

Dell EMC Data Domain Virtual Edition

Version 4.0

Installation and Administration Guide

302-005-025

REVISION 03

Copyright © 2012-2018 Dell Inc. or its subsidiaries. All rights reserved.

Published September 2018

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS-IS." DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.

Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

Preface	7	
Chapter 1	Introduction to DD VE	9
	Revision history.....	10
	What is DD VE?.....	10
	DD VE in the Cloud or on Premise: Overview.....	10
Chapter 2	DD VE on Premise	13
	DD VE capabilities	14
	DD VE management.....	15
	Supported virtual environments.....	15
	Provision physical storage.....	16
	Raw physical capacity needed	16
	Overview of SSD cache tier.....	17
	SSD cache tier—Supported models and their SSD requirements....	
	18	18
	Rules for adding storage to the cache tier.....	18
	SSD cache tier considerations.....	18
	SSD cache tier—Using the DAT tool to meet SSD IOPS	
	requirements.....	19
	Managing SSD cache tier in DD VE (GUI).....	20
	Managing SSD cache tier in DD VE CLI.....	22
	Performance Monitoring	23
	Configuration of other resources	24
	Configuration requirements for DD Cloud Tier support	26
	How to Install the DD VE on Premise.....	27
	Performing the DD VE Download	28
	Performing the DD VE Installation	28
	Initial virtual machine configuration.....	38
	Upgrading DD VE.....	40
	Powering on the virtual machine	40
	Define the Data Domain System Information for Your Site	41
	Initial System Configuration	43
	DD VE storage guidelines.....	43
	Configuring DD VE in Data Domain System Manager.....	44
	Provisioning the storage with the CLI.....	48
	Configure the System for Data Access	49
	Configuration of optional software and internal licenses.....	50
	Optional Additional System Configuration	50
	Overview of Deploying the DD VE on KVM Hypervisor.....	51
	Configuration of other resources	51
	Deploying the DD VE on the KVM Hypervisor.....	52
	Deploying the DD VE on KVM with a reference script.....	52
	Deploying the DD VE on KVM using the Virtual Machine Manager....	
	54	
	Deploying the DD VE on KVM Hypervisor Best Practices.....	55

Chapter 3	DD VE in the Cloud -- All providers	57
	Supported cloud environments.....	58
	DD VE capabilities	58
	DD VE in the Cloud -- DD VE on S3 Storage.....	60
	Overview of DD VE on S3 Storage.....	60
	Deploying DD VE on AWS.....	60
	AWS System Configuration requirements	60
	AWS Storage Size Specifications	61
	Deploy the DD VE in AWS with Cloud Formation Template	61
	Configuring the DD VE on S3 Storage.....	63
	Creating an S3 bucket.....	63
	AWS role-based access for S3 object store access.....	64
	Configuring DD VE using the Data Domain System Manager interface.....	67
	Configuring DD VE using initial configuration or config setup.....	69
	Configuring DD VE using the Command Line Interface.....	71
	System Headswap.....	75
	System Recovery.....	76
	Instance Upgrade and Capacity Expansion.....	78
	DD VE on Azure hot blob storage.....	78
	Azure System Configuration requirements	78
	Specifications for DD VE in Azure	79
	Creating DD VE from Azure Marketplace.....	79
	Adding disks in Azure.....	80
	Resizing the DD VE instance in Azure	83
	Capacity expansion in Azure	84
	Configuring DD VE on Azure hot blob storage using the Data Domain System Manager interface.....	85
	Configuring DD VE in Azure using the Command Line Interface.....	87
	System Headswap - Azure.....	92
	System Recovery - Azure.....	93
	VMware Cloud on AWS	94
	VMC System Configuration requirements	95
	DD VE Deployment in VMware Cloud on AWS	95
	Configuring DD VE in VMC using the Data Domain System Manager interface.....	98
	DD VE in VMC configuration using CLI.....	98
	VMware Cloud on AWS commands.....	99
	DD VE in VMC System HeadSwap	99
	DD VE in VMC System Recovery.....	99
	Initial System Configuration	99
	Provisioning the storage with the CLI	100
	Configuring DD VE in Data Domain System Manager.....	102
	Configure the System for Data Access	108
Chapter 4	DD VE Administration - on Premise and in Cloud	109
	Served Licensing Model for DD VE.....	110
	Adding virtual storage	110
	Disk (Spindle) Group Configuration	111
	Optional Additional System Configuration	112
	Extensions to DD OS for DD VE.....	112
	Storage performance evaluation.....	113
	perf.....	114
	system vresource.....	115
	DD VE-only commands.....	115

System Recovery CLI.....	118
Modified DD OS Commands.....	120
Performance Troubleshooting	122
Unsupported DD OS Commands	123
Upgrading DD VE.....	128
Define the Data Domain System Information for Your Site	128
Setting Up NTP Time Synchronization	130
Configuration of optional software and internal licenses.....	131
Migrating DD VE.....	131
Appendix A DD VE Cloud Workflow and Recommendations	133
Overview of DD VE on Block Storage.....	134
Configuring DD VE on block storage using the DDSM interface...	134
Network infrastructure setup.....	134
Ports for inbound traffic.....	136
Ports for outbound traffic.....	138
Network setup recommendations.....	139
Setting Up NTP Time Synchronization in AWS.....	140

CONTENTS

Preface

As part of an effort to improve its product lines, we periodically release revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Purpose

This manual describes how to install, configure, and administer Data Domain Virtual Edition (DD VE) systems.

Audience

This manual is intended for use by both system administrators and general users of Data Domain Virtual Edition.

Related documentation

The following publications and websites provide additional information:

- *Data Domain Operating System Release Notes*
- *Data Domain Operating System Initial Configuration Guide*
This manual explains configuration steps that are common to hardware and virtual Data Domain systems.
- *Data Domain Operating System OS Command Reference Guide*
This manual explains how to administer Data Domain systems from the command line.
- *Data Domain Operating System OS Administration Guide*
This manual explains how to administer Data Domain systems with the System Manager graphical user interface.
- *Data Domain Boost for OpenStorage Administration Guide*
This manual explains how to use the DD Boost protocol for data transfer between backup software and Data Domain systems.
- *Avamar, Data Domain and NetWorker Compatibility Guide*: <http://compatibilityguide.emc.com:8080/CompGuideApp/>
This website lists Avamar and NetWorker software support for DD VE.

Where to get help

We support, product, and licensing information can be obtained as follows:

Product information

For documentation, release notes, software updates, or information about products, go to Online Support at <https://support.emc.com>.

Technical support

For technical support of this release of DD VE, go to Online Support at <https://support.emc.com>.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to DPAD.Doc.Feedback@emc.com.

CHAPTER 1

Introduction to DD VE

This chapter includes the following topics:

- [Revision history](#)..... 10
- [What is DD VE?](#)..... 10
- [DD VE in the Cloud or on Premise: Overview](#)..... 10

Revision history

Table 1 Data Domain Virtual Edition 4.0 Installation and Administration Guide revision history

Revision	Date	Description
01	July 2018	Initial Publication of DD VE 4.0 with DD OS 6.1.2.5 release.
02	July 2018	AWS content updates.
03	September 2018	Azure content updates.

What is DD VE?

Data Domain Virtual Edition (DD VE) is a software-only protection storage appliance: a virtual deduplication appliance that provides data protection for entry, enterprise and service provider environments. Like any Data Domain system, DD VE is always paired with backup software.

DD VE runs the Data Domain Operating System (DD OS), and provides the DD OS command line interface (CLI) and the Data Domain System Manager graphical user interface (GUI) for performing all system operations.

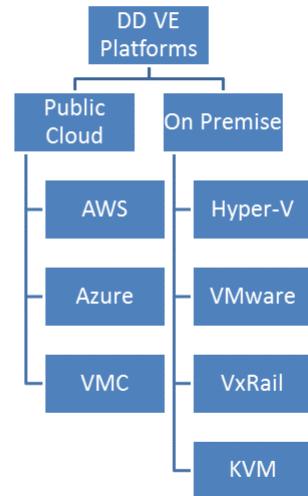
Data Domain Virtual Edition maintains the core Data Domain features that differentiate it as the industry-leading protection storage. This includes high-speed, variable length deduplication for a 10 - 30x reduction in storage requirements, unparalleled data integrity to ensure reliable recovery, and seamless integration with leading backup and archiving applications.

DD VE also comes with DD Boost, which speeds backups by 50%, DD Encryption for enhanced security of data, and DD Replicator, which enables network efficient replication for faster time-to-DR readiness.

DD VE runs on two platforms, on-premises or in the public cloud. On-premises, DD VE supports VMware, Hyper-V, KVM, and VxRail. DD VE also runs in the Amazon Web Services (AWS), Azure, and VMware Cloud on AWS public cloud platforms. For more information about the features and capabilities of Data Domain systems (both physical and virtual), see the *Data Domain Operating System Administration Guide*.

DD VE in the Cloud or on Premise: Overview

The release of Data Domain Virtual Edition (DD VE) 4.0 (with DD OS 6.1.2.5) has in the cloud and on-premise capabilities. Customers now have the ability to deploy DD VE in the cloud to protect their applications in the cloud. For example, customers can back up and replicate data within the cloud or do so on-premise. Customers can replicate data between on-premise and cloud or between different regions in the cloud.



In the Cloud

DD VE enables data protection in the cloud for applications running in the cloud. The DD VE is deployed in the cloud and backup applications can use any existing backup protocols such as DDBlack or BoostFS to backup and restore data to and from the DD VE.

Data replication options that are available include:

- Replication in the cloud among DD VE instances in the same cloud. Replication may be among different available zones and regions.
- Replication in a hybrid cloud environment with your on-premise and cloud environments. Such replications among the DD VEs may be bi-directional.
- Mtree replication and managed file replication are supported, but directory replication and collection replication are not supported.

DD VE runs in Amazon Web Services (AWS), Azure, and VMC to provide these in the cloud capabilities. DD VE runs the Data Domain Operating System (DD OS), and provides the DD OS command line interface (CLI) and the Data Domain System Manager graphical user interface (GUI) for performing all system operations.

On premise

DD VE provides the capabilities of a physical Data Domain system in a virtual machine template for VMware ESXi, Microsoft Hyper-V, or KVM hypervisor on supported Linux distributions.

Data Domain Virtual Edition maintains the core Data Domain features that differentiate it as the industry-leading protection storage. This includes high-speed, variable length deduplication for a 10 - 30x reduction in storage requirements, unparalleled data integrity to ensure reliable recovery, and seamless integration with leading backup and archiving applications.

Note

For more information about the features and capabilities of Data Domain systems (both physical and virtual), see the *Data Domain Operating System Administration Guide*.

CHAPTER 2

DD VE on Premise

This chapter covers the requirements for DD VE on premise including Hyper-V, KVM, VMware ESXi or VxRail 4.0. This chapter includes the following topics:

• DD VE capabilities	14
• DD VE management	15
• Supported virtual environments	15
• Provision physical storage	16
• Overview of SSD cache tier	17
• Performance Monitoring	23
• Configuration of other resources	24
• Configuration requirements for DD Cloud Tier support	26
• How to Install the DD VE on Premise	27
• Define the Data Domain System Information for Your Site	41
• Initial System Configuration	43
• Configuration of optional software and internal licenses	50
• Optional Additional System Configuration	50
• Overview of Deploying the DD VE on KVM Hypervisor	51
• Configuration of other resources	51
• Deploying the DD VE on the KVM Hypervisor	52

DD VE capabilities

DD VE provides the capabilities of a physical Data Domain system in a virtual machine template for VMware ESXi, Microsoft Hyper-V, or KVM. DD VE is available in the following resource configurations:

- Up to 500 GB (evaluation version only)
- Up to 4 TB
- Up to 8 TB
- Up to 16 TB
- Up to 32 TB
- Up to 48 TB
- Up to 64 TB
- Up to 96 TB

Actual DD VE capacity is available in 1 TB increments starting at 1 TB, and up to 96 TB.

The following sections list supported and unsupported Data Domain protocols and features in DD VE.

Supported Data Domain protocols

- CIFS
- NFS
- Data Domain Boost (DD Boost) over IP
- Data Domain Boost (DD Boost) FS

Supported Data Domain features

- DD Boost managed file replication (MFR)
- Encryption
- Data Domain System Manager GUI for DD VE management
- Secure multitenancy (SMT) with Network Isolation Support in 6.0
- Data Domain Cloud Tier (supported in 16 TB, 64 TB, and 96 TB configurations)
- Hadoop Application Agent
- KMIP
- More restricted IPtables settings
- Managed file replication and Mtree replication
- Retention Lock Governance Edition is supported on DD VE (both on-prem and in the cloud)
- KVM
- VMware
- Hyper-V (Hyper-V HA)

Please see the DD OS Administration Guide, DD Boost OST Guide, DD Boost for Partner Integration Administration Guide for additional information on the supported protocols and features above.

Unsupported Data Domain features

- DD Boost over FC
- Extended retention
- DD High Availability (HA)--however, VMware and Hyper-V HA are supported
- NDMP
- VTL
- Collection replication
- Directory replication
- Instant access
- Retention Lock Compliance Edition

DD OS commands related to these unsupported features, and commands for hardware features that are not applicable to a virtual machine, are not supported on the DD VE platform.

DD VE management

Use the VMware vSphere client software, the VMware vSphere Web client, or Hyper-V Manager to install the DD VE and define its virtual hardware: CPUs, memory, network interfaces, and virtual disks.

The DDSM GUI provides a configuration wizard to guide you through the steps required to configure DD OS after the virtual appliance is installed.

Once the DD VE is configured and running, you can access the system console to run DD OS commands by using the VMware vSphere Web client, or Hyper-V Manager. You can also administer the DD VE by using a terminal emulator or `ssh` command line to use the command-line interface.

The default login credentials for the DD VE instance are:

- Username: `sysadmin`
- Password: `changeme` or the password specified during deployment

Note

The system may panic if an RSA DPM client certificate is within 15 days of expiring.

Supported virtual environments

DD VE is supported in the following virtual environments:

- Microsoft Windows Server 2012 R2 with Hyper-V.
- KVM hypervisor on supported Linux distributions
- VMware ESXi servers, either standalone or managed by VMware vCenter, versions 5.5, 6.0, and 6.5 with the corresponding versions of the VMware vSphere client application.

Note

DD VE supports virtual hardware versions of virtual machines up to the latest version of the ESXi in use, and minimum of version 10. Consult VMware latest documentation for any virtual hardware version upgrades.

Table 2 Supported Virtual Hardware

ESXi Version	Up to Virtual Hardware version
ESXi 6.5	13
ESXi 6.0	11
ESXi 5.5	10

Note

The OpenVMTools are pre-installed on the DD VE image. When you update the DD OS software on the DD VE from within the DD OS environment, any necessary updates to the OpenVMTools get installed automatically. It is not possible to update OpenVMTools on the virtual machine from outside DD OS.

The hypervisor user should have only read-only privileges on the data center object where the server or cluster hosting the DD VE instance resides.

For information about compatibility with more recent versions of VMware or Microsoft products, visit the support portal at <https://support.emc.com>.

Provision physical storage

Provision storage on the ESXi or Windows Hyper-V server to host the DD VE instance.

- Provision VMware storage
- Provision Microsoft storage

Raw physical capacity needed

The table below shows the raw capacity needed to get the desired usable capacity. For raw capacities not shown in the table, use the same raw capacity in TiB as the usable capacity.

For example:

1. To get 40 TB usable capacity, you need to provision 40 TiB.
2. For 5 TB capacity with 4 TB configuration, provision 5 TiB.

Important: The capacity in vCenter or Hyper-V manager is in TiB. When you create a virtual disk in vCenter of 1 TB, a storage capacity of 1 TiB is allocated.

Note

This table does not apply to configurations with DD Cloud Tier.

Table 3 Raw physical capacity requirements

Usable Capacity (TB)	Raw Capacity (GiB) at each configuration						
	4	8	16	32	48	64	96
0.5	620	622	620	632	694	752	846
1	1114	1116	1114	1126	1188	1245	1340
2	2100	2101	2100	2112	2174	2232	2326
3	3072	3072	3072	3098	3161	3218	3312
4	4096	4096	4096	4096	4148	4205	4299
5		5120	5120	5120	5120	5191	5286
6		6144	6144	6144	6144	6144	6272
7		7168	7168	7168	7168	7168	7259
8			8192	8192	8192	8192	8245
9				9216	9216	9216	9216

Table 4 Conversions for raw physical capacity

GiB	TB
1	0.001074
TB	TiB
1	0.909495

Overview of SSD cache tier

SSD cache tier for DD VE systems creates caches for file system metadata using flash technologies. The SSD Cache is a low latency, high input/output operations per second (IOPS) cache to accelerate metadata and data access.

Note

- Supported only for DD VE on ESX, it is not supported for DD VE on HyperV or KVM.
- The minimum software version required is DD OS 6.1.2.5.

The SSD cache tier provides the SSD cache storage for the file system. The file system draws the required storage from the SSD cache tier without active intervention from the user. Caching the file system metadata on SSDs improves I/O performance for both traditional and random workloads:

- For traditional workloads, offloading random access to metadata from HDDs to SSDs allows the hard drives to accommodate streaming write and read requests.
- For random workloads, SSD cache provides low latency metadata operations, which allows the HDDs to serve data requests instead of cache requests.

Read cache on SSD improves random read performance by caching frequently accessed data. Writing data to vNVRAM combined with low latency metadata

operations to drain the vNVRAM faster improve random write latency. The absence of cache does not prevent file system operation, it only impacts file system performance.

When the cache tier is first created, a file system restart is only required if the cache tier is being added after the file system is running. Additional cache can be added to a live system, without the need to disable and enable the file system (file system restart is not needed). If the file system has not been created or if the file system is not enabled, creating or enabling the file system after adding the cache tier will activate the SSD cache.

Note

- If encryption is enabled for the active tier, data cache and DM cache will not be enabled.
- One specific condition in regard to SSDs, when the number of spare blocks remaining gets close to zero, the SSD enters a read-only condition. When a read-only condition occurs, DD OS treats the drive as read-only cache and sends an alert.

SSD cache tier—Supported models and their SSD requirements

It is important that customers review the SSD requirements for their specific DD VE model. Note that SSD cache tier is not supported for configurations less than 16 TB.

Table 5 SSD models with Requirements

DD VE Model	SSD Requirement
DD VE 16 TB	160 GB
DD VE 32 TB	320 GB
DD VE 48 TB	480 GB
DD VE 64 TB	640 GB
DD VE 96 TB	960 GB

Rules for adding storage to the cache tier

Review the guidelines for adding storage.

- Adding storage to the cache tier requires a CAPACITY license in the system.
- Only add SSD type disks to the cache tier.
- The minimum disk size for adding to the cache tier is 100 GB.
- The tier is under provisioned when the filesystem is enabled; the system alerts the user to add the required cache capacity.
- If the tier is over provisioned; the cache tier storage add fails.

SSD cache tier considerations

Be aware of the following items for SSD cache:

- One specific condition with regard to SSDs is when the number of spare blocks remaining gets close to zero, the SSD enters a read only condition. When a read

only condition occurs, DD OS treats the drive as read-only cache and sends an alert.

- When SSDs are deployed within a controller, those SSDs are treated as internal root drives. They display as enclosure 1 in the output of the `storage show all` command.
- Manage individual SSDs with the `disk` command the same way HDDs are managed.
- Run the `storage add tier cache` command to add an individual SSD or SSD enclosure to the SSD cache tier.
- The SSD cache tier space does not need to be managed. The file system draws the required storage from the SSD cache tier and shares it among its clients.
- The `filesys create` command creates an SSD volume if SSDs are available in the system.

Note

If SSDs are added to the system later, the system should automatically create the SSD volume and notify the file system. SSD Cache Manager notifies its registered clients so they can create their cache objects.

- If the SSD volume contains only one active drive, the last drive to go offline will come back online if the active drive is removed from the system.

SSD cache tier—Using the DAT tool to meet SSD IOPS requirements

Before you begin

It is recommended to run DAT tool before creating file system and after adding SSD to cache tier.

Note

Running the DAT tool after you create the file system will not provide the usual numbers. DAT will perform read only test, if the disks are utilized by the file system.

Procedure

1. Start an SSD cache benchmark test only. See example below where dev3 is a cache disk on 16TB DD VE.

```
disk benchmark start cache dev3
This will take about 5 minutes to complete.
Are you sure? (yes|no) [no]: yes

ok, proceeding.
Checking devices, please wait.
Benchmark test 1 started, use 'disk benchmark watch' to monitor
its progress.

Results:
disk benchmark show
Last benchmark test 1 was completed.
Devices:           dev3 dev3 dev3 dev3 dev3 dev3
Start Time:       2018/06/22 17:23:29
Duration (hh:mm:ss): 00:04:55

Cache Write      Cache Read      Cache Write      Cache Read
Random IOPS      Random IOPS      Random Latency (ms)  Random Latency
(ms)           -----          -----          -----          -----
```

```

----- 102 24092 6.75
0.98
-----
----- This set of devices is suitable for use in a 16 TiB
file system.

2. Start an SSD cache test along with the data disk. See example below where
dev3 is a cache disk and dev4 , dev5 is a data disk on a 16TB DD VE.

disk benchmark start dev4+dev5 cache dev3
This will take about 10 minutes to complete.
Are you sure? (yes|no) [no]: yes

ok, proceeding.

Checking devices, please wait.
Benchmark test 30 started, use 'disk benchmark watch' to
monitor its progress.

Results:
disk benchmark show

Checking devices, please wait.
Benchmark test 30 was completed.
Devices: dev4+dev5 dev3+dev3 dev3+dev3 dev3+dev3
dev3+dev3 dev3+dev3 dev3+dev3
Start Time: 2018/06/22 15:46:01
Duration (hh:mm:ss): 00:09:40

Sequential Read Random Read Random
vNVRAM
Throughput (MiB/s) IOPS Latency (ms)
Write IOPS
-----
----- -----
----- 2322 1244
5.50 n/a -----
----- -----
----- Cache Write Cache Read Cache Write Cache Read
Random IOPS Random IOPS Random Latency (ms) Random Latency
(ms)
-----
----- -----
----- 102 24544 6.79 0.99
-----
----- This set of devices is suitable for use in a 16 TiB file system.

```

Managing SSD cache tier in DD VE (GUI)

Methods to add or remove SSD cache tiers

- **Configuration Wizard while creating file system** SSDs can be added or removed to cache tier using the wizard to while creating the file system. Maintenance > System > Configure System > Configuration Wizard > File System > Configure

Cache Tier

DELL EMC | Data Domain System Manager

Home    

Health

Data Management

Replication

Protocols

Hardware

Administration

Maintenance

System

Model Number	DD3300
Version	6.1.2.1-5948
System Uptime	System has
System Serial No	123456789C
Chassis Serial No	CN7792174

Configuration Wizard

The DD File System will be created without support for Active Tier.

Licenses

Network

File System

- Configure Active Tier
- Configure Cloud Tier
- Configure Cache Tier

Create File System Summary

System Settings

DD Boost Protocol

CIFS Protocol

NFS Protocol

Upload Upgrade Package 

File Name
No record found.
Items Selected: 0

Upgrade packages may be downloaded.

Upgrade History

Version
6.1.2.1-594885
6.1.2.1-594497

Count: 2

Capacity Expansion History

- **Configuration Wizard from start page** Launch the wizard from the starting page to add or remove SSD storage for cache tier
Data Management > File System > Create > Create File System > Configure Cache Tier

The screenshot shows the Dell EMC Data Domain System Manager interface. The left sidebar contains navigation links: Home, Health, Data Management, File System, Replication, Protocols, Hardware, Administration, and Maintenance. The 'File System' link is selected, highlighted in blue. The main content area is titled 'File System' and shows 'System Partitions'. A 'Create' button is visible. The 'Create File System' dialog is open, showing 'Configure Cache Tier (Step 3 of 5)'. It lists 'Addable Storage' with a table where 'dev5' is selected (checked). An 'ADD TO TIER' button is available. Below it, the 'Cache Tier' section shows a table with 'No devices found.' A 'BACK' button is at the bottom right.

DELL EMC Data Domain System Manager

File System

SUMMARY DD ENCRYPTION CHARTS

Create

System Partitions

Partition
/ddvar
/ddvar/core

Create File System

Configure Cache Tier (Step 3 of 5)

Addable Storage

Device	Size
dev5	160.00 GiB

ADD TO TIER REMOVE FROM TIER

Cache Tier

Device	Size
No devices found.	

BACK

- **With or without a file system** Launch page to add or remove SSD in cache tier.
Hardware > Storage > Overview > Cache Tier > Configure >

Configure Cache Tier **The screenshot illustrates adding the selected storage to cache tier.**

The screenshot shows the Dell EMC Data Domain System Manager interface. The left sidebar has a 'Storage' section selected. The main area is titled 'Storage' and shows 'OVERVIEW', 'ENCLOSURES', and 'DISKS' tabs. Under 'OVERVIEW', the 'Cache Tier' section is expanded, showing a 'Configure' button and a 'Disks in Use' table with the message 'No disk groups found.' A callout box on the right says 'Configure Cache Tier Status' and 'Add Disk dev5 to Cache Tier... In Progress' with a green progress icon.

Managing SSD cache tier in DD VE CLI

The following CLI are used to manage SSD cache tier for the DD VE.

Table 6 SSD cache tier CLI

Command	Description
<code># storage add tier cache dev3</code>	Adds storage to the cache tier. <code># storage add tier cache dev3</code> Checking storage requirements...done Adding dev3 to the cache tier...done Updating system information...done dev3 successfully added to the cache tier.
<code># storage remove dev3</code>	Removes storage from the cache tier. <code># storage remove dev3</code> Removing dev3...done Updating system information...done dev3 successfully removed.
<code># storage show tier cache</code>	Displays the storage configured for cache tier. <code># storage show tier cache</code> Cache tier details: Device Device Device Group Size

Table 6 SSD cache tier CLI (continued)

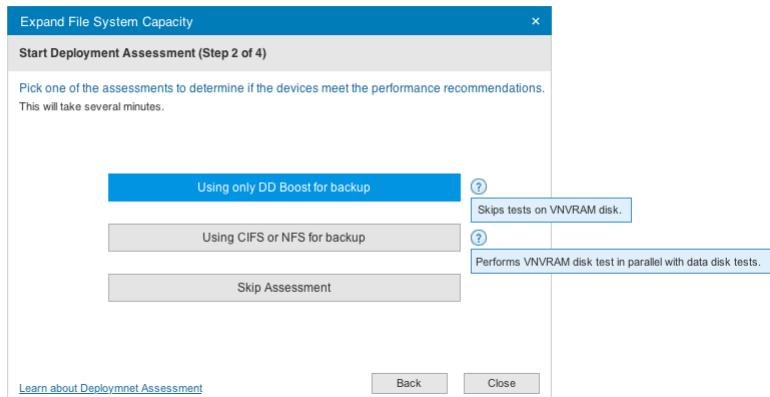
Command	Description							
	<pre> dg0 3 320.0 GiB ----- ----- Spindle Devices Count Total Size Group ----- ----- ----- ----- 1 3 1 320.0 GiB ----- ----- ----- -----</pre> <p>Current cache tier size: 0.3 TiB</p> <p>Storage addable devices:</p> <table> <thead> <tr> <th>Device Type</th> <th>Device</th> <th>Device</th> <th>Size</th> </tr> </thead> </table>				Device Type	Device	Device	Size
Device Type	Device	Device	Size					

Performance Monitoring

Data Domain recommends that you enable the performance monitoring features of the DD VE instance. If you ever need to troubleshoot a DD VE performance problem, you should begin by using the performance monitoring software to detect and resolve any performance problems on the physical storage layer.

Data Domain provides a deployment assessment tool (DAT) which runs on DD VE. It measures the underlying I/O performance and determines the size of the file system. Dat may be used to scan the available physical storage to determine if the storage meets the DD VE requirements. [Table 7](#) on page 25 lists the required physical storage specifications for DD VE.

DD VE 3.1/DD OS 6.1 provides DAT testing for the Boost protocol only-and will skip vNVRAM values to enhance DD VE performance. You may access the DAT using CLIs or from the GUI console of the DD VE (see below):

Figure 1 Expand File System Capacity Screenshot

DAT test results

1. After a serial benchmark test, DAT will parse the list of serial log files from the parallel log file. Then, DAT will open all the serial log files one by one and will parse the performance values for the device and vNVRAM. At the end, DAT will print the average output values of all the devices tested in serial.

2. After a parallel benchmark test, DAT will parse the list of device and vNVRAM log file from the parallel log file. Then, DAT will open the device log file and vNVRAM log file one by one and will parse the performance values. At the end, DAT will print the average output value of all the devices tested in parallel.

Note

The `with-vnvramp` option should be used if you are primarily using NFS to write backups. Some users use Boost to perform backups and then use NFS to get read access, if so, the use of the flag is not required.

Configuration of other resources

This section discusses resources other than storage.

CPU resources

For VMware environment, [Initial virtual machine configuration](#) on page 38 lists the CPU requirements. For Windows environment, the CPU reservation is configured as percentage, “Virtual Machine Reserve” should be set to 100%.

Note

Do not reduce system memory after you have created the file system in DD OS. This makes the file system unusable.

Network adapters

DD VE can support up to eight virtual network adapters.

For VMware environments, the ova package creates two VMXNET3 virtual network adapters by default. DHCP will be configured automatically on these two interfaces inside the DD VE. DHCP can be configured manually on any additional interfaces.

For Windows environments, DHCP will be configured automatically for up to two network interfaces. DHCP can be configured manually on any additional interfaces.

Disk controllers

For VMware environments, DD VE supports up to four VMware Paravirtual SCSI Controllers. Other types of SCSI controllers are not supported.

For Windows environments, DD VE supports up to four Microsoft SCSI controllers.

One SCSI Controller is configured by default. The maximum number of disks for each controller is 15 for vSphere and 64 for Hyper-V. If the environment requires more than the maximum number of disks, you may add extra SCSI HBA controllers to the DD VE system, but do not change the HBA controller type from the type of the first HBA controller. If you make changes accidentally, power off the virtual machine and restore the original settings.

Using resource pools and vApp containers (VMware only)

If you put DD VE systems into resource pools or vApp containers, do not override the default memory and CPU resource allocation settings. The DD VE virtual machine will fail to boot up and report an `insufficient resource` message if it cannot satisfy the minimum resource requirements shown in the next table.

Table 7 DD VE resource reservations

Resources		Up to 4 TB	Up to 8 TB	Up to 16 TB	16 TB to 32 TB	Up to 48 TB	Up to 64 TB	Up to 96 TB
Computing resources	CPU	2 x GHz vCPU		4 x GHz vCPU			8 x GHz vCPU	
	Memory	6 GB	8 GB	16 GB	24 GB	36 GB	48 GB	64 GB
	Shares	Normal						
	Limit	Unlimited						
Underlying storage requirements	Random IOPS	160	320	650	1280	1920	2560	3200
	Random I/O latency	14 ms						
	Sequential throughput	40 MB/s	80 MB/s	160 MB/s	320 MB/s	480 MB/s	640 MB/s	960 MB/s
	RAID	RAID 5/6 or similar fault tolerance storage						
	SCSI controllers	For Hyper-V and ESX: Up to 4 SCSI controllers. For KVM: Virtio SCSI.						
	vNVRAM simulation file size	512 MB	512 MB	512 MB	1 GB	1 GB	1 GB	2GB
		Configurations with DD Cloud Tier support has the same vNVRAM size as the corresponding ones without DD Cloud Tier.						
	System disks	<ul style="list-style-type: none"> 250 GB root disk 10 GB vNVRAM disk 						
		<p>Note</p> <p>The root disk and vNVRAM disk are required to deploy DD VE.</p>						
Data disks		<ul style="list-style-type: none"> The minimum first data disk size: 500GiB for 64TB, Cloud 64TB, 96TB, and Cloud 96TB; 200GiB for all other configurations All subsequent data disks: at least 100 GB <p>Whenever possible, use disks that are larger than the minimum required disk sizes. The maximum capacity of DD VE is defined by the DD VE license, and the maximum virtual disk size supported by the hypervisor.</p>						
		<p>Note</p> <p>The usable capacity available on a data disk is less than the capacity specified when the disk was created because of overhead requirements. DD VE storage guidelines on page 43 provides additional details about DD VE storage overhead requirements.</p>						
Network adapters		Up to 8 network adapters						

The system displays an error message if you attempt to configure a higher capacity with fewer memory and CPU resources than the amounts listed in the table above. To check these settings, use the Resources tab of the Virtual Machine Properties dialog box in the vSphere client, or the `system vresource show requirements` command.

