

Composing Dell EMC PowerEdge modular infrastructure using VMware Cloud Foundation

Accelerate your cloud deployment with integrated composability

Abstract

This whitepaper provides an overview of the integration of VMware Cloud Foundation and Dell EMC PowerEdge MX7000 server with instructions on how to compose compute and storage resources in PowerEdge MX7000 to accelerate your cloud deployment.

December 2019

Revisions

Date	Description
December 2019	Initial release

Acknowledgements

This paper was produced by the following:

Author: Peter Giulietti and Shiva Katta

Support: VMware Cloud Foundation development team and Dell Open Manage development team

Other: Ramya D R and Vasanth D, IDD team

The information in this publication is provided “as is.” Dell Inc. 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 software described in this publication requires an applicable software license.

Copyright © <12/05/2019>Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
Executive summary.....	5
1 Introduction.....	6
1.1 Audience and Scope	6
2 Overview.....	7
3 Deployment requirements	8
3.1 Connectivity	8
3.2 Network services	8
3.3 Networking requirements.....	9
3.4 Management domain.....	9
3.5 Storage enclosure mode	9
4 Validated components.....	10
4.1 Hardware components	10
4.2 Software components.....	10
4.3 Firmware and Drivers	11
4.4 Using the VMware Compatibility Guide.....	11
4.4.1 Viewing BIOS and compatible ESXi versions for system and server platform	11
5 Hardware overview.....	12
5.1 Dell EMC PowerEdge MX7000	12
5.2 Dell EMC PowerEdge MX9002m management module	12
5.3 Dell EMC PowerEdge MX740c compute sled	13
5.4 Dell EMC PowerEdge MX5016s storage sled.....	13
5.5 Dell EMC PowerEdge MX5000s SAS switch	13
6 Composability workflow.....	15
6.1 Composability layer	15
6.1.1 Connection with chassis groups	16
6.1.2 Signed SSL certificate	16
6.1.3 Connecting composability layer.....	17
6.2 Viewing Available Resources	18
6.2.1 Resource Zone	19
6.2.2 Compute Systems	19
6.2.3 Storage – SSD.....	19

6.2.4 Storage – HDD	19
6.2.5 Network Interfaces.....	19
6.3 Identifying composable resources.....	20
6.3.1 Resource selection	20
6.4 Compose resources	20
6.5 Commission servers	21
6.6 Decompose resources.....	23

Executive summary

Business IT departments are challenged to deploy cloud infrastructure in an increasingly complex environment. New hardware and software provide greater capabilities, but with greater complexity. Multiple management interfaces are required to configure and deploy your compute, storage and software to provide a singular solution.

Many challenges in IT departments can be resolved through management software, but management interfaces are limited to their own scope (compute, storage, networking or software). The Hyperconverged Infrastructure (HCI) addresses virtualization, software-defined storage, and software-defined networking. This HCI must still exist on top of hardware that is generally not addressed in the HCI management and deployment process.

The VMware Cloud Foundation integration with the Dell EMC PowerEdge MX solution and the Dell EMC OpenManage Modular edition software, can combine the configuration of hardware and software into one easy-to-manage tool. By integrating the efficiencies of the Dell EMC chassis group management paradigm into VMware's streamlined HCI deployment, you can increase time savings and precision. This allows IT departments to deploy HCI infrastructure faster and more reliably.

Deployment of HCI resources is only part of the picture. In our ever-changing IT landscape, it is important to quickly and accurately disassemble and redeploy hardware and software resources. The complex composition of shared storage resources allocated to multiple compute resources requires attention when dismantling hardware.

The integration of the Dell EMC OpenManage Modular capabilities into VMware's software-defined data center (SDDC) Manager provides a single interface to administer hardware and software deployment, management, and recovery. SDDC Manager provides the link between the hardware configuration and the software-defined datacenter running on top of it. SDDC Manager can also track the modular hardware used to deploy a cluster, so that when it is time to dismantle and redeploy the cluster hardware it can be done with the confidence that other running systems will not be negatively impacted.

