

Data Domain Virtual Edition on Premises

Version DD VE 4.0 with DD OS 6.2.0.10

Installation and Administration Guide

302-005-338 REV 02

Copyright © 2016-2019 Dell Inc. or its subsidiaries. All rights reserved.

Published March 2019

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

	Revision history	5
	Preface	7
Chapter 1	Introduction to DD VE	9
	What is DD VE?.....	10
	DD VE capabilities	10
Chapter 2	Deploy the DD VE	13
	DD VE management.....	14
	DD VE management.....	14
	Supported virtual environments.....	14
	Provision physical storage.....	15
	Raw physical capacity needed	15
	Overview of SSD cache tier.....	16
	SSD cache tier—Supported models and their SSD requirements....	17
	Rules for adding storage to the cache tier.....	17
	SSD cache tier considerations.....	17
	SSD cache tier—Using the DAT tool to meet SSD IOPS requirements.....	18
	Managing SSD cache tier in DD VE (GUI).....	19
	Managing SSD cache tier in DD VE CLI.....	22
	Configuration requirements for DD Cloud Tier support	22
	Install the DD VE on Premises.....	24
	Performing the DD VE Download	24
	Performing the DD VE Installation	25
	Initial virtual machine configuration.....	34
	Upgrade DD OS.....	36
	Powering on the virtual machine	36
	Define the Data Domain System Information for Your Site	37
	Overview of Deploying the DD VE on KVM Hypervisor.....	38
	Configuration of other resources	39
	Deploying the DD VE on the KVM Hypervisor.....	40
	Deploying the DD VE on KVM with a reference script.....	40
	Deploying the DD VE on KVM using the Virtual Machine Manager....	42
	Deploying the DD VE on KVM Hypervisor Best Practices.....	42
Chapter 3	DD VE Initial Configuration	45
	Initial System Configuration	46
	DD VE storage guidelines.....	46
	Configuring DD VE in Data Domain System Manager.....	47
	Provisioning the storage with the CLI.....	51
	Configure the System for Data Access	52
	Performance Monitoring	53

	Configuration of other resources	54
	Configuration of optional software and internal licenses.....	55
	Optional Additional System Configuration	56
Chapter 4	DD VE Administration	57
	DD VE Licensing.....	58
	Adding virtual storage	59
	Disk (Spindle Group) Configuration	60
	Optional Additional System Configuration	60
	Extensions to DD OS for DD VE.....	61
	Storage performance evaluation.....	61
	perf.....	63
	system vresource.....	63
	DD VE-only commands.....	64
	Modified DD OS commands.....	66
	Performance Troubleshooting	68
	Unsupported DD OS Commands	69
	Upgrade DD OS.....	74
	Define the Data Domain System Information for Your Site	75
	Setting Up NTP Time Synchronization	76
	Configuration of optional software and internal licenses.....	77
	Migrating DD VE.....	77
	Best Practices for DD VE on VMware and DD VE on HyperV.....	77

Revision history

Table 1 DD VE 4.0 On Premises Installation and Administration Guide revision history

Revision	Date	Description
02	March 2019	Editorial updates
01	December 2018	Initial Publication (with DD OS 6.2.0.5)

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:

- [What is DD VE?](#) 10
- [DD VE capabilities](#) 10

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.

DD VE 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 types of platforms, on premises or in the cloud. On premises, DD VE supports VMware, Hyper-V, KVM, and VxRail. In the cloud, DD VE also runs in the Amazon Web Services (AWS) (cloud and gov cloud), Azure (cloud and gov cloud), VMware Cloud on AWS cloud platforms, and Google Cloud Platform (GCP). 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 capabilities

Resource configurations depend on your DD VE configuration. For capabilities for cloud configurations within the admin guide for your specific cloud provider, see .

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

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 premises and in the cloud)

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

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 that are 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.

CHAPTER 2

Deploy the DD VE

• DD VE management.....	14
• DD VE management.....	14
• Supported virtual environments.....	14
• Provision physical storage.....	15
• Overview of SSD cache tier.....	16
• Configuration requirements for DD Cloud Tier support	22
• Install the DD VE on Premises.....	24
• Define the Data Domain System Information for Your Site	37
• Overview of Deploying the DD VE on KVM Hypervisor.....	38
• Configuration of other resources	39
• Deploying the DD VE on the KVM Hypervisor.....	40

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.

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.

Note

Ensure the disks provisioned in VMware are configured as "Independent Persistent Disks".

- 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 8 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					
	8	16	32	48	64	96
0.5	622	620	632	694	752	846
1	1116	1114	1126	1188	1245	1340
2	2101	2100	2112	2174	2232	2326
3	3072	3072	3098	3161	3218	3312
4	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	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

The minimum software version required is HyperV and KVM is DD OS 6.2.x.x.

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. In this case `force` option should be used in the `storage add tier cache` CLI command like `storage add tier cache <dev> force`.
- 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
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          IOPS              Latency (ms)
Throughput (MiB/s)
Write IOPS
-----
5.50      2322      n/a      1244
-----

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

Note

For KVM Hypervisor, instead of using the GUI, users must add devices to cache tier for DD VEs running on KVM Hypervisor by using the CLI command `storage add tier cache force dev[]`.

- **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

The screenshot shows the Dell EMC Data Domain System Manager interface. The left sidebar contains navigation links: Home, Health, Data Management, Replication, Protocols, Hardware, Administration, Maintenance (selected), System, Support, and Logs. The main content area is titled 'Configuration Wizard' and is divided into several sections:

- System Information:** A table showing system details:

Model Number	DD3300
Version	6.1.2.1-5948
System Uptime	System has
System Serial No	1234567890
Chassis Serial No	CN7792174
- Upgrade Packages Available:** A section with an 'UPLOAD UPGRADE PACKAGE' button and a table for file names. It shows 'No record found.' and 'Items Selected: 0'. Below this is a message: 'Upgrade packages may be downloaded from the Dell EMC website.'