Configuration requirements for DD Cloud Tier support

To configure DD Cloud Tier for each DD VE, refer to The DD OS Administration Guide. This section discusses resources for DD Cloud Tier support for each DD VE.

Table 8 DD VE resource reservations

Configuration	CPU	Memory (GiB)	Minimum Metadata Tier Size(GiB)	Metadata Tier Throughput(MB /S)	Metadata Tier IOPS	Metadata Tier Latency (ms)
16 TB	4	32	500	160	640	14
64 TB	8	60	500	640	2560	14
96 TB	8	80	500	960	3200	14

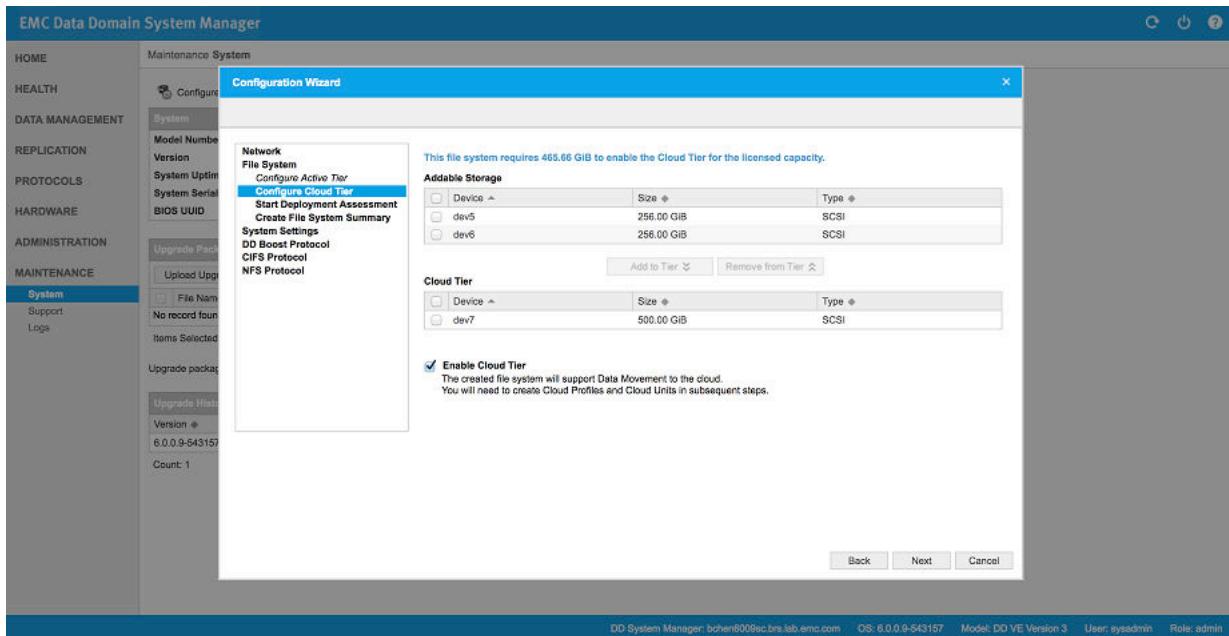
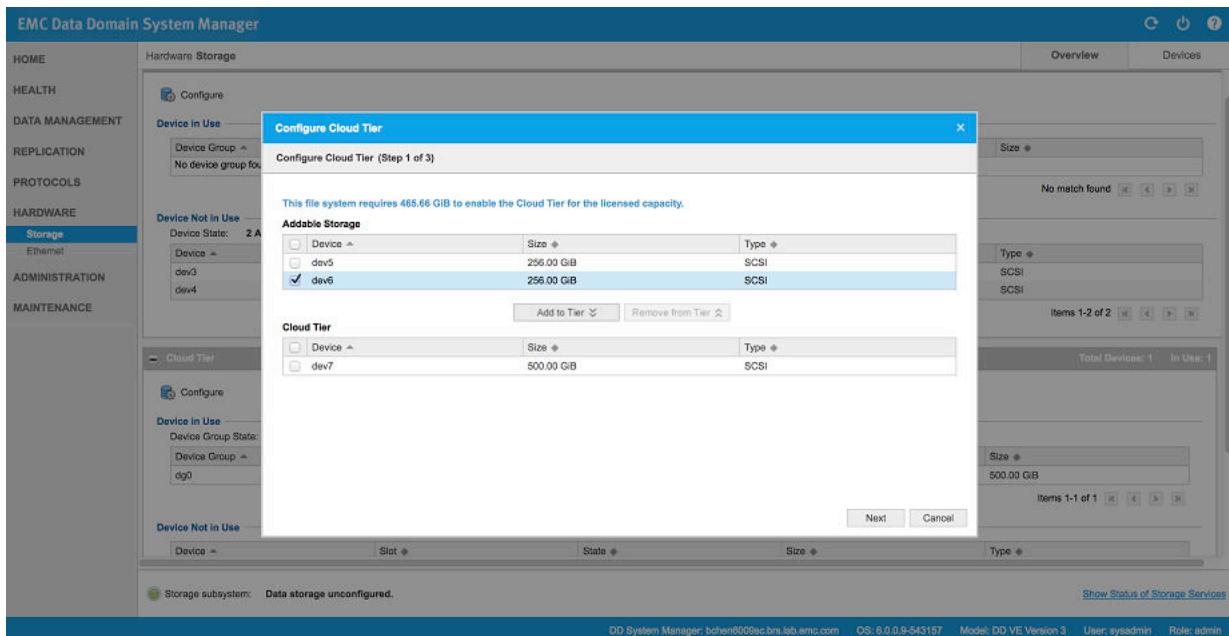
The minimum metadata size is a hard limit. We recommend users start with 1 TB metadata tier and use 1 TB as incremental size.

Table 9 DD Cloud Tier Meta Data Size for DD VE

CU Size(TiB)	1~16	16~32	32~48	48~64	64~80	80~96	96~112	112~128	128~144	144~160	160~176	176~192
MD Size(TiB)	1	2	3	4	5	6	7	8	9	10	11	12

*The CU (TiB) size includes both cloud units.

The table above is a breakdown of recommended metadata tier size for the corresponding cloud unit (CU) size range. Refer to the *EMC Data Domain Operating System 6.0 Administration Guide* for additional information.

Figure 2 DD Cloud Tier Configuration Wizard for DD VE**Figure 3** Configure DD Cloud Tier for DD VE

The images above contain DD Cloud Tier screenshots for DD VE. Refer to the *EMC Data Domain Operating System 6.0 Administration Guide* for additional DD Cloud Tier information.

How to Install the DD VE on Premise

DD VE runs on two platforms, on premise or in the public cloud. On premise, DD VE supports VMware, Hyper-V, VxRail, and KVM. For more information about the features and capabilities of Data Domain systems (both physical and virtual), see the *Data Domain Operating System Administration Guide*.

Performing the DD VE Download

DD VE is packaged as a zip file that contains a virtual machine template (.ova for VMware, or .vhd for Microsoft) file. The zip file is available from EMC Online Support at <https://www.emc.com/products-solutions/trial-software-download/data-domain-virtual-edition.htm>. There are separate zip packages for VMware and Microsoft Hyper-V available for download. Customer should choose the package that best suits the customer's environment.

Note

The 500 GiB evaluation version is also available from the EMC Trial Download page.

All capacity configurations are available from the same zip file.

DD VE requires a minimum of 760 GiB for configurations of 64TB, Cloud 64TB, 96TB, and cloud 96TB. The minimum size of the first data disk is 500 GiB. DD VE requires a minimum of 460 GiB of available storage to deploy. The 460 GiB breaks down as follows:

- System disk: 250 GiB
- vNVRAM disk: 10 GiB
- First data disk: 200 GiB

Performing the DD VE Installation

Before you begin

Be sure you have downloaded the DD VE template, or know its URL, as described in [Performing the DD VE Download](#) on page 28.

Installation procedures are included for:

- [Installing on a VMware ESXi Server](#) on page 28
- [Installing through a VMware vCenter Server](#) on page 29
- [Installing on a Microsoft Windows Server 2012 R2 or Windows Server 2016 with Hyper-V Server](#) on page 30
- You can also use `ovftool` to install the software from the command line. For help, see the VMware documentation.

Installing on a VMware ESXi Server

[Table 10](#) on page 28 lists the information required to deploy the DD VE instance on a VMware ESXi server.

Table 10 Installing DD VE on a VMware ESXi server

Installation step	Description
Username and password for the ESXi server.	Specify the credentials in the vSphere Client to log into the ESXi server.
Launch virtual machine deployment wizard.	Use the VMware deployment wizard to deploy the DD VE instance.
Choose the deployment method.	Deploy from a local file, or deploy from a network location.

Table 10 Installing DD VE on a VMware ESXi server (continued)

Installation step	Description
Review the VM details.	Review the details to this point, and proceed if they look correct.
Review the End User License Agreement (EULA).	Accept the EULA.
Specify a name for the DD VE virtual machine.	This name identifies the virtual machine on the VMware server; it does not become a host name on your LAN.
Choose a datastore to host the DD VE instance.	Select the datastore where the DD VE instance will reside. For best performance, Data Domain recommends that you use a dedicated datastore, one that is not shared by other virtual machines.
Review and complete the deployment.	Review the deployment summary and finish the wizard.
Configure the virtual machine	See Initial Virtual Machine Configuration with the vSphere Client .

The hypervisor documentation provides additional details.

Installing through a VMware vCenter Server

[Table 11](#) on page 29 lists the information required to deploy the DD VE instance on a VMware vCenter server.

Note

When using version 5.5 of the vSphere web client to install DD VE, the system displays the following warning: The OVF package contains extra configuration options, which possess a potential security risk. Review the extra configuration options below and accept to continue the deployment. Select **Accept extra configuration options** to continue.

Table 11 Installing DD VE on a VMware vCenter server

Installation step	Description
Username and password for the vCenter server.	Specify the credentials in the vSphere Client to log into the vCenter server.
Launch virtual machine deployment wizard.	Use the VMware deployment wizard to deploy the DD VE instance.
Choose the deployment method.	Deploy from a local file, or deploy from a network location.
Review the VM details.	Review the details to this point, and proceed if they look correct.
Review the End User License Agreement (EULA).	Accept the EULA.

Table 11 Installing DD VE on a VMware vCenter server (continued)

Installation step	Description
Specify a name for the DD VE virtual machine.	This name identifies the virtual machine on the VMware server; it does not become a host name on your LAN.
Select an Inventory Location.	Select the inventory location, or data center to assign the DD VE instance to a host or cluster.
Select a host or cluster.	Choose a host or cluster in the specified inventory location or data center where the DD VE instance will reside.
Choose a datastore to host the DD VE instance.	Select the datastore where the DD VE instance will reside. For best performance, Data Domain recommends that you use a dedicated datastore, one that is not shared by other virtual machines.
Choose the format for the virtual disks.	Data Domain recommends Thick Provision Lazy Zeroed to provide the best balance of performance and deployment time. Thick Provision Eager Zeroed provides the best performance, but takes a long time to deploy.
Review and complete the deployment.	Review the deployment summary and finish the wizard.
Configure the virtual machine	See Initial Virtual Machine Configuration with the vSphere Client .

The hypervisor documentation provides additional details.

Installing on a Microsoft Windows Server 2012 R2 or Windows Server 2016 with Hyper-V Server

[Table 11](#) on page 29 lists the information required to deploy the DD VE instance on a Windows server.

Note

There are three ways to perform this installation: creating a VM, running the powershell script to install DD VE on Hyper-V manager machine, or running the powershell for MS System Center.

Installing on a Microsoft Windows Server 2012 R2 or Windows Server 2016 with Hyper-V Server by creating a VM

Table 12 Hyper-V Installation for Windows via VM

Installation step	Description
Username and password for the Windows server.	Specify the credentials to log into the Windows server.
Launch virtual machine deployment wizard.	Use the Hyper-V deployment wizard to deploy the DD VE instance.

Table 12 Hyper-V Installation for Windows via VM (continued)

Installation step	Description
Specify a name for the DD VE virtual machine.	This name identifies the virtual machine on the Windows server; it does not become a host name on your LAN.
Specify the amount of memory.	Assign memory to the virtual machine.
Configure networking.	Connect the DD VE virtual machine to the Hyper-V networking switch.
Select the virtual disk.	Select the .vhd file that contains the DD VE instance.
Configure the virtual machine	See Initial Virtual Machine Configuration .

The hypervisor documentation provides additional details.

Installing on a Microsoft Windows Server 2012 R2 or Windows Server 2016 with Hyper-V Server via powershell script for Hyper-V Manager

This installation script deploys DD VE on Hyper-V.

Syntax: C:\ddve-hyperv-0.6000.11.0-524942\ddve-installer.ps1 [-VMName] <String> [-Configuration] <String> [[-VirtualMachinePath] <String>] [[-VirtualHardDiskPath] <String>] [-Force] [<CommonParameters>]

Example: C:\PS>ddve-installer.ps1 -VMName DDVE -Configuration 4TB -VirtualMachinePath C:\DDVE -VirtualHardDiskPath C:\DDVE

Table 13 Hyper-V Installation for Windows Powershell Script

Parameter	Description
-VMName <String>	Specify the name of DD VE virtual machine
-Configuration <String>	Specify the configuration of the DD VE. This parameter accepts one of these values 4TB, 8TB, 16TB, 32TB, 48TB, 64TB, 96TB.
-VirtualMachinePath <String>	Specify the directory to store files for the DD VE virtual machine. This parameter is optional and wildcard characters are supported for this parameter value. If this parameter were not specified in command line, Hyper-V setting would be used for this parameter value.
-VirtualHardDiskPath <String>	Specify the directory to store virtual hard disks for the DD VE. This parameter is optional and wildcard characters are supported for this parameter value. If this parameter were not specified in command line, Hyper-V setting would be used for this parameter value.
-Force [<SwitchParameter>]	
<CommonParameters>	This cmdlet supports the common parameters: Verbose, Debug, ErrorAction, ErrorVariable,

Table 13 Hyper-V Installation for Windows Powershell Script (continued)

Parameter	Description
	WarningAction, WarningVariable, OutBuffer, PipelineVariable, and OutVariable. For more information, see about_CommonParameters , available from the Microsoft website.

Please see ddve-installer-help.txt for additional information.

Installing on a Microsoft Windows Server 2012 R2 or Windows Server 2016 with Hyper-V Server via powershell script for MS System Center

This installation script deploys DD VE on System Center Virtual Machine Manager (SCVMM).

Syntax: C:\ddve-hyperv-0.6000.11.0-524942\ddve-installer-sc.ps1 [-VMName] <String> [-Configuration] <String> [-SCVMMServer] <String> [-SCVMHost] <String> [-SCVMNetwork] <String> [[-VirtualMachinePath] <String>] [[-VirtualHardDiskPath] <String>] [[-NetworkAdapterCount] <Int32>] [[-IPAddress] <String>] [[-Gateway] <String>] [[-Netmask] <String>] [[-DnsServer1] <String>] [[-DnsServer2] <String>] [<CommonParameters>]

Example: C:\PS>.\ddve-installer-sc.ps1 -VMName DDVE -Configuration 4TB -SCVMMServer localhost -SCVMHost osdev-ucs30d -SCVMNetwork mktest-vmnet -NetworkAdapterCount 3

Table 14 Hyper-V Installation for Windows Powershell Script for MS system Center

Parameter	Description
-VMName <String>	Specify the name of DD VE virtual machine.
-Configuration <String>	Specify the configuration of the DD VE. This parameter accepts one of these values 4TB, 8TB, 16TB, 32TB, 48TB, 64TB, 96TB, Cloud16TB, Cloud64TB, Cloud96TB.
-SCVMMServer <String>	Specify the system center VMM server name.
-SCVMHost <String>	Specify the Hyper-V host where DD VE will be deployed.
-SCVMNetwork <String>	Specify a VM Network.
-VirtualMachineHostname <String>	Specify the hostname of the DD VE virtual machine. This parameter is optional and wildcard characters are supported for this parameter value. If this parameter is not specified in command line, localhost would be used for this parameter value.
-VirtualMachinePath <String>	Specify the directory to store files for the DD VE virtual machine. The VirtualMachinePath folder must exist on host SCVMHost. If this parameter is not specified in command line, Hyper-V setting would be used for this parameter value.

Table 14 Hyper-V Installation for Windows Powershell Script for MS system Center (continued)

Parameter	Description
-VirtualHardDiskPath <String>	Specify the directory to store virtual hard disks for the DD VE. The VirtualHardDiskPath folder must exist on host SCVMHost. If this parameter is not specified in command line, Hyper-V setting would be used for this parameter value.
-NetworkAdapterCount <Int32>	Specify the number of network adapters to be added to DD VE. If this parameter is not specified, 2 networks adapters are added. This parameter accepts value in range 1 to 8.
-IPAddress <String>	Specify IP address.
-Gateway <String>	Specify gateway IP address.
-Netmask <String>	Specify netmask.
-DnsServer1 <String>	Specify first DNS server IP address.
-DnsServer2 <String>	Specify second DNS server IP address.
<CommonParameters>	This cmdlet supports the common parameters: Verbose, Debug, ErrorAction, ErrorVariable, WarningAction, WarningVariable, OutBuffer, PipelineVariable, and OutVariable. For more information, see about_CommonParameters, available from the Microsoft website.

Please see ddve-installer-sc-help.text for additional information.

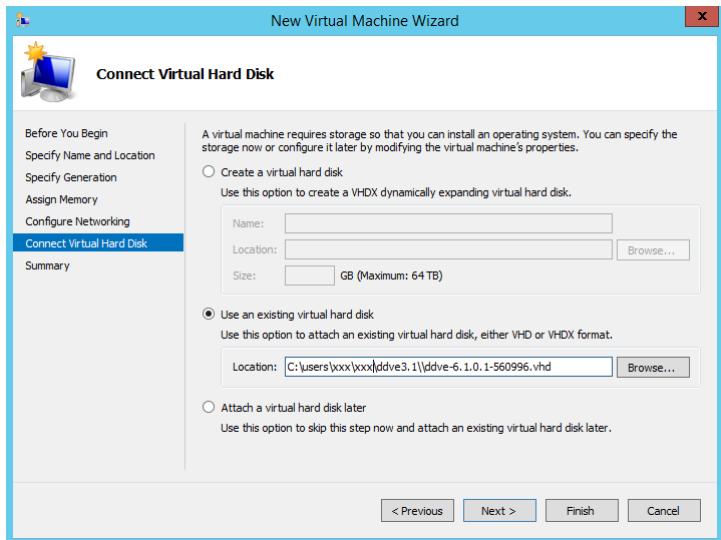
Installing the DD VE on Hyper-V using the GUI

The following highlights the key steps from the New Virtual Machine Wizard to install and configure the DD VE on Hyper-V.

Procedure

1. On the Hyper-V Manager screen, select the managed hyper-v server on which you wish to deploy the new DD VE.
2. Launch the New Virtual Machine Wizard.
3. Click “Next” to begin creating a DD VE virtual machine with a custom configuration.
4. On the “Specify Name and Location” page, to choose a name and location for this virtual machine. In the “Name:” section, enter: a name that will be easily identified for this virtual machine in the “<ENTER DDVE NAME HERE>” cell. Click “Next”.
5. On the “Specify Generation” page, choose “Generation 1”.
6. On the “Assign Memory” page, Enter the exact memory size (a multiple of 512 GB) required memory for the DD VE capacity you are creating. Refer to [Initial virtual machine configuration](#). Then click “Next”.
7. On the “Configure Networking” page, select a configured virtual switch to attach to the new DD VE. Select “Next”.

8. On the “Connect Virtual Hard Disk”, select Use an “existing virtual hard disk” and enter the path to the .vhd file extracted from the DD VE OS download.



Verify your configuration and select “Finish”.

9. On the “Summary” page review your configuration and select “Finish” to deploy your new DD VE VM.

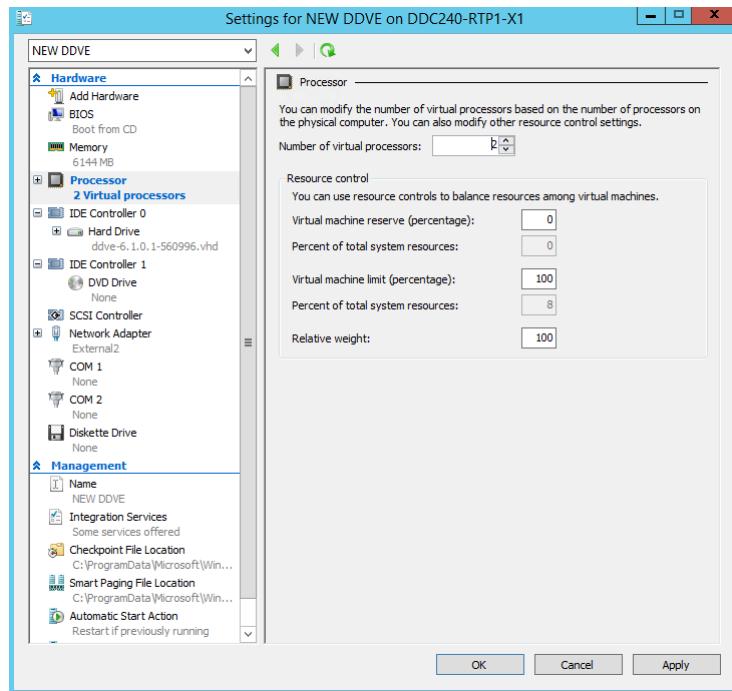
You will see the successful completion message such as, “You have successfully completed the New Virtual Machine Wizard. You are about to create the following virtual machine.”

Configuring the DD VE on Hyper-V using the GUI

The following highlights the key steps from the New Virtual Hard Disk Wizard to configure the newly installed DD VE on Hyper-V.

Procedure

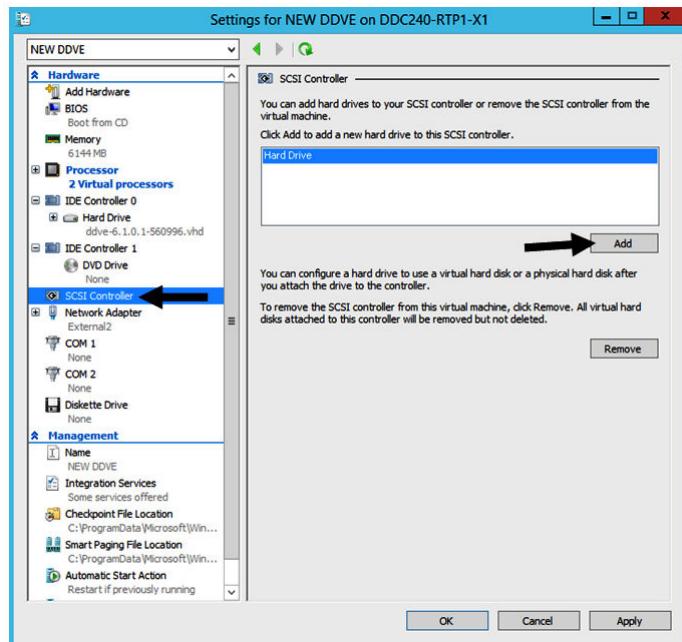
1. Once the VM is deployed, it will be displayed in the Hyper-V manager on the under “virtual Machines” list.
2. Select the new DD VE VM and select the Settings tab to begin configuration of CPU, vNVRAM, and storage.
3. Select the “add hardware” > “processor” tab and add the correct amount of CPU as required for your DD VE capacity. Refer to [Initial virtual machine configuration](#) to make your selection then click “Apply”.



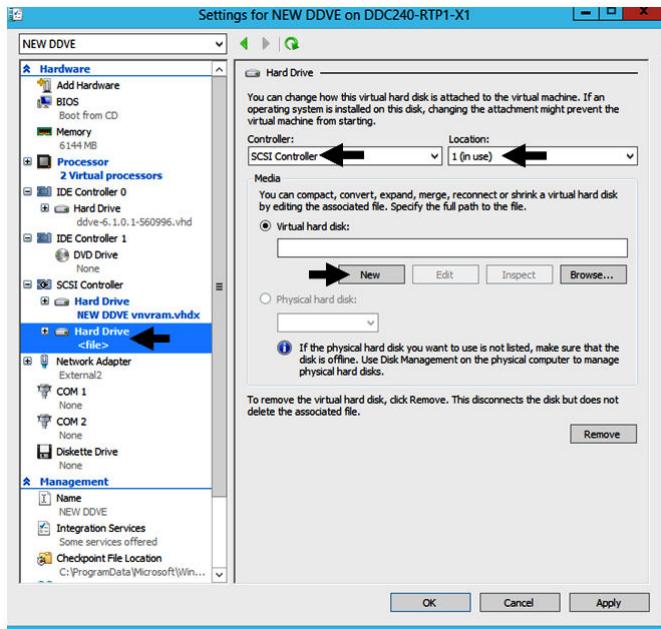
Note

Do not add any restrictions to CPU resources.

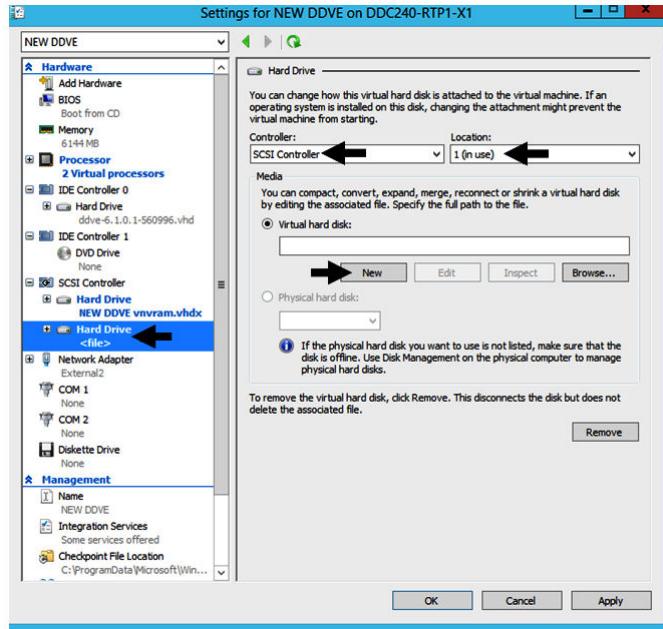
4. Begin disk configuration by adding a 10 GB virtual disk for vram. All disks should be thick provisioned (lazy zero).
5. Select Add hardware and create the first disk. Select the first SCSI controller under IDE controller 1 and click the Add button, then select "New Disk" tab. Then the New Virtual Hard Disk Wizard will open



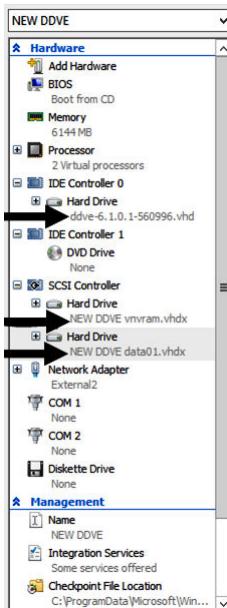
6. Attach the first 10GB virtual disk (vNVRAM) to the first SCSI controller under IDE controller 1. vNVRAM must be set to location zero, then hit "New" button to launch disk configuration wizard.



7. Select the disk format to use for the virtual hard disk on the "Choose Disk Format" page. Select "VHDX".
8. Select the type of disk on the "Choose Disk Type" page. "Fixed size" offers the greatest performance, however you may select "Dynamically expanding" instead. Select Next.
9. On the Specify Name and Location page, name the first disk as a vnram disk. this ensures that this vnram disk resides on the highest performing physical storage.
 - For example, for Name enter "NEW DDVE vnram.vhdx"
 - For Location, enter "D:\PremiumRAIDStorage\Hyper-V\Virtual Hard Disks \
"
 - Then select Next.
10. On the Configure Disk page, select the option "Create a new blank virtual hard disk". Enter 10 for the "size" to create a 10G B disk for vnram. verify your configuration then select Finish.
11. On the Completing the New Virtual hard Disk Wizard summary page. You will see a successful completion message and a window that says "Creating the new virtual hard disk". Verify the configuration then click Next.
12. Next add more virtual disks for backup data. Use the same disk creation wizard to create up to 14 data disks for user data. The data disks should be attached to SCSI controller beginning at location 1.



13. On the Specify Name and Location page, specify the name and location of the virtual hard disk file.
 - Name: New DDVE data01.vhdx
 - Location: D:\PremiumRAIDstorage\Hyper-V\Virtual Hard Disks\
14. Continue within disk creation wizard. Keep in mind that using several smaller disks can offer better overall performance than fewer larger disks.
 - Select "Create a new blank virtual hard disk"
 - Size: 500 GB
15. At this stage, you have created the minimum configuration for a functioning DD VE:
 - a. OS disk now attached at IDE controller zero.
 - b. vNVRAM disk attached at the first SCSI controller : location zero.
 - c. Data storage disk attached at first scsi controller : location one.
 - d. For larger capacity DD VE, additional scsi controllers can be added, however no performance gain is expected.
16. Select your new DD VE and select Start. Then configure the new VM as you would for any other DD appliance.



Initial virtual machine configuration

The DD VE template does not include any storage, so you need to add data disks to the system. The procedure in this section explains how to add the disks before you start the virtual machine.

Note

You can add the first or additional virtual data disks while the virtual machine is running, provided that you do not also need to add more virtual memory to support the additional disks. DD VE supports virtual disk hot-plugging, but not CPU, memory, HBA card or NIC card hot-plugging.

Depending on the amount of disk space, you may also need to add memory to the virtual machine. The next table shows the supported storage configurations and their virtual CPU and memory requirements.

Table 15 Initial virtual machine configuration

Hardware configuration		Storage capacity range (TB)							
		Up to 500 GB	Up to 4 TB	Up to 8 TB	Up to 16 TB	Up to 32 TB	Up to 48 TB	Up to 64 TB	Up to 96 TB
CPU	Topology	1 socket with 2 cores			1 socket with 4 cores			1 socket with 8 cores	
	Reservation	2 x 1.5 GHz			4 x 1.5 GHz			8 x 1.5 GHz	
Memory	Topology	6 GB		8 GB	16 GB	24 GB	36 GB	48 GB	64 GB
	Reservation								

The system displays an error message if you attempt to configure a higher capacity with fewer memory and CPU resources than the amounts listed in the table above.

The `system vresource show requirements` command lists the virtual resources available on the host.

Adding NICs

When initially deployed, DD VE is provisioned with two VMXNET3 NICs which can be configured as required. Additional NIC cards can be added up to a maximum of:

- DD VE 3.1: 8 NICs total

Note

The following applies specifically to VMware:

- Additional NICs can only be of type VMXNET3. NICs of type VMXNET2 and E1000(E) can be added to the DD VE virtual machine, but will not be visible or usable within DD OS on the appliance.
- DD VE does not support hot add of NICs. A VMXNET3 NIC can be added while DD VE is powered on, the NIC will not be visible or useable within the DD OS on the appliance until the DD VE appliance is restarted.
- DD VE does not support hot remove of NICs. Any attempt to remove a NIC while the DD VE appliance is powered on will cause ESXi/vSphere to report the following error: The guest operating system did not respond to a hot-remove request for device ethernet3 in a timely manner. The DD VE appliance must be powered off before NICs can be removed.