1 Introduction

This document provides instructions on how to compose Dell EMC PowerEdge MX 7000 components for additional workload domains. This guide is a follow on to the Dell EMC VMware Cloud Foundation for PowerEdge Deployment Guides. The deployment of a VMware Cloud Foundation Management Domain is same as the previous guides except that the firmware and software versions have been updated to the latest versions. For more information, see the *Dell EMC VMware Cloud Foundation for PowerEdge MX7000 Deployment Guide* available at www.support.dell.com.

1.1 Audience and Scope

Dell EMC recommends that you have the following prerequisite knowledge and skills:

- Knowledge on Dell EMC PowerEdge MX7000 products including the location of buttons, cables, and components in the hardware.
- Functional knowledge of the components as described in the Dell EMC owner's manuals for the products.
- Knowledge on VMware products, components, and the features of VMware Cloud Foundation.
- Best practices on managing data center infrastructure components such as server, storage, networking, and environmental consideration such as power and cooling.

This guide assumes that you have an existing and functional management domain in VMware Cloud Foundation. The management domain is installed on either Dell EMC PowerEdge modular or rack servers. For more information about deploying a Cloud Foundation Management Domain, see the *Dell EMC VMware Cloud Foundation deployment guides* available at www.support.dell.com.

The scope of this document includes only the specific hardware and software components mentioned in this guide and excludes the existing infrastructure components. Dell EMC does not take responsibility for any issues that may be caused to an existing infrastructure during deployment.

2 Overview

The ability to compose the PowerEdge MX7000 modular platform provides a mechanism to create additional workload domains using the resources required for a workload.

VMware in conjunction with Dell EMC have integrated the compute and storage composability of Dell EMC modular infrastructure into VMware Cloud Foundation SDDC Manager. With VMware Cloud Foundation version 3.9.0 SDDC Manager, you can inventory, allocate, and configure compute and storage resources for deployment as workload domains.

The key benefits of composed Dell EMC PowerEdge modular platforms are:

- Single management interface to deploy and manage both hardware and software in a Cloud Foundation environment.
- Ability to identify the compute resources depending on the workload specification.
- Ability to view the processor, memory and disk resources available in all PowerEdge MX740c and PowerEdge MX840c compute sleds.
- Ability to assign modular storage resources to compute sleds.
- Ability to track the resource availability.

3 Deployment requirements

Verify that the requirements are met before proceeding with the deployment process.

3.1 Connectivity

The SDDC Manager must be able to communicate with the PowerEdge MX9002m management modules of all PowerEdge MX7000 chassis used to create the workload domain. The workload domain generally uses a different vSAN and vMotion network, which is different from the management domain. However, if required, the management domain and workload domain can share the same management and VXLAN networks.

3.2 Network services

The network services for PowerEdge MX7000 composability must be same as the network services used for VMware Cloud Foundation deployment. The different network services are:

- Domain Name Service (DNS)
- Dynamic Host Control Protocol (DHCP)
- Network Time Service (NTP)

3.3 Networking requirements

Ensure that the network communication is established between the SDDC Controller of the management domain and the PowerEdge MX9002m management modules on the target PowerEdge MX7000 chassis. The vSAN, vMotion, and VXLAN networks that can be shared with the management domain or segregated based on the need of the intended workload domain.

In the following figure, the cabling is shown for Dell EMC PowerEdge MX9002m management modules within three Dell EMC PowerEdge MX7000. The cabling can be extended up to 20 Dell EMC PowerEdge MX7000 chassis if the first and last chassis connect to the management work.




Figure 1 Cabling for MX9002m management module

3.4 Management domain

Ensure that the VMware Cloud Foundation management domain running version 3.9.0 is used to deploy a workload domain. The management domain and selected compute sleds must meet the requirements of vSAN. For more information, see [VMware Compatibility Guide](#).

3.5 Storage enclosure mode

The storage sleds support **Drive-Assigned** and **Enclosure-Assigned** modes. The MX composability function requires configuring the **Drive-Assigned** mode in the storage sleds. These assignment modes must be configured on each storage sled in your chassis group.

