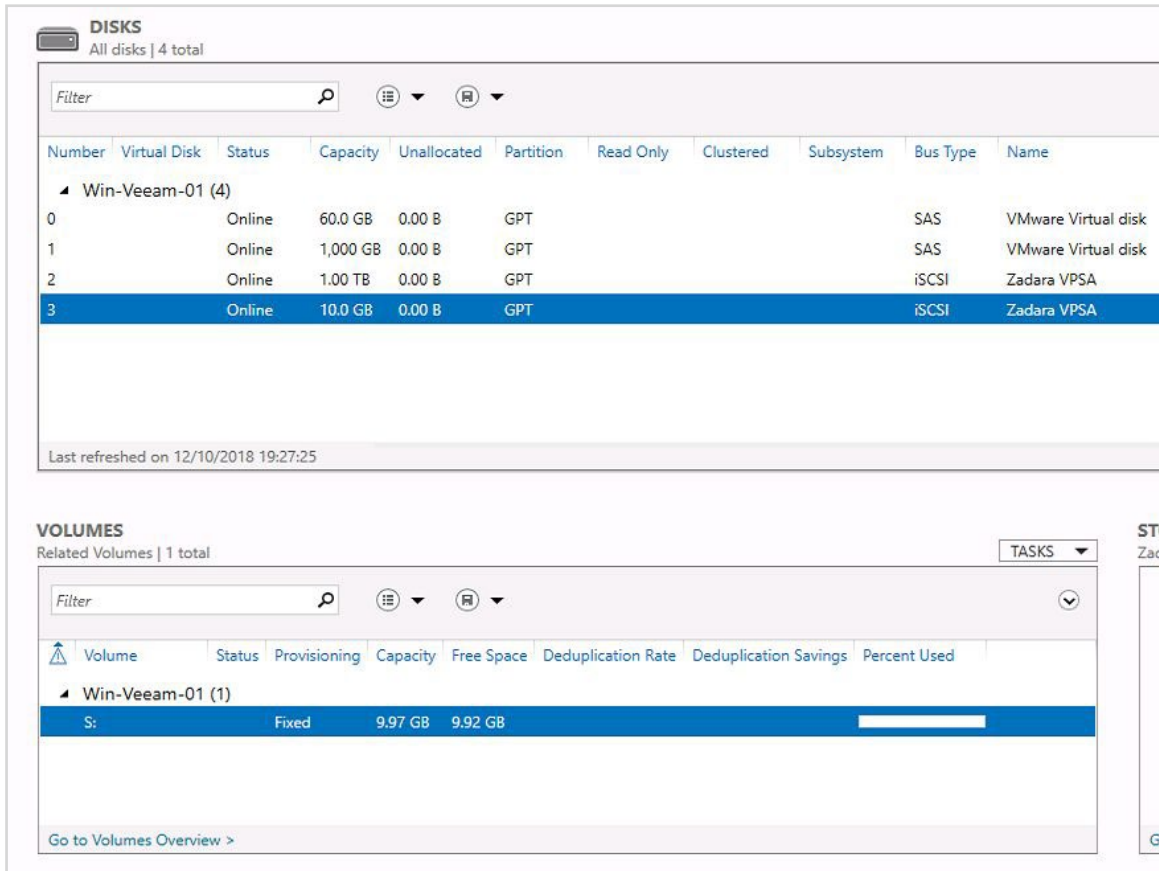
* Shadow copy ID = {3f40063f-f9f4-4fcf-81d5-20290d663c3c}           %VSS_SHADOW_1%
- Shadow copy set: {5891544f-28e5-407b-8026-1d3347dce18f}       %VSS_SHADOW_SET%
- Original count of shadow copies = 1
- Original volume name: \\?\Volume{f7265667-489c-4b27-83b5-b692b14dfcac}\ [volume not on this machine]
- Creation time: 12/10/2018 17:35:31
- Shadow copy device name: \\?\Volume{35fa8b4e-cd81-11e8-9c66-005056875b46}
- Originating machine: Zadara-WIN-DC1.zadara.local
- Service machine: Zadara-WIN-DC1.zadara.local
- Not exposed
- Provider ID: {6a111feb-1ea9-496e-ab0e-b85ff8a42e46}
- Attributes: Transportable No_Auto_Release Persistent Hardware Imported

Number of shadow copies listed: 1


DISKSHADOW> expose {3f40063f-f9f4-4fcf-81d5-20290d663c3c} S:
The shadow copy was successfully exposed as S:\.