Instance Upgrade and Capacity Expansion

For instance upgrade refer to: [Changing the Instance Type](#)

For capacity expansion refer to: [Raw Physical Capacity Needed](#)

Setting Up NTP Time Synchronization

By default, NTP is disabled on the DD VE system. If you need to enable NTP on the DD VE, follow these steps:

Note

Skip this task if you are going to join the DD VE to an Active Directory domain. Because the Windows domain controller obtains the time from an external source, NTP must be configured. See the Microsoft documentation on how to configure NTP for the Windows operating system version or service pack that is running on your domain controller. After joining the domain, the system time is periodically synchronized with the domain controller time. When the host joins the Active Directory, the DD VE displays a warning if multiple time sources are in use.

Later, while performing initial configuration of the DD VE system, enable NTP by selecting the appropriate options from the configuration wizards. If you do not use the wizards to perform initial configuration, you can use the `ntp enable` command on the DD OS command line. Enabling NTP with the `ntp enable` command automatically disables synchronizing the time on the guest to the host time.

To reenable synchronizing the guest time to the host time, run the `ntp disable` command.

Note

NTP is disabled by default. The `ntp reset` command also deactivates NTP on the guest.

For more information about AWS time synchronization, see [AWS NTP Time Synchronization](#).

Upgrading DD VE

Upgrading from DD VE 3.0 and 3.1

DD VE 4.0 uses DD OS 6.1.2.5 Refer to the *Data Domain Operating System 6.1.2 Administration Guide* and the *Data Domain Virtual Edition Installation and Administration Guide* for additional information.

Upgrading DD VE to a higher capacity

If the higher capacity does NOT need additional resources (refer to [Table 7](#) on page 25), follow these steps.

1. Add the needed hard disks for the new capacity
2. Configure the newly added data disks using the CLI command `storage add dev tier active<device ID>` (Or, use DD SM GUI)
3. Expand the file system using the CLI command `filesys expand`

If the higher capacity will require DD VE to have higher resources, follow these steps to upgrade DD VE to a higher capacity.

1. Shutdown DD VE.
2. Switch to the appropriate instance type and refer to [DD VE capabilities table](#) for cloud providers and resource configuration sizes.
 - AWS
 - Azure
3. Add the needed hard disks for the new capacity
4. Power on the DD VE
5. Add the license for the new capacity
6. Configure the newly added data disks using the CLI command `storage add dev tier active<device ID>`
7. Expand the file system using the CLI command `filesys expand`

Powering on the virtual machine

If the installation is successful, you should be able to power on the DD VE virtual machine and log into the system.

Procedure

1. From the Hyper-V or VMware, power on the DD VE virtual machine.

Note

There may be a delay of several minutes until the DD OS prompt appears, depending on your hardware and configuration.

2. Optionally, open the virtual machine console to view the boot and initialization process. You should see the CLI prompt to log in for a successful boot.

3. Note the IP Address assigned to the system by DHCP and shown in the previous figure.

You can use this address to configure or administer the system outside the hypervisor.

After you finish

The next step is the initial system configuration in DD OS. See the *Data Domain Operating System Initial Configuration Guide* for detailed instructions.

Note

To shut down the DD VE virtual machine, shut down the guest operating system from the DDSH with the command `system poweroff` or `system reboot` for reboot. Do not reset or power off the DD VE virtual machine, which will perform a hard reset of the system rather than an orderly shutdown. Currently, the Guest OS shutdown and Guest OS reboot features in the hypervisor also cannot guarantee an orderly shutdown and reboot.

The hypervisor documentation provides additional details.

Define the Data Domain System Information for Your Site

An installation requires information unique to your site. Before starting the installation, provide values for the system information listed below.

Note

Data Domain recommends that you print the tables in this section and record the information. Be sure to enter the serial number correctly to avoid DD VE issues.

Table 16 System Setup Worksheet for DD VE

Information	Your Values
A fully qualified host name for the system:	
The DNS domain name:	
A default gateway IP address (if you are not using DHCP):	
DNS server IP addresses (if you are not using DHCP): <ul style="list-style-type: none"> • Primary • Secondary • Tertiary 	
If you will enable CIFS access, enter the information for your CIFS authentication method: <ol style="list-style-type: none"> 1. For Workgroup authentication: <ul style="list-style-type: none"> • Workgroup name: 	

Table 16 System Setup Worksheet for DD VE (continued)

Information	Your Values
<ul style="list-style-type: none"> • Backup user name: • Password: <p>2. For Active Directory authentication:</p> <ul style="list-style-type: none"> • Realm name: • Domain admin name: • Password 	
Host name from which to administer the system:	
Administrator's email address (or admin group alias):	
Mail server (SMTP) host name:	
Hypervisor server name:	
(Optional) Physical location of the hypervisor server:	
Time zone name (default is US/Pacific):	
Serial number (SN) provided to you by Data Domain:	
Virtual machine unique ID (after initial configuration, use the system show serialno command to display this ID):	

Use this table to enter Ethernet connectivity information. By default, DHCP is enabled on both ports.

Table 17 Ethernet Connectivity Worksheet

Ethernet Connectivity	Enable	Use DHCP	IP Address (if no DHCP)	Netmask (if no DHCP)
ethV0				
ethV1				
ethV2				
ethV3				
ethV4				
ethV5				
ethV6				
ethV7				

Initial System Configuration

You can connect to the system to perform the initial system configuration with the DDSM Configuration Wizard or manually using the CLI.

DHCP is enabled on the DD VE system by default. If the DHCP service is available, the DD VE system will receive IP addresses from the DHCP server.

Note

DHCP is only activated automatically for the first network interface card (NIC) which is built into the virtual machine template. Any extra NICs must be configured manually by following instructions here <https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-network-interface-vm>.

Using the CLI

Access the CLI by using `ssh` or a terminal emulator to access the DD OS command line. The CLI configuration utility contains four sections: Network, eLicense, System, and DD Boost.

Using the GUI

Access DDSM by entering the IP address of the DD VE into a web browser, and logging in. The GUI Configuration Wizard contains six sections: Networking, File System, System Settings, DD Boost, CIFS, and NFS.

[Provisioning the storage with the CLI](#) on page 100 describes how to configure the DD VE manually with the CLI instead of using the configuration utility.

DD VE storage guidelines

Be aware of the following virtual storage guidelines for DD VE before deploying a DD VE instance.

DD VE licenses are sold in Terabytes (TB), but VMware disk sizes are actually measured in Tebibytes (TiB). 1 TB is equal to 0.97 TiB or 1.02 TB is equal to 1 TiB. Because of this, DD VE allows a 10% buffer to reach the licensed capacity measured in TB. Additionally, Hyper-V manager also uses GiB for "GB" in its GUI.

Table 18 TB to TiB equivalents

Licensed capacity in TB	Licensed capacity in TiB	Maximum capacity with buffer in TB
4 TB	3.6 TiB	4.4 TB
8 TB	7.3 TiB	8.8 TB
16 TB	14.6 TiB	17.6 TB
32 TB	29.2 TiB	35.2 TB
48 TB	43.8 TiB	52.8 TB
64 TB	58.4 TiB	70.4 TB
96 TB	87.6 TiB	105.6 TB

Individual virtual disks are subject to overhead that reduce their amount of usable capacity to amounts lower than their specified capacity.

- The first data disk is subject to 120 GB of base overhead, with 5.6% of the remaining capacity reserved for RAID-on-LUN
- All subsequent data disks are subject to 5.6% overhead reserved for RAID-on-LUN

Table 19 Virtual disk overhead calculations

Disk	Overhead calculation	Usable capacity examples
First data disk (200 GB or more)	$(\text{Total capacity} - 120 \text{ GB}) * 0.944$	<ul style="list-style-type: none"> • 200 GB disk: 75.5 GB • 300 GB disk: 169.9 GB • 400 GB disk: 264.3 GB • 500 GB disk: 358.7 GB
All subsequent data disks (100 GB or more)	Total capacity * 0.944	<ul style="list-style-type: none"> • 100 GB disk: 94.4 GB • 200 GB disk: 188.8 GB • 300 GB disk: 283.2 GB • 400 GB disk: 377.6 GB • 500 GB disk: 472 GB

Configuring DD VE in Data Domain System Manager

DD VE licensing and configuration can be accomplished through the Configuration Wizard in Data Domain System Manager. After the initial installation of a DD VE instance, the Configuration Wizard automatically appears after the licensing screen on the first launch of DDSM.

Note

The DAT is not supported for cloud DD VE.

Enter the DD VE virtual machine IP address into a web browser to launch Data Domain System Manager. Log in with the following credentials:

- Username: sysadmin
- Password: changeme

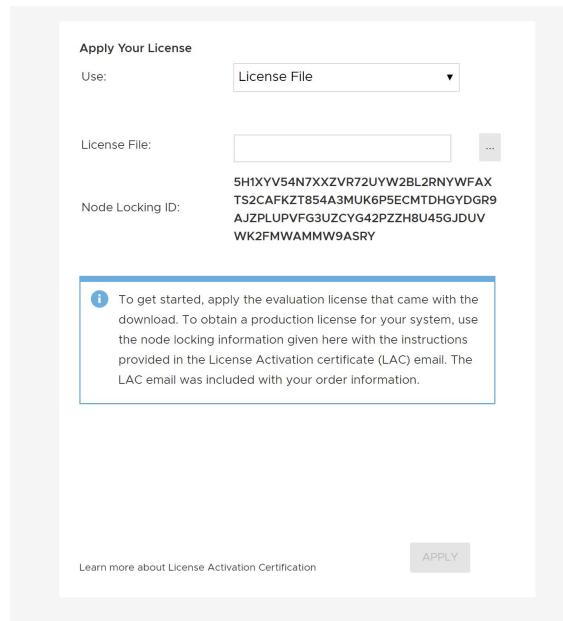
DD VE licensing

The **Apply Your License** window is the first screen that appears when DDSM is launched for the first time. The DD VE instance is locked until a license file is applied.

Click **Browse**, locate the license file for a purchased capacity license or the evaluation license included with the DD VE download, then click **Apply**.

Note

If you begin the configuration with the evaluation license, but wish to purchase a license later, you will need the Node Locking ID for the DD VE instance. Click **Administration > Licenses** to view the Node Locking ID.

Figure 4 DD VE Node Locking ID

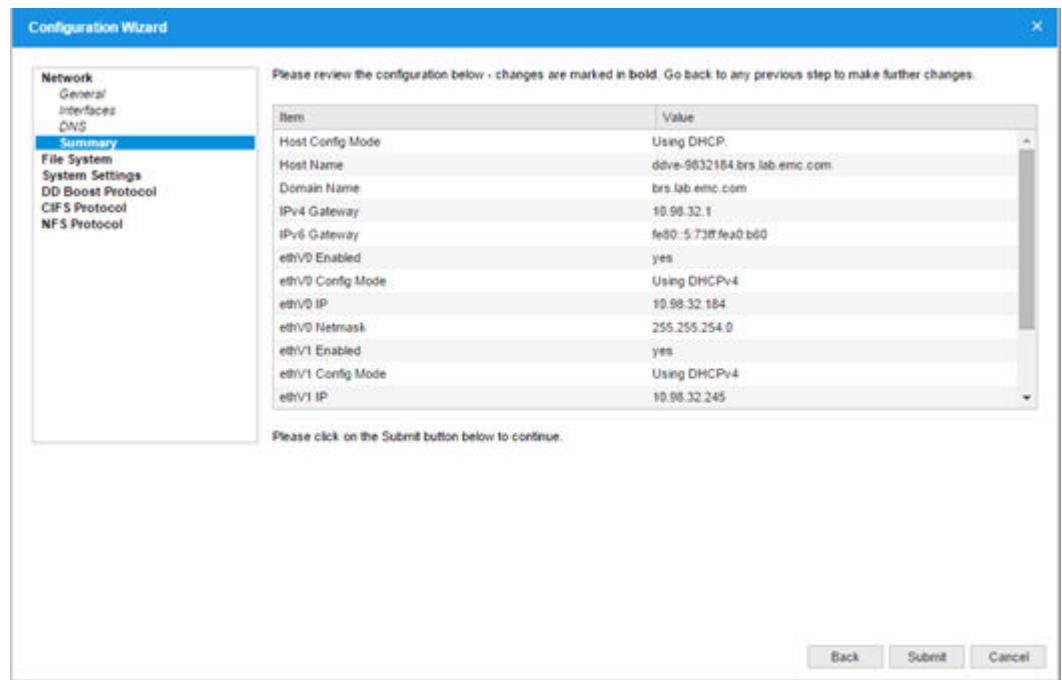
Note

When you obtain the original license file name the server, do not enter the comma in the license file name. DD OS will not accept the name if the comma is used.

DD VE configuration

After applying the DD VE license, the Configuration Wizard begins automatically. The wizard assists in configuring the following aspects of the DD VE:

- Networking
 - DHCP or manual settings
 - Virtual interface ethV0 and ethV1 configuration
 - DHCP or manual DNS configuration

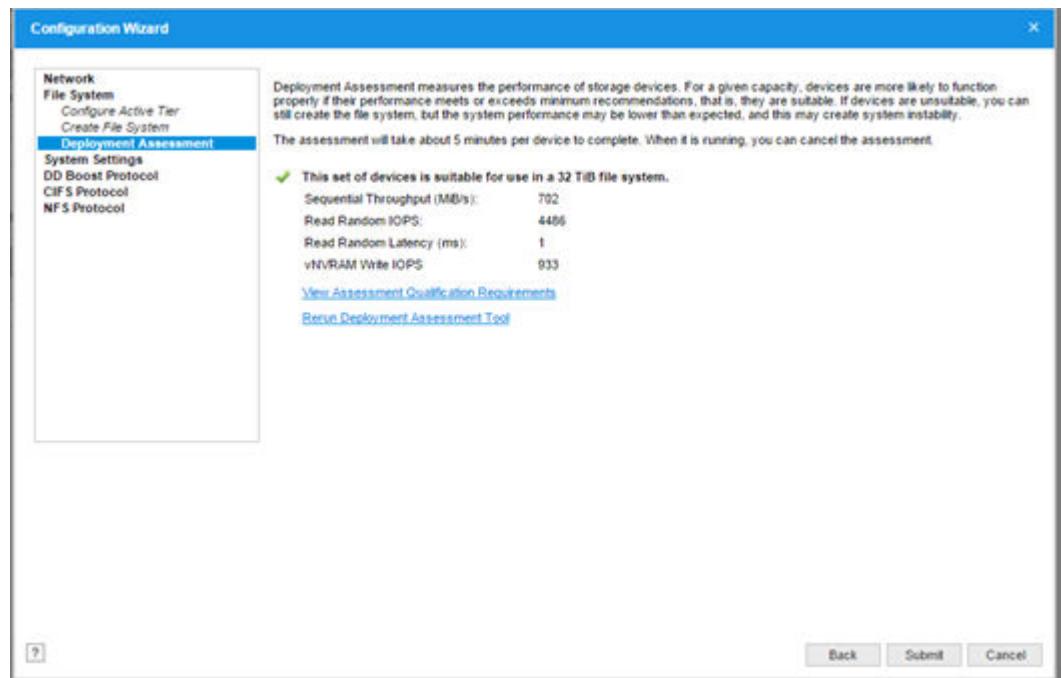
Figure 5 Configuration Wizard - Network

- File system

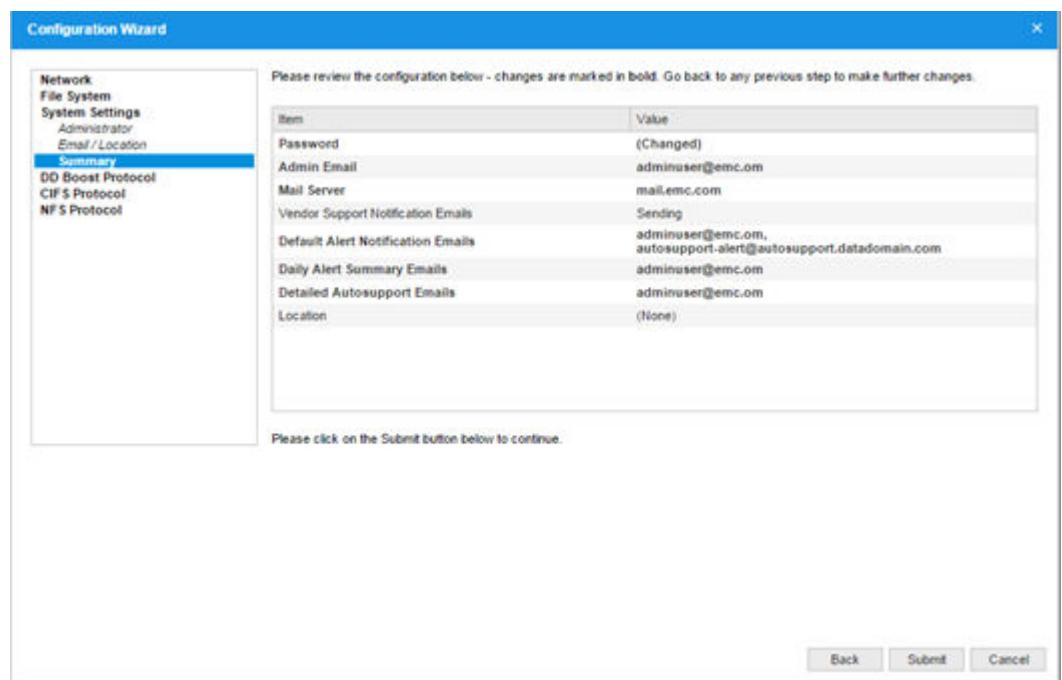
Note

DD VE supports a maximum of six MTrees active at a given time, however up to 100 MTrees can be created on DD VE.

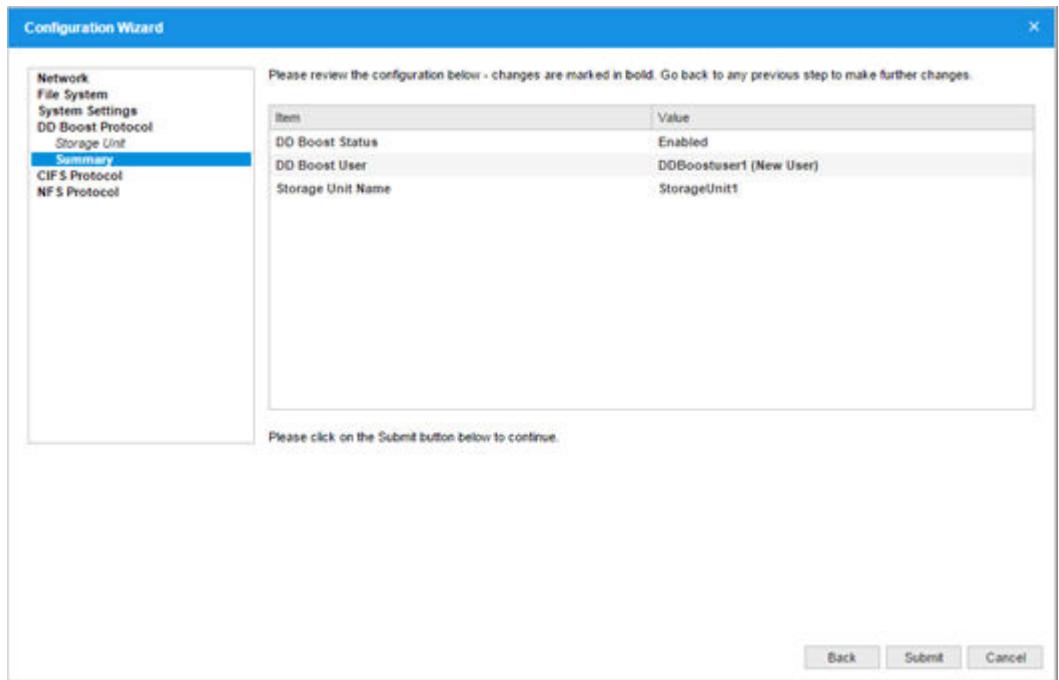
- Create virtual storage devices
- Optionally enable the DDFS automatically after creating it

Figure 6 Configuration Wizard - File System

- System settings
 - Update the sysadmin password
 - Optionally configure alert and autosupport email settings

Figure 7 Configuration Wizard - System Settings

- DD Boost
 - Create a Boost storage-unit, and assign a user ID to own it

Figure 8 Configuration Wizard - DD Boost Protocol

Provisioning the storage with the CLI

Before you begin

See [Disk \(Spindle\) Group Configuration](#) on page 111.

Procedure

1. Log into the system with the user name of `sysadmin`.

The default password is `changeme`.

At the first login, use the `elicense` command to add a DD VE license.

2. Type control-C to exit the configuration utility.
3. Confirm that virtual disk `dev3` exists and has the expected size:

`# disk show hardware`

The output should include a line similar to the following example:

`dev3 VMware Virtual disk 1.0 (unknown) 256.00 GiB SAS n/a`
 The first two virtual disks (`dev1` and `dev2`) are used for the system software and cannot be used for storage. The `disk show state` command shows `System Dev` for these system disks.

4. Choose to run DAT tool test (optional):

`#disk benchmark start dev3`

This command starts the DAT tool test.

`#disk benchmark watch`

You can monitor the test's progress by entering this command.

`#disk benchmark show`

Once the test is complete, you can use this command to see the test's result.

5. Add the storage disk to the active storage tier:

`# storage add dev3`

If you are adding more than one virtual disk, repeat the storage add command for each disk. For guidelines on specifying the optional spindle-group argument, see [Configuring Disk \(Spindle\) Groups](#).

6. Create the file system:

```
# filesys create
```

The "filesys create" may take longer to complete if the hypervisor's storage is slow and does not meet the criteria.

7. Enable the file system:

```
# filesys enable
```

After you finish

You can now complete the initial system configuration. See [Completing Initial Configuration with the Command-Line](#).

Completing Initial Configuration with the Command-Line

Procedure

1. Enter the `config setup` command to start the configuration utility.
2. When prompted, enter the Name, and Domain Name of the system.
3. When prompted, configure the initial IP port.

You can:

- Choose DHCP.
- Enter a static IP address and Net Mask.

4. When prompted, either exit the configuration utility and continue configuring the system using the Graphical User Interface, or continue using the CLI configuration utility.

The list entries in the utility can be comma-separated, space-separated, or both.

- At each prompt, enter a value, OR
- Enter a question mark (?) for more details, OR
- Press Enter to accept the value displayed in braces.

Follow the configuration utility instructions for entering appropriate values. At the end of each configuration section, you can choose to: `Save`, `Cancel`, or `Retry` (restart the input entry as the beginning of the current section).

Note

If you need to enable NTP, you can do so with the configuration utility.

Configure the System for Data Access

The DD VE system provides the DD Boost protocol. You need to configure one or more protocols for data access, depending on your environment. You also need to configure the clients for accessing the DD VE with the protocol of your choice.

If you did not configure data access with the configuration wizard, use the instructions in this section.

DD Boost (DD VE includes the DD Boost for cloud or on premise)

For setting up the Data Domain DD Boost feature, see the *Data Domain Boost for OpenStorage Administration Guide* or *Data Domain Boost for Partner Integration Administration Guide* available at <https://support.emc.com>.

Application Integration

For information about how to integrate the Data Domain system with backup software, see the documentation for the applicable application at the Data Domain Integration Documentation section on the Data Domain Support web site <https://support.emc.com>.

Configuration of optional software and internal licenses

If you need to configure optional software features, you need to install and activate those licenses before you configure those features. See **DD VE capabilities** on page 14 for information about features and licenses that are available to for DD VE. A separate license is required for DD Cloud Tier.

Information about installing licenses and configuring optional software can be found in the *Data Domain Administration Guide*. Refer to the applicable *Data Domain Operating System Release Notes* for the most up-to-date information on product features, software updates, software compatibility guides, and information about our products, licensing, and service. Access the latest documents at <https://support.emc.com>.

Optional Additional System Configuration

See the *Data Domain Operating System Initial Configuration Guide* for help performing typical but optional initial system configuration tasks. Below is a summary of the DD OS CLI commands for some common tasks.

Note

Any system command that accepts a list, such as a list of IP addresses, accepts entries separated by either commas or spaces. See the *Data Domain Operating System Command Reference Guide* for command details.

Add users to the email list that reports system problems:

```
# alerts notify-list add group-name
```

Add users to the system report email list:

```
# autosupport add {alert-summary|asup-detailed} emails email-list
```

Enable FTP or TELNET:

```
# adminaccess enable {ftp|telnet}
```

Add remote hosts to use FTP:

```
# adminaccess ftp add <host list>
```

Add a user:

```
# user add name [role {admin|user}]
```

Change a user's password:

```
# user change password username
```

To enable remote management, refer to the *Data Domain Operating System Administration Guide* for details.

To Shut Down The System:

```
# system poweroff
```

Overview of Deploying the DD VE on KVM Hypervisor

DD VE can be run as a virtual machine on top of the KVM Hypervisor which provides a full virtualization solution for Linux. This information helps you to install DD VE on hypervisor on supported Linux distributions.

Supported Linux distributions

DD VE on KVM only supports Intel-based processors. The following Linux distributions are supported by DD VE.

Table 20 Supported Linux distributions

Linux distribution	Version
CentOS	7-1611
Red Hat	7.2, 7.3
SUSE	12-SP2
Ubuntu	14.04 LTS Trusty, 16.04 LTS xenia

Note

Use the proper libvirt version provided by the supported Linux distributions. DD VE is only supported on KVM libvirt version 1.2.2 or newer. For the Ubuntu 16.04 distribution, the supported libvirt package is the 1.3.4 version.

Prerequisites

- Install the most recent version updates for your Linux distribution.
- Create a bridged network or open vSwitch for KVM

Configuration of other resources

DD VE resource reservations for DD VE on KVM

Table 21 DD VE resource reservations

Resources		Up to 4 TB	Up to 8 TB	Up to 16 TB	16 TB to 32 TB	Up to 48 TB	Up to 64 TB	Up to 96 TB
Computing resources	CPU	2 x vCPU				4 x vCPU		
	Memory	6 GB	8 GB	16 GB	24 GB	36 GB	48 GB	64 GB
Underlying storage requirements	Random IOPS	160	320	650	1280	1920	2560	3200
	Random I/O latency	14 ms						
	Sequential throughput	40 MB/s	80 MB/s	160 MB/s	320 MB/s	480 MB/s	640 MB/s	960 MB/s
	RAID	RAID 5/6 or similar fault tolerance storage						

Table 21 DD VE resource reservations (continued)

Resources	Up to 4 TB	Up to 8 TB	Up to 16 TB	16 TB to 32 TB	Up to 48 TB	Up to 64 TB	Up to 96 TB				
SCSI controllers	Virtio SCSI										
vNVRAM simulation file size	512 MB			1 GB		2GB					
	Configurations with DD Cloud Tier support has the same vNVRAM size as the corresponding ones without DD Cloud Tier.										
System disks	<ul style="list-style-type: none"> 250 GB root disk 10 GB vNVRAM disk 										
	<p>Note</p> <p>The root disk and vNVRAM disk are required to deploy the DD VE.</p>										

Note

We recommend that you do not over commit CPU and memory resources on the KVM host with DD VEs.

DD VE on KVM Stream Counts**Table 22** Stream counts for each capacity

Configuration	Write Stream	Read Stream	Repl Source	Repl Dest	Mixed Stream	Max Mtree
4 TB	16	4	15	20	20	6
8 TB	20	16	20	20	30	6
16 TB	45	30	45	45	60	6
32 TB	90	50	90	90	90	14
48 TB	90	50	90	90	90	14
64 TB	90	50	90	90	90	32
96 TB	180	50	90	180	180	32

Deploying the DD VE on the KVM Hypervisor

Includes the information required to deploy the DD VE instance on KVM hypervisor.

Deploying the DD VE on KVM with a reference script

The DD VE for KVM image has reference scripts to deploy the DD VE image on KVM hypervisor and adds the data disks to a deployed DD VE. The scripts can be modified

to run in your environment and are included in the tar.gz file you downloaded for the DD VE on KVM installation.

Procedure

1. Extract the tar.gz file. `tar -xvf ddve-kvm-0.6120.12.0-563880.tar.gz` (or `tar -xvf ddve-kvm-0.6120.12.0-xxxxxx.tar.gz` where "xxxxxx" is the build number).
2. Change directory to the extracted folder.
3. To deploy a DD VE alone without any data disks, run as below example (the example below will create a 16TB configuration DD VE with name ddve-test on br0 network interface with the DD VE root disk and vNVRAM disk on '/kvm-root' directory) `./kvm-ddve-installer.sh -n ddve-test -r /kvm-root -c 16TB -b br0`

4. **Note**

```
[root@ddqa-r730-d05 ddve-kvm-6.1.2.5-595467]# ./kvm-ddve-installer.sh -h
Distribution: rhel Version:7.3. The host version check done.
Basic Validation done. Usage: ./kvm-ddve-installer.sh [options]
```

Where options are:

- a. -n - Specify the virtual machine name (default will be the name of the DD VE build).
- b. -r - Specify the root disk full path (default will be the folder where installer script executed).
- c. -c - Specify the configuration 4TB, 8TB, 16TB, 32TB, 48TB, 64TB, 96TB, Cloud16TB, Cloud64TB, and Cloud96TB (default will be 4TB).
- d. -b - Specify the bridge name (default will be br0).
- e. -p - Specify provision type for the data disk created on NFS. By default, the thin provisioning disk is created on NFS.
- f. -s - Specify the disk size in TB or GB (when option -s is specified, options -d and -x are mandatory).
- g. -x - Specify the number of data disks (when option -x is specified, options -d and -s are mandatory).
- h. -d - Specify the path where the data disks will be stored (when option -s is specific, options -s and -x are mandatory).

Note

These three options are to be used as a set: -s -x or -d. If any one of these options is used, the other two in the set are required.

- i. -w -Wait for IP address. The IP address of the DD VE will be displayed after deployment after a 5-minute wait for it to become available.
- j. -h - Help message.
5. To deploy a DD VE with data disk, run as in the example below (This example will create 4x4 TB data disk on datastore1 along with the deployment of a 16TB configuration DD VE). `./kvm-ddve-installer.sh -n ddve-test -r /kvm-root -c 16TB -b br0 -x 4 -s 4TB -d /datastore1`

6. If you want to add a data disk to the existing DD VE, run as below example (the example below will create a 4x2TB data disks on datastore2 for the DD VE with the name ddve-8tb). `./kvm-add-disk.sh -n ddve-8tb -x 4 -s 2TB -d /datastore2`

Note

`kvm-ddve-installer.sh` can only be run once per extraction since it will convert the `qcow2` into `raw` format. Any subsequent execution of the script will cause failures because `qcow2` will no longer exist.

Results

The script will complete following operations:

1. Generate the VM domain XML file `config.xml`.
2. Create a new VM domain named "ddve-test".
3. Convert the `.qcow2` root disk in the `tar.gz` package to `.raw` format for better performance.
4. Attach the root disk to SCSI 0:0.
5. Create a 10G raw disk named "ddve-test-vnvr" and attach the disk to SCSI 0:1.
6. Setup CPU and memory configuration for the DD VE VM as per DD VE sizing guidelines.
7. Setup network configuration for the DD VE instance.
 - a. Create 2 virtual network interfaces.
 - b. Use "bridge" as the interface type.