Note: You must not configure the **Enclosure-Assigned** mode on the storage sleds. If the **Enclosure-Assigned** mode is configured in the storage sled, then the assignable drives will not be listed, and the drives will be unusable.

4 Validated components

4.1 Hardware components

This solution is validated using the following hardware components:

Manufacturer	Model	Description	Specifications
Dell EMC	PowerEdge MX7000	Modular chassis	NA
Dell EMC	PowerEdge MX740c	Compute sled	2x Xeon Gold processor, 256 GB RAM, Cache drives
Dell EMC	PowerEdge MX5016s	Storage sled	12 Gbps SAS capacity drives
Dell EMC	PowerEdge MX5000s	SAS Fabric IOM	12 Gbps SAS
Dell EMC	HBA330 MX	Disk controller – internal drives	NA
Dell EMC	HBA330 MMZ	Disk controller – PowerEdge MX5016s drives	NA
Dell EMC	BOSS MX	Boot / OS device	2x 256 GB M.2, RAID 01 vDisk
QLogic	QL41232HMKR	Network Interface Card	Slot Mezz 1A and 1B 2 port (each) x 25GbE

In VMware Cloud Foundation, configuring vSAN disk groups require the following protocol for the drive population:

- Configure identical drive in each target host.
- Use the same size for all cache drive, and same size for all capacity drives.
- The number of capacity drives in a host is divided by the number of cache drives. (That is, when you divide the capacity drives with the number of cache drive, the result must be a whole number.)

The management cluster can be deployed either on a rack-based cluster or modular-based cluster. In this example, the management cluster was deployed on PowerEdge R740xd rack servers.

Note: Cloud Foundation automatically configures vSAN disk groups, based on the internal drives and drives assigned from the PowerEdge MX5016s storage sled.

Note: Dell EMC supports the latest BIOS/firmware versions that are available at www.support.dell.com.

4.2 Software components

This solution is validated using the following software products. The validated software versions are certified from VMware for compatibility with VMware Cloud Foundation. To confirm, verify the compatibility in the [VMware Compatibility Guide](#). See [Using the VMware Compatibility Guide](#).

Software Product	Version	Build
Dell Customized VMware vSphere ESXi	6.7.0 Update 3	14320388
VMware Cloud Foundation	3.9.0	14866160

By using the Dell customized VMware ESXi installation media, the appropriate drivers will be installed on the hosts.

4.3 Firmware and Drivers

This solution is validated using the following product firmware and driver versions. The validated firmware and drivers are certified from VMware for compatibility with VMware Cloud Foundation. To confirm, verify the compatibility in the [VMware Compatibility Guide](#). For more information on how to use, see [Using the VMware Compatibility Guide](#).

Product	Firmware Version
PowerEdge MX7000 chassis	1.10.00
PowerEdge MX740c BIOS	2.3.10
iDRAC	3.36.36.36
QLogic 2x25 GbE and QL41232HMKR	14.10.07
HBA 330 MX	16.17.00.05
HBA 330 MMZ	16.17.00.05
PowerEdge MX5016s storage sled	2.40, A00-00
PowerEdge MX5000s storage IOM	1.0.9.6

4.4 Using the VMware Compatibility Guide

The VMware Compatibility Guide provides information about the versions of certain types of firmware and drivers, which are certified to be compatible with vSphere and vSAN. These include server platform, vSAN disk controllers, and network interface cards. For more information on other components, see www.support.dell.com.

The list of software and firmware versions are certified by VMware for compatibility with Cloud Foundation. For the list of validated software components, see [Software Component](#).

For the list of validated Firmware and Drivers, see [Firmware and Drivers](#).

For further confirmation, check the compatibility in the VMware Compatibility Guide.

The [Viewing BIOS and compatible ESXi version for system and server platform](#) section describes how to check compatibility using the VMware Compatibility Guide and to view the device specification, view driver and firmware versions.