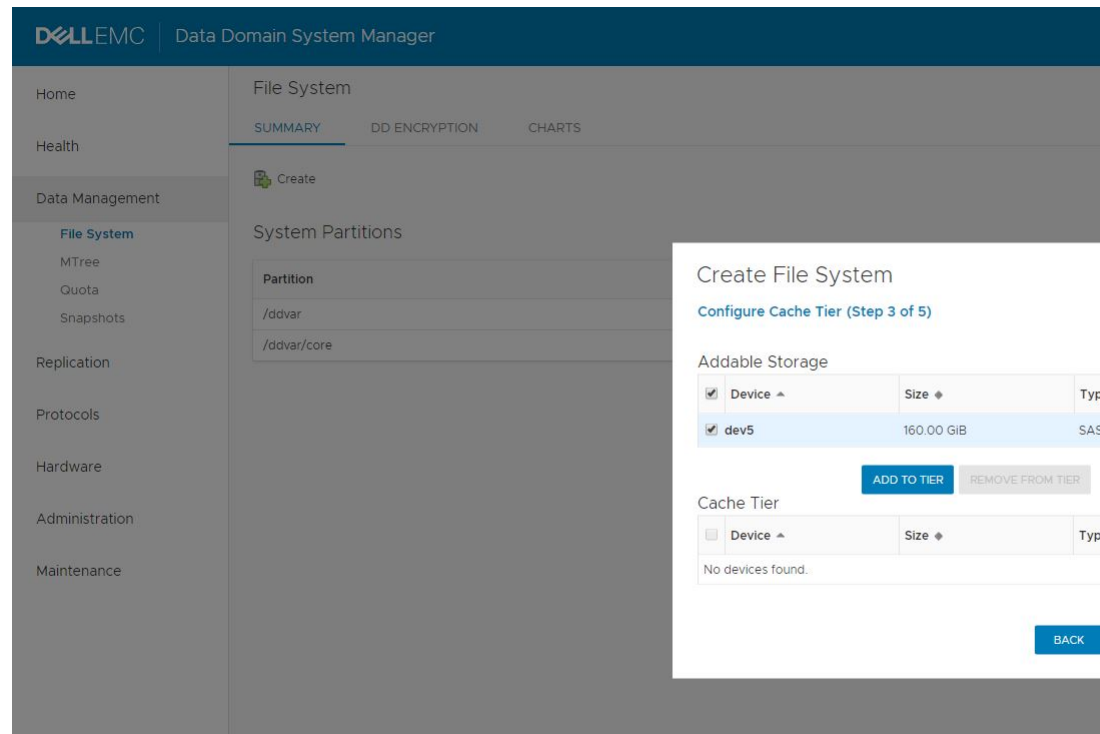
File Name
No record found.
- Upgrade History:** A table showing previous upgrades:

Version
6.1.2.1-594885
6.1.2.1-594497

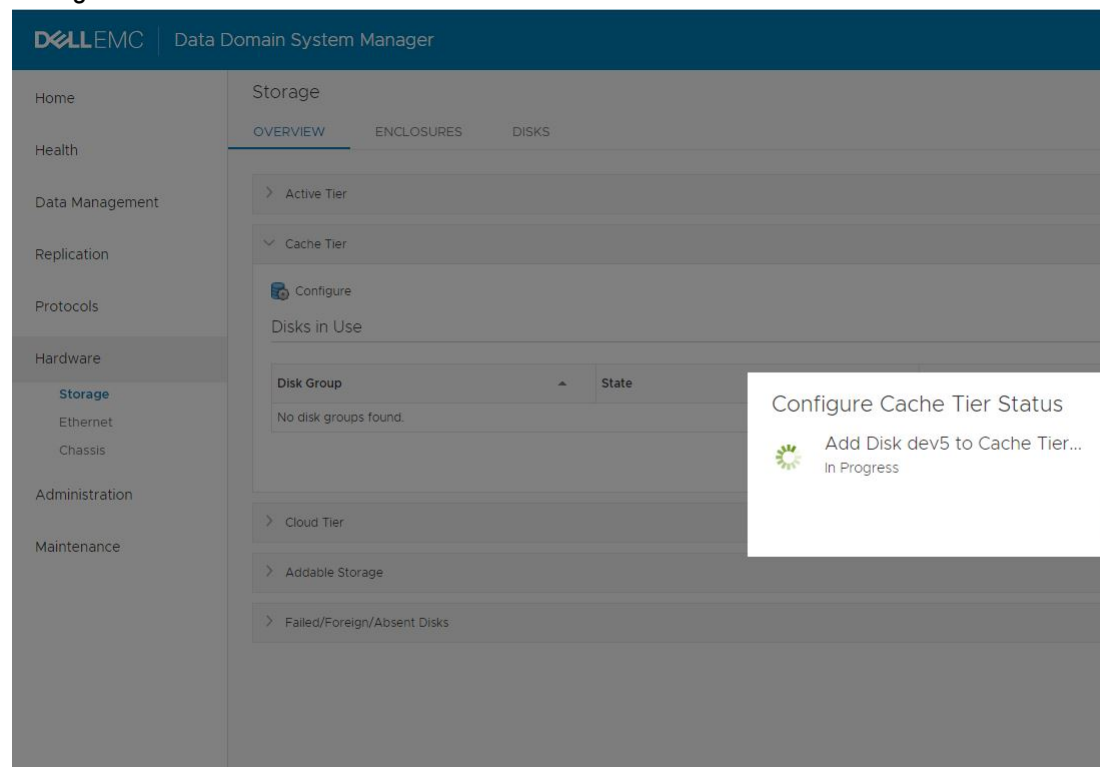
Count: 2
- Configuration Wizard Steps:** A vertical list of steps: Licenses, Network, File System (selected), Configure Active Tier, Configure Cloud Tier, Configure Cache Tier, Create File System Summary, System Settings, DD Boost Protocol, CIFS Protocol, and NFS Protocol.
- File System Configuration:** A section on the right showing configuration options:
 - Active Tier:** Size: 7.5 TiB (approximate)
 - Cache Tier:** Size: 160.0 GiB (approximate)
 - ☒ Enable file system after creation

- **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



- **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.



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
# storage add tier cache dev3	Adds storage to the cache tier. <pre># storage add tier cache dev3 Checking storage requirements...done Adding dev3 to the cache tier...done Updating system information...done dev3 successfully added to the cache tier.</pre>
# storage remove dev3	Removes storage from the cache tier. <pre># storage remove dev3 Removing dev3...done Updating system information...done dev3 successfully removed.</pre>
# storage show tier cache	Displays the storage configured for cache tier. <pre># storage show tier cache Cache tier details: Device Device Device Group Size ----- dg0 3 320.0 GiB ----- Spindle Devices Count Total Size Group ----- 1 3 1 320.0 GiB ----- Current cache tier size: 0.3 TiB Storage addable devices: Device Device Device Type Size</pre>

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 7 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

Table 7 DD VE resource reservations (continued)

Configuration	CPU	Memory (GiB)	Minimum Metadata Tier Size (GiB)	Metadata Tier Throughput (MB /S)	Metadata Tier IOPS	Metadata Tier Latency (ms)
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 8 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.