Deploying the DD VE on KVM using the Virtual Machine Manager

The DD VE for KVM image deploys the DD VE image on the KVM hypervisor using `virt-manager`. The steps are included below to create a DD VE VM and then to configure the DD VE.

The following is an example configuration to deploy 4 TB DD VE.

Procedure

1. Create a DD VE VM
 - a. Open Virtual Machine Manager.
 - b. Click "File", select "New Virtual Machine".
 - c. Select "Import existing disk image", click "Forward".
 - d. Enter the existing storage path, click "Browse..." to select the root disk file you will use.
 - e. Choose an operating system type and version. Select "Linux", "Red Hat Enterprise Linux 6.4", click "Forward".
 - f. Enter the VM name. Make the "Customize configuration before install" selected. Click "Finish".
2. Configure the DD VE
 - a. Click "Processor" on the left plane. Set the "Current/Maximum allocation" to 2. Set the Model to "Hypervisor Default". Select the "Manually set CPU topology" option and set Sockets 1, Cores 2, Threads 1. Click "Apply".

- b. Click "Add Hardware" button, select "Controller", Type "SCSI" and Model "VirtIO SCSI", click "Finish".
- c. Click "Add Hardware" button, select "Network". Network source "Bridge..." (The bridge name on the host), Device Model "virtio". Click "Finish".
- d. Click "Begin Installation".

Deploying the DD VE on KVM Hypervisor Best Practices

- Shutdown the VM before you make any changes to VM settings with "virsh edit." Some settings will not be reflected after a VM reboot--the changes will only be in effect after a VM shutdown.
- We recommend you sync to NTP for the KVM host. By default for DD VE on KVM, the KVM clock will be enabled. Power on and reboot the system, the DD VE will sync with the KVM hypervisor.
- For the Ubuntu 16.04 distribution, please update libvirt package to libvirt 1.3.4 version.
- Be sure to verify the directory permissions on KVM host before trying to deploy the DD VE. If directory permissions are incorrect, you may encounter these error messages: "native: could not open disk image" or "Could not open ... Permission Denied." Be aware of this especially when mounting to a remote directory. Make sure the directory permissions are correct for both the qemu user and current user during the deployment.

Note

The qemu user needs search permissions all the way up the path of the directory tree.

For example,

1. The directory /data-san1 was created with the owner/group as root/root.
2. When /dev/mapper/mpathb1 /data-san1 was mapped, the owner/group changed to user1/user1. The DD VE failed to power on due to this permission issue.
3. While mapping the datastore, the command `chown -R root:root /data-san1` was issued. The DD VE was then successfully deployed.

Creating a bridged network interface using Virtual Manager

This information is for those who wish to use the Virtual Manager. Each supported Linux distribution may differ in terms of configuration files used and the configuration via the command line interface (CLI). If you wish to use the CLI, please refer to the Linux vendor documentation for your Linux distribution for additional information.

Procedure

1. By default KVM will use "Usermode Networking" where NAT is performed on traffic through the host interface to the outside network. As a best practice, however, a Bridge Network should be used for DD VE. this allows external hosts to access the guest VM directly--guest VMs are connected directly to the host network
2. Verify the bridging kernel is installed by running this command (an error message should result stating that module is already in the kernel). `modprobe --first-time bridge`

3. Find out the name of the physical interface on your host, e.g., used in the network bridge and verify that it is physically connected and functional. You can use the GUI to determine this or the CLI. CLI commands that are helpful include:
 - `ifconfig` - lists the interfaces on the host.
 - `ethtool <int name>` - displays details about the interface .
4. Open virt-manager (Virtual Manager).
5. Go to Edit>Connection Details.
6. In Connection Details, click on the "Network Interfaces" tab.
7. On the QEM/KVM Connection Details screen, click the "+" button at the bottom of the screen to start the add interface wizard.
8. On the Configure Network Interface screen, ensure the "Bridge" option is chosen from the 'interface type' drop down menu, then click the forward button.
9. Ensure the bridge interface name is "br0", and select "onboot" from the 'start mode' drop down menu.
10. Verify the bridge interface has the correct settings for IP address. If you want to change any of the IP settings click the "configure" button for IP settings.
11. On the IP Configuration screen, click through the IPv4 and the IPv6 tabs and set the settings you desire. Click ok once done. You can configure the bridge interface to pick up IP from DHCP or you can statically configure the IP if you know the network settings.
12. Click on the "configure" button for the bridge settings.
13. Uncheck the "Enable STP" check box, then click "Ok".
14. Select the physical interface that will be a member of this bridge. Make sure only one interface is selected. Click "finish" when done.
15. Reboot or restart network services. To restart network services, open terminal and enter:

```
systemctl restart network
```

16. Verify bridge interface is up after network service restarts:

```
ifconfig br0
```

17. You may also verify the bridge interface in the virt-manager GUI: Edit>Connection details>Network interfaces>br0.
18. Connection is complete.

CHAPTER 3

DD VE in the Cloud -- All providers

This chapter covers the following topics:

• Supported cloud environments	58
• DD VE capabilities	58
• DD VE in the Cloud -- DD VE on S3 Storage	60
• Deploying DD VE on AWS	60
• DD VE on Azure hot blob storage	78
• VMware Cloud on AWS	94
• Initial System Configuration	99

Supported cloud environments

DD VE supports the following cloud environments.

- Amazon Web Services (AWS)
- Microsoft Azure
- VMware Cloud on AWS (VMC)

For information about compatibility with more recent versions of Azure products, visit the support portal at <https://support.emc.com>

DD VE capabilities

DD VE provides the capabilities of a cloud Data domain system using the following cloud providers and resource configuration sizes.

Table 23 DD VE Cloud Providers and Resource Configuration Size

Cloud Provider	Type	Resource Configuration Size
Amazon Web Services (AWS)	DD VE on Block Storage	up to 16 TB
AWS	DD VE on S3 Storage	<ul style="list-style-type: none"> • up to 7 TB • up to 96 TB <p>Note Actual DD VE capacity is available in 1 TB increments starting at 1 TB, and up to 15 TB.</p>
Azure	DD VE on Block Storage	<ul style="list-style-type: none"> • DD VE on Block Storage: up to 15TB <p>Note Actual DD VE capacity is available in 1 TB increments starting at 512 GB, and up to 8TB or 16TB respectively for 8TB and 16TB type.</p>
Azure	DD VE on Hot Blob Storage	<ul style="list-style-type: none"> • Actual DD VE capacity is available up to 96 TB • Metadata consumes block storage of 10

Table 23 DD VE Cloud Providers and Resource Configuration Size (continued)

Cloud Provider	Type	Resource Configuration Size
		percent of licensed capacity

The following sections list supported and unsupported Data Domain protocols and features in DD VE.

Supported Data Domain protocols

- Data Domain Boost (DD Boost) over IP
- Data Domain Boost (DD Boost) FS

Supported Data Domain features

- DD Boost managed file replication (MFR)
- Encryption
- MTree replication
- Data Domain System Manager GUI for DD VE management
- Secure multitenancy (SMT) with Network Isolation Support in 6.0
- DD Boost for Big Data
- Key Management Interoperability Protocol (KMIP)
- More restricted IPTables settings
- In both AWS and Azure cloud, DD VE supports two types of data storage:
 - For AWS
 - DD VE on S3 Storage
 - DD VE on Block Storage
 - For Azure
 - DD VE on Hot Blob Storage
 - DD VE on Block storage.

Note

DD VE 4.0 supports these replication capabilities:

- Managed file replication and Mtree replication
- Replication across availability zones and regions
- Bidirectional replication between on-prem and AWS

Please see the DD OS Administration Guide, DD Boost OST Guide, DD Boost for Partner Integration Administration Guide for additional information on the supported protocols and features above.

Unsupported Data Domain features

- Cloud Tier
- Gov/C2S cloud regions

- DD Boost over FC
- Extended retention
- DD High Availability (HA)--however, VMware and Hyper-V HA are supported
- NDMP
- VTL
- Collection replication
- Directory replication
- Instant access
- Retention Lock Compliance Mode
- CIFS and NFS for backups

DD OS commands related to these unsupported features, and commands for hardware features that are not applicable to a virtual machine, are not supported on the DD VE platform.

DD VE in the Cloud -- DD VE on S3 Storage

Overview of DD VE on S3 Storage

DD VE on S3 storage provides enterprise customers and service providers who are running applications in the public cloud with a dedupe data protection appliance that provides object storage efficiency and ease of management. Supported cloud platforms include Amazon Web Services (AWS) and Microsoft Azure. It is important to note that DD VE on S3 is deployed in these environments differently.

DD VE on S3 Storage supports:

- Backup/restore using DD VE on S3 storage data into cloud object store while DD VE is running in the cloud
- DD SM to configure, manage, and monitor DD VE in AWS and Azure with DD VE on S3 Storage
- VMware Cloud (VMC) in AWS environments (VMware Cloud is an AWS offering)

Deploying DD VE on AWS

DD VE on S3 Storage provides a data protection solution that enables customers to protect their operational data in the cloud, to backup/restore the DD VE on S3 Storage data into cloud object store, while the DD VE is running in the cloud. This section describes first-time setup procedures, and includes how to manage and monitor the DD VE in an AWS environment.

Note

For DD VE on Block Storage see [Configuring DD VE on block storage using the DD SM interface](#)

AWS System Configuration requirements

These are the system configuration requirements for configuring the DD VE on S3 Storage on AWS.

System Configuration Requirements for AWS Meta data disk type: GP2

Table 24 AWS System Requirements

Instance type	M4.xlarge	M4.2xlarge	M4.4xlarge
CPU	4	8	16
Memory (GiB)	16	32	64
System Disk	250 GiB GP2 Root disk	250 GiB GP2 Root disk	250 GiB GP2 Root disk
	10 GiB GP2 vNVRAM disk	10 GiB GP2 vNVRAM disk	10 GiB GP2 vNVRAM disk
Maximum File System Capacity	16 TB	32 TB	96 TB

AWS Storage Size Specifications

These are the system configuration requirements for configuring the DD VE on S3 Storage on AWS.

System Configuration Requirements for AWS Meta data disk type: GP2

Table 25 Storage Size Specifications

Configuration	Instance Type	Block Storage Volumes			
		Root Disk	vNVRAM Disk	Metadata Disks	Object Storage Capacity
16 TB	M4.xlarge	GP2/250GiB	GP2/10GiB	GP2 (1-2 TiB Disk)	0-16 TB
32 TB	M4.2xlarge	GP2/250GiB	GP2/10GiB	GP2 (1-4 TiB Disk)	0-32 TB
96 TB	M4.4xlarge	GP2/250GiB	GP2/10GiB	GP2 (1-10 TiB Disk)	0-96 TB

Deploy the DD VE in AWS with Cloud Formation Template

- The DD VE on S3 Storage can be configured on AWS using Cloud Formation Template Deployment option.
- Refer to [Network Setup Recommendations](#) for subnet and security group settings.
- Create IAM role and key pair to attach to the instance during deployment.

Cloud Formation Template Deployment

Procedure

- Go to AWS Marketplace.
- Search for “Data Domain Virtual” in the search bar for “AMI and SaaS”.

3. “Choose Dell EMC Data Domain Virtual Edition (DD VE) v4.0” and click on “Continue to Subscribe”

4. Click on “Continue to Configuration”
5. Select the fulfillment option to launch a DD VE instance. “CloudFormation” is the recommended option. Select the “Region” in which to deploy the DD VE. Then click “Continue to Launch”
6. Review the configuration details, select “Launch the Cloud Formation” template, then click “Launch”.

7. The template URL is populated. Click “Next”.
8. The following values need to be populated in order to create the stack.
 - Stack name
 - DD VE Capacity: pick any capacity from the drop down list. The recommended amount of metadata storage and the instance type will be attached automatically by the template, based on the selected capacity.
 - DD VE name tag
 - IAM Role for S3 access: type in the correct IAM role to be attached to the DD VE
 - Key pair: select an existing key pair from the drop down list
 - Subnet ID
 - Security Group ID
9. Proceed to create instance once the values are populated.

10. Once the stack creation is complete, verify the EC2 instance is running.

Overview	Outputs	Resources	Events	Template	Parameters	Tags	Stack Policy	Change Sets	Rollback Triggers
Filter by: Status ▾									
Search events									
2018-07-17	Status	Type	Logical ID	Status Reason					
12:10:19 UTC-0700	CREATE_COMPLETE	AWS::CloudFormation::Stack	ddve-selvab1-ge-cft-deploy-test						
12:11:08 UTC-0700	CREATE_COMPLETE	AWS::EC2::VolumeAttachment	MountPointMetadataDisk1						
12:11:08 UTC-0700	CREATE_COMPLETE	AWS::EC2::VolumeAttachment	MountPointNvram						
12:11:08 UTC-0700	CREATE_IN_PROGRESS	AWS::EC2::VolumeAttachment	MountPointMetadataDisk1						
12:10:50 UTC-0700	CREATE_IN_PROGRESS	AWS::EC2::VolumeAttachment	MountPointNvram						
12:10:35 UTC-0700	CREATE_IN_PROGRESS	AWS::EC2::VolumeAttachment	MountPointMetadataDisk1						
12:10:35 UTC-0700	CREATE_IN_PROGRESS	AWS::EC2::VolumeAttachment	MountPointNvram						
12:10:32 UTC-0700	CREATE_COMPLETE	AWS::EC2::Volume	MetadataDisk1						
12:10:31 UTC-0700	CREATE_COMPLETE	AWS::EC2::Volume	Nvram						
12:10:15 UTC-0700	CREATE_IN_PROGRESS	AWS::EC2::Volume	MetadataDisk1						
12:10:15 UTC-0700	CREATE_IN_PROGRESS	AWS::EC2::Volume	Nvram						
12:10:14 UTC-0700	CREATE_IN_PROGRESS	AWS::EC2::Volume	MetadataDisk1						
12:10:14 UTC-0700	CREATE_IN_PROGRESS	AWS::EC2::Volume	Nvram						
12:10:10 UTC-0700	CREATE_COMPLETE	AWS::EC2::Instance	DdveEc2Instance						
12:09:49 UTC-0700	CREATE_IN_PROGRESS	AWS::EC2::Instance	DdveEc2Instance						
12:09:48 UTC-0700	CREATE_IN_PROGRESS	AWS::EC2::Instance	DdveEc2Instance						
12:09:42 UTC-0700	CREATE_IN_PROGRESS	AWS::CloudFormation::Stack	ddve-selvab1-ge-cft-deploy-test	User Initiated					

Configuring the DD VE on S3 Storage

This section provides best practices for the DD VE on S3 Storage configuration .

Before you begin

- Refer to the recommended network setup recommendations.
- Metadisks should be added in 1TiB increments. The recommended metadata storage is 10% of the total capacity.
- The bucket and IAM role need to created. Refer to [bucket](#) and [role](#) instructions.

The DD VE can be configured using one of the following options

- [Graphical User Interface \(GUI\)](#)
- [Command Line Interface \(CLI\)](#)

Note

The role must be attached to DD VE instance, before configuring object store feature in DD VE using its CLI or GUI.

After you finish

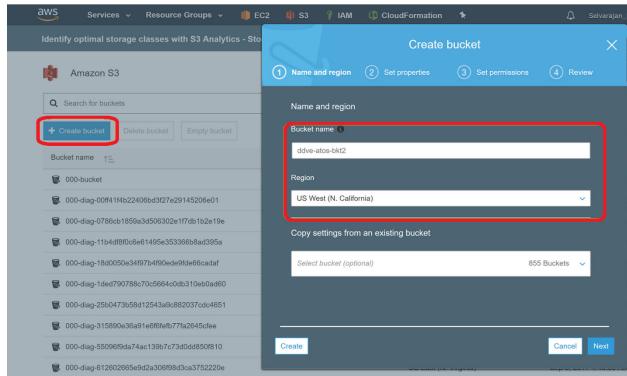
The following sections describe the steps to configure the DD VE using [GUI](#) or [CLI](#).

Creating an S3 bucket

Create a bucket in S3 and make note of the bucket name. The bucket name will be used in the IAM policy template and also used to create the cloud profile on the DD VE.

Note

Bucket should be created in the same region as the DD VE instance. The length of the bucket name should be no more than 48 characters.



AWS role-based access for S3 object store access

Object store in AWS uses role-based access for S3 access. This section provides the steps to achieve this access by creating and attaching the IAM role to the DD VE. The DD VE then fetches the access credentials to access the S3 bucket. The access credentials are rotated periodically by the AWS infrastructure. The new credentials are automatically fetched by the DD VE just before the old credentials expire.

Before you begin

In order to create the IAM role and the policy associated with the role, the AWS user should have the necessary IAM privileges. The following are some of the IAM privileges/actions that are required to create and attach the IAM role.

```
"iam:AddRoleToInstanceProfile",
"iam:AttachRolePolicy",
"iam:CreateRole",
"iam:DeleteRole",
"iam:DeleteRolePolicy",
"iam:DetachRolePolicy",
"iam:GetRole",
"iam:GetRolePolicy",
"iam>ListRolePolicies",
"iam>ListRoles",
"iam:PassRole",
"iam:RemoveRoleFromInstanceProfile",
"iam:UpdateRolePolicy",
"iam:CreateInstanceProfile",
"iam:PutRolePolicy",
"iam:DeleteInstanceProfile"
```

Once you have the necessary privileges as an AWS, continue creating the role-based for S3 object store access as follows.

Create the policy for S3 bucket access

Procedure

1. Sign in to the AWS Management Console and open the IAM Service Console.
2. In the navigation pane of the IAM console, choose “Policies” and then click “Create policy” button.
3. In the “Create policy” web page, select the tab “JSON”
 - a. Replace the text under the JSON tab with the following content.

Note

Substitute “my-bucket-name” with the name of the bucket that you have created for the DD VE.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3>ListBucket",
        "s3GetObject",
        "s3PutObject",
        "s3DeleteObject"
      ],
      "Resource": [
        "arn:aws:s3:::my-bucket-name",
        "arn:aws:s3:::my-bucket-name/*"
      ]
    }
  ]
}
```

- b. Verify this information on your screen then click the “Review policy” button.
- c. Provide the name and description of your choice.

Create policy

1 2

Review policy

Name* Use alphanumeric and '+=_@-' characters. Maximum 128 characters.

Description Maximum 1000 characters. Use alphanumeric and '+=_@-' characters.

Summary

Summary			
Service	Access level	Resource	Request condition
S3	Limited: List, Read, Write	Multiple	None

* Required

Cancel Previous **Create policy**

- d. Follow the steps in the next section to create the role and attach the policy you have just created to the role.

Create the role for S3 bucket access

Procedure

1. In the navigation pane of the IAM console, choose “Roles” and then click the “Create role” button.
2. On the “Create role” page,
 - a. For the “Select type of trusted entity” option, select “AWS service”
 - b. For the “Choose the service that will use this role” option, select “EC2”
 - c. Then click “Next Permissions” to advance to the next section.
3. On the “Attach permissions policies” page, search for the policy that you created in the previous section, such as “ddve-s3-access-policy”. Then select the check box for that policy. Click the “Next review” button to advance to the next section.

Create role

1 2 3

Attach permissions policies

Choose one or more policies to attach to your new role.

[Create policy](#) [Refresh](#)

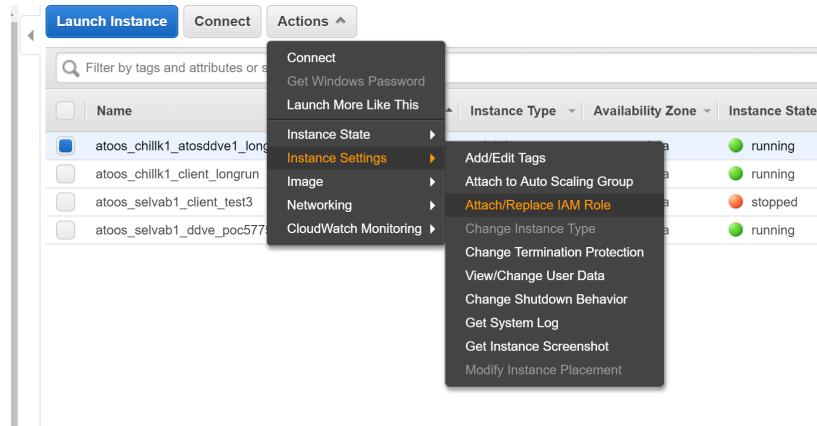
Filter: Policy type		Q ddve-s3-access-policy	Showing 1 result
	Policy name	Attachments	Description
<input checked="" type="checkbox"/>	ddve-s3-access-policy	0	The policy defines the actions that are permitted by DD VE o...

* Required

[Cancel](#) [Previous](#) [Next: Review](#)

4. For the “Review” section, provide the name for the role. then click the “Create role” button.

5. Select the role you created during the DD VE deployment or attach the role after the DD VE is deployed.



Note

The role must be attached to the DD VE instance, before you can configure the object store feature in DD VE using the CLI or UI.

Configuring DD VE using the Data Domain System Manager interface

DD VE can be configured in AWS using one of the Data Domain System Manager (DD SM) interface options: GUI or CLI.

Before you begin

- Recommended metadata storage is 10% of the total capacity.
- The default password is the instance-id.
- To configure DD VE on **Block Storage**, refer to [Configuring DD VE on block storage using the DDSM interface](#).

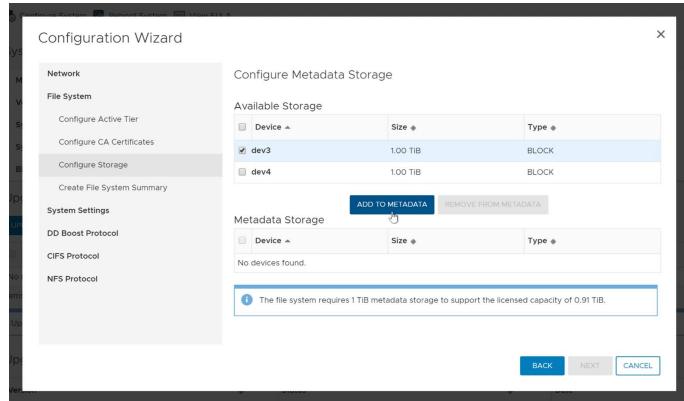
Configuring DD VE using the DD SM interface ([http/](http://)[https](https://))

The GUI configuration wizard helps go through the DD VE on S3 Storage configuration and file system creation on DD.

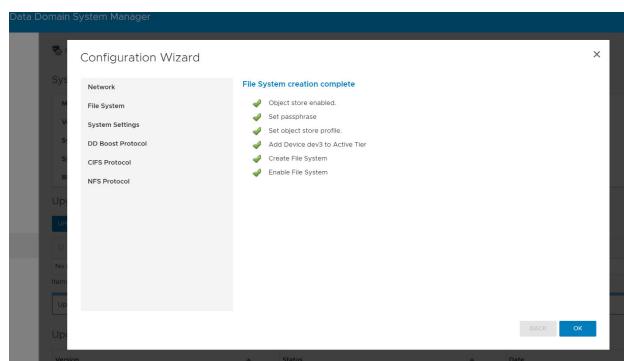
Procedure

1. Log in to the Data Domain System Manager.
2. For the “Apply your license” step, select one of the three license types available on the drop down menu:
 - Pre-install Evaluation: (500GB)
 - License File: Node locked license (unserved mode)
 - License Server: Served mode license
3. Then click “Apply”/Accept the End User License Agreement (EULA) by clicking “I accept the terms of EULA”.
4. Next, the “Configuration Wizard” will guide you through the DD VE on S3 Storage configuration and file system creation on the DD VE. Select “File System” and click “Yes”.

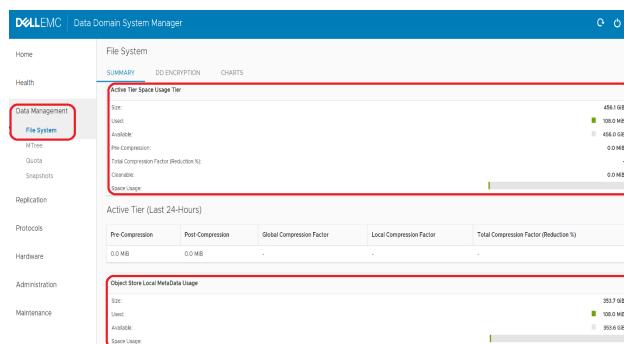
5. Select “Configure Active Tier”. Then, select “Enable Object Store”checkbox to configure the DD VE on S3 Storage system.
6. Enter the passphrase. (**Be sure to make note of this passphrase, you will need it again later.**) Enter the S3 bucket name created in the same region as the DD VE instance. Refer to the bucket and role creation instructions.
7. Import the Baltimore CyberTrust Root certificate to communicate with AWS S3 Object Store.
8. Add the metadata storage.



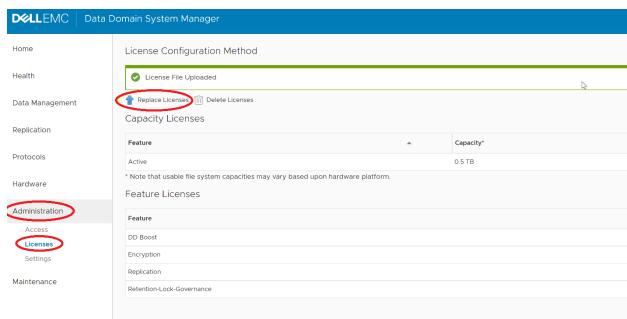
9. Review the summary and Click “Submit” to create the file system and enable it.
10. Review “File System Creation Complete” page.



11. The File System section under the Data Management tab has the space usage and availability details for the S3 Object storage as well as the local metadata storage.



12. To configure or update the elicense on the DD VE, click “Replace licenses” button in the Licenses page as shown in this image.



13. To relaunch the configuration wizard, select “System” under “Maintenance”, then “Configure System”

After you finish

Note

Set up NTP time synchronization by following the steps in [Setting up NTP Time Synchronization in AWS](#).

Configuring DD VE using initial configuration or config setup

Before you begin

Recommended Metadata storage is 10% of the total capacity.

DD VE can be configured as a DD VE on S3 Storage system using one of the following options

Configuring DD VE using initial configuration wizard or the “config setup”

You can add an eLicense by selecting “yes” for eLicense configuration. After the wizard ends, you need to add the devices.

Procedure

1. The eLicense wizard will prompt you, as shown below. Enter “yes” for each option.

Press any key then hit enter to acknowledge the receipt of EULA information:

Press any key then hit enter to acknowledge the receipt of EULA information:

Press any key then hit enter to acknowledge the receipt of EULA information:

Press any key then hit enter to acknowledge the receipt of EULA information:

Press any key then hit enter to acknowledge the receipt of EULA information:

Press any key then hit enter to acknowledge the receipt of EULA information:

Press any key then hit enter to acknowledge the receipt of EULA information:

Press any key then hit enter to acknowledge the receipt of EULA information:

Press any key then hit enter to acknowledge the receipt of EULA information:

Press any key then hit enter to acknowledge the receipt of EULA information:

Enter new password:

Re-enter new password:

Passwords matched.

Do you want to configure system using GUI wizard
(yes|no) : yes

```
Network Configuration
Configure Network at this time (yes|no): yes

System Configuration
Configure System at this time (yes|no): yes

Storage object-store profile
Configuration
Configure Storage object-store profile at this
time (yes|no): yes

Do you want to enable object store (yes|no): yes

A passphrase needs to be set on the system.
Enter new passphrase:
Re-enter new passphrase:
Config object store
DD VE is running in AWS. Role-based access will be used
to access s3.
Enter the bucket name: ddve-cloud-automation-
ncali2

Object-store endpoint needs the Baltimore
CyberTrust Root certificate to be imported.
Do you want to import that certificate with
below fingerprint?
D4:DE:20:D0:5E:66:FC:53:FE:1A:50:88:2C:78:DB:
28:52:CA:E4:74 (yes|no): yes

Pending Object Store Settings
Bucket name: ddve-cloud-automation-ncali2

Do you want to save these settings (Save|Cancel|
Retry): Save
The passphrase is set

Successfully set object store profile.

Configuration complete.
sysadmin@ip-10-1-1-170#
```

2. Create the file system using # filesys create

Screen will show the following:

```
A filesystem of approximate size 352.33 GiB will be created.
Do you want to continue? (yes|no) [yes]:
ok, continuing.

This will take 5 - 10 minutes.

Provisioning storage...
#####
Initializing filesystem...
#####
snapshot schedules deleted

You now have a freshly initialized filesystem.
Enable the filesystem using 'filesys enable'.
```

3. Enable file system using #filesys enable

Screen will show the following:

```
Please wait.....  
The filesystem is now enabled.
```

Configuring DD VE using the Command Line Interface

Before you begin

Recommended Metadata storage is 10% of the total capacity.

DD VE can be configured using the CLI option.

Configuring the DD VE using the CLI

Procedure

1. Log in the DD VE instance via SSH using the sysadmin account and password and instance-id. Or you may use SSH using key-pair.
 - When logging in for the first time, you will be asked to change the password. Enter the new password. The initial configuration wizard will start.
 - As shown in the screen capture below, you can add the eLicense by answering “yes” to the eLicense configuration prompt. This wizard can also be used to configure the network.

```
$ ssh -l sysadmin <DDVE ip address>
The authenticity of host '***.*.*.*' (***.*.*.*' can't be
established.
ECDSA key fingerprint is SHA256:evoXXGRgCzp/
tmrtWRIAeOWLpI7ym0q9mwTBWh9J2bs.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '***.*.*.*' (ECDSA) to the list of
known hosts.
EMC Data Domain Virtual Edition
Password:
Password:
Welcome to Data Domain OS *****

-----
Press any key then hit enter to acknowledge the receipt of EULA
information:
Press any key then hit enter to acknowledge the receipt of EULA
information: q
Do you want to configure system using GUI wizard (yes|no) [no]: y

Network Configuration
    Configure Network at this time (yes|no) [no]: y

eLicenses Configuration
    Configure eLicenses at this time (yes|no) [no]: y

System Configuration
    Configure System at this time (yes|no) [no]: y

Storage object-store profile Configuration
    Configure Storage object-store profile at this time
(yes|no) [no]: y

Configuration complete.
```

2. To update the eLicense on DD VE, copy license file to /ddvar and use the file name as follows: # elicense update <filename.lic>.

```
# elicense update atos_cap_96_TB.lic

Existing licenses:

Capacity licenses:
```

```

## Feature Capacity Type State
Expiration Date Note
-- -----
-- -----
1 CAPACITY 0.45 TiB unexpired evaluation active n/a
-- -----
-- -----
Feature licenses:
## Feature Count Type
State Expiration Date Note
-- -----
-- -----
1 REPLICATION 1 unexpired evaluation
active n/a
2 DDBOOST 1 unexpired evaluation
active n/a
3 RETENTION-LOCK-GOVERNANCE 1 unexpired evaluation
active n/a
4 ENCRYPTION 1 unexpired evaluation
active n/a
-- -----
-- -----
New licenses:
Capacity licenses:
## Feature Capacity Type State
Expiration Date Note
-- -----
-- -----
1 CAPACITY 87.31 TiB permanent (int) active n/a
-- -----
-- -----
Feature licenses:
## Feature Count Type State
Expiration Date Note
-- -----
-- -----
1 DDBOOST 1 permanent (int) active n/a
2 ENCRYPTION 1 permanent (int) active n/a
3 REPLICATION 1 permanent (int) active n/a
-- -----
-- -----
** New license(s) will overwrite all existing
license(s).

Do you want to proceed? (yes|no) [yes]: yes

eLicense(s) updated.

```

Use the # elicense show command to verify.