4.4.1 Viewing BIOS and compatible ESXi versions for system and server platform

1. Open [VMware Compatibility Guide](#).
2. Expand What are you looking for and then select Systems/Servers.
3. In the **Keyword** field, enter the model of the server. For example, PowerEdge MX740c.
4. Click Update and View Results.
5. In the **Server Device and Model Information** section, select the model.
The BIOS and the compatible ESXi versions for the selected model are displayed.

Note: Only the minimum supported versions of server BIOS are displayed. For information about the latest firmware and driver, see the [VMware vSAN VCG](#).

Follow the same steps from 1 to 5 for information about drivers and firmware of the PowerEdge MX9002m Management Module, and other devices.

5 Hardware overview

This section provides information about the following hardware components that are part of the composable infrastructure. This section does not cover MX IOMs (IO modules) because this solution is not specific to any network.

- Dell EMC PowerEdge MX7000
- Dell EMC PowerEdge MX740c compute sled
- Dell EMC PowerEdge MX5016s storage sled
- Dell EMC PowerEdge MX9002m management module
- Dell EMC PowerEdge MX5000s SAS switch

For more information about the hardware components, see the *Dell EMC VMware Cloud Foundation for PowerEdge MX7000 Deployment Guide* available at www.support.dell.com.

5.1 Dell EMC PowerEdge MX7000

With kinetic architecture and Agile management, the PowerEdge MX portfolio dynamically configures compute, storage, and fabric, increases team effectiveness, and accelerates operations. The responsive design delivers the innovation and longevity customers of all sizes need for their IT and digital business transformations. Key features of PowerEdge MX7000 include:

- 7U modular enclosure with eight slots that can accommodate 2S single-width or four 4S double-width compute sleds and 12 Gbps single-width storage sleds.
- 25 Gb Ethernet, 12 Gb SAS, and 32 Gb Fiber channel I/O options.
- Three I/O network fabrics—two for general use and one for storage only; each with redundant modules.
- Multi-chassis networking up to 10 chassis.
- Single management point for compute, storage, and networking.
- High-speed technology connections, now and into the future, with no mid-plane upgrade.

Note: If there are PowerEdge MX5016s storage sleds in the lead chassis, the storage sleds in the lead chassis must be placed in different slots than storage sleds in any of the member chassis. For example, if you have four compute sleds and four storage sleds in every chassis of a chassis group; place the four storage sleds in the lead chassis in slots 1 to 4 and then place the storage sleds in slots 5 to 8 of each of the member chassis.

5.2 Dell EMC PowerEdge MX9002m management module

The Dell EMC PowerEdge MX9002m management module controls the complete chassis power, cooling, and hosts the OpenManage Enterprise-Modular (OME-M) console. Two external 1G-BaseT Ethernet ports are provided to enable management connectivity and to connect more PowerEdge MX7000 chassis into a single logical chassis. The PowerEdge MX7000 chassis supports two PowerEdge MX9002m management modules for redundancy.

The PowerEdge MX9002m management module is also the target of the PowerEdge MX7000 composability process. These management modules communicate with all hardware deploy in that chassis. It is the management module that inventories and configures the hardware components.

5.3 Dell EMC PowerEdge MX740c compute sled

Dell EMC PowerEdge MX740c is a two-socket, full-height, single-width compute sled that offers high performance and scalability. It is ideal for dense virtualization environments and serves as a foundation for collaborative workloads. The PowerEdge MX7000 chassis supports up to eight PowerEdge MX740c sleds (if no other sleds are used, such as PowerEdge MX5016s storage sleds)

- 24 DIMM slots of DDR4 memory
- Up to six SAS or SATA SSD or hard drive and NVMe PCIe SSDs
- Boot device options such as BOSS-S1
- Two PCIe mezzanine card slots for connecting to network Fabric A and B
- One PCIe mini-mezzanine card slot for connecting to storage Fabric C
- iDRAC9 with Lifecycle Controller