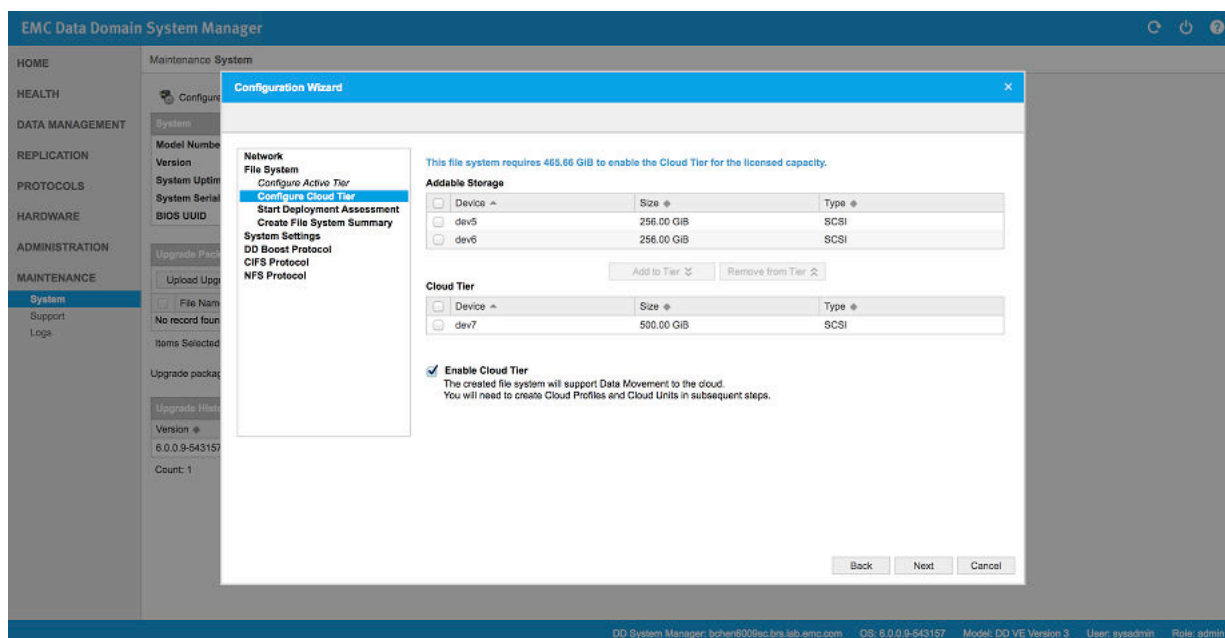
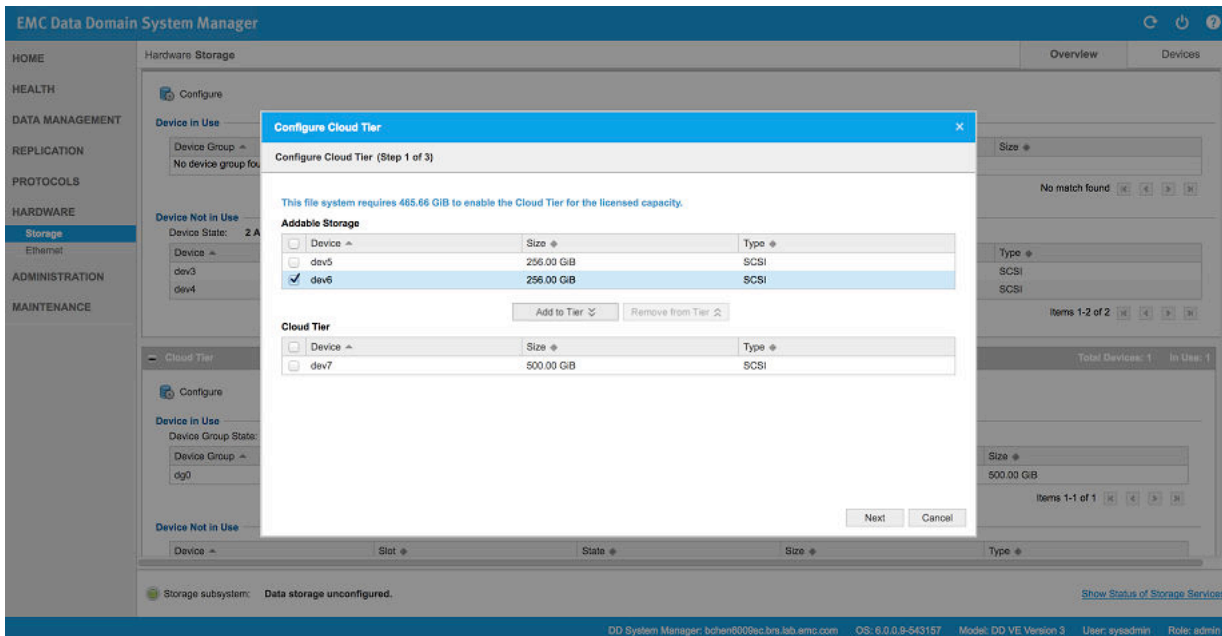
Figure 1 DD Cloud Tier Configuration Wizard for DD VE

Figure 2 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.

Install the DD VE on Premises

DD VE runs on two platforms, on premises or in the public cloud. On premises, 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 24.

Installation procedures are included for:

- [Installing on a VMware ESXi Server](#) on page 25
- [Installing through a VMware vCenter Server](#) on page 26
- [Installing on a Microsoft Windows Server 2012 R2 or Windows Server 2016 with Hyper-V Server](#) on page 27
- 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 9](#) on page 25 lists the information required to deploy the DD VE instance on a VMware ESXi server.

Table 9 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.
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 10 on page 26 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 10 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.
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.

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

Installation step	Description
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 10](#) on page 26 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 11 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.
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 8TB -VirtualMachinePath C:\DDVE -VirtualHardDiskPath C:\DDVE`

Table 12 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 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, 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 -SCVMMServer localhost -SCVMHost osdev-ucs30d -SCVMNetwork mktest-vmnet -NetworkAdapterCount 3`

Table 13 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 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.
-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

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

Parameter	Description
	about_CommonParameters, available from the Microsoft website.