DISKSHADOW>
```

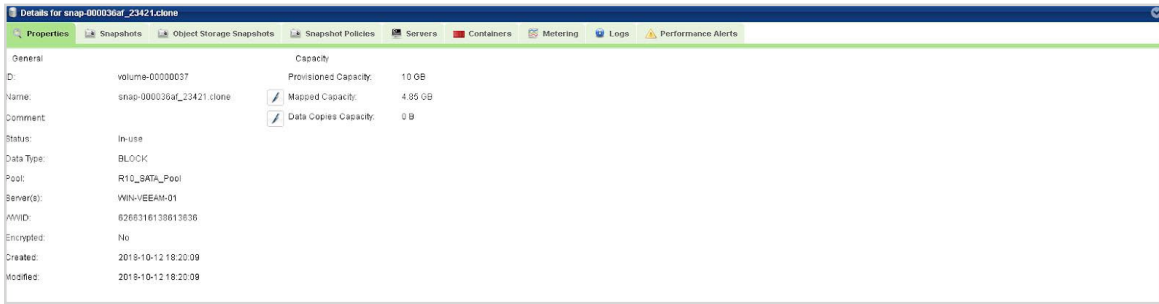
If we now check the Windows Disk Manager we can see the disk has mounted.



Windows Explorer and we can see the files present at the time of the snapshot on the original host.



Back to our VPSA Array GUI and we can see we have a new volume with the name of the snapshot, starting with its ID, some unique identifier, and ending in ".clone".



As can be seen, this clone is consuming no additional space, yet it is a completely read / write independent binary copy of the source volume. Only the differences ("deltas") between the source volume and the clone, caused by modifications of either the source or clone volumes after clone creation, will consume capacity. This instant availability and "only store the differences" approach results in massive space, time and cost savings.

A nice feature of DiskShadow is you can use an input file to automate the above commands with a scripting language like PowerShell. This will work for any VSS Aware Application, SQL Server, Exchange, Oracle etc. on a Windows Server. For other applications such as Linux you can use application tools such as RMAN to place the Database into hot standby mode and use the Zadara API or the Python module zadarapy to do the automation.

## VEEAM / ZADARA VSS HW PROVIDER INTEGRATION

The point of all of this is that, with the Veeam Backup and Recovery Agent installed on a host, it is possible to do a server-free backup using a transportable snapshot script to present and mount a Windows application on a backup server. Doing so reduces the load on production servers during backup hours. Details on utilising pre- and post-scripts with Veeam update 4 can be found here [https://helpcenter.veeam.com/docs/backup/vsphere/backup\\_job\\_vss\\_scripts\\_vm.html?ver=95u4](https://helpcenter.veeam.com/docs/backup/vsphere/backup_job_vss_scripts_vm.html?ver=95u4).

Details on scripting DiskShadow can be found here <https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/diskshadow>.

Here is a simple script for the application to create a transportable snapshot:

```
> Diskshadow -s myVSSCreateScript.txt
```

```
#Zadara Diskshadow transportable file myVSSCreateScript.txt #  
reset  
set context persistent  
set option transportable  
set metadata "z:\meta\ESX_Guest_Snapshot_Meta.cab"  
begin backup  
add volume e:  
create  
  
end backup  
#end of myVSSCreateScript.txt
```

To import on the Veeam Backup Server or any other server:

```
> Diskshadow -s myVSSImportScript.txt
```

```
#Zadara Diskshadow Import VSS file myVSSImportScript.txt #  
reset  
load metadata "z:\meta\ESX_Guest_Snapshot_Meta.cab"  
import  
expose %VSS_SHADOW_1% S:  
  
#end myVSSImportScript.txt
```

**Note:** A clone presented to a Windows host is treated as read-only by that Windows host, even though it is read / write as far as the Zadara storage is concerned. To gain write access from Windows, break VSS snapshot with the following command before presenting the volume.

```
> break writable %VSS_SHADOW_1%
```

Once the volume is no longer required you can unpresent it and remove the shadow. This will remove the volume from the VPSA also, unless the volume is also mapped to another host.

```
> Diskshadow -s myRemoveVSSClone.txt
```

```
#Zadara Diskshadow remove Clone file myRemoveVSSClone.txt #  
unexpose S:  
delete shadows %VSS_SHADOW_1%  
reset  
exit  
  
#end myRemoveVSSClone.txt
```

# SUMMARY

In this document, we have demonstrated deploying Zadara enterprise storage-as-a-service with VMware Cloud on AWS can deliver a simple-to-operate and highly available hybrid cloud. By combining the flexibility and performance of Zadara with the native hybrid cloud capabilities of VMware Cloud, we enable simplified workload migrations between VMware Clouds, wherever located, without modifying applications and without compromising security or performance.

Having established our “basic” hybrid cloud, we then extended it, adding three key capabilities:

- Enterprise data protection and availability with Veeam Availability Suite 9.5
- VDI, including Windows-native self-service data recovery, with VMware Horizon 7 View
- High availability for Windows enterprise applications with Windows Failover Server Cluster and application-consistent snapshots

Zadara is uniquely positioned to deliver both technical and commercial benefits to organisations looking to implement VMware in AWS, Backup Solutions such as Veeam, Failover Clusters, VDI, Hybrid Cloud, utilise Disaster Recovery as a Service and many more solutions.

**zadara**

Zadara is enterprise storage made easy.  
Any data type. Any protocol. Any location.

Contact us at:

[www.zadara.com](http://www.zadara.com)

[info@zadara.com](mailto:info@zadara.com)