5.4 Dell EMC PowerEdge MX5016s storage sled

The PowerEdge MX5016s storage sled delivers scale-out, shared storage within the PowerEdge MX architecture. The PowerEdge MX5016s sled provides customizable 12 GB/s direct-attached SAS storage with up to 16 SAS hard drives or SSDs. Both the PowerEdge MX740c and the PowerEdge MX840c compute sleds can share drives with the PowerEdge MX5016s sled using the PowerEdge MX5000s SAS module. Internal server drives can be combined with up to seven PowerEdge MX5016s sleds in one chassis for extensive scalability. The PowerEdge MX7000 chassis supports up to seven PowerEdge MX5016s storage sleds.

Note: If there are MX5016s storage sleds in the lead chassis the storage sleds in the lead chassis must be placed in different slots than storage sleds in any of the member chassis. For example: If you there are four compute sleds and four storage sleds in every chassis of a chassis group; place the four storage sleds in the lead chassis in slots 1 to 4 and then place the storage sleds in slots 5 to 8 of each of the member chassis.

Note: SATA and NVMe devices are not supported in the PowerEdge MX5016s storage sled (it is SAS only). All three drive types are supported as local drives in the PowerEdge MX740c and PowerEdge MX840c compute sleds.

The MX5016s storage sleds support the Drive-Assigned and Enclosure-Assigned modes. The MX composability function requires configuring the Drive-Assigned mode in the MX5016s storage sleds. These assign modes must be configured on each MX5016s storage sled in your chassis group.

Note: Do not configure the Enclosure-Assigned mode. If the Enclosure-Assigned mode is configured in the MX5016s storage sled, then the assignable drives will not be listed, and the drives will be unusable.

5.5 Dell EMC PowerEdge MX5000s SAS switch

The Dell EMC PowerEdge MX5000s SAS module supports x4 SAS internal connections to all eight front-facing slots in the PowerEdge MX7000 chassis. The PowerEdge MX5000s uses T10 SAS zoning to provide multiple SAS zones or domains for the compute sleds. Storage management is done using the OME-Modular console.

PowerEdge MX5000s provides Fabric C SAS connectivity to each compute and one or more PowerEdge MX5016s storage sleds. Compute sleds connect to the PowerEdge MX5000s using either SAS Host Bus Adapters (HBA) or a PowerEdge RAID Controller (PERC) in the mini-mezzanine PCIe slot.

The PowerEdge MX5000s switches are deployed as redundant pairs to offer multiple SAS paths to the individual SAS disk drives. The PowerEdge MX7000 chassis supports redundant PowerEdge MX5000s in Fabric C.

6 Composability workflow

Following is the workflow to compose the compute and storage resources within a PowerEdge MX7000 chassis:

1. Connect to the composability layer of the PowerEdge MX7000 chassis group. See [Connecting Composability Layer](#).
2. View the available resources in zones. See [View Available Resources](#).
3. Identify and select the composable resources. See [Identifying composable resources](#).
4. Compose the resources. See [Compose Resources](#).
5. Commission the servers. See [Commission servers](#).