```

# elicense show
System locking-id:
V4MXYV1S7R6VZVRWW6T9JTMPPBZEGY4CL25FSPX775WJC8GM6P57YKTD
HGYDGR9AJZ4Y66CSH152YJRS6UPHFUZ2PP6VATMY2FMWSSKKZ8SHD

System software-id: Not available
Instance software-id: Not available

Licensing scheme: EMC Electronic License Management System
(ELMS) node-locked mode

Capacity licenses:
## Feature Capacity Type State
Expiration Date Note
-- -----
-- -----

```

```

-----  -----
1  CAPACITY  87.31 TiB  permanent (int)  active  n/a
-----  -----
-----  -----
Feature licenses:
##  Feature      Count      Type      State
Expiration Date  Note
-----  -----
-----  -----
1  REPLICATION  1  permanent (int)  active  n/a
2  DDBOOST      1  permanent (int)  active  n/a
3  ENCRYPTION   1  permanent (int)  active  n/a
-----  -----
-----  -----
License file last modified at : 2018/05/07 18:56:36.

```

3. Enable object store with # storage object-store enable

```

# storage object-store enable
Object-store is enabled.

```

4. Create the Object store

- a. **System Passphrase** is required to encrypt the object store credentials. It will also will be used to encrypt keys if file system encryption is enabled. If the passphrase has already been set, user will not be prompted to enter passphrase.
- b. For AWS, the **Baltimore CyberTrust Root** certificate is needed to communicate with object store and should be imported for the profile creation to succeed.

```

# storage object-store profile set
# storage object-store profile set
A passphrase needs to be set on the system.
Enter new passphrase: <enter-passphrase-string-meeting-
requirements>
Re-enter new passphrase: <re-enter-passphrase-string>
Passphrases matched.
The passphrase is set
DD VE is running in AWS. Role-based access will be used to
access s3.
Enter the bucket name: <name-of-the-bucket>
Object-store endpoint needs the Baltimore CyberTrust
Root certificate to be imported.
Do you want to import that certificate with below
fingerprint?
D4:DE:20:D0:5E:66:FC:53:FE:1A:50:88:2C:78:DB:
28:52:CA:E4:74  (yes|no) [yes]:

```

Profile is set.

5. Add the storage using # storage add tier active dev3

```

# storage add tier active dev3
Checking storage requirements...done
Adding dev4 to the active tier...done
Updating system information...done
dev4 successfully added to the active tier.

```

Multiple devices can also be added as metadata storage using the following CLI command. This will be useful in when adding dev4, dev5, and dev6 to the DD VE:

```
VE: # storage add tier active dev4-6
# storage add tier active dev4-6

Checking storage requirements...done
Adding dev4 to the active tier...done

Updating system information...done

dev4 successfully added to the active tier.

Checking storage requirements...
done
Adding dev5 to the active tier...done

Updating system information...done

dev5 successfully added to the active tier.

Checking storage requirements...
done
Adding dev6 to the active tier...done

Updating system information...done

dev6 successfully added to the active tier.
```

Note

Use the command below to see the disks that are attached.

```
# storage show all
# storage show all
Active tier details:
Device      Device      Device
Group          Size
-----
(availble)    3          500.0 GiB
(availble)    5          200.0 GiB
-----
Spindle     Devices     Count     Total Size
Group
-----
1           5           1          200.0 GiB
3           3           1          500.0 GiB
-----
Active tier maximum capacity: 0.5 TiB
```

```
Storage addable devices:
Device      Device      Device
Type          Size
-----
(unknown)    4          100.0 GiB
-----
```

6. Create the file system# filesystem create

```
# filesystem create
A filesystem of approximate size 2.71 TiB will be created.
Do you want to continue? (yes|no) [yes]: yes
```

```

ok, continuing.

This will take 5 - 10 minutes.

Provisioning storage...
#####
[100%]

Initializing filesystem...
#####
[100%]

snapshot schedules deleted

You now have a freshly initialized filesystem.
Enable the filesystem using 'filesys enable'.

```

7. Enable file system# filesys enable

```

# filesys enable
Please wait.....
The filesystem is now enabled.

```

After you finish

Note

Set up NTP time synchronization by following the steps in [Setting up NTP Time Synchronization in AWS](#).

System Headswap

This section describes how the system headswap command recovers DD VE with head unit failure in AWS.

To perform system headswap, vNVRAM disk and Metadata disks from system A (original system) should be available, and they will be attached to the new instance B. If either vNVRAM disk or any metadata disk is not available, the command “system recovery from object-store” should be used instead

Procedure

1. Create instance B with Head Unit (root disk only) with the same instance type as the original one.
2. Detach the vNVRAM and Meta-data storage from the broken head unit.
3. Attach the vNVRAM and Meta-data storage above to instance B Head Unit.
4. Set system passphrase

Note

Please set the passphrase to match with system A, otherwise, headswap will fail to proceed.

```

# system passphrase set
Enter new passphrase:
Re-enter new passphrase:
Passphrases matched.
The passphrase is set.

```

5. **Note**

Before executing the headswap command, please make sure that the system A is powered off . This step is required to detach the bucket from system A and make it available to be attached with system B.

Execute system headswap

Note

System will reboot during the headswap process

```
# system headswap
This command returns the system back to its prior operational
conditions. The system will be rebooted before
resuming normal operations.

** If system passphrase was set on the old head, you will
   need to do one of the following after headswap completes:
      - unlock the filesystem           if you have encrypted
data, or
      - set the system passphrase      if you don't have
encrypted data
Are you sure? (yes|no) [no]: yes

ok, proceeding.

Please enter sysadmin password to confirm 'system headswap':
Restoring the system configuration, do not power off /
interrupt process ...
Broadcast message from root (Mon Apr 30 13:44:10 2018):

The system is going down for reboot NOW!
```

6. Check filesys status **after the headswap process is complete.**

```
# filesys status
The filesystem is enabled and running.
```

System Recovery

This section describes how the system recovery command recovers DD VE with head unit, vNVRAM disk, Metadata disk, on failure .

Before you begin

The system recovery command recovers DD VE with head unit, vNVRAM disk, metadata disk failure, or any combination of the three. However, if both vNVRAM disk and Metadata disks are available, then the `system headswap` command should be used instead.

Procedure

1. Create instance B with the same configuration as instance A, including instance type, metadata disk capacity.
2. Enable object-store

```
# storage object-store enable
Object-store is enabled.
```

3. Set object-store profile

a. Set the passphrase to match with system A, otherwise, the recovery will fail to proceed.

b. Set the same s3 bucket name from system A.

```
# storage object-store profile set
A passphrase needs to be set on the system.
Enter new passphrase: <enter-passphrase-string-meeting-
requirements>
Re-enter new passphrase: <re-enter-passphrase-string>
Passphrases matched.
The passphrase is set
DD VE is running in AWS. Role-based access will be used to
access s3.
    Enter the bucket name: <name-of-the-bucket>
    Object-store endpoint needs the Baltimore CyberTrust
Root certificate to be imported.
    Do you want to import that certificate with below
fingerprint?
    D4:DE:20:D0:5E:66:FC:53:FE:1A:50:88:2C:78:DB:
28:52:CA:E4:74 (yes|no) [yes]:
Profile is set.
```

```
# storage object-store profile set
```

c. Follow rest of CLI prompts.

4. Add EBS volumes to the active tier

Note

Add EBS volumes to match or exceed the capacity of system A .

```
# storage add dev3
Object-store is not enabled. Filesystem will use block storage
for user data.
    Do you want to continue? (yes|no) [no]: yes
    Checking storage requirements...done
    Adding dev3 to the active tier...done
    Updating system information...done
    dev3 successfully added to the active tier.
```

5. Run system recovery precheck

```
# system recovery precheck from object-store
Recovery precheck passed. Use start command to start the
recovery.
```

6. Execute the recovery

```
# system recovery start from object-store
System recovery has started. Use status command to check the
status.
```

7. Check the status with recovery status

```
# system recovery status
System recovery is running: stage 2 of 6 (attaching object-
store)
```

Note

The system will reboot during the recovery process.

8. Check filesystem status after the recovery process completed.

```
# filesystem status
The filesystem is enabled and running.
```

Instance Upgrade and Capacity Expansion

For instance upgrade refer to: [Changing the Instance Type](#)

For capacity expansion refer to: [Raw Physical Capacity Needed](#)

DD VE on Azure hot blob storage

DD VE on Azure hot blob storage provides a data protection solution that enables customers to protect their operational data in the cloud, to backup/restore the active tier's data into cloud object store, while the DD VE is running in the cloud. This section describes first-time setup procedures, and includes how to manage and monitor DD VE in an Azure environment. (This section describes first-time setup procedures, and includes how to manage and monitor DD VE in an Azure environment.)

Note

For DD VE on Block Storage see [Configuring DD VE on Block Storage using the DDSM Interface](#) Azure Deployment Options:

- Template deployment
- Marketplace deployment

Azure System Configuration requirements

These are the system configuration requirements for configuring DD VE on Azure.

System Configuration Requirements for Azure Meta data disk type: Standard

Table 26 Azure System Requirements

Instance Type	DD VE Capacity	#vCPU Memory
Standard_F8	16 TB	8, 16 GB
Standard_D4_v2	32 TB	8, 28 GB
Standard_D16_v3	96 TB	16,64 GB

Table 27 Azure System Requirements (cont.)

Instance type	Standard_F8	Standard_D4_v2	Standard_D16_v3
CPU	4	8	16
Memory (GiB)	16	28	64
System Disk	250 GiB Standard Root disk	250 GiB Standard Root disk	250 GiB Standard Root disk
	10 GiB Standard NVRAM disk	10 GiB Standard NVRAM disk	10 GiB Standard NVRAM disk

Table 27 Azure System Requirements (cont.) (continued)

Instance type	Standard_F8	Standard_D4_v2	Standard_D16_v3
Storage capacity	16 TB	32 TB	96 TB

Specifications for DD VE in Azure

The tables show the instance types and storage volumes needed by DD VE in Azure. Standard HDD storage is used for all volumes. Azure DD VE has a different max capacity for optimizing the instance type cost.

Specifications for DD VE on Azure Hot Blob Storage

Table 28 Azure Hot Blob Storage Infrastructure requirements

Configuration	Instance Type	Block Storage Volumes			Object Storage Capacity	Network Interface
		Root Disk	vNVRAM Disk	Metadata Disk		
16TB	Standard_F8	Standard/ 250GB	Standard/ 10GB	Standard (1 - 2TB)	0-16TB	Default = 1
32TB	Standard_D4_V2	Standard/ 250GB	Standard/ 10GB	Standard (1 - 4TB)	0-32TB	
96TB	Standard_D16_V3	Standard/ 250GB	Standard/ 10GB	Standard (1 - 10TB)	0-96TB	

Table 29 Azure Hot Blob Storage Stream Counts

Configuration	Write Stream	Read Stream	Repl Source	Repl Dest	Mixed Stream	Max Mtree
16TB	45	30	45	45	60	6
32TB	90	50	90	90	90	14
96TB	180	50	90	180	180	32

Note

Host cache will not be supported for DD VE on Azure hot blob storage.

Creating DD VE from Azure Marketplace

DD VE is available in the Azure Marketplace. This section lists the steps for deploying DD VE from the Azure Marketplace.

Creating the DD VE from the Marketplace

Procedure

1. Login to the Azure portal
2. Search for Dell EMC to find Data Domain Virtual Edition in Azure Place.
3. Begin the deployment.

4. Configure basic settings.
 - Name: the name for DD VE. (Maximum length of the name is 10 characters)
 - VM disk type: only HDD is supported
 - Username: must be “sysadmin”
 - Authentication type: Both SSH publicKey and Password authentication are supported.
 - SSH public key: Copy and paste the SSH publicKey. Change the default password of “sysadmin” at the first login
 - Password: the password for sysadmin
 - Subscription: specify a subscription
 - Resource group: specify or create a resource group
5. Choose VM size: Select the VM size according to the DD VE Specification table based on the DD VE capacity.
6. Configure option features
 - Storage: Use managed disks or un-managed disks
 - Storage account: Specify the storage amount
 - Network: Specify the virtual network, subnet, network security group and public IP address. (Deployment of DD VE in private subnet and leave the public IP address as “None” is recommended for security consideration)
 - Extensions: no extensions
 - High availability: None
 - Monitoring
 - Boost Diagnostics: Enabled
 - Guest OS Diagnostics: Disabled
 - Diagnostics Storage Account: specify the storage account in which to save the diagnostics logs
7. Review the summary: Review the configuration summary for creating the DD VE and make changes if necessary.
8. Select the “Create ”button. The Azure portal will start the DD VE deployment.

Adding disks in Azure

Make sure you have enough licensed capacity available to add new capacity to DD VE. When adding additional capacity, make sure the DD VE instance can support the new capacity. If the new capacity is more than the DD VE instance supported capacity, please upgrade the DD VE instance.

Note

Azure provides two types of disk storage: Premium and Standard. DD VE only supports Standard HDD as metadata disks (DD VE on block storage) or metadata disks (DD VE on hot blob storage). Also, you won't need to add a vNVRAM disk. After deploying the DD VE in Azure, the 10 GB vNVRAM disk will be automatically created.

New storage for the DD VE must meet the following requirements:

- **DD VE on hot blob storage**
 - The minimum size of the first metadata disk is 1 TB.
 - The minimum size of any subsequent metadata disks is 1 TB.
- **DD VE on block storage**
 - The first data disk is 512 GB.
 - The size of subsequent data disks is 1 TB.

After you finish

To add additional storage in the future, follow the requirements above. It is not necessary to shut down the virtual machine before adding storage.

Note

The virtual disk cannot be resized. Create a new virtual disk to add additional storage to the virtual machine.

Adding unmanaged disks or managed disks in Azure

Use these procedures to add unmanaged or managed disks to DD VE in Azure.

Before you begin

Make sure you have enough licensed capacity available to add new capacity to DD VE. When adding additional capacity, make sure the DD VE instance can support the new capacity. If the new capacity is more than the DD VE instance's supported capacity, please upgrade the DD VE instance.

Note

Azure provides two types of disk storage :Premium and Standard. DD VE only supports Standard HDD as data disks (DD VE on block storage) or metadata disks(DD VE on hot blob storage). After deploying the DD VE in Azure, the vNVRAM disk will be automatically created.

Note the following important information:

- For template-based deployments, only unmanaged disks are supported.
- Use these same procedures to add additional storage in the future.
- It is not necessary to shut down the VM before adding storage.
- The virtual disk cannot be resized. Users must create a new virtual disk to add additional storage to the VM.

New Storage for the DD VE on object storage must meet the following requirements:

- The minimum size of the first metadata disk is 1 TB.
- The minimum size of any subsequent data disks is 1 TB.

Procedure

1. Allocate and attach unmanaged disks to DDVE:
 - a. Log into the Azure portal.
 - b. Add the name of the DD VE VM.
 - c. Navigate to the right pane, and click **SETTINGS > Disks**.

Figure 9 Login to the Azure portal and select SETTINGS > Disks

d. Click the **Add data disk** button.

e. A pop up window displays. Add the following:

- A name for the data (DD VE on block storage) or metadata (DD VE on hot blob storage) disk
- Specify the **Source Type** as **New (empty disk)**
- Specify the **Account type** as **HDD**, and insert a size up to 1024 GiB
- Navigate to the **Storage container** file path.
- Enter the **Storage blob** name

2. Click **Save** (with disk icon) in the top left corner of the page to add the data disk.

3. Allocate and attach managed disks to DD VE:

a. Log into the Azure portal.

b. Add the name of the DD VE VM.

c. Navigate to the right pane, and click **SETTINGS > Disks**.

d. Click the **Add data disk** button.

e. Choose **Create disk** from the drop-down menu (red box).

f. A pop up window displays. Add the following:

- A name for the data (DD VE on block storage) or metadata (DD VE on hot blob storage) disk
- Specify the **Resource group**
- Specify the **Location**
- Specify the **Availability zone**

- Specify the **Account type**
- Specify the **Source Type**
- Specify the **Size**

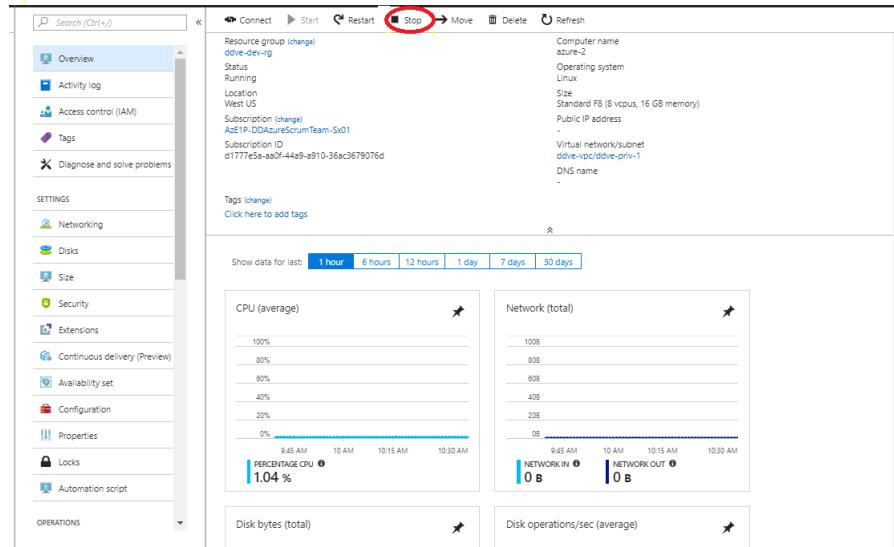
4. Click the **Create** button.
5. Click **Save** (with disk icon) in the top left corner of the page to add the data disk.

Resizing the DD VE instance in Azure

Follow these steps to resize the DD VE virtual machine instance.

Procedure

1. Stop the current DD VE instance.



2. Resize the VM.

- a. **Settings:** In the Azure portal, click the DD VE instance settings option.
- b. **Size:** Select the new size to be upgraded for the DD VE instance.

Note

The DD VE instance type may only be upgraded in this direction:
Standard_F8>Standard_D4_v2>Standard_D16_v3

RECOMMENDED	SKU	TYPE	COMPUTE ...	VCPUS	GB RAM	DATA DISKS	MAX IOPS	LOCAL SSD	Premium ...	ADDITIONAL ...	USD/MONTH
	D2_v3	Standard	General purpos	2	8	4	2x500	50 GB	HDD		\$87.05
	D4_v3	Standard	General purpos	4	16	8	4x500	100 GB	HDD		\$174.10
	D8_v3	Standard	General purpos	8	32	16	8x500	200 GB	HDD		\$348.19
	D16_v3	Standard	General purpos	16	64	32	16x500	400 GB	HDD		\$696.38
	D32_v3	Standard	General purpos	32	128	32	32x500	800 GB	HDD		\$1,392.77
	D1_v2	Standard	General purpos	1	3.5	4	2x500	50 GB	HDD		\$47.76
	D2_v2	Standard	General purpos	2	7	8	4x500	100 GB	HDD		\$95.23
	D3_v2	Standard	General purpos	4	14	16	8x500	200 GB	HDD		\$191.95
	D4_v2	Standard	General purpos	8	28	32	16x500	400 GB	HDD		\$383.16
	D5_v2	Standard	General purpos	16	56	64	32x500	800 GB	HDD		\$766.32
	D11_v2	Standard	Memory optimi	2	14	8	4x500	100 GB	HDD		\$121.27
	D12_v2	Standard	Memory optimi	4	28	16	8x500	200 GB	HDD		\$242.54
	D13_v2	Standard	Memory optimi	8	56	32	16x500	400 GB	HDD		\$485.09
	D14_v2	Standard	Memory optimi	16	112	64	32x500	800 GB	HDD		\$970.18
	D15_v2	Standard	Memory optimi	20	140	64	40x500	1000 GB	HDD		\$1,213.46

3. The Azure portal indicates the DD VE VM has been successfully resized. Click the Start button to launce the resized DD VE instance.

VM details for ddve-1

Computer name: ddve-2
Operating system: Linux
Size: Standard F8 (8 vcpus, 16 GB memory)
Public IP address: 52.177.111.11
Virtual network/subnet: ddve-vpc/ddve-priv-1
DNS name: ddve-2.eastus.cloudapp.azure.com

Metrics (last 1 hour):

- CPU (average): 1.04 %
- Network (total): 0 B
- Disk bytes (total): 0 B
- Disk operations/sec (average): 0

Capacity expansion in Azure

Upgrading DD VE to a higher capacity

If the higher capacity is supported by the current DD VE configuration, follow these steps.

1. Add the needed hard metadisks for the new capacity.
2. Configure the newly added data disks using the CLI command `storage add tier active dev<device id>(Or, use DD SM GUI).`
3. Expand the file system using the CLI command `filesys expand`.

If the higher capacity is not supported by the VM size configuration of the current DD VE, you must upgrade to larger instance type from the current VM size to a higher VM size (see the tables in [Specifications for DD VE in Azure](#) on page 79).

Configuring DD VE on Azure hot blob storage using the Data Domain System Manager interface

DD VE can be configured on Azure Hot Blob Storage using one of the Data Domain System Manager (DD SM) interface options: GUI or CLI.

Before you begin

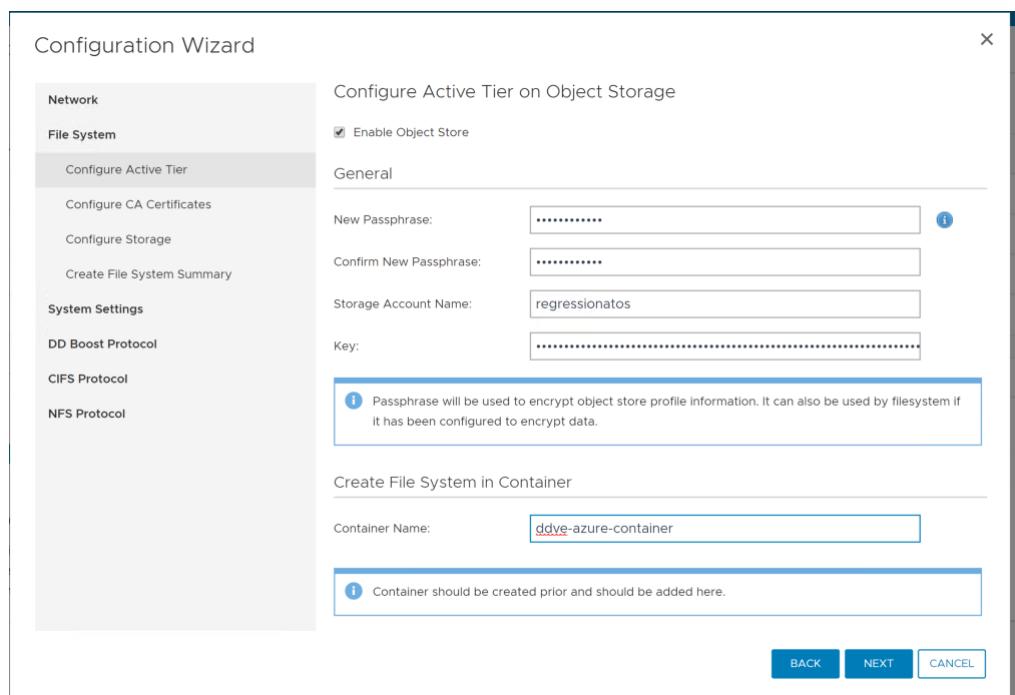
Recommended metadata storage is 10% of the total capacity.

Configuring DD VE on Azure hot blob storage using the DD SM interface (http/https)

The UI configuration wizard helps go through the Azure hot blob storage configuration and file system creation on DD.

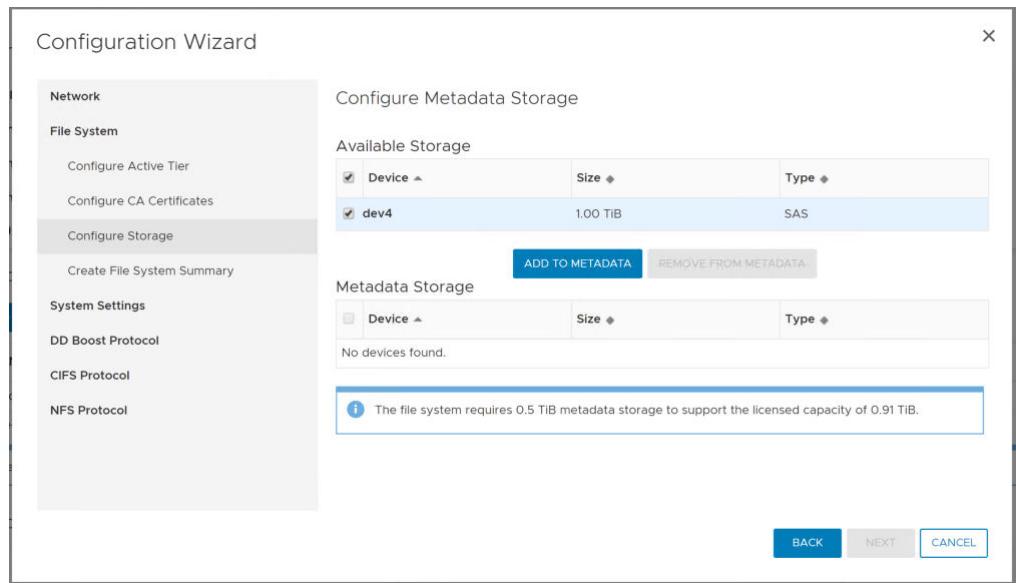
Procedure

1. Log in to the Data Domain System Manager with the account name: sysadmin and the password specified at deployment.
2. For the “Apply your license” step: select “Pre-installed evaluation license” with 500 GiB. Then click “Apply”.
3. Accept the End User License Agreement (EULA) by clicking “I accept the terms of EULA”
4. Next, the “Configuration Wizard” will guide you through the Azure hot blob storage configuration and file system creation on the DD VE. Select “File System” and click “Yes”.
5. Select “Configure Active Tier”. Then, select “Enable Object Store”checkbox to configure the Azure hot blob storage system.
6. Enter the passphrase, account name, key, and container name. (The container may be created through the Azure portal.)

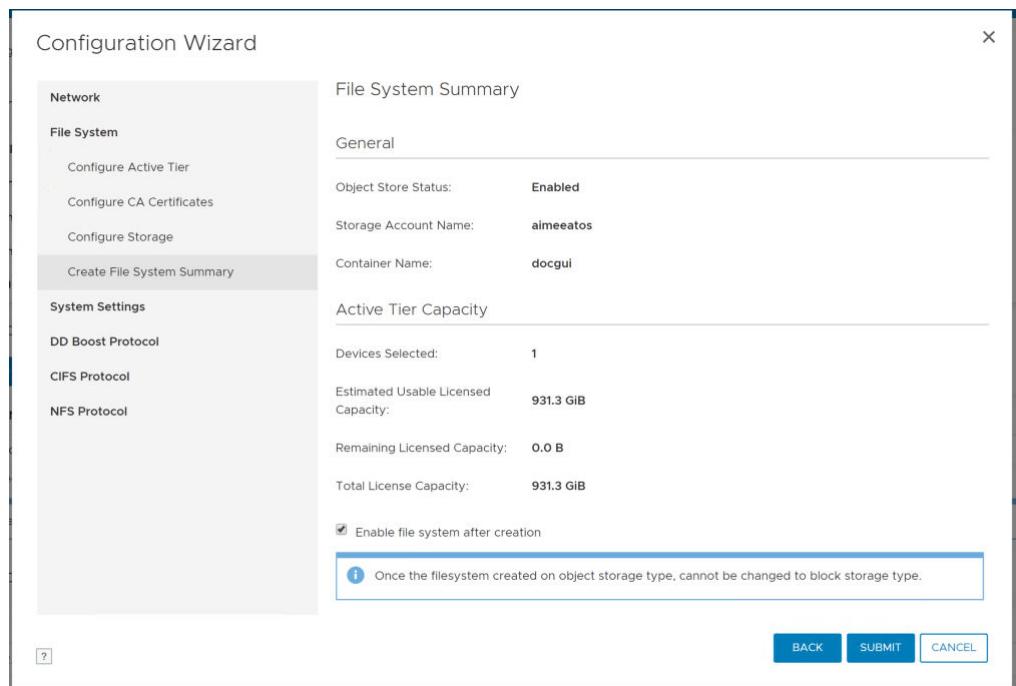


7. Import the Baltimore CyberTrust Root certificate to communicate with Azure Object Store.

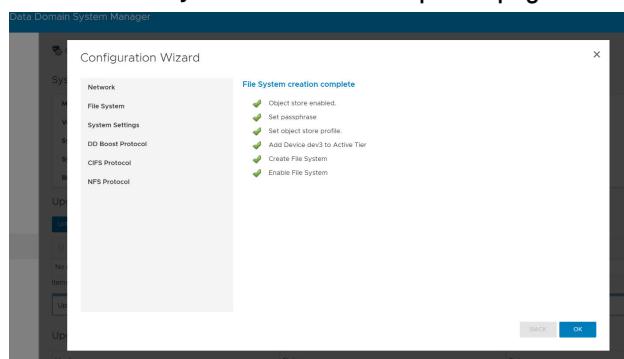
8. Add the metadata storage.



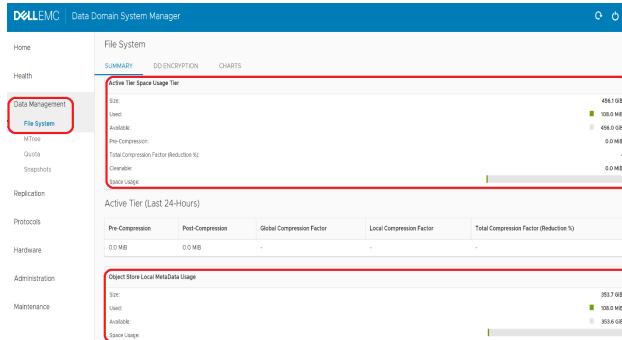
9. Review the summary and Click "Submit" to create the file system and enable it.



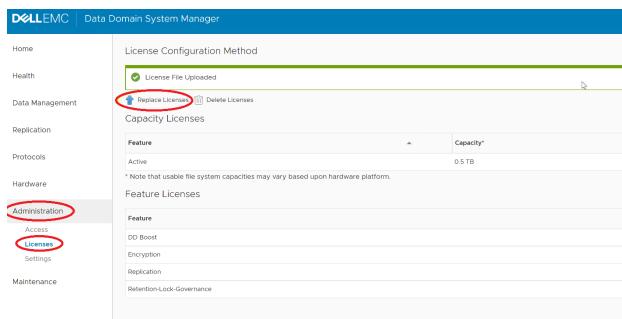
10. Review "File System Creation Complete" page.



11. The File System section under the Data Management tab has the space usage and availability details for the hot blob storage as well as the local metadata storage.



12. To configure or update the eLicense on the DD VE, click “Replace licenses” button in the Licenses page as shown in this image.



13. To relaunch the configuration wizard, select “System” under “Maintenance”, then “Configure System”

Configuring DD VE in Azure using the Command Line Interface

Before you begin

DD VE can be configured using the CLI option.

Configuring the DD VE using the CLI

Procedure

1. Log in to the DD VE instance via SSH using the sysadmin account and password you specified in the template.
 - When logging in for the first time, you will be asked to change the password. Enter the new password. The initial configuration wizard will start.
 - As shown in the screen capture below, you can add the eLicense by answering “yes” to the eLicense configuration prompt. This wizard can also be used to configure the network.