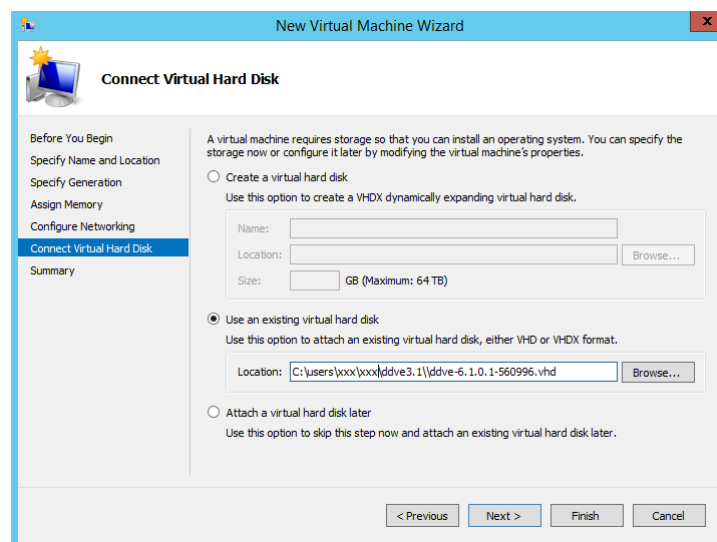
Please see ddve-installer-sc-help.text for additional information.

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.

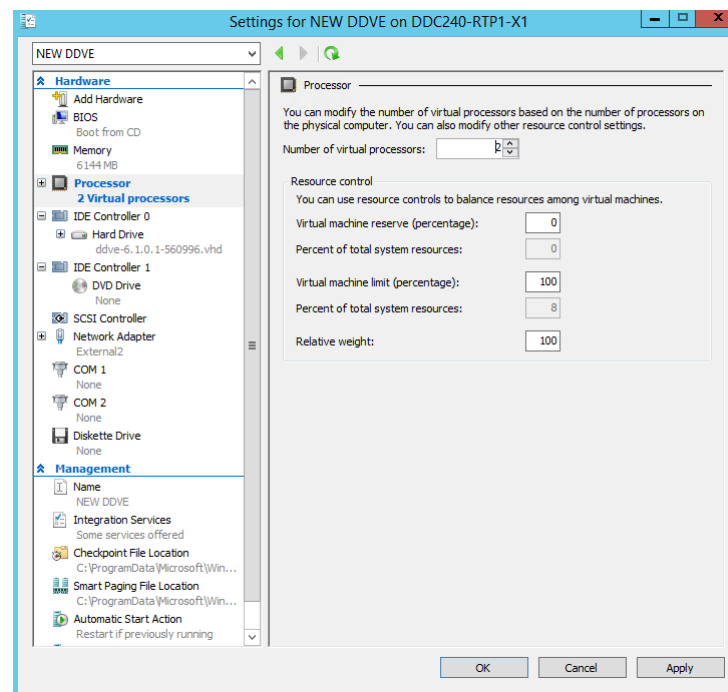
You will see the successful completion message such as, “You have successfully completed the New Virtual Machine Wizard. You are about to create the following virtual machine.”

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".



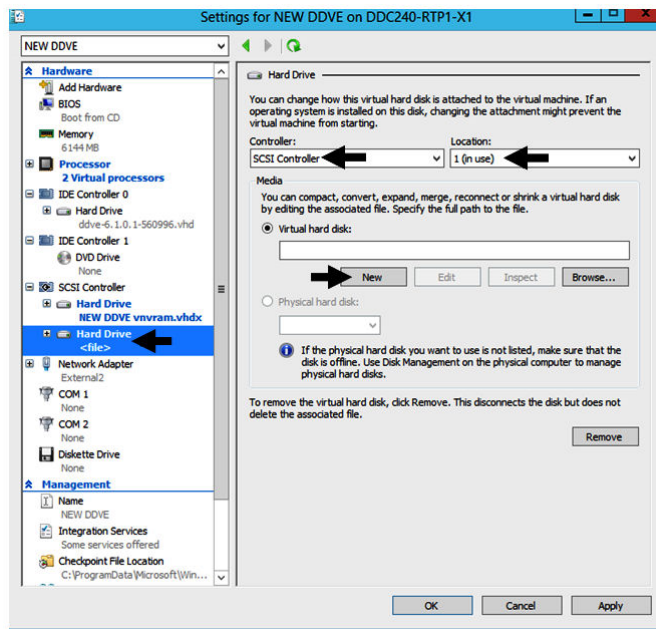
Note

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

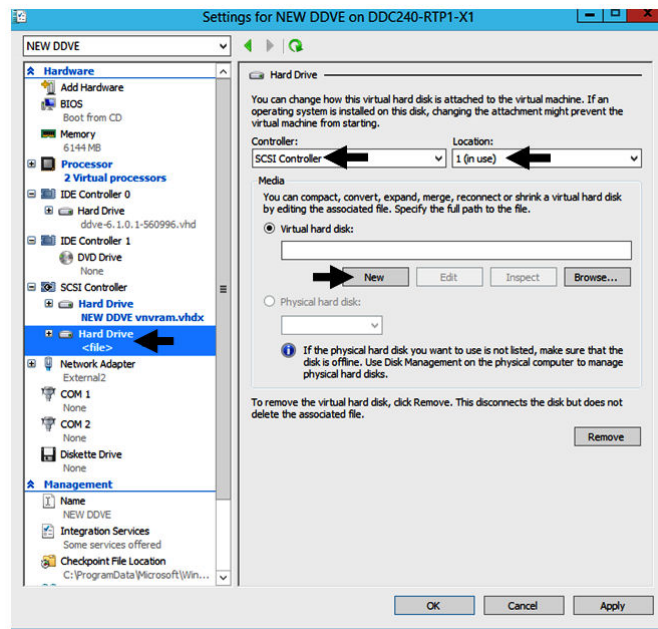


6. Attach the first 10GB virtual disk (vNVRAM) to the first SCSI controller under IDE controller 1. vNVRAM must be set to location zero, then hit "New" button to launch disk configuration wizard.




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 vram. 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 14 Initial virtual machine configuration

Hardware configuration		Storage capacity range (TB)						
		Up to 500 GB	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	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.
-

Setting Up NTP Time Synchronization

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 cloud provider 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.

.

Upgrade DD OS

The Data Domain Operating System can be upgraded using the rpm package file. For more information, refer to the Data Domain Operating System 6.2 Administration Guide.

DD VE system upgrade for higher capacity

1. Shutdown the DD VE using the command `system poweroff`
2. Upgrade the CPU and memory resources and add additional metadata disks that are required for the new configuration as per the following:

Table 15 Upgrade requirements

Instance Type (custom)	#vCPU	Memory	DD Storage Capacity	Metadata disks (num. of disks x size of each disk)
custom-4-16384	4	16 GiB	Up to 16 TB	2 x 1024 GiB
custom-8-32768	8	32 GiB	Up to 32 TB	4 x 1024 GiB
custom-16-65536	16	64 GiB	Up to 96 TB	10 x 1024 GiB

3. Power on the DD VE
4. Add the license for the new capacity
5. Configure the newly added metadata disks using the CLI command `storage add dev tier active<device ID>`
6. Expand the file system using the CLI command `filesystem 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 unique VM 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 	

Table 16 System Setup Worksheet for DD VE (continued)

Information	Your Values
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:	
Region: 1. Zone 2. VPC 3. subnet	
Serial number (SN) provided to you by Data Domain:	
Virtual machine unique ID (after initial configuration, use the <code>system show serialno</code> command to display this ID):	

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

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				

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 18 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 19 DD VE resource reservations

Resources		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			8 x vCPU	
	Memory	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						
	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 disk10 GB vNVRAM disk						

Table 19 DD VE resource reservations (continued)

Resources		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
		Note The root disk and vNVRAM disk are required to deploy the DD VE.					

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 20** Stream counts for each capacity

Configura tion	Write Stream	Read Stream	Repl Source	Repl Dest	Mixed Stream	Max Mtree
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 -hDistribution: rhel Version:7.3. The host version check done.
Basic Vaidation 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 8TB, 16TB, 32TB, 48TB, 64TB, 96TB, Cloud16TB, Cloud64TB, and Cloud96TB.
 - 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-vnvrn" 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 8TB 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 pane. 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 Initial Configuration

• Initial System Configuration	46
• Performance Monitoring	53
• Configuration of other resources	54
• Configuration of optional software and internal licenses	55
• Optional Additional System Configuration	56

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](#) 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 21 TB to TiB equivalents

Licensed capacity in TB	Licensed capacity in TiB	Maximum capacity with buffer in 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 22 Virtual disk overhead calculations

Disk	Overhead calculation	Usable capacity examples
First data disk (200 GB or more)	(Total capacity - 120 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 3 DD VE Node Locking ID

Apply Your License

Use: License File ▼

License File: ...

Node Locking ID: 5H1XYV54N7XXZVR72UYW2BL2RNYWFAX
TS2CAFKZT854A3MUK6P5ECMTDHGYDGR9
AJZPLUPVFG3UZCYG42PZZH8U45GJDUV
WK2FMWAMMW9ASRY

i To get started, apply the evaluation license that came with the download. To obtain a production license for your system, use the node locking information given here with the instructions provided in the License Activation certificate (LAC) email. The LAC email was included with your order information.

[Learn more about License Activation Certification](#) **APPLY**

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 4 Configuration Wizard - Network

Please review the configuration below - changes are marked in bold. Go back to any previous step to make further changes.

Item	Value
Host Config Mode	Using DHCP
Host Name	ddve-9632184.brs.lab.emc.com
Domain Name	brs.lab.emc.com
IPv4 Gateway	10.98.32.1
IPv6 Gateway	fe80::573ff:fea0:b60
ethV0 Enabled	yes
ethV0 Config Mode	Using DHCPv4
ethV0 IP	10.98.32.184
ethV0 Netmask	255.255.254.0
ethV1 Enabled	yes
ethV1 Config Mode	Using DHCPv4
ethV1 IP	10.98.32.245

Please click on the Submit button below to continue.


Back Submit Cancel

- File system

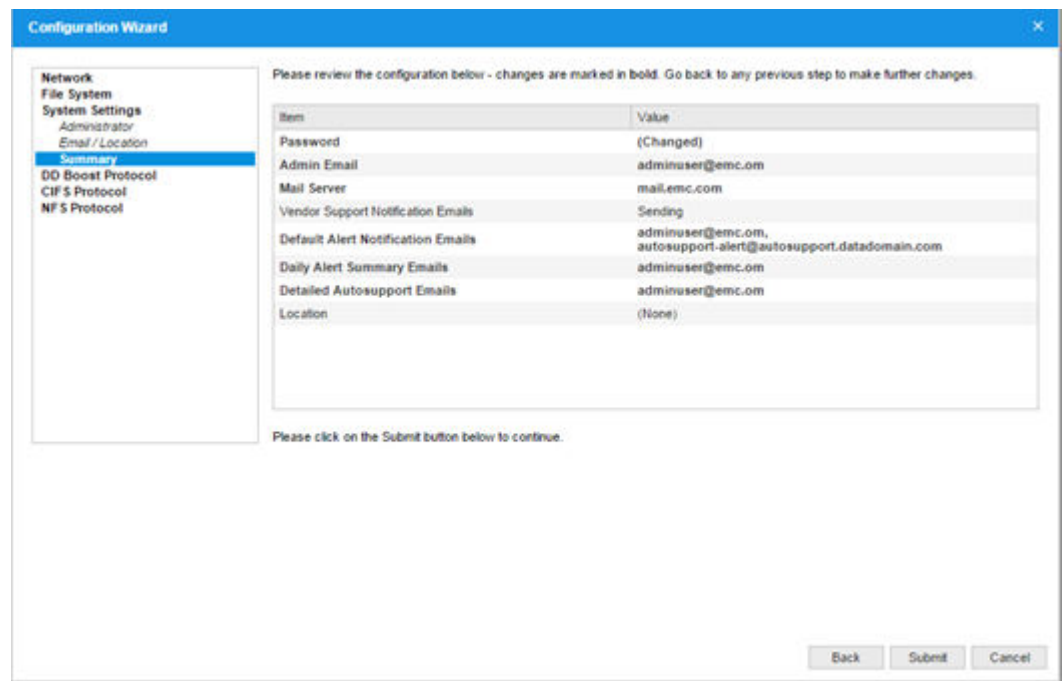
Note

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 5 Configuration Wizard - File System

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

Figure 6 Configuration Wizard - System Settings

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

Figure 7 Configuration Wizard - DD Boost Protocol

Provisioning the storage with the CLI

Before you begin

See [Disk \(Spindle Group\) Configuration](#) on page 60.

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 virtual disks (`dev1` and `dev3`) 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 premises)

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>.

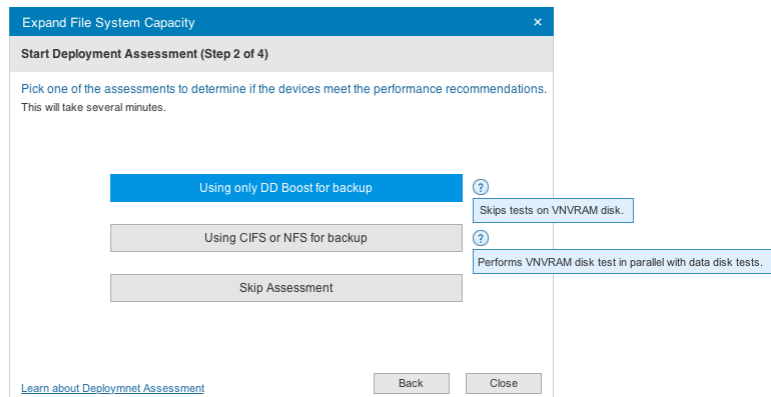
Performance Monitoring

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 23](#) on page 54 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 8 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-vnvram` 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 34 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 23 DD VE resource reservations