```
$ ssh -l sysadmin <DDVE ip address>
The authenticity of host '***.***.***.*** (**.**.**.**)' can't be
established.
ECDSA key fingerprint is SHA256:evoXXGRgCzp/
tmrtWRIAeOWLpI7ymOq9mwTBwH9J2bs.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '***.***.***.***' (ECDSA) to the list of
known hosts.
EMC Data Domain Virtual Edition
```

```
-----  
Password:  
Password:  
Welcome to Data Domain OS *****  
-----  
Press any key then hit enter to acknowledge the receipt of EULA  
information:  
Press any key then hit enter to acknowledge the receipt of EULA  
information: q  
Do you want to configure system using GUI wizard (yes|no) [no]:  
  
Network Configuration  
    Configure Network at this time (yes|no) [no]:  
  
eLicenses Configuration  
    Configure eLicenses at this time (yes|no) [no]:  
  
System Configuration  
    Configure System at this time (yes|no) [no]:  
  
Storage object-store profile Configuration  
    Configure Storage object-store profile at this time  
(yes|no) [no]:  
  
Configuration complete.
```

2. To update the eLicense on DD VE, copy license file to /ddvar and use the file name as follows: # elicense update <filename.lic>.

```
# elicense update atos_cap_96_TB.lic

Existing licenses:

Capacity licenses:
##  Feature    Capacity  Type          State
Expiration Date  Note
---  -----  -----  -----  -----
1    CAPACITY  0.45 TiB  unexpired evaluation  active  n/a
---  -----  -----  -----  -----
-----  ----

Feature licenses:
##  Feature          Count  Type
State    Expiration Date  Note
---  -----  -----  -----
1    REPLICATION      1  unexpired evaluation
active  n/a
2    DDBOOST          1  unexpired evaluation
active  n/a
3    RETENTION-LOCK-GOVERNANCE  1  unexpired evaluation
active  n/a
4    ENCRYPTION        1  unexpired evaluation
active  n/a
---  -----  -----  -----
-----  ----

New licenses:

Capacity licenses:
##  Feature    Capacity  Type          State
Expiration Date  Note
---  -----  -----  -----  -----
1    CAPACITY  87.31 TiB  permanent (int)  active  n/a
---  -----  -----  -----  -----
-----  ----

Feature licenses:
##  Feature  Count  Type          State
```

	Expiration Date	Note				
1	DDBOOST	1	permanent	(int)	active	n/a
2	ENCRYPTION	1	permanent	(int)	active	n/a
3	REPLICATION	1	permanent	(int)	active	n/a

Use the # elicense show command to verify.

3. Enable object store with # storage object-store enable

```
# storage object-store enable  
Object-store is enabled.
```

4. Create the Object store

- a. **System Passphrase** is required to encrypt the object store credentials. It will also be used to encrypt keys if file system encryption is enabled. If the passphrase has already been set, user will not be prompted to enter passphrase.

b. **Account Name** if there is no any account existed, create one first. Account kind is blob storage.

c. **Primary Key storage** accounts>settings>access keys.

d. **Container Name** create a container under storage account.

e. **Baltimore Cyber Trust Root** certificate is needed to communicate with object store and should be imported for the profile creation to succeed.

```
# storage object-store profile set
# storage object-store profile set
A passphrase needs to be set on the system.
Enter new passphrase: <enter-passphrase-string-meeting-
requirements>
Re-enter new passphrase: <re-enter-passphrase-string>
Passphrases matched.
The passphrase is set
    Enter the account name: <name-of-the-account-name>
    Enter the primary key: <name-of-the-primary-key>
    Enter the container name: <name-of-the-container-name>

    Object-store endpoint needs the Baltimore CyberTrust
Root certificate to be imported.
    Do you want to import that certificate with below
fingerprint?
    D4:DE:20:D0:5E:66:FC:53:FE:1A:50:88:2C:78:DB:
28:52:CA:E4:74  (yes|no) [yes]: yes

    Profile is set.
```

5. Add the storage using # storage add tier active dev4

```
# storage add tier active dev4
Checking storage requirements...done
Adding dev4 to the active tier...done

Updating system information...done

dev4 successfully added to the active tier.
```

Multiple devices can also be added as metadata storage using the following CLI command. This will be useful in when adding dev4, dev5, and dev6 to the DD VE:

```
# storage add tier active dev4-6
# storage add tier active dev4-6

Checking storage requirements...done
Adding dev4 to the active tier...done

Updating system information...done

dev4 successfully added to the active tier.
```

```
Checking storage requirements...
done
Adding dev5 to the active tier...done

Updating system information...done

dev5 successfully added to the active tier.
```

```
Checking storage requirements...
done
Adding dev6 to the active tier...done

Updating system information...done

dev6 successfully added to the active tier.
```

Note**Use the command below to see the disks that are attached.**

```
# storage show all
# storage show all
Active tier details:
Device      Device      Device
Group          Size
-----
(availble)   4          1023.0 GiB
(availble)   5          1023.0 GiB
(availble)   6          1023.0 GiB
-----
Spindle     Devices     Count     Total Size
Group
-----
2          4          1          1023.0 GiB
3          5          1          1023.0 GiB
4          6          1          1023.0 GiB
-----
Current active tier size: 2.9 TiB
Active tier maximum capacity: 35.2 TiB**
** The maximum capacity supported by system memory.

Capacity License:
License     Total      Used      Remaining
-----
CAPACITY   87.31 TiB   2.70 TiB   84.61 TiB
-----
```

6. Create the file system# filesys create

```
# filesys create
A filesystem of approximate size 2.71 TiB will be created.
Do you want to continue? (yes|no) [yes]: yes
ok, continuing.

This will take 5 - 10 minutes.

Provisioning storage...
#####
[100%]

Initializing filesystem...
#####
[100%]

snapshot schedules deleted

You now have a freshly initialized filesystem.
Enable the filesystem using 'filesys enable'.
```

7. Enable file system# filesys enable

```
# filesys enable
Please wait.....
The filesystem is now enabled.
```

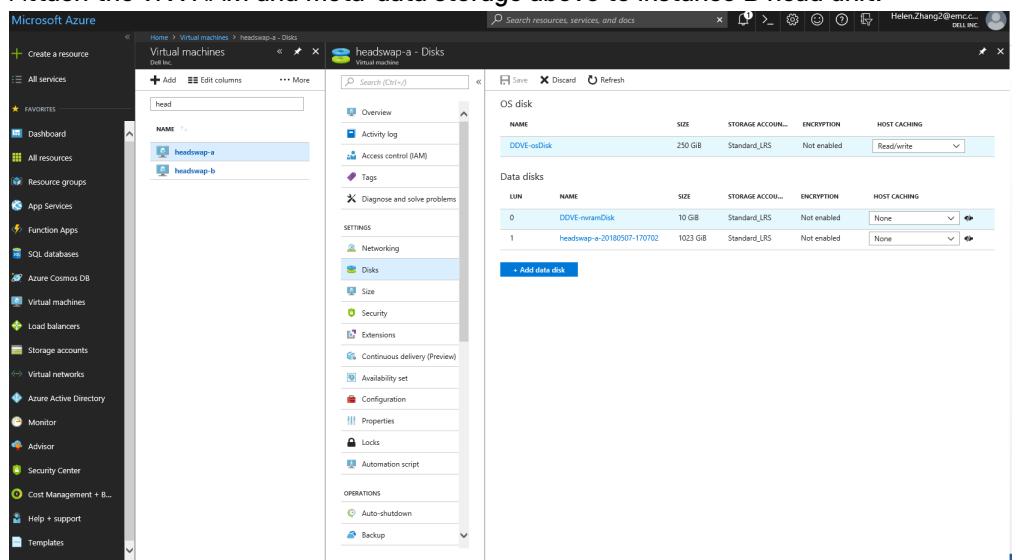
System Headswap - Azure

This section describes how the system headswap command recovers DD VE with head unit failure in Azure.

To perform system headswap, vNVRAM disk and metadata disks from system A should be available, and they will be attached to the new instance B. If either vNVRAM disk or any metadata disk is not available, the `system recovery` operation from object-store should be used instead.

Procedure

1. Create instance B with head unit (root disk only) with the same instance type as the original one.
2. Detach the vNVRAM and meta-data storage from the broken head unit.
3. Attach the vNVRAM and meta-data storage above to instance B head unit.



4. Set system passphrase.

Note

Please set the passphrase to match with system A, otherwise, headswap will fail to proceed.

```
# system passphrase set
Enter new passphrase:
Re-enter new passphrase:
Passphrases matched.
The passphrase is set.
```

- 5.

Note

Before executing the headswap operation, please make sure that the system A is powered off. This step is required to detach the bucket from system A and make it available to be attached with system B.

Execute system headswap.

Note**System will reboot during the headswap process.**

```
# system headswap
This command returns the system back to its prior operational
conditions. The system will be rebooted before
resuming normal operations.

** If system passphrase was set on the old head, you will
   need to do one of the following after headswap completes:
      - unlock the filesystem           if you have encrypted
data, or
      - set the system passphrase      if you don't have
encrypted data
Are you sure? (yes|no) [no]: yes

ok, proceeding.

Please enter sysadmin password to confirm 'system headswap':
Restoring the system configuration, do not power off /
interrupt process ...
Broadcast message from root (Mon Apr 30 13:44:10 2018):
The system is going down for reboot NOW!
```

6. Check filesys status after the headswap process is complete.

```
# filesys status
The filesystem is enabled and running.
```

System Recovery - Azure

This section describes how the system recovery command recovers DD VE with head unit, vNVRAM disk, Metadata disk, on failure .

Before you begin

The system recovery command recovers DD VE with head unit, vNVRAM disk, metadata disk failure, or any combination of the three. However, if both vNVRAM disk and Metadata disks are available, then the `system headswap` command should be used instead.

Procedure

1. Create instance B with the same configuration as instance A, including instance type, metadata disk capacity.
2. Enable object-store

```
# storage object-store enable
```

3. Set object-store profile

- a. Set the passphrase to match with system A, otherwise, the recovery will fail to proceed.
- b. Set the same storage account/container name from system A.
- c. Follow rest of CLI prompts.

```
# storage object-store profile set
A passphrase needs to be set on the system.
Enter new passphrase: <enter-passphrase-string-meeting-
```

```
requirements>
Re-enter new passphrase: <re-enter-passphrase-string>
Passphrases matched.
The passphrase is set
    Enter the account name: <name-of-the-storage-account>
    Enter the primary key:
    Enter the container name: <name-of-the-container-name>

    Object-store endpoint needs the Baltimore CyberTrust
Root certificate to be imported.
    Do you want to import that certificate with below
fingerprint?
    D4:DE:20:D0:5E:66:FC:53:FE:1A:50:88:2C:78:DB:
28:52:CA:E4:74 (yes|no) [yes]:  
Profile is set.
```

4. Add data disk

Note

Add data disk with the capacity to match or exceed the capacity of system A .

```
# storage add dev4
```

5. Run system recovery precheck

```
# system recovery precheck from object-store
```

6. Execute the recovery

```
# system recovery start from object-store
```

7. Check the status with recovery status

```
# system recovery status
```

Note

The system will reboot during the recovery process.

8. Check filesystem status after the recovery process completed.

```
# filesys status
```

VMware Cloud on AWS

DD VE in VMC (VMware Cloud on AWS) provides a data protection solution that enables customers to protect their operational data in the cloud, to backup/restore data into cloud object store, while the DD VE is running in the cloud. This section

describes first-time setup procedures, and includes how to manage and monitor DD VE in the VMC environment.

VMC System Configuration requirements

These are the system configuration requirements for VMC (VMware Cloud on AWS) configuration.

System Configuration Requirements for Meta data disk type: Standard/VSAN

Note

The instance type mentioned in the table below is only for logical significance.

Table 30 VMC System Requirements

Instance Type	DD VE Capacity	#vCPU Memory
Standard_VMC_16	16 TB	4, 16 GB
Standard_VMC_32	32 TB	4, 24 GB
Standard_VMC_96	96 TB	8, 64 GB

DD VE Deployment in VMware Cloud on AWS

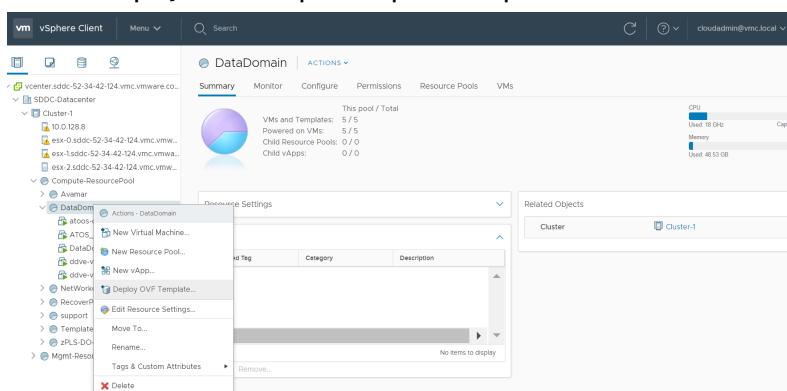
Before you begin

For deploying DD VE in VMware Cloud on AWS (VMC) on S3 object store:

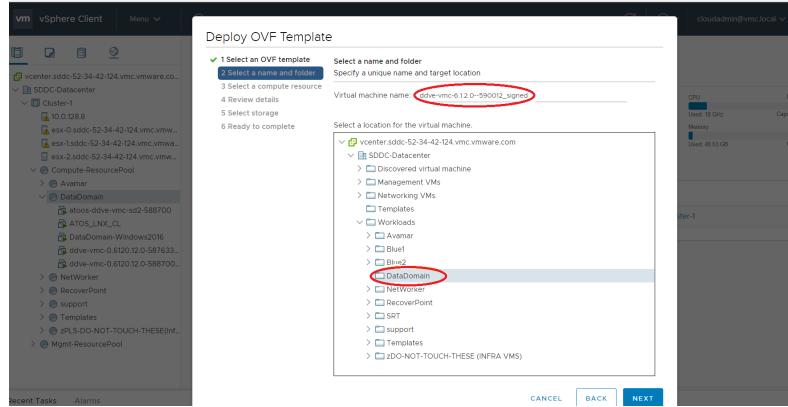
- You need an AWS account linked to your VMWare cloud account.
- The SDDC in VMC will be connected to an AWS account during creation. The subnet selected within the AWS account should be in the same region as the SDDC.
- Create the AWS S3 bucket, used by the DD VE, in the same region as the SDDC and within the same AWS account.
- The S3 traffic from VMC must be routed internally toward the AWS infrastructure. (During the SDDC in VMC setup, you should have already linked your AWS/VPC subnet account to the VMC account.) Secondly, within AWS VPC make sure to create the S3 endpoint so that the object store traffic routed within AWS infrastructure.
- Refer to VMC documentation for additional details.

Procedure

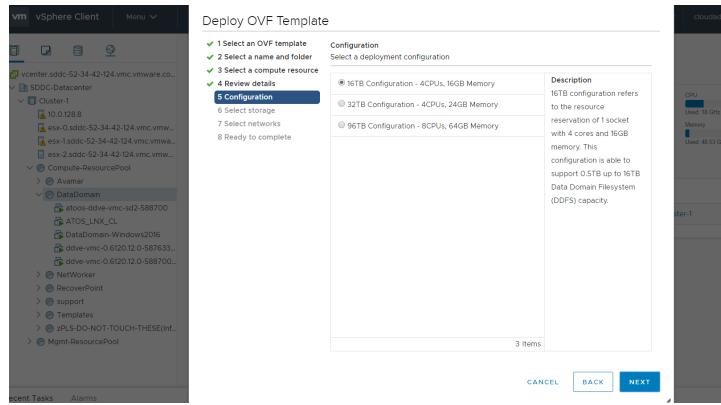
1. Select “Deploy OVF Template” option to upload the OVA file.



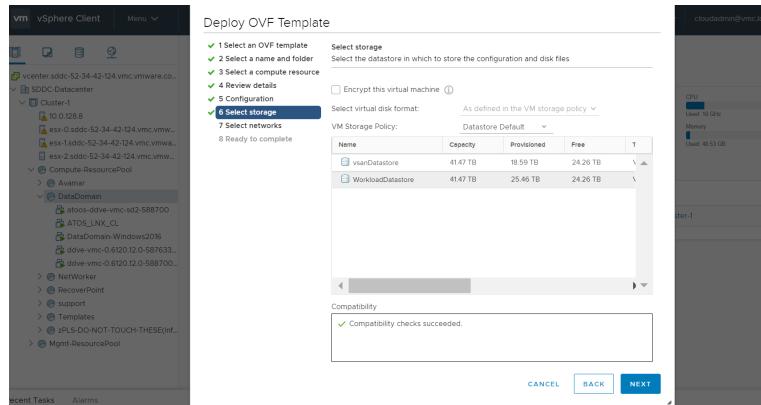
2. Choose the OVA file.
3. Create AWS S3 bucket in the same region as the SDDC.



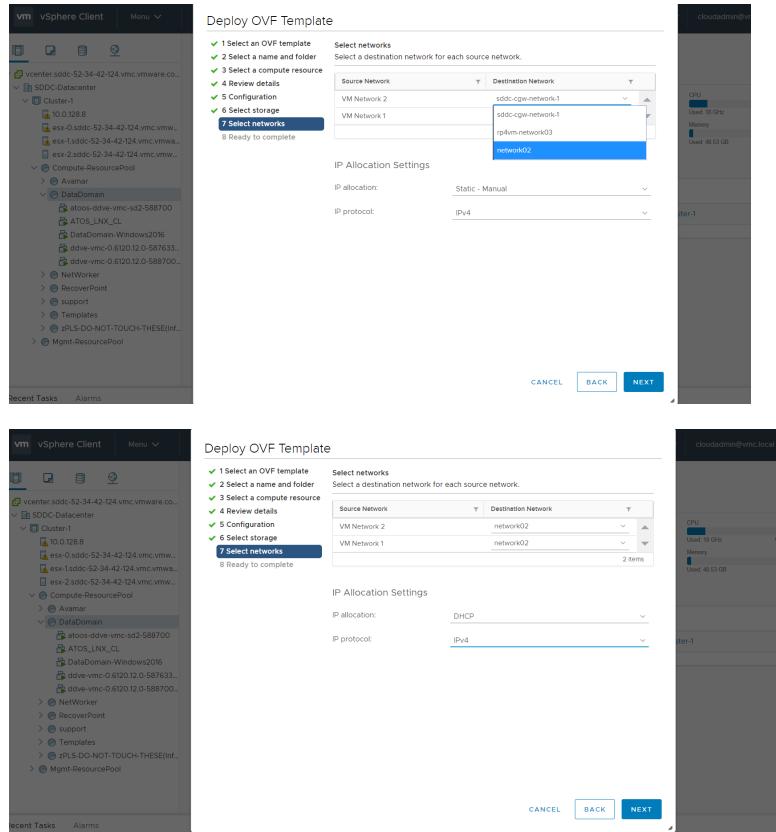
4. Select the compute resource pool. It may be one of the nodes in the cluster or a pre-configured resource pool. Ensure the compatibility checks succeed.
5. Review and verify the template details.
6. Choose the deployment configuration.



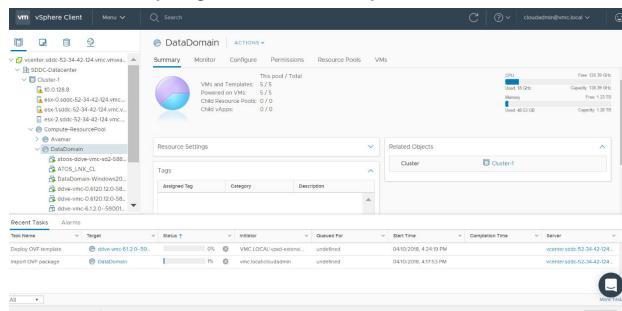
7. Select the storage for the metadata disks.



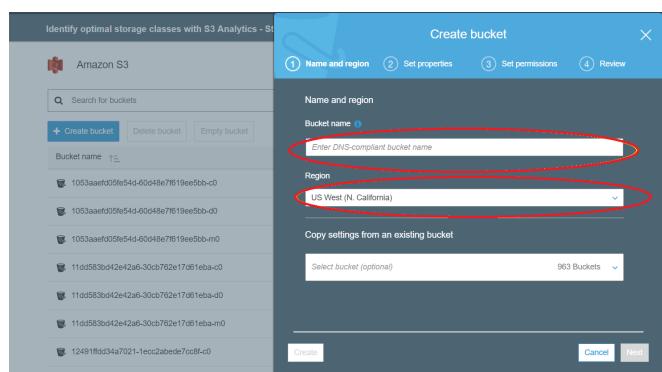
8. Select network.



9. Review the configuration details.
10. Monitor the progress of the deployment under the recent tasks panel.



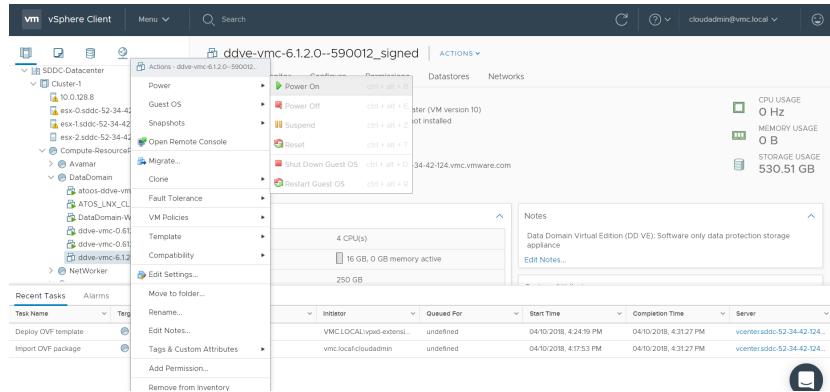
11. While the DD VE instance is being deployed from the AWS console.
 - a. Create AWS S3 bucket
 - b. Be sure to enter the bucket name
 - c. Select the same region as the VMC region



Note

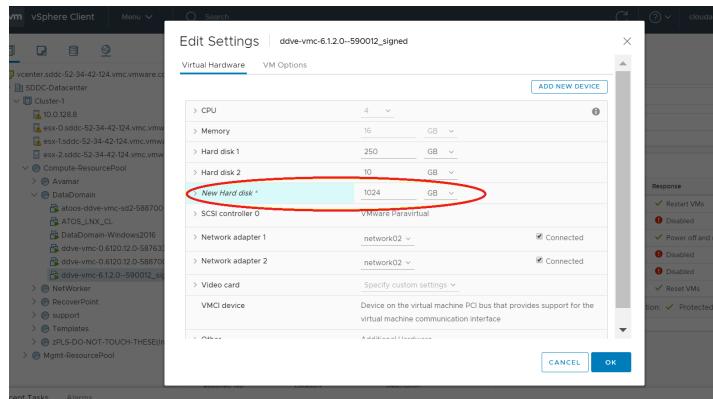
To avoid additional costs and potential performance issues, please ensure the S3 bucket and DD VE in VMC are available in the same region.

12. Power on the DD VE on the VMC console.



13. On the settings page, select “Add New Device”.

14. Add the meta data storage.



Configuring DD VE in VMC using the Data Domain System Manager interface

This section describes how to configure the DD VE in VMC using the DD SM interface ([http/https://](http://https://)).

Before you begin

Please refer to [Configuring DD VE using the Data Domain System Manager interface](#)

DD VE in VMC configuration using CLI

This section describes how to configure the DD VE in VMC.

Before you begin

Procedure

1. Log into the DD VE.
2. Enable object store using `storage object-store enable`.
3. Create profile by providing the passphrase and access credentials, region, and bucket. Use `storage object-store profile set`.

4. Add metadata tier #storage add dev3
5. Create file system #filesys create
6. Enable file system # filesys enable

VMware Cloud on AWS commands

VMware Cloud on AWS (VMC) enables AWS to run applications in vSphere-based cloud environments while making use of AWS services. These CLI commands have been modified for the interaction with VMC. These commands are not supported on physical Data Domain systems.

Table 31 DD VE-only commands

Command	Description
storage object-store profile set	<p>Set the access credentials/profile information. Role required: admin. When this command runs for the first time:</p> <pre>#storage object-store profile set Enter access key: Enter secret key: Enter region: Enter bucket name: Profile is set.</pre>
storage object-store profile show	<p>Show details of the object-storage profile. Role required: admin</p> <pre>#storage object-store profile show Region: <region> End Point: <end-point> Bucket Name: <bucket-name> Provider: AWS (VMware Cloud)</pre>

DD VE in VMC System HeadSwap

This section describes how the system headswap command recovers the DD VE in VMC with head unit failure.

Refer to [System Headswap](#)

DD VE in VMC System Recovery

This section describes how the system recovery command recovers DD VE with head unit, vNVRAM disk, Metadata disk, or failure .

Refer to [System Recovery](#)

Initial System Configuration

You can connect to the system to perform the initial system configuration with the DDSM Configuration Wizard or manually using the CLI.

DHCP is enabled on the DD VE system by default. If the DHCP service is available, the DD VE system will receive IP addresses from the DHCP server.

Note

DHCP is only activated automatically for the first network interface card (NIC) which is built into the virtual machine template. Any extra NICs must be configured manually by following instructions here <https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-network-interface-vm>.

Using the CLI

Access the CLI by using `ssh` or a terminal emulator to access the DD OS command line. The CLI configuration utility contains four sections: Network, eLicense, System, and DD Boost.

Using the GUI

Access DDSM by entering the IP address of the DD VE into a web browser, and logging in. The GUI Configuration Wizard contains six sections: Networking, File System, System Settings, DD Boost, CIFS, and NFS.

[Provisioning the storage with the CLI](#) on page 100 describes how to configure the DD VE manually with the CLI instead of using the configuration utility.

Provisioning the storage with the CLI

Before you begin

When configuring DD VE on S3 storage for AWS, verify that the maximum length of the bucket name or the container name when configuring the DD VE on Azure hot blob storage, ensuring it does not exceed 48 characters. Setup will not complete if the names exceed 48 characters. Refer to [How to find Amazon instance ID](#).

Procedure

1. Login to the system with the following credentials:
 - User name: `sysadmin`
 - Password:
 - **AWS default password:** `instance-id`
 - **Azure default password:** `changeme`

Note

At the first login prompt, run the `elicense` command to add a DD VE license.

2. Type **Ctrl-C** to exit the configuration utility.
3. **Use these steps to provision storage for active tier on block storage:**
 - a. Confirm that `virtualdisk dev3` exists, and shows the expected size:

```
# disk show hardware
```

The output includes a line similar to the following example:

```
# dev3 VMware Virtual disk 1.0 (unknown) 256.00 GiB SAS n/a
```

Note

The first two virtual disks (dev1 and dev2) are used for the system software and not for storage. Use the `disk show state` command to display System Dev details for these system disks.

b. Add the storage disk to the active storage tier:

```
# storage add dev3
```

If adding more than one virtual disk, repeat **step b.** for each virtual disk.

c. Create the file system:

```
# filesys create
```

Note

The `filesys create` command may take longer to complete if the hypervisor's storage is slow and does not meet the criteria.

d. Enable the file system:

```
# filesys enable
```

4. Use these steps to provision object storage:

a. Enable object store capability on DD VE:

```
# storage object-store enable
```

b. Set up the object storage profile:

```
# storage object-store profile set
```

It is important to note that:

- **For AWS:** DD VE automatically configures role-based S3 access.
- **For Azure:** The user supplies the storage account name, primary key, container name, and certificate.

c. Add the metadata disk to the system:

```
# storage add dev3 tier active
```

If adding more than one virtual disk, repeat **step c.** for each virtual disk.

d. Create the file system:

```
# filesys create
```

Note

The `filesys create` command may take longer to complete if the hypervisor's storage is slow and does not meet the criteria.

e. Enable the file system:

```
# filesys enable
```

Configuring DD VE in Data Domain System Manager

DD VE licensing and configuration can be accomplished through the Configuration Wizard in Data Domain System Manager. After the initial installation of a DD VE instance, the Configuration Wizard automatically appears after the licensing screen on the first launch of DDSM.

Note

The DAT is not supported for cloud DD VE.

Enter the DD VE virtual machine IP address into a web browser to launch Data Domain System Manager. Log in with the following credentials:

- Username: sysadmin
- Default password: changeme
 - **AWS:** Default password is the instance-id
 - **Azure:** Default password is defined in the template, or is specified from the input of Marketplace deployment.

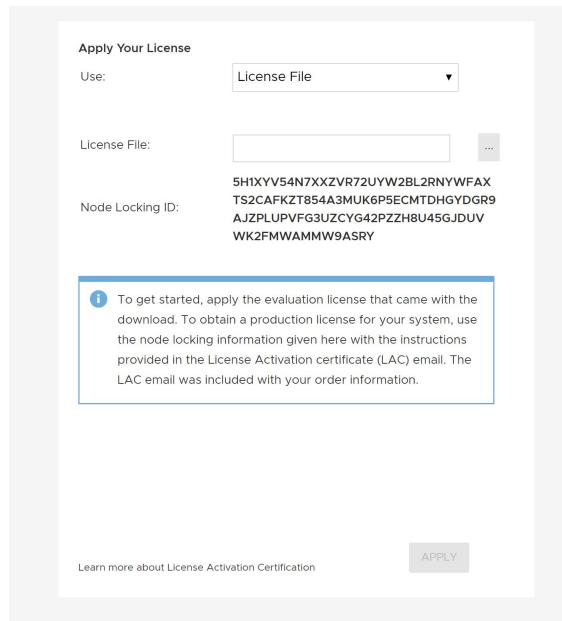
DD VE licensing

The **Apply Your License** window is the first screen that appears when DDSM is launched for the first time. The DD VE instance is locked until a license file is applied.

Click **Browse**, locate the license file for a purchased capacity license or the evaluation license included with the DD VE download, then click **Apply**.

Note

If you begin the configuration with the evaluation license, but wish to purchase a license later, you will need the Node Locking ID for the DD VE instance. Click **Administration > Licenses** to view the Node Locking ID.

Figure 10 DD VE Node Locking ID

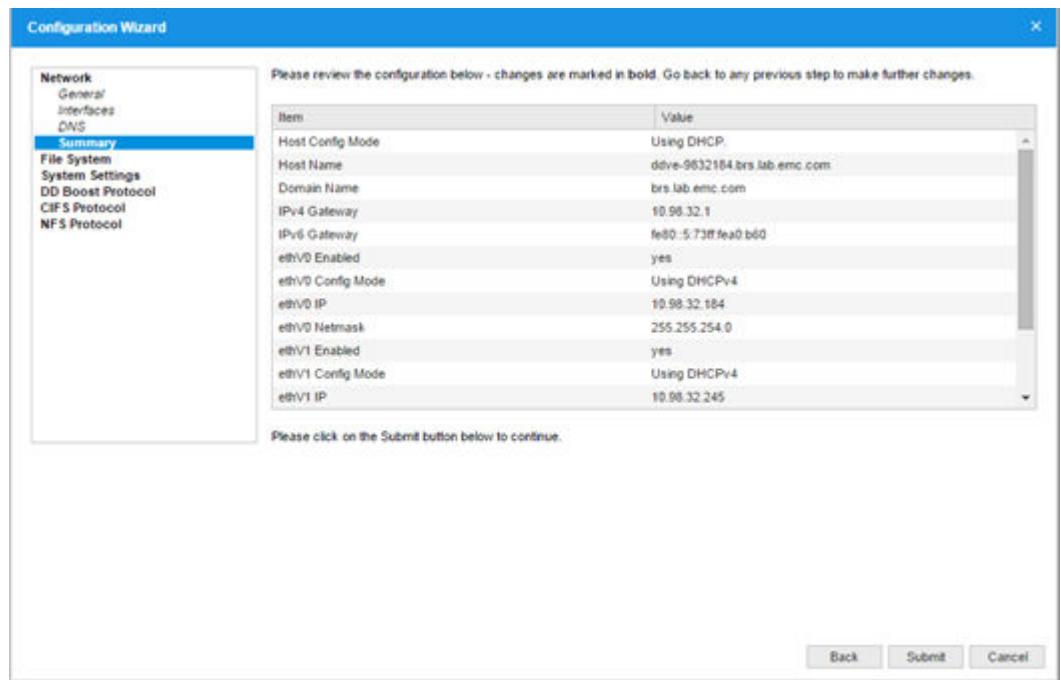
Note

When you obtain the original license file name the server, do not enter the comma in the license file name. DD OS will not accept the name if the comma is used.

DD VE configuration

After applying the DD VE license, the Configuration Wizard begins automatically. The wizard assists in configuring the following aspects of the DD VE:

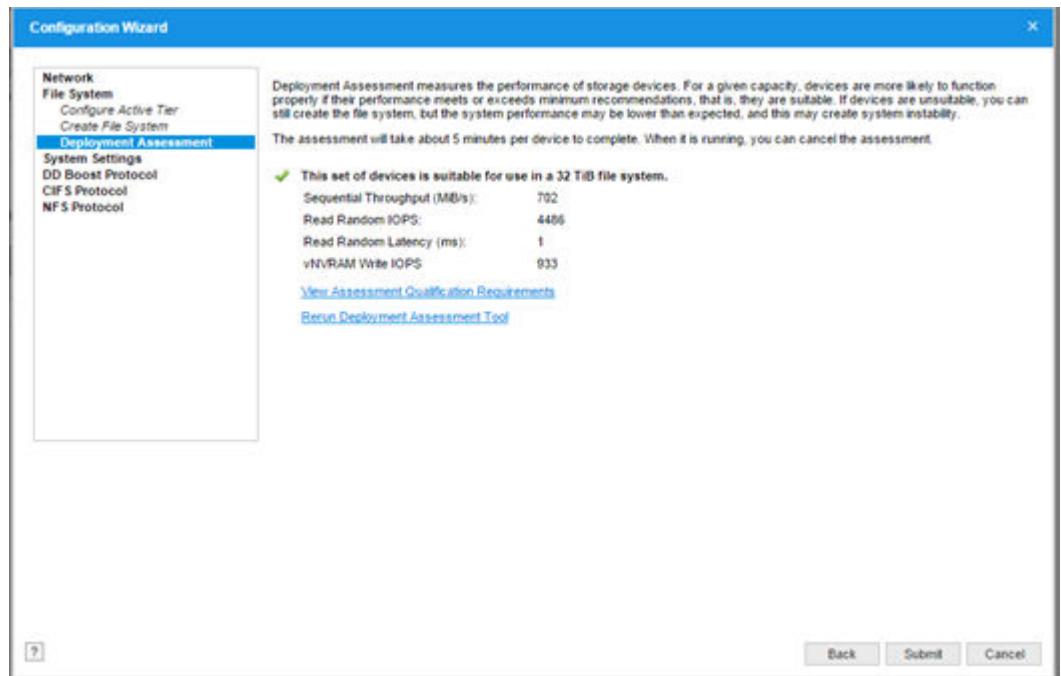
- Networking
 - DHCP or manual settings
 - Virtual interface ethV0 and ethV1 configuration
 - DHCP or manual DNS configuration

Figure 11 Configuration Wizard - Network

- File system

Note

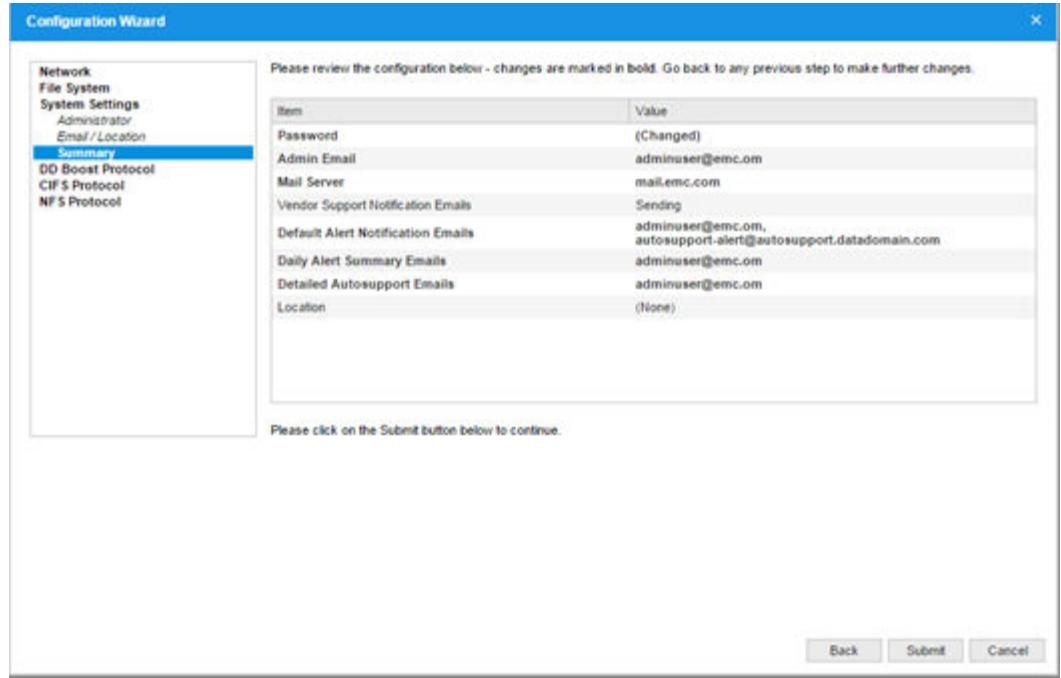
DD VE supports a maximum of six MTrees active at a given time, however up to 100 MTrees can be created on DD VE.

Figure 12 Configuration Wizard - File System

- System settings
 - Update the sysadmin password

- Optionally configure alert and autosupport email settings

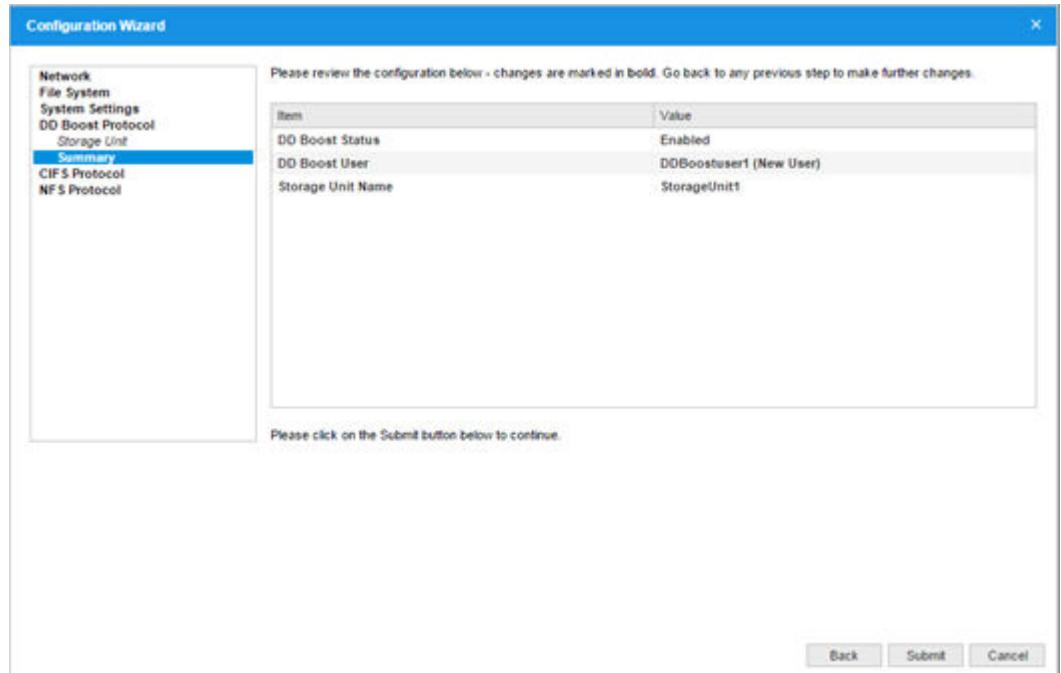
Figure 13 Configuration Wizard - System Settings



- DD Boost

- Create a Boost storage-unit, and assign a user ID to own it

Figure 14 Configuration Wizard - DD Boost Protocol



Enabling the object store feature with DD SM

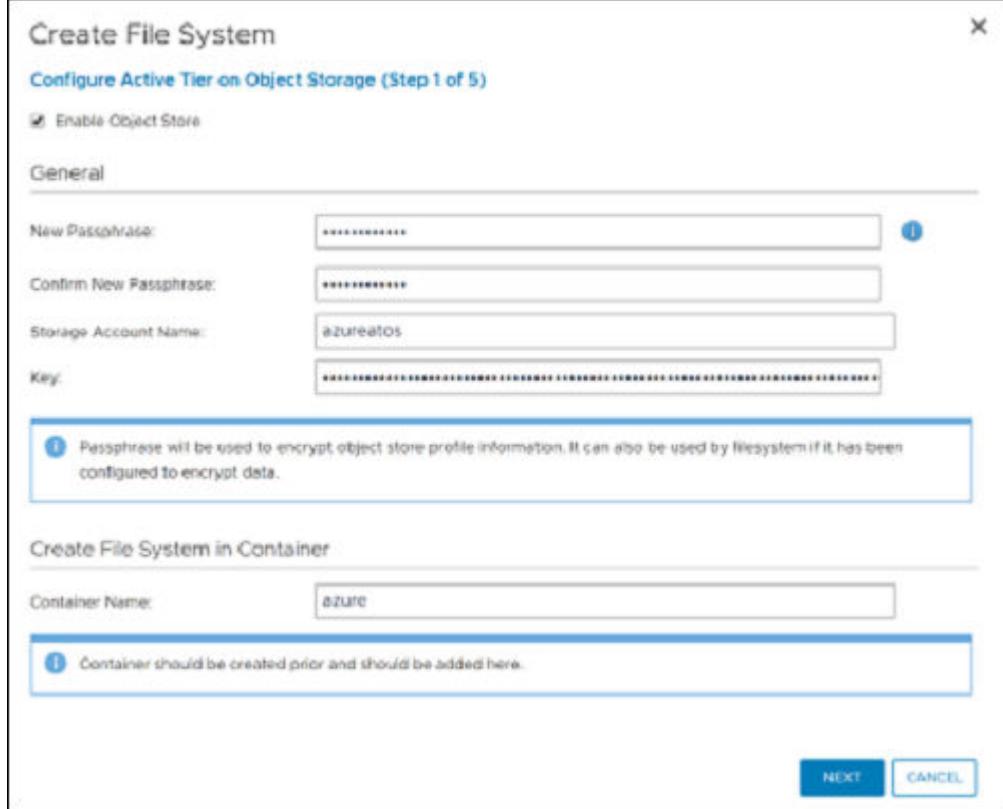
Use these steps to configure object store. Alternatively, see the next section to configure block storage.

Note

Before you begin, you will need the container name that was created using the Azure portal.

1. Navigate to **DD SM Configuration Wizard** main menu.
2. Select **File System** > **Create File System**.
3. Select the **Enable Object Store** checkbox. If **Enable Object Store** is not enabled, block storage is configured (see the next section for steps).
4. Enter the following information under **General** as shown:
 - New Passphrase
 - Storage Account Name
 - Key
 - Container Name
5. Click **Next**.

Figure 15 Enable Object Store on Azure (example)



6. The **Manage CA Certificates** pane displays. Click **+ Add**, and click **Yes** to automatically import the CA Certificates.
7. Click **Next**.
8. The **Configure the Object Store Cache** pane displays. Select the appropriate devices to add the addable cache storage and the active tier cache storage.
9. Click **Next**.
10. The **File System Summary** displays. Select the **Enable the system after creation** checkbox.

11. Click **Finish** to create the File System.

Creating the File System takes several minutes to complete.

12. Verify that the File System confirmation prompts indicate these six tasks:

- Object store enabled
- Set passphrase
- Set object store profile
- Add Device dev3
- Create File System
- Enable File System

Enabling the block store feature with DD SM

Use these steps to configure block storage. Alternatively, see the previous section to configure object storage.

1. Navigate to DD SM Configuration Wizard main menu.
2. Select **File System > Create File System**.
3. Do not select the **Enable Object Store**, the checkbox is unchecked by default.
4. Click **NEXT**.
5. Add data disks as shown.

Figure 16 Configure Block Storage Cache

Create File System

Configure Object Storage Cache (Step 3 of 5)

Configure Block Storage Cache

Addable Storage

<input type="checkbox"/>	Device	Size	Type
No devices found.			

ADD TO TIER REMOVE FROM TIER

Active Tier

<input type="checkbox"/>	Device	Size	Type
<input type="checkbox"/>	dev4	500.00 GiB	SAS

Estimated Active Tier Capacity

Estimated Usable Licensed Capacity: 0.14 TiB Remaining Licensed Capacity: 0.32 TiB

Total Licensed Capacity: 0.45 TiB

BACK NEXT CANCEL

6. Select **Enable file system after creation**, and click **Finish**.

7. Click **Close**.

Configure the System for Data Access

The DD VE system provides the DD Boost protocol. You need to configure one or more protocols for data access, depending on your environment. You also need to configure the clients for accessing the DD VE with the protocol of your choice.

If you did not configure data access with the configuration wizard, use the instructions in this section.

DD Boost (DD VE includes the DD Boost for cloud or on premise)

For setting up the Data Domain DD Boost feature, see the *Data Domain Boost for OpenStorage Administration Guide* or *Data Domain Boost for Partner Integration Administration Guide* available at <https://support.emc.com>.

Application Integration

For information about how to integrate the Data Domain system with backup software, see the documentation for the applicable application at the Data Domain Integration Documentation section on the Data Domain Support web site <https://support.emc.com>.

CHAPTER 4

DD VE Administration - on Premise and in Cloud

This chapter covers the following topics:

• Served Licensing Model for DD VE	110
• Adding virtual storage	110
• Optional Additional System Configuration	112
• Extensions to DD OS for DD VE	112
• DD VE-only commands	115
• System Recovery CLI	118
• Modified DD OS Commands	120
• Performance Troubleshooting	122
• Unsupported DD OS Commands	123
• Upgrading DD VE	128
• Define the Data Domain System Information for Your Site	128
• Setting Up NTP Time Synchronization	130
• Configuration of optional software and internal licenses	131
• Migrating DD VE	131

Served Licensing Model for DD VE

DD VE licensing

DD VE 4.0 features the Served Licensing Model for DD VE which provides the solution for managing licensing for the deployment of DD VE(s). This solution is only available for virtual systems, not physical systems at this time. The sales ordering process will remain the same. Licenses are retrieved, by the customer, from the Software Licensing Central (SLC) portal. This allows you to deploy the license server software (the hardware server is not provided) by downloading this license, loading it into the license server, and configuring the DD VE to talk to the license server.

Refer to the applicable *Data Domain Operating System Release Notes* for the most up-to-date information on product features, software updates, software compatibility guides, and information about products, licensing, and service.

How to configure: the served licensing model

Table 32 Existing license server information on license server supported platforms

Environment	System
vCloud	Linux
AWS	Linux
	Windows
Azure	Linux
	Windows

Note

The CLIs eLicense checkout and elicense checkin are used to obtain licenses from the DD VE.

- If you experience an "invalid key magic" issue after a headswap, set the passphrase on the new DD VE, then perform the `headswap ddboost user revoke token-access sysadmin`.
- If you experience a certificate authentication issue after a headswap, detach and re-attach the DD from the AV-server. The AV-server will then regenerate the certificate and import it to DD.

Note

If the system has already been configured to use license server and some served licenses have already been checked out, after the system upgrade, those licenses may go into grace period and some unexpected alerts may be generated. You can run `telnet <license-server-ip> <license-server-port>` to verify the connection between the DD VE and the license server. If the connection is good, licenses become "served" again in 24 hours.

Adding virtual storage

Additional virtual storage can be added using the GUI or the CLI.

Note

DD VE does not support the use of VMware RDM volumes.

Note

It is not possible to extend a virtual disk if it has already been used by the file system. Instead, expand the storage by adding a new virtual disk.

Using the GUI

- In DDSM, click **Hardware > Storage > Configure Storage** to add additional devices to the DD VE active tier.
- Since the GUI cannot detect AWS storage, you will need to configure the storage from the Data Domain command line.

Note

If no addable devices exist, create new virtual disks as described in [Instance Upgrade and Capacity Expansion](#) on page 39.

After adding the storage, click **Data Management > File System > Expand Capacity > Next** to launch the DAT to verify the hardware supporting the virtual storage meets the requirements for DD VE.

Using the CLI

When you add a new virtual data disk to an existing DD OS file system, use the `filesys expand` command instead of the `filesys create` command. For instructions and restrictions, see [Initial Virtual Machine Configuration with the vSphere Client](#).

Disk (Spindle) Group Configuration

DD VE 3.1 supports 16 spindle-groups. We recommend that virtual disks from the same storage be configured with same spindle-group number. Virtual disks with different storage should be configured with a different spindle-group number. However, on DD VE in AWS or Azure, by default, disks are assigned with different spindle-groups. The best practice is NOT to assign spindle-group manually.

Note

DD VE 3.1 supports up to 58 virtual disks on VMware and supports up to 120 virtual disks on Hyper-V.

Load and capacity balancing in DD OS depends on correct assignment of virtual disks to disk (spindle) groups. You do this by specifying the optional `spindle-group` argument to the `storage add` command. The command syntax is:

```
# storage add devdisk-id spindle-group 1
```

For example, if three disks are configured on DD VE, `dev3` and `dev4` are from the same storage, and `dev5` is from a different storage.

```
# storage add dev3 spindle-group 1
# storage add dev4 spindle-group 1
# storage add dev5 spindle-group 2
```

Note

`storage add` command does not support multiple devices in one command line. As a workaround you may:

- Use

```
# storage add dev3,dev4,dev5
```

or

- Use

```
# storage add dev3-5
```

Optional Additional System Configuration

See the *Data Domain Operating System Initial Configuration Guide* for help performing typical but optional initial system configuration tasks. Below is a summary of the DD OS CLI commands for some common tasks.

Note

Any system command that accepts a list, such as a list of IP addresses, accepts entries separated by either commas or spaces. See the *Data Domain Operating System Command Reference Guide* for command details.

Add users to the email list that reports system problems:

```
# alerts notify-list add group-name
```

Add users to the system report email list:

```
# autosupport add {alert-summary|asup-detailed} emails email-list
```

Enable FTP or TELNET:

```
# adminaccess enable {ftp|telnet}
```

Add remote hosts to use FTP:

```
# adminaccess ftp add <host list>
```

Add a user:

```
# user add name [role {admin|user}]
```

Change a user's password:

```
# user change password username
```

To enable remote management, refer to the *Data Domain Operating System Administration Guide* for details.

To Shut Down The System:

```
# system poweroff
```

Extensions to DD OS for DD VE

Several DD OS commands are supported on the DD VE platform only. This section describes these commands.

Storage performance evaluation

Manage virtual disk performance benchmark tests.

Storage performance can be evaluated in two ways:

- With the DAT in DDSM
- With the `disk benchmark` command in the DD OS CLI

DAT

In addition to being part of the DDSM Configuration Wizard, DAT can be run as part of the process of adding new devices to expand the file system on the DD VE instance.

Note

Running DAT before expanding the file system requires the presence of at least one device to add to the active storage tier. Also, benchmark should not be run when virtual disks are absent.

In DDSM, click **Data Management > File System > Expand Capacity**.

Click **Configure** to add devices to the active tier, or click **Next** if you have already added the devices to the active tier.

Run the DAT to analyze the underlying storage performance.

disk benchmark

disk benchmark requirements

Displays the currently configured recommended performance characteristics by disk capacity.

`disk benchmark start <dev-list>`

Start a performance benchmark test on one or more data disks in series or in parallel.

- Specify `dev[3-5]+dev7+dev[10-12]` to test the specified devices in sequence, one after the other.
- Specify `dev(3-5) dev7 dev(10-12)` to test the specified devices in parallel.

Example 1

Test `dev3`. When that test finishes, test `dev4`.

```
# disk benchmark start dev[3-4]
```

Test `dev3` and `dev4` in parallel.

```
disk benchmark start dev(3-4)
```

Start two series of tests in parallel. The two series of tests are `dev3` followed by `dev4`, and `dev5` followed by `dev6`.

```
# disk benchmark start dev(3-4) dev(5-6)
```

`disk benchmark show {[[detailed] test-id] | all | requirements}`
 Print disk performance benchmark test results and report a recommended capacity for the tested configuration. With no arguments, the command prints information about the most recent test for every disk. Specify a `test-id` to see information about a single test. Specify `all` to see a table of previous and currently running tests. Specify `requirements` to see a table of performance goals.

Note

The **disk benchmark show** command shows the results of the test of storage performance on the host system to determine which DD VE capacity configurations can be supported on the host.

disk benchmark start

Conducts a test of storage performance.

disk benchmark stop

Stop all running tests.

disk benchmark watch

Display the ongoing results of all tests in progress.

disk benchmark (for SSD cache performance)

disk benchmark start <dev-list> cache <dev-list>

Start a performance benchmark test on one or more cache devices. Where **cache <dev-list>** is an optional argument.

Note

The **<dev-list>** argument for cache devices is listed in parallel. Serial testing of cache devices is restricted, and not supported.

Example 2

Cache device in parallel (supported):

```
# disk benchmark start dev5+dev6+dev7 cache dev3:dev4
```

Cache device in serial (not supported):

```
# disk benchmark start dev5+dev6+dev7 cache dev3+dev4
```

disk benchmark show requirements <cache>

Displays the currently configured recommended performance characteristics for cache capacity. Where **<cache>** is an optional argument.

perf

Collect and show DD VE performance statistics.

perf disable trace event-regexp [module {default | ddfs}]
Disable tracing of specified events.

perf enable trace event-regexp [module {default | ddfs}]
Enable tracing of the specified events.

perf start histogram [module {default | ddfs}]

Start collecting performance histograms. This command may reduce performance marginally.

perf start stats

Start printing statistics. This command may reduce performance marginally.

perf start trace [allow-wrap] [module {default | ddfs}]

Start tracing events. This command may reduce performance marginally.

perf status trace event-regexp [module {default | ddfs}]

Shows whether tracing is enabled or disabled for the specified events.

`perf stop histogram histogram-filename [module {default | ddfs}]`
Stop collecting histograms and write the collected histograms to the specified file.

`perf stop stats`
Stop printing statistics.

`perf stop trace trace-filename [module {default | ddfs}]`
Stop tracing events and write the collected traces to the specified file.

system vresource

Display details about the virtual CPU and memory resources on the DD VE host.
Display the cache tier requirements.

`system vresource show [current | requirements]`
Display details about the virtual CPU and memory resources on the DD VE host.

Example 3

Display the cache tier requirements:

```
# ddsh -s system vresource show requirements
```

Sample output:

Active Tier Capacity (TB)	Cloud Tier Capacity (TB)	Cache Tier Capacity (GB)	vCPU Count	Memory (GiB)
4	n/a	n/a	2	6
8	n/a	n/a	2	8
16	n/a	160	4	16
32	n/a	320	4	24
48	n/a	480	4	36
64	n/a	640	8	48
96	n/a	960	8	64
16	32	160	4	32
64	128	640	8	60
96	192	960	8	80

DD VE-only commands

The following commands only work on DD VE, and are not supported on physical Data Domain systems.

Table 33 DD VE-only commands

Command	Description
<code>disk benchmark [[detailed] <test-id>]</code>	Create a disk benchmark test, and specify a unique ID for the test. Use the <code>detailed</code> option to collect more advanced information.
<code>disk benchmark show all</code>	List all the disk benchmark tests stored on the system.

Table 33 DD VE-only commands (continued)

Command	Description
disk benchmark show requirements	Displays the physical storage requirements for running DD VE.
disk benchmark show requirements <cache>	Displays the currently configured recommended performance characteristics for cache capacity.
disk benchmark start <dev-list>	Start disk benchmarking tests on the specified device or group of devices. For the <dev-list> parameter: <ul style="list-style-type: none"> Specify dev[3-5]+dev7+dev[10-12] to test the specified devices in sequence, one after the other. Specify dev(3-5) dev7 dev(10-12) to test the specified devices in parallel.
disk benchmark start <dev-list> cache <dev-list>	Start a performance benchmark test on one or more cache devices.
disk benchmark stop	Stop all disk benchmarking.
disk benchmark watch	Displays information about an in-progress disk benchmarking test, including the test being run, the device being tested, and the percent complete. This command blocks the system until the test completes, or the user types Ctrl + C .
elicense checkout feature-license <feature-name-list>	Allows user to check out the features of licenses for License Server installation
elicense checkout capacity-license <feature-name> value <n> {TB GB}	Allows user to check out the capacity of licenses for License Server installation. Here is sample output: <pre>sysadmin@localhost# elic checkout capacity-license capacity value 10 TB Checking out CAPACITY license will also checkout available feature licenses. An addition 10 TB CAPACITY license will be checked out. 10 TB additional CAPACITY license has been checked out. License(s) have been checked out for REPLICATION, DDBOOST, ENCRYPTION. Total 10 TB CAPACITY license is now available on this system.</pre>
elicense checkin {<feature-name-list> all}	Allows user to check in features for licences for License Server installation

Table 33 DD VE-only commands (continued)

Command	Description
elicense license-server set server {<ipaddr> <hostname>} port <port-number>	
elicense license-server reset	Returns DD VE to factory license settings.
elicense license-server show	
net hosts add	Two DD VEs in different regions cannot resolve each other's hostname. Run this command to add a host list entry.
	Note For VNET to VNET connection between different regions in Azure, see Microsoft.com.
storage object-store enable	Enables the object-store feature for DD VE.
storage object-store profile set	Is used to setup new profile for object-store.
storage object-store profile show	This CLI lists the object-store profile information set on the DD VE.
storage object-store profile status	This CLI lists the object-store profile information set on the DD VE.
system vresource show [requirements]	Displays the file system capacity, the number of virtual CPUs, and the amount of memory assigned to the virtual machine running the DD VE instance. The requirements option displays the physical storage requirements for DD VE.
vserver config set	DD VE supports the hypervisor's functionality to collect performance statistics from the hypervisor. These performance statistics can be used to troubleshoot the DD VE performance problems. To do that, users need to specify the vServer information (hostname or IP address) and the credential information(username and password). The vServer can be a vCenter server, an ESXi host for vSphere, a Hyper-V server, or an SVCMM server for Hyper-V. Once this information is configured, DD VE will collect performance statistics from the vServer every 5 minutes.
vserver config reset	Reset the vServer credentials for DD VE to their default values.
vserver config show	Display the vServer credentials for DD VE.

System Recovery CLI

The following system recovery commands are only applicable for the DD VE platform running on object store. These CLIs include:

1. system recovery precheck from object-store
2. system recovery start from object-store
3. system recovery status

Table 34 Object Store Command Descriptions

Command	Description
system recovery precheck from object-store	<p>This command checks if system configuration satisfies the requirement of system recovery. The same check will also be run for command system recovery start from object-store</p> <ul style="list-style-type: none"> • system recovery precheck from object-store Role required: admin # system recovery precheck from object-store Recovery precheck passed. Use start command to start the recovery. • Failure Cases <ul style="list-style-type: none"> ▪ Object-store is not enabled. # system recovery precheck from object-store **** Cannot run precheck: object-store is not enabled. ▪ Profile is not configured # system recovery precheck from object-store **** Cannot run precheck: object-store profile is not configured. ▪ Object store is not configured# system recovery precheck from object-store **** Cannot run precheck: object-store is not configured. ▪ Platform configuration doesn't match the original. # system recovery precheck from object-store Precheck found the following issues: <ol style="list-style-type: none"> 1. DD VE version version does not match the original version version 2. Instance type <i>instance</i> does not match the original instance type <i>instance</i> 3. Passphrase does not match the original passphrase 4. Active tier capacity <i>n</i> GiB is smaller than the original capacity <i>m</i> GiB 5. The object-store <i>name</i> does not have valid filesystem data 6. The filesystem already exists 7. The system recovery is already in progress
system recovery start from object-store	<p>This command starts system recovery from object-store. Since precheck is run again before recovery is actually started, all failure</p>

Table 34 Object Store Command Descriptions (continued)

Command	Description
	<p>cases for system recovery start from object-store also apply for this command.</p> <ul style="list-style-type: none"> • system recovery start from object-store Role required: admin# system recovery start from object-store System recovery has started. Use status command to check the status. • Failure cases# system recovery start from object-store Precheck found the following issues: <ol style="list-style-type: none"> 1. DD VE version <i>version</i> does not match the original version <i>version</i> 2. Instance type<i>instance</i> does not match the original instance type<i>instance</i> 3. Passphrase does not match the original passphrase 4. Active tier capacity<i>n</i> GiB is smaller than the original capacity <i>m</i> GiB 5. The object-store <i>name</i> does not have valid filesystem data 6. The filesystem already exists 7. The system recovery is already in progress**** Failed to start system recovery.
system recovery status	<p>This command shows the current system recovery status.</p> <ul style="list-style-type: none"> • system recovery status Role required: anyone# system recovery status System recovery is running: stage x of 6 (<stage name>). where <stage name> := [starting attaching object-store formatting active tier restoring configurations rebooting system restoring filesystem] • Cases <ul style="list-style-type: none"> ▪ Recovery has never run# system recovery status System recovery has never run. ▪ Recovery has completed # system recovery status System recovery completed on <date time>. where <date time> format is, for example, "Tue Feb 1 15:37:32 2018". ▪ Fail to create volume # system recovery status **** System recovery did not complete: failed to format active tier. ▪ Fail to restore configurations# system recovery status **** System recovery did not complete: failed to restore system configurations from object-store.

Table 34 Object Store Command Descriptions (continued)

Command	Description
	<ul style="list-style-type: none"> Fail to restore filesystem # system recovery status **** System recovery did not complete: failed to restore the filesystem.

Modified DD OS Commands

The behavior of the following commands has been modified on the DD VE platform:

Table 35 Modified DD OS Commands

Command	Changes
alert	The tenant-unit parameter is not supported.
compression	The tenant-unit parameter is not supported.
config setup show	Arguments for configuring features not available in DD VE have been removed.
ddboost clients show active	The tenant-unit parameter is not supported.
ddboost file-replication show active	The tenant-unit parameter is not supported.
ddboost file-replication show detailed-file-history	The tenant-unit parameter is not supported.
ddboost file-replication show file-history	The tenant-unit parameter is not supported.
ddboost option reset	The fc parameter is not supported.
ddboost option show	The fc parameter is not supported.
ddboost storage-unit create	The tenant-unit parameter is not supported.
ddboost storage-unit modify	The tenant-unit parameter is not supported.
ddboost storage-unit show	The tenant-unit parameter is not supported.
ddboost streams show active	The tenant-unit parameter is not supported.
ddboost streams show history	The tenant-unit parameter is not supported.
disk rescan	The <enlclosure-ID>.<disk-ID> parameter is not supported.

Table 35 Modified DD OS Commands (continued)

Command	Changes
disk show state	DD VE system disks show the System Dev state.
disk show stats	The DD VE format for this command is disk show stats [dev <n>]
disk status	The Spare row has been removed from the output. The System row has been added.
enclosure show all	The [<enclosure>] parameter is not supported.
enclosure show controllers	The [<enclosure>] parameter is not supported.
enclosure show cpus	The [<enclosure>] parameter is not supported.
enclosure show io-cards	The [<enclosure>] parameter is not supported.
enclosure show memory	The [<enclosure>] parameter is not supported.
filesys encryption keyes delete	The [tier {active archive} archive-unit <unit-name>] parameter is not supported.
filesys encryption keys show	The [tier {active archive} archive-unit <unit-name>] parameter is not supported.
filesys fastcopy	The [retention-lock] parameter is supported with DD VE 4.0. Retention lock governance mode is supported for DD VE on premise. Retention lock compliance mode is not supported for any DD VE.
filesys show compression	The [tier {active archive} archive-unit <unit-name>] parameter is not supported.
filesys show space	The [tier {active archive} archive-unit <unit-name> arcjove-unit {all <unit-name>}] parameter is not supported.
mtree create	The tenant-unit parameter is not supported.
mtree list	The tenant-unit parameter is not supported.
mtree show compression	The tenant-unit and tenant-unit parameters are not supported.

Table 35 Modified DD OS Commands (continued)

Command	Changes
mtree show performance	The <code>tenant-unit</code> parameter is not supported.
net create interface	The <code><virtual-ifname></code> parameter is not supported.
net destroy	The <code><virtual-ifname></code> parameter is not supported.
perf	The <code>vtl</code> option is not supported on any <code>perf</code> command.
storage add	The <code>tier</code> , <code>enclosure</code> , and <code>disk</code> parameters are not supported.
storage remove	The <code>tier</code> , <code>enclosure</code> , and <code>disk</code> parameters are not supported.
storage show	The <code>archive</code> option is not supported.
system show stats	NVRAM statistics are not reported, because DD VE systems do not have physical NVRAM.
quota	The <code>tenant-unit</code> parameter is not supported.
replication	MTree replication is the only type of replication supported.
snapshot	The <code>tenant-unit</code> parameter is not supported.

Performance Troubleshooting

You can check DD VE performance statistics with the native tools available in VMware vCenter or ESXi, or Microsoft Hyper-V.

You can also use the `perf show` and `disk benchmark` commands in DD OS to monitor and benchmark performance. See [Extensions to DD OS for DD VE](#) on page 112 for information about these commands.

CPU Performance

The two key statistics for CPU performance are:

- CPU usage: CPU usage as a percentage during the interval
- CPU ready: the percentage of time that the virtual machine was ready, but could not get scheduled to run on the physical CPU. This counter might not be displayed by default.

If these counters are high, there may be a performance problem on the hypervisor host.

Memory Performance

The key statistic for memory performance is memory swapping,: the current amount of guest physical memory swapped out to the virtual machine's swap file. This value

should always be zero. If it is not, there is a memory resource contention on the hypervisor host, which is likely to have a severe impact on DD VE performance.

Virtual Disk Performance

The key statistics for virtual disk performance are:

- I/O throughput: a decrease in these values indicates a performance issue.
- I/O latency: an increase in read and write latency values indicates a performance problem.

Failed commands: an increase in the average number of outstanding read and write requests indicates a performance problem.

Note

The controls and names for these statistics and counters are quite different in some versions of the vSphere client, and between VMware and Microsoft hypervisors. See your hypervisor documentation for help displaying these statistics in your environment.

Unsupported DD OS Commands

The following DD OS commands and command options are not supported on the DD VE platform.

Table 36 Unsupported Commands and Command Options

Unsupported Command or Command Option	Notes
adminaccess https generate certificate	Deprecated. Use <code>adminaccess certificate generate</code> instead.
alerts add	Deprecated. Use <code>alerts notify-list add</code> instead.
alerts del	Deprecated. Use <code>alerts notify-list del</code> instead.
alerts notify-list option set <i>group-name</i> tenant-alert-summary {enabled disabled}	
alerts notify-list option reset <i>group-name</i> tenant-alert-summary	
alerts reset	Deprecated. Use <code>alerts notify-list reset</code> instead.
alerts show alerts-list	Deprecated. Use <code>alerts notify-list show</code> instead.
alerts test	Deprecated. Use <code>alerts notify-list test</code> instead.
archive	
authorization	
autosupport display	Deprecated. Use <code>autosupport show report</code> instead.
autosupport reset support-list	Deprecated. Use <code>autosupport reset { all alert-summary asup-detailed support-notify }</code> instead.
autosupport show support-list	Deprecated. Use <code>autosupport show { all asup-detailed alert-summary support-notify }</code> instead.

Table 36 Unsupported Commands and Command Options (continued)

Unsupported Command or Command Option	Notes
cifs set authentication nt4	Deprecated. Use <code>cifs set authentication active-directory</code> instead.
cluster	
ddboost fc	
ddboost option reset fc	
ddboost option show fc	
ddboost show image-duplication	Deprecated. Use <code>ddboost file-replication show</code> instead.
ddboost user option set <i>user</i> default-tenant-unit <i>tenant-unit</i>	
ddboost user option reset <i>user</i> [default-tenant-unit]	
disk add devdisk-id [spindle-group 1-16]	Deprecated. Use <code>storage add</code> instead.
disk add enclosure enclosure-id	Deprecated. Use <code>storage add</code> instead.
disk benchmark start	Not supported by DDVE in cloud
disk benchmark show	Not supported by DDVE in cloud
disk benchmark stop	Not supported by DDVE in cloud
disk benchmark watch	Not supported by DDVE in cloud
disk expand	Deprecated. Use <code>storage add</code> instead.
disk failenclosure-id.disk-id	
disk multipath	
disk port	
disk rescan [enclosure-id.disk-id]	
disk show detailed-raid-info	Deprecated. Use <code>disk show state</code> and <code>storage show</code> instead.
disk show failure-history	
Disk show performance	Not supported by DDVE in cloud
disk show raid-info	Deprecated. Use <code>disk show state</code> and <code>storage show</code> instead.
disk show reliability-data	
disk disk show stats	Not supported by DDVE in cloud
disk unfail	
enclosure beacon	
enclosure show all [enclosure]	This command is supported, but not with the <i>enclosure</i> argument.

Table 36 Unsupported Commands and Command Options (continued)

Unsupported Command or Command Option	Notes
enclosure show chassis	
enclosure show controllers <i>enclosure</i>	This command is supported, but not with the <i>enclosure</i> argument.
enclosure show cpus [<i>enclosure</i>]	This command is supported, but not with the <i>enclosure</i> argument.
enclosure show fans	
enclosure show io-cards [<i>enclosure</i>]	This command is supported, but not with the <i>enclosure</i> argument.
enclosure show memory [<i>enclosure</i>]	This command is supported, but not with the <i>enclosure</i> argument.
enclosure show nvram	
enclosure show powersupply	
enclosure show summary	
enclosure show temperature-sensors	
enclosure show topology	
enclosure test topology	
filesys archive	
filesys clean update-stats	Deprecated. Use <code>filesys show space</code> instead.
filesys encryption	
filesys encryption passphrase change	Deprecated. Use <code>system passphrase change</code> instead.
filesys retention-lock	Deprecated. Use <code>mtree retention-lock</code> instead.
filesys show compression tier	The <code>tier</code> option is not supported.
filesys show history	Deprecated. Use <code>filesys show compression daily</code> instead.
ha create	Not supported by DDVE in cloud
ha destroy	Not supported by DDVE in cloud
ha status	Not supported by DDVE in cloud
ha failover	Not supported by DDVE in cloud
ha online	Not supported by DDVE in cloud
ha offline	Not supported by DDVE in cloud
license	The license commands are not supported because DD VE uses new <code>elicense</code> commands.
mtree show compression <i>mtree_path tier</i>	
net aggregate	
net config <i>ifname</i> type cluster	

Table 36 Unsupported Commands and Command Options (continued)

Unsupported Command or Command Option	Notes
net create interface <i>virtual-ifname</i>	
net create interface <i>physical-ifname</i> vlan <i>vlan-id</i>	
net create virtual <i>vethid</i>	
net destroy <i>virtual-ifname</i>	
net destroy <i>vlan-ifname</i>	
net failover	
net modify <i>virtual-ifname</i> bonding {aggregate failover}	
net set portnaming	
ndmp	
ndmpd	
nfs option disable report-replica-as-writable	Deprecated. Use filesys option disable report-replica-as-writable instead.
nfs option enable report-replica-as-writable	Deprecated. Use filesys option enable report-replica-as-writable instead.
nfs option reset report-replica-as-writable	Deprecated. Use filesys option reset report-replica-as-writable instead.
nfs option show report-replica-as-writable	Deprecated. Use filesys option show report-replica-as-writable instead.
perf * module vtl	
san	
shelf migration start	Not supported by DDVE in cloud
shelf migration status	Not supported by DDVE in cloud
shelf migration suspend	Not supported by DDVE in cloud
shelf migration resume	Not supported by DDVE in cloud
shelf migration precheck	Not supported by DDVE in cloud
shelf migration option	Not supported by DDVE in cloud
shelf migration finalize	Not supported by DDVE in cloud
shelf migration show history	Not supported by DDVE in cloud
snapshot add schedule <i>name</i> [<i>days days</i>] time <i>time</i> [, <i>time...</i>] [<i>retention period</i>]	Deprecated. Use snapshot schedule create instead.
snapshot add schedule <i>name</i> [<i>days days</i>] time <i>time</i> every <i>mins</i> [<i>retention period</i>]	Deprecated. Use snapshot schedule create instead.

Table 36 Unsupported Commands and Command Options (continued)

Unsupported Command or Command Option	Notes
snapshot add schedule name [days days] time time-time [every hrs mins] [retention period]	Deprecated. Use <code>snapshot schedule create</code> instead.
snapshot del schedule {name all}	Deprecated. Use <code>snapshot schedule destroy</code> instead.
snapshot modify schedule name {[days days] time time [,time...] [retention period]}	Deprecated. Use <code>snapshot schedule modify</code> instead.
snapshot modify schedule name {[days days] time time every {mins none} [retention period]}	Deprecated. Use <code>snapshot schedule modify</code> instead.
snapshot modify schedule name {[days days] time time-time [every {hrs mins none}] [retention period]}	Deprecated. Use <code>snapshot schedule modify</code> instead.
snapshot reset schedule	Deprecated. Use <code>snapshot schedule reset</code> instead.
snapshot show schedule	Deprecated. Use <code>snapshot schedule show</code> instead.
storage add enclosure enclosure-id	
storage add disk enclosure-id.disk-id	
storage remove enclosure enclosure-id	
storage remove disk enclosure_id.disk-id	
system firmware	
system option set console	
system retention-lock	
system sanitize	
system show anaconda	
system show controller-inventory	
system show nvram	
system show nvram-detailed	
system show oemid	
system upgrade continue	
user	
user change priv	Deprecated, with no replacement.
vserver config set host	Not supported by DDVE in cloud
vserver config reset	Not supported by DDVE in cloud
vserver config show	Not supported by DDVE in cloud
vserver config perf-stats start	Not supported by DDVE in cloud
vserver config perf-stats stop	Not supported by DDVE in cloud

Table 36 Unsupported Commands and Command Options (continued)

Unsupported Command or Command Option	Notes
vserver config perf-stats status	Not supported by DDVE in cloud
vtl lunmask	Deprecated. Use vtl group instead.
vtl lunmask add	Deprecated. Use vtl group add instead.
vtl lunmask del	Deprecated.
vtl lunmask show	Deprecated. Use vtl group show instead.

Upgrading DD VE

Upgrading from DD VE 3.0 and 3.1

DD VE 4.0 uses DD OS 6.1. 2.5 Refer to the *Data Domain Operating System 6.1.2 Administration Guide* and the *Data Domain Virtual Edition Installation and Administration Guide* for additional information.

Upgrading DD VE to a higher capacity

If the higher capacity does NOT need additional resources (refer to [Table 7](#) on page 25), follow these steps.

1. Add the needed hard disks for the new capacity
2. Configure the newly added data disks using the CLI command `storage add dev tier active<device ID>`(Or, use DD SM GUI)
3. Expand the file system using the CLI command `filesys expand`

If the higher capacity will require DD VE to have higher resources, follow these steps to upgrade DD VE to a higher capacity.

1. Shutdown DD VE.
2. Switch to the appropriate instance type and refer to [DD VE capabilities table](#) for cloud providers and resource configuration sizes.
 - AWS
 - Azure
3. Add the needed hard disks for the new capacity
4. Power on the DD VE
5. Add the license for the new capacity
6. Configure the newly added data disks using the CLI command `storage add dev tier active<device ID>`
7. Expand the file system using the CLI command `filesys expand`

Define the Data Domain System Information for Your Site

An installation requires information unique to your site. Before starting the installation, provide values for the system information listed below.

Note

Data Domain recommends that you print the tables in this section and record the information. Be sure to enter the serial number correctly to avoid DD VE issues.

Table 37 System Setup Worksheet for DD VE

Information	Your Values
A fully qualified host name for the system:	
The DNS domain name:	
A default gateway IP address (if you are not using DHCP):	
DNS server IP addresses (if you are not using DHCP): <ul style="list-style-type: none"> • Primary • Secondary • Tertiary 	
If you will enable CIFS access, enter the information for your CIFS authentication method: <ol style="list-style-type: none"> 1. For Workgroup authentication: <ul style="list-style-type: none"> • Workgroup name: • Backup user name: • Password: 2. For Active Directory authentication: <ul style="list-style-type: none"> • Realm name: • Domain admin name: • Password 	
Host name from which to administer the system:	
Administrator's email address (or admin group alias):	
Mail server (SMTP) host name:	
Hypervisor server name:	
(Optional) Physical location of the hypervisor server:	
Time zone name (default is US/Pacific):	
Serial number (SN) provided to you by Data Domain:	
Virtual machine unique ID (after initial configuration, use the system show serialno command to display this ID):	

Use this table to enter Ethernet connectivity information. By default, DHCP is enabled on both ports.

Table 38 Ethernet Connectivity Worksheet

Ethernet Connectivity	Enable	Use DHCP	IP Address (if no DHCP)	Netmask (if no DHCP)
ethV0				
ethV1				
ethV2				
ethV3				
ethV4				
ethV5				
ethV6				
ethV7				

Setting Up NTP Time Synchronization

By default, NTP is disabled on the DD VE system. If you need to enable NTP on the DD VE, follow these steps:

Note

Skip this task if you are going to join the DD VE to an Active Directory domain. Because the Windows domain controller obtains the time from an external source, NTP must be configured. See the Microsoft documentation on how to configure NTP for the Windows operating system version or service pack that is running on your domain controller. After joining the domain, the system time is periodically synchronized with the domain controller time. When the host joins the Active Directory, the DD VE displays a warning if multiple time sources are in use.

Later, while performing initial configuration of the DD VE system, enable NTP by selecting the appropriate options from the configuration wizards. If you do not use the wizards to perform initial configuration, you can use the `ntp enable` command on the DD OS command line. Enabling NTP with the `ntp enable` command automatically disables synchronizing the time on the guest to the host time.

To reenable synchronizing the guest time to the host time, run the `ntp disable` command.

Note

NTP is disabled by default. The `ntp reset` command also deactivates NTP on the guest.

For more information about AWS time synchronization, see [AWS NTP Time Synchronization](#).

Configuration of optional software and internal licenses

If you need to configure optional software features, you need to install and activate those licenses before you configure those features. See [DD VE capabilities](#) on page 14 for information about features and licenses that are available to for DD VE. A separate license is required for DD Cloud Tier.

Information about installing licenses and configuring optional software can be found in the *Data Domain Administration Guide*. Refer to the applicable *Data Domain Operating System Release Notes* for the most up-to-date information on product features, software updates, software compatibility guides, and information about our products, licensing, and service. Access the latest documents at <https://support.emc.com>.

Migrating DD VE

The virtual machine running DD VE supports live migration and cold migration in VMware vCenter environments. The virtual machine running DD VE supports live migration in Hyper-V environments. DD VE supports live migration and cold migration.

Note

After changing the virtual host, verify the network adapters are connected with the correct network label, otherwise the virtual machine will not be able to acquire an IP address. The system generates a warning when a virtual machine host does not have the network label available.

Note

DD VE uses dynamic mac address on Hyper-V platform. When you perform DD VE migration on Hyper-V, the mac address may change. Use DHCP so that IP address will change. However, if you want to keep the MAC address, you can configure DD VE with static mode before migration. For additional information, see *Hyper-V and Dynamic MAC Address Regeneration* at <https://blogs.msdn.microsoft.com> and *Understanding MAC Address Behavior During Hyper-V Live Migration* at <http://www.virtualizationadmin.com>

APPENDIX A

DD VE Cloud Workflow and Recommendations

• Overview of DD VE on Block Storage	134
• Network infrastructure setup	134
• Ports for inbound traffic	136
• Ports for outbound traffic	138
• Network setup recommendations	139

Overview of DD VE on Block Storage

DD VE on block storage provides enterprise customers and service providers who are running applications in the public cloud with a dedupe data protection appliance that provides object storage efficiency and ease of management. Supported cloud platforms include Amazon Web Services (AWS) and Microsoft Azure. It is important to note that DD VE on block storage is deployed in these environments differently.

DD VE on block storage supports:

- Backup/restore using active tier's data into cloud block storage while DD VE is running in the cloud
- DD SM to configure, manage, and monitor DD VE in AWS and Azure on block storage
- DD MC for multisystem management of DD VE systems in AWS and Azure with on block storage

Configuring DD VE on block storage using the DDSM interface

DD VE can be configured as an active tier on block storage system on AWS or Azure using one of the DDSM interface option.

Configuring DD VE on block storage using the DDSM interface ([http/https](http://))

Procedure

1. The GUI configuration wizard helps go through the active tier configuration and file system creation on DD VE. Login with the account name: sysadmin and the password specified at deployment.
2. Proceed without selecting the “Enable Object Store” checkbox to configure active tier on block storage. Click “Next”.
3. Add the block storage attached to the DD VE to the active tier.
4. Review the summary and select “Submit” to create the file system and enable it.
5. The File System section under the Data Management tab has the space usage and availability details for block storage.
6. To configure or update the eLicense on the DD VE, select “Replace License”s on the Licenses page.
7. To relaunch the configuration wizard, select “Configure System” under “Maintenance”, then select “Configure System”.

Network infrastructure setup

This section describes security group restrictions for AWS.

Security groups

The security groups restrict access to an instance based on

1. Port
2. IP range
3. Security group (its own or another)

Inbound control

The security groups are stateful which means that the responses to the inbound traffic will be allowed to go out regardless of outbound rules. The following are the inbound ports that are allowed for DD VE.

Table 39 DD VE Inbound Ports

Port	Service	Description
TCP 22	SSH	Used for SSH (CLI) access and for configuring DD VE.
TCP 443	HTTPS	Used for DDSM (GUI) access and for configuring DD VE.
TCP 2049	DD Boost/NFS	Main port used by NFS - can be modified using the <code>nfs set server-port</code> command which requires SE mode.
TCP 2051	Replication/DD Boost/ Optimized Duplication	Used only if replication is configured (run <code>replication show config</code> on Data Domain system to determine). This port can be modified using <code>replication modify</code> .
TCP 3009	SMS (system management)	Used for managing a system remotely using Data Domain System Manager. This port cannot be modified. This port is used only on Data Domain systems running DD OS 4.7.x or later. This port will also need to be opened if you plan to configure replication from within the Data Domain System Manager, as the replication partner needs to be added to the Data Domain System Manager.

Depending on the protocol that is used to backup data to DD VE, additional ports will be allowed with inbound security group rules.

Note

For the complete list of inbound ports that are used by DD VE, please refer to [Inbound Ports Table](#). Depending on the protocol used the respective ports shall be allowed.

Outbound control

As stated earlier the security groups are stateful, which means that if a request is allowed to be sent out of a DD VE, its responses will be allowed regardless of inbound rules. The following are the outbound ports that shall be allowed for DD VE.

Table 40 DD VE Outbound Ports

Port	Service	Description
UDP 123	NTP	Used by the Data Domain system to synchronize to a time server.
TCP 443	HTTPS	Used for DD VE to be able to communicate with Object store (S3).
TCP 2049	DD Boost/NFS	Main port used by NFS - can be modified using the <code>nfs set server-port</code> command which requires SE mode.
TCP 2051	Replication/DD Boost/ Optimized Duplication	Used only if replication is configured (run <code>replication show config</code> on Data Domain system to determine). This port can be modified using <code>replication modify</code> .

Table 40 DD VE Outbound Ports (continued)

Port	Service	Description
TCP 3009	SMS (system management)	Used for managing a system remotely using Data Domain System Manager. This port cannot be modified. This port is used only on Data Domain systems running DD OS 4.7.x or later. This port will also need to be opened if you plan to configure replication from within the Data Domain System Manager, as the replication partner needs to be added to the Data Domain System Manager.

Depending on the other applications/services that are being used, additional ports shall be allowed.

Note

For the complete list of outbound ports that are used by DD VE, please refer to [Outbound ports table](#). Depending on the protocol used the respective ports shall be allowed.

Ports for inbound traffic

The following are the ports that are used by the Data Domain system for inbound traffic.

Table 41 Ports Used by Data Domain System for Inbound Traffic

Port	Service	Note
TCP 21	FTP	Port is used for control only if FTP is enabled (run 'adminaccess show' on the Data Domain system to determine if this is the case).
TCP 22	SSH	Port is used only if SSH is enabled (run 'adminaccess show' on the Data Domain system to determine if this is the case).
TCP 23	Telnet	Port is used only if Telnet is enabled (run 'adminaccess show' on the Data Domain system to determine if this is the case).
TCP 80	HTTP	Port is used only if HTTP is enabled (run 'adminaccess show' on the Data Domain system to determine if this is the case).
TCP 111	DDBOOST/ NFS (portmapper)	Used to assign a random port for the mountd service used by NFS and DDBOOST. Mountd service port can be statically assigned.
UDP111	DDBOOST/ NFS (portmapper)	Used to assign a random port for the mountd service used by NFS and DDBOOST. Mountd service port can be statically assigned.
UDP 123	NTP	Port is used only if NTP is enabled on the Data Domain system. Run <code>ntp status</code> to determine if this is the case.

Table 41 Ports Used by Data Domain System for Inbound Traffic (continued)

Port	Service	Note
UDP 137	CIFS (NetBIOS Name Service)	Port used by CIFS for NetBIOS name resolution.
UDP 138	CIFS (NetBIOS Datagram Service)	Port used by CIFS for NetBIOS Datagram Service.
TCP 139	CIFS (NetBIOS Session Service)	Port used by CIFS for session information.
UDP 161	SNMP (Query)	Port is used only if SNMP is enabled. Run 'snmp status' to determine if this is the case.
TCP 389	LDAP	LDAP server listens on this port for any LDAP client request. By Default it uses TCP.
TCP 443	HTTPS	Port is used only if HTTPS is enabled (run <code>adminaccess show</code> on the Data Domain system to determine if this is the case).
TCP 445	CIFS (Microsoft-DS)	Main port used by CIFS for data transfer.
TCP 2049	DD Boost / NFS	Main port used by NFS. Can be modified via the 'nfs set server-port' command. Command requires SE mode.
TCP 2051	Replication / DD Boost / Optimized Duplication	Port is used only if replication is configured on the Data Domain system. Run <code>replication show config</code> to determine if this is the case. This port can be modified via the <code>replication modify</code> command.
TCP 2052	NFS Mountd / DD BOOST / Optimized Duplication	Main port used by NFS MOUNTD
TCP 3009	SMS (System Management)	Port is used for managing a system remotely using Web Based GUI DD EM (Data Domain Enterprise Manager). This port cannot be modified. This port is only used on Data Domain systems running DD OS 4.7.x or later. This port will also need to be opened if you plan to configure replication from within the Data Domain GUI interface, as the replication partner needs to be added to the DD Enterprise Manager.
TCP 5001	iPerf	Port is default used by <code>iperf</code> . To change the port, it requires <code>-p</code> option from <code>se iperf</code> or <code>port</code> option from the <code>net iperf</code> command. The remote side must listen on the new port.
TCP 5002	Congestion-checker	Port is default used by <code>congestion-checker</code> , when it runs <code>iperf</code> . To change the port the new port needs to be specified in the <code>port</code> option of the <code>net congestion-check</code> command. The remote side must also be listen on the new port. It is available only for DD OS 5.2 and above.

Ports for outbound traffic

The following are the ports that are used by the Data Domain system for outbound traffic.

Table 42 Ports Used by Data Domain System for Outbound Traffic

Port	Service	Note
TCP 20	FTP	Port is used for data only if FTP is enabled (run <code>adminaccess show</code> on the Data Domain system to determine if this is the case).
TCP 25	SMTP	Used by the Data Domain system to send email autosupports and alerts.
UDP/TCP 53	DNS	Port is used by Data Domain system to perform DNS lookups when DNS is configured. Run <code>net show dns</code> to review DNS configuration.
TCP 80	HTTP	Used by Data Domain system for uploading log files to Data Domain Support via the <code>support upload</code> command.
UDP 123	NTP	Used by the Data Domain system to synchronize to a time server.
UDP 162	SNMP (Trap)	Used by the Data Domain system to send SNMP traps to SNMP host. Use <code>snmp show trap-hosts</code> to see destination hosts and <code>snmp status</code> to display service status.
TCP 443	HTTPS	Port is used for communicating with Object store (S3).
UDP 514	Syslog	Used by the Data Domain system to send syslog messages, if enabled. Use 'log host show' to display destination hosts and service status.
TCP 2051	Replication / OST / Optimized Duplication	Used by Data Domain system only if replication is configured. Use <code>replication show config</code> to determine if this is the case.
TCP 3009	SMS (System Management)	Port is used for managing a system remotely using Web Based GUI DD EM (Data Domain Enterprise Manager). This port cannot be modified. This port is only used on Data Domain systems running DD OS 4.7.x or later. This port will also need to be opened if you plan to configure replication from within the Data Domain GUI interface, as the replication partner needs to be added to the DD Enterprise Manager.
TCP 5001	iPerf	Port is default used by iPerf. To change the port, it requires -p option from <code>se iperf</code> or <code>port</code> option from the <code>net iperf</code> command. And the remote side must listen on the new port.

Table 42 Ports Used by Data Domain System for Outbound Traffic (continued)

Port	Service	Note
TCP 5002	Congestion-checker	Port is default used by congestion-checker, when it runs iperf. To change the port the new port needs to be specified in the port option of the <code>net congestion-check</code> command. The remote side must also be able to listen on the new port. It is available only for DD OS 5.2 and above.
TCP 27000	Avamar client communications with Avamar server	Avamar client network hosts.
TCP 27000	Avamar server communications with Replicator target server (Avamar proprietary communication)	Required if server is used as replicator source.
TCP 28001	Avamar client communications with administrator server	Avamar clients required.
TCP 28002	Administrator server communications with Avamar client	Optional for browsing clients and cancelling backups from Avamar administrator management console.
TCP 29000	Avamar client Secure Sockets Layer (SSL) communications with Avamar server	Avamar clients required.
TCP 29000	Avamar server SSL communications with Replicator target server	Required if server is replicator source.

Network setup recommendations

VPC Architecture

We recommend you use public or private subnet architecture to deploy the DD VE in private subnet. It will secure the DD VEs (VMs) with the appropriate use of various VPC components such as route tables, access control lists, security groups, etc.

Public IP address

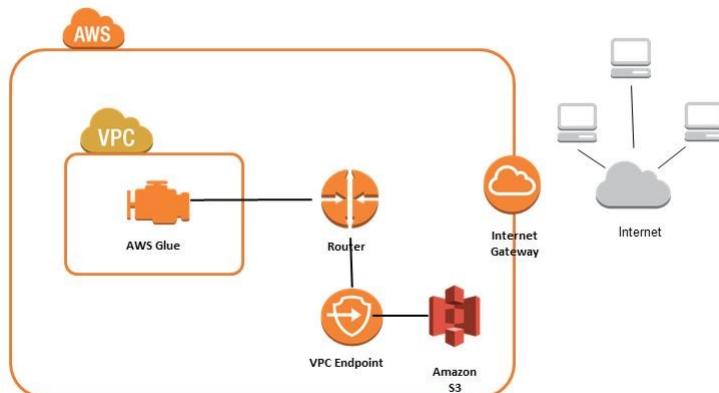
Due to security considerations and in order to protect the DD VE from potential attacks over open internet, the DD VE MUST NOT be exposed using Public IP directly over internet. It is highly recommended that you use VPN connections between different geographical regions (VPCs). For example, the replication between different VPCs, different cloud regions, cloud to on-premise and vice versa can be used via the secure VPN connection.

Object store connectivity

The DD VE object store feature needs connectivity to its object storage, such as to the S3 bucket. The object store communication is over https, so the outbound security group setting must allow communication over port 443. There are different

ways to enable DD VE connectivity to the object store. Out of the following three we recommend only the third option (Using VPC endpoint).

- Using the public IP from the public subnet: should not be used
- Using NAT (Network Address Translation): If the private subnet is configured to use NAT, then DD VE will be able to communicate to object store over NAT.
- We strongly recommend using VPC endpoint for accessing the Amazon S3. It does not require the DD VE to have a public IP address to communicate to S3, it uses the private IP address instead. (In this case, an internet gateway, NAT, or virtual private gateway are not needed to access S3). This method also allows the traffic to the S3 endpoint to stay within the Amazon network and will be routed internally to S3.



Note

- Refer to [Role based access for S3 object store](#) for configuring the DD VE to access the S3 bucket securely.
- The S3 bucket that was created for DD VE use, MUST be in the same region where DD VE is running.
- For information see [Amazon AWS documentation](#).

Setting Up NTP Time Synchronization in AWS

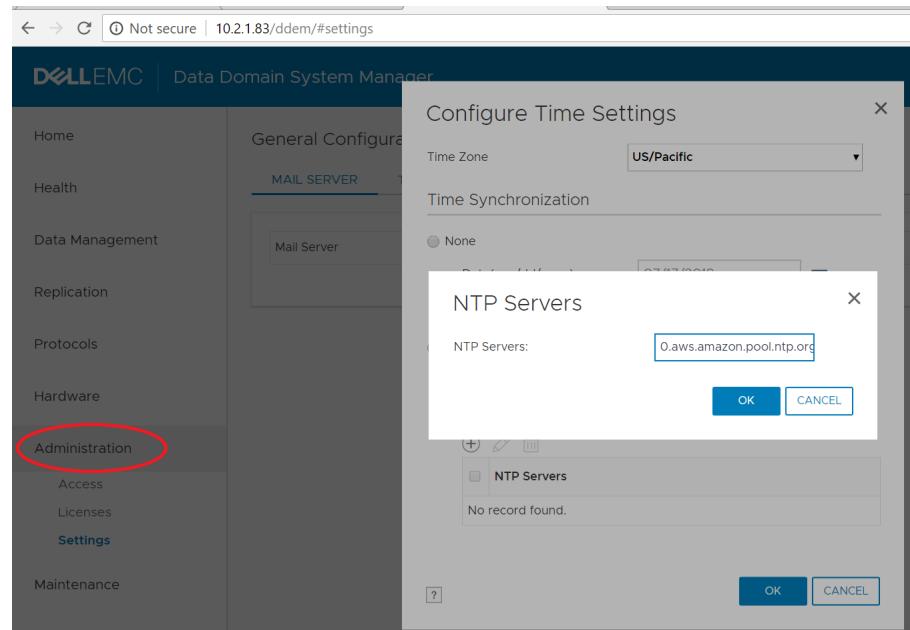
It is important for the DD VE's time to be properly synchronized. Any drift in time might impact the object store communication, system time is one component required for secure communication. Therefore, NTP must be configured for the DD VE that is running in AWS. While performing initial configuration of the DD VE system, enable NTP and configure the NTP server. If you do not use the CLI configuration wizard to perform initial configuration, you can use the `ntp enable` command on the DD OS command line. According to AWS documentation use the following details to configure NTP if you do not have your own NTP server.

```
server 0.amazon.pool.ntp.org
```

Follow the steps below to configure NTP on the DD VE (using GUI)

Procedure

1. Select the settings under the Administration tab.
2. Select “Configure Time Settings” from the drop down menu of “More Tasks”.
3. Select the “Manually Configure” option under NTP and add the NTP servers as 0.amazon.pool.ntp.org



4. Run the following commands to configure NTP on the DD VE (using CLI)
 - `ntp add timeserver 0.amazon.pool.ntp.org`
 - `ntp enable`
 - `ntp sync`