Resources		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	8 GB	16 GB	24 GB	36 GB	48 GB	64 GB
	Shares	Normal					

Table 23 DD VE resource reservations (continued)

Resources		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
	Limit	Unlimited					
Underlying storage requirements	Random IOPS	320	650	1280	1920	2560	3200
	Random I/O latency	14 ms					
	Sequential throughput	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	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 <hr/> Note The root disk and vNVRAM disk are required to deploy DD VE.					
	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 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. <hr/> Note 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 46 provides additional details about DD VE storage overhead requirements.					
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 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](#) 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
```


CHAPTER 4

DD VE Administration

This chapter covers the following topics:

- [DD VE Licensing](#).....58
- [Adding virtual storage](#) 59
- [Optional Additional System Configuration](#)60
- [Extensions to DD OS for DD VE](#)..... 61
- [DD VE-only commands](#).....64
- [Modified DD OS commands](#)..... 66
- [Performance Troubleshooting](#)68
- [Unsupported DD OS Commands](#) 69
- [Upgrade DD OS](#).....74
- [Define the Data Domain System Information for Your Site](#)75
- [Setting Up NTP Time Synchronization](#)76
- [Configuration of optional software and internal licenses](#)..... 77
- [Migrating DD VE](#).....77
- [Best Practices for DD VE on VMware and DD VE on HyperV](#)..... 77

DD VE Licensing

DD VE licensing may be via

- Served Licensing
- File based license

Licensing for DD VE is based on capacity, with the minimum purchased capacity being 1 TB and going up in 1 TB increments. There are no differences in the available features and functionality between any of the available resource configurations.

DD VE Served 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 licensing model is useful if you have multiple DD VE instances in your environment. 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.

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. Please save the filename with a hyphen or underscore instead of a comma.

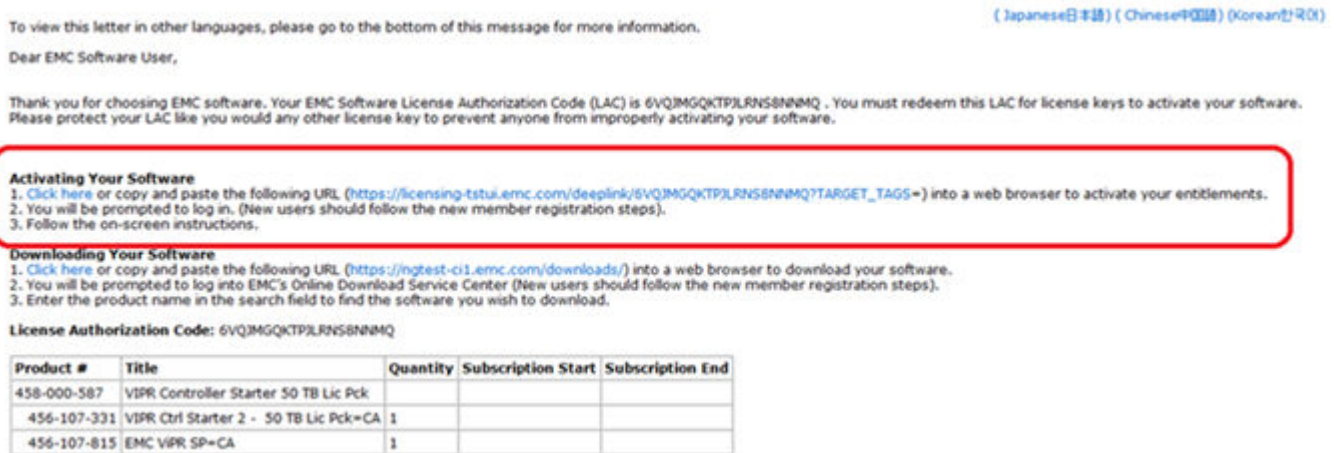
[Figure 9](#) on page 59 shows a sample email generated by the Software Licensing Central portal system.

<https://support.emc.com/servicecenter/license/> provides additional information about software licensing. If you cannot find your License Authorization e-mail, contact your account representative or support.

- 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

File based license

DD VE is licensed through the Software Licensing Central portal. When you purchase DD VE, you receive an email which email contains a license authorization code (LAC) to redeem for the DD VE software license. Follow the instructions in the email to create the license for the DD VE instance.

Figure 9 Sample LAC email

How to configure: the served licensing model

Table 24 Existing license server information on license server supported platforms

Environment	System
vCloud	Linux

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 DD VE was attached to an AV-server and you experienced 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.

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 DD SM, click **Hardware > Storage > Configure Storage** to add additional devices to the DD VE active tier.

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

Note

DD VE 3.1 and above support 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
```

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 DD SM
- With the `disk benchmark` command in the DD OS CLI

DAT

In addition to being part of the DD SM 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

Example 1 (continued)

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-regex [module {default | ddfs}]
```

Disable tracing of specified events.

```
perf enable trace event-regex [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-regex [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.

```
system vresource show [current | requirements]
```

```
# system vresource show requirements
Active Tier      Cloud
Tier
Capacity (TB)    Capacity
(TB)
-----
16               n/a   custom-4-16384 (Only block storage is
supported)
16               n/a   custom-4-16384
32               n/a   custom-8-32768
96               n/a   custom-16-65536
-----
** The maximum allowed system capacity for active tier on block
storage is 16 TB
```

DD VE-only commands

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

Table 25 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.
<code>disk benchmark show requirements</code>	Displays the physical storage requirements for running DD VE.
<code>disk benchmark show requirements <cache></code>	Displays the currently configured recommended performance characteristics for cache capacity.
<code>disk benchmark start <dev-list></code>	Start disk benchmarking tests on the specified device or group of devices. For the <code><dev-list></code> parameter: <ul style="list-style-type: none"> Specify <code>dev[3-5]+dev7+dev[10-12]</code> to test the specified devices in sequence, one after the other. Specify <code>dev(3-5) dev7 dev(10-12)</code> to test the specified devices in parallel.
<code>disk benchmark start <dev-list> cache <dev-list></code>	Start a performance benchmark test on one or more cache devices.
<code>disk benchmark stop</code>	Stop all disk benchmarking.
<code>disk benchmark watch</code>	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 <code>Ctrl + C</code> .
<code>elicense checkout feature-license <feature-name-list></code>	Allows user to check out the features of licenses for License Server installation
<code>elicense checkout capacity-license <feature-name> value <n> {TB GB}</code>	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 willl also checkout available feature licenses. An addition</pre>

Table 25 DD VE-only commands (continued)

Command	Description
	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.
<code>elicense checkin {<feature-name-list> all}</code>	Allows user to check in features for licences for License Server installation
<code>elicense license-server set server {<ipaddr> <hostname>} port <port-number></code>	
<code>elicense license-server reset</code>	Returns DD VE to factory license settings.
<code>elicense license-server show</code>	
<code>filesystem show space tier active local-metadata</code>	Displays the usage for the metadata storage .
<code>net hosts add</code>	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 .
<code>storage object-store enable</code>	Enables the object-store feature for DD VE.
<code>storage object-store disable</code>	Disables the object-store feature for DD VE.
<code>storage object-store profile set</code>	Configures the object-store access profile.
<code>storage object-store profile show</code>	Displays the object-store access profile.
<code>storage object-store profile status</code>	This CLI lists the object-store profile information set on the DD VE.
<code>system vresource show [requirements]</code>	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 <code>requirements</code> option displays the physical storage requirements for DD VE.
<code>vserver config set</code>	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

Table 25 DD VE-only commands (continued)

Command	Description
	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.
<code>vserver config reset</code>	Reset the vServer credentials for DD VE to their default values.
<code>vserver config show</code>	Display the vServer credentials for DD VE.

Modified DD OS commands

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

Table 26 Modified DD OS commands

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

Table 26 Modified DD OS commands (continued)

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

Table 26 Modified DD OS commands (continued)

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

- `perf show`
- `disk benchmark`

See [Extensions to DD OS for DD VE](#) on page 61 for information about 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 27 Unsupported Commands and Command Options

Unsupported Command or Command Option	Notes
<code>adminaccess https generate certificate</code>	Deprecated. Use <code>adminaccess certificate generate</code> instead.
<code>alerts add</code>	Deprecated. Use <code>alerts notify-list add</code> instead.
<code>alerts del</code>	Deprecated. Use <code>alerts notify-list del</code> instead.
<code>alerts notify-list option set group-name tenant-alert-summary {enabled disabled}</code>	
<code>alerts notify-list option reset group-name tenant-alert-summary</code>	
<code>alerts reset</code>	Deprecated. Use <code>alerts notify-list reset</code> instead.
<code>alerts show alerts-list</code>	Deprecated. Use <code>alerts notify-list show</code> instead.
<code>alerts test</code>	Deprecated. Use <code>alerts notify-list test</code> instead.

Table 27 Unsupported Commands and Command Options (continued)

Unsupported Command or Command Option	Notes
<code>archive</code>	
<code>authorization</code>	
<code>autosupport display</code>	Deprecated. Use <code>autosupport show report</code> instead.
<code>autosupport reset support-list</code>	Deprecated. Use <code>autosupport reset { all alert-summary asup-detailed support-notify }</code> instead.
<code>autosupport show support-list</code>	Deprecated. Use <code>autosupport show { all asup-detailed alert-summary support-notify }</code> instead.
<code>cifs set authentication nt4</code>	Deprecated. Use <code>cifs set authentication active-directory</code> instead.
<code>cluster</code>	
<code>ddboost fc</code>	
<code>ddboost option reset fc</code>	
<code>ddboost option show fc</code>	
<code>ddboost show image-duplication</code>	Deprecated. Use <code>ddboost file-replication show</code> instead.
<code>ddboost user option set user default-tenant-unit <i>tenant-unit</i></code>	
<code>ddboost user option reset user [default-tenant-unit]</code>	
<code>disk add devdisk-id [spindle-group 1-16]</code>	Deprecated. Use <code>storage add</code> instead.
<code>disk add enclosure enclosure-id</code>	Deprecated. Use <code>storage add</code> instead.
<code>disk benchmark start</code>	Not supported by DDVE in cloud
<code>disk benchmark show</code>	Not supported by DDVE in cloud
<code>disk benchmark stop</code>	Not supported by DDVE in cloud
<code>disk benchmark watch</code>	Not supported by DDVE in cloud
<code>disk expand</code>	Deprecated. Use <code>storage add</code> instead.
<code>disk failenclosure-id.disk-id</code>	
<code>disk multipath</code>	
<code>disk port</code>	
<code>disk rescan [enclosure-id.disk-id]</code>	
<code>disk show detailed-raid-info</code>	Deprecated. Use <code>disk show state</code> and <code>storage show</code> instead.
<code>disk show failure-history</code>	
<code>Disk show performance</code>	Not supported by DDVE in cloud

Table 27 Unsupported Commands and Command Options (continued)

Unsupported Command or Command Option	Notes
<code>disk show raid-info</code>	Deprecated. Use <code>disk show state</code> and <code>storage show</code> instead.
<code>disk show reliability-data</code>	
<code>disk disk show stats</code>	Not supported by DDVE in cloud
<code>disk unfail</code>	
<code>enclosure beacon</code>	
<code>enclosure show all [enclosure]</code>	This command is supported, but not with the <i>enclosure</i> argument.
<code>enclosure show chassis</code>	
<code>enclosure show controllers enclosure</code>	This command is supported, but not with the <i>enclosure</i> argument.
<code>enclosure show cpus [enclosure]</code>	This command is supported, but not with the <i>enclosure</i> argument.
<code>enclosure show fans</code>	
<code>enclosure show io-cards [enclosure]</code>	This command is supported, but not with the <i>enclosure</i> argument.
<code>enclosure show memory [enclosure]</code>	This command is supported, but not with the <i>enclosure</i> argument.
<code>enclosure show nvram</code>	
<code>enclosure show powersupply</code>	
<code>enclosure show summary</code>	
<code>enclosure show temperature-sensors</code>	
<code>enclosure show topology</code>	
<code>enclosure test topology</code>	
<code>filesystem archive</code>	
<code>filesystem clean update-stats</code>	Deprecated. Use <code>filesystem show space</code> instead.
<code>filesystem encryption</code>	
<code>filesystem encryption passphrase change</code>	Deprecated. Use <code>system passphrase change</code> instead.
<code>filesystem retention-lock</code>	Deprecated. Use <code>mtree retention-lock</code> instead.
<code>filesystem show compression tier</code>	The <code>tier</code> option is not supported.
<code>filesystem show history</code>	Deprecated. Use <code>filesystem show compression daily</code> instead.
<code>ha create</code>	Not supported by DDVE in cloud
<code>ha destroy</code>	Not supported by DDVE in cloud
<code>ha status</code>	Not supported by DDVE in cloud

Table 27 Unsupported Commands and Command Options (continued)

Unsupported Command or Command Option	Notes
<code>ha failover</code>	Not supported by DDVE in cloud
<code>ha online</code>	Not supported by DDVE in cloud
<code>ha offline</code>	Not supported by DDVE in cloud
<code>license</code>	The <code>license</code> commands are not supported because DD VE uses new <code>elicense</code> commands.
<code>mtree show compression mtree_path tier</code>	
<code>net aggregate</code>	
<code>net config ifname type cluster</code>	
<code>net create interface virtual-ifname</code>	
<code>net create interface physical-ifname vlan vlan-id</code>	
<code>net create virtual vethid</code>	
<code>net destroy virtual-ifname</code>	
<code>net destroy vlan-ifname</code>	
<code>net failover</code>	
<code>net modify virtual-ifname bonding {aggregate failover</code>	
<code>net set portnaming</code>	
<code>ndmp</code>	
<code>ndmpd</code>	
<code>nfs option disable report-replica-as-writable</code>	Deprecated. Use <code>filesystems option disable report-replica-as-writable</code> instead.
<code>nfs option enable report-replica-as-writable</code>	Deprecated. Use <code>filesystems option enable report-replica-as-writable</code> instead.
<code>nfs option reset report-replica-as-writable</code>	Deprecated. Use <code>filesystems option reset report-replica-as-writable</code> instead.
<code>nfs option show report-replica-as-writable</code>	Deprecated. Use <code>filesystems option show report-replica-as-writable</code> instead.
<code>perf * module vtl</code>	
<code>san</code>	
<code>shelf migration start</code>	Not supported by DDVE in cloud
<code>shelf migration status</code>	Not supported by DDVE in cloud
<code>shelf migration suspend</code>	Not supported by DDVE in cloud
<code>shelf migration resume</code>	Not supported by DDVE in cloud
<code>shelf migration precheck</code>	Not supported by DDVE in cloud

Table 27 Unsupported Commands and Command Options (continued)

Unsupported Command or Command Option	Notes
<code>shelf migration option</code>	Not supported by DDVE in cloud
<code>shelf migration finalize</code>	Not supported by DDVE in cloud
<code>shelf migration show history</code>	Not supported by DDVE in cloud
<code>snapshot add schedule name [days days] time time [,time...] [retention period]</code>	Deprecated. Use <code>snapshot schedule create</code> instead.
<code>snapshot add schedule name [days days] time time every mins [retention period]</code>	Deprecated. Use <code>snapshot schedule create</code> instead.
<code>snapshot add schedule name [days days] time time-time [every hrs mins] [retention period]</code>	Deprecated. Use <code>snapshot schedule create</code> instead.
<code>snapshot del schedule {name all}</code>	Deprecated. Use <code>snapshot schedule destroy</code> instead.
<code>snapshot modify schedule name {[days days] time time [,time...] [retention period]}</code>	Deprecated. Use <code>snapshot schedule modify</code> instead.
<code>snapshot modify schedule name {[days days] time time every {mins none} [retention period]}</code>	Deprecated. Use <code>snapshot schedule modify</code> instead.
<code>snapshot modify schedule name {[days days] time time-time [every {hrs mins none}] [retention period]}</code>	Deprecated. Use <code>snapshot schedule modify</code> instead.
<code>snapshot reset schedule</code>	Deprecated. Use <code>snapshot schedule reset</code> instead.
<code>snapshot show schedule</code>	Deprecated. Use <code>snapshot schedule show</code> instead.
<code>storage add enclosure enclosure-id</code>	
<code>storage add disk enclosure-id.disk-id</code>	
<code>storage remove enclosure enclosure-id</code>	
<code>storage remove disk enclosure_id.disk-id</code>	
<code>system firmware</code>	
<code>system option set console</code>	
<code>system retention-lock</code>	
<code>system sanitize</code>	
<code>system show anaconda</code>	
<code>system show controller-inventory</code>	
<code>system show nvram</code>	
<code>system show nvram-detailed</code>	
<code>system show oemid</code>	
<code>system upgrade continue</code>	
<code>user</code>	

Table 27 Unsupported Commands and Command Options (continued)

Unsupported Command or Command Option	Notes
<code>user change priv</code>	Deprecated, with no replacement.
<code>vserver config set host</code>	Not supported by DDVE in cloud
<code>vserver config reset</code>	Not supported by DDVE in cloud
<code>vserver config show</code>	Not supported by DDVE in cloud
<code>vserver config perf-stats start</code>	Not supported by DDVE in cloud
<code>vserver config perf-stats stop</code>	Not supported by DDVE in cloud
<code>vserver config perf-stats status</code>	Not supported by DDVE in cloud
<code>vtl lunmask</code>	Deprecated. Use <code>vtl group</code> instead.
<code>vtl lunmask add</code>	Deprecated. Use <code>vtl group add</code> instead.
<code>vtl lunmask del</code>	Deprecated.
<code>vtl lunmask show</code>	Deprecated. Use <code>vtl group show</code> instead.

Upgrade DD OS

The Data Domain Operating System can be upgraded using the rpm package file. For more information, refer to the Data Domain Operating System 6.2 Administration Guide.

DD VE system upgrade for higher capacity

1. Shutdown the DD VE using the command `system poweroff`
2. Upgrade the CPU and memory resources and add additional metadata disks that are required for the new configuration as per the following:

Table 28 Upgrade requirements

Instance Type (custom)	#vCPU	Memory	DD Storage Capacity	Metadata disks (num. of disks x size of each disk)
custom-4-16384	4	16 GiB	Up to 16 TB	2 x 1024 GiB
custom-8-32768	8	32 GiB	Up to 32 TB	4 x 1024 GiB
custom-16-65536	16	64 GiB	Up to 96 TB	10 x 1024 GiB

3. Power on the DD VE
4. Add the license for the new capacity
5. Configure the newly added metadata disks using the CLI command `storage add dev tier active<device ID>`
6. Expand the file system using the CLI command `filesystem 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 29 System Setup Worksheet for DD VE

Information	Your Values
A unique VM 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"> For Workgroup authentication: <ul style="list-style-type: none"> • Workgroup name: • Backup user name: • Password: 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:	
Region:	

Table 29 System Setup Worksheet for DD VE (continued)

Information	Your Values
1. Zone 2. VPC 3. subnet	
Serial number (SN) provided to you by Data Domain:	
Virtual machine unique ID (after initial configuration, use the <code>system show serialno</code> command to display this ID):	

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

Table 30 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

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 cloud provider 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.

.

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](#) 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>

Best Practices for DD VE on VMware and DD VE on HyperV

For additional information, see DD VE on *VMware Best Practice Guide* at <https://community.emc.com/docs/DOC-59452> and *DD VE on HyperV Best Practice Guide* at <https://community.emc.com/docs/DOC-59453>.