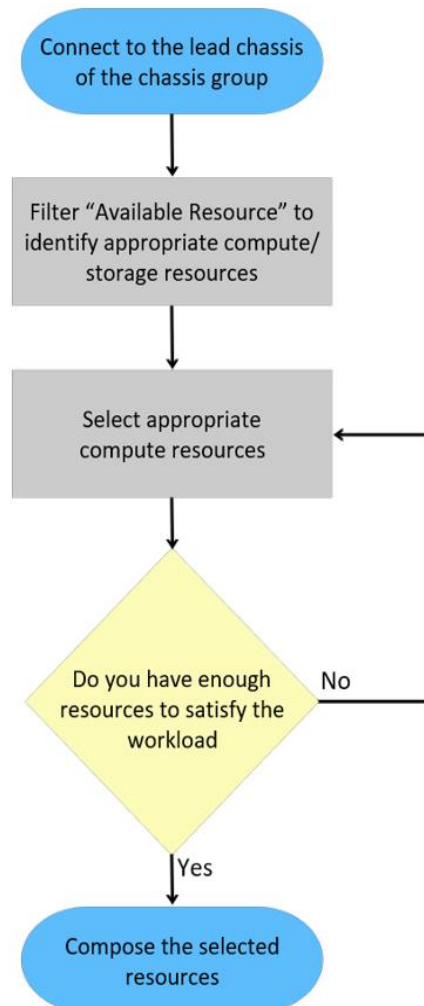


Figure 2 Composability workflow

6.1 Composability layer

The composability layer is accessed from the SDDC Manager interface. Within the OME Modular system, there are software interfaces that allow SDDC Manager to connect to the PowerEdge MX7000 chassis group, and then to the individual compute sled iDRAC interfaces. These software interfaces are referred to as the Composability Layer. The connection between SDDC Manager and the target chassis group is established and used to inventory all resources in the chassis group. The OME Modular interface reports the compute

and storage resources from the entire group to the SDDC Manager. After the information is gathered, the SDDC Manager displays the resources. The information about these resources will then be displayed in SDDC Manager.

6.1.1 Connection with chassis groups

PowerEdge MX7000 chassis are logically combined into chassis groups. Each chassis group comprises a lead chassis and member chassis. The multi-chassis chassis group concept is designed to improve the one-to-many management capabilities of a larger number of PowerEdge MX7000 chassis. SDDC Manager establishes a connection to a lead chassis and obtain details about all the resources in both lead and member chassis that are part of the multi-chassis chassis group. The resources in the entire chassis group can be composed.

Chassis groups are dynamic. Chassis can be added or removed from chassis groups without any impact on the remaining chassis. Other than the connectivity to the management modules, the chassis in a chassis group can be networked independently or combined to a group of chassis.

Note: If there are PowerEdge MX5016s storage sleds in the lead chassis, the storage sleds in the lead chassis must be placed in different slots than storage sleds in any of the member chassis. For example, if there are four compute sleds and four storage sleds in every chassis of a chassis group, place the four storage sleds in the lead chassis in slots 5 to 8, and then place the storage sleds in slots 1 to 4 on the member chassis.

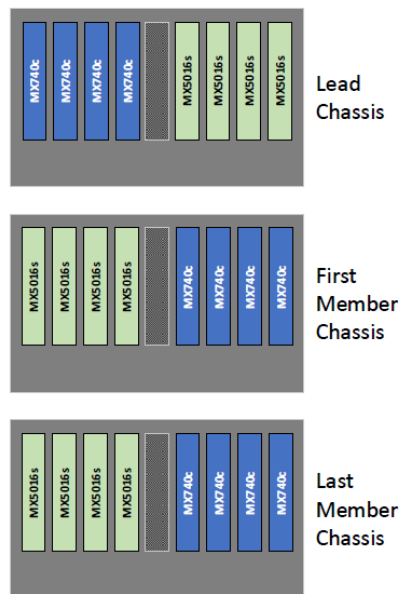


Figure 3 Lead chassis storage

For more information on creating and managing multi-chassis chassis groups, see the *Dell EMC VMware Cloud Foundation for PowerEdge MX7000 Deployment Guide* available at www.support.dell.com.

6.1.2 Signed SSL certificate

For the connection to be established, you must also generate a signed SSL certificates and apply on each PowerEdge MX7000 nodes. Because the FQDN is included as distinguished name in the certificate, you must have a matching DNS entry for each PowerEdge MX7000.

All certificate operations on the PowerEdge MX7000 are handled through the OME Modular interface.

1. Log in to the OME Modular web interface of PowerEdge MX7000.
2. From the main menu, select **Application Settings** and select **Security**.
3. From the **Security** menu, select the **Certificates** tab.
4. Click **Generate a Certificate Signing Request** and upload the signed certificate when it is available.

6.1.3 Connecting composability layer

To connect to PowerEdge MX7000, do the following:

1. Log in to the SDDC Manager web interface.
2. On the left pane, click Administration → Composable Infrastructure.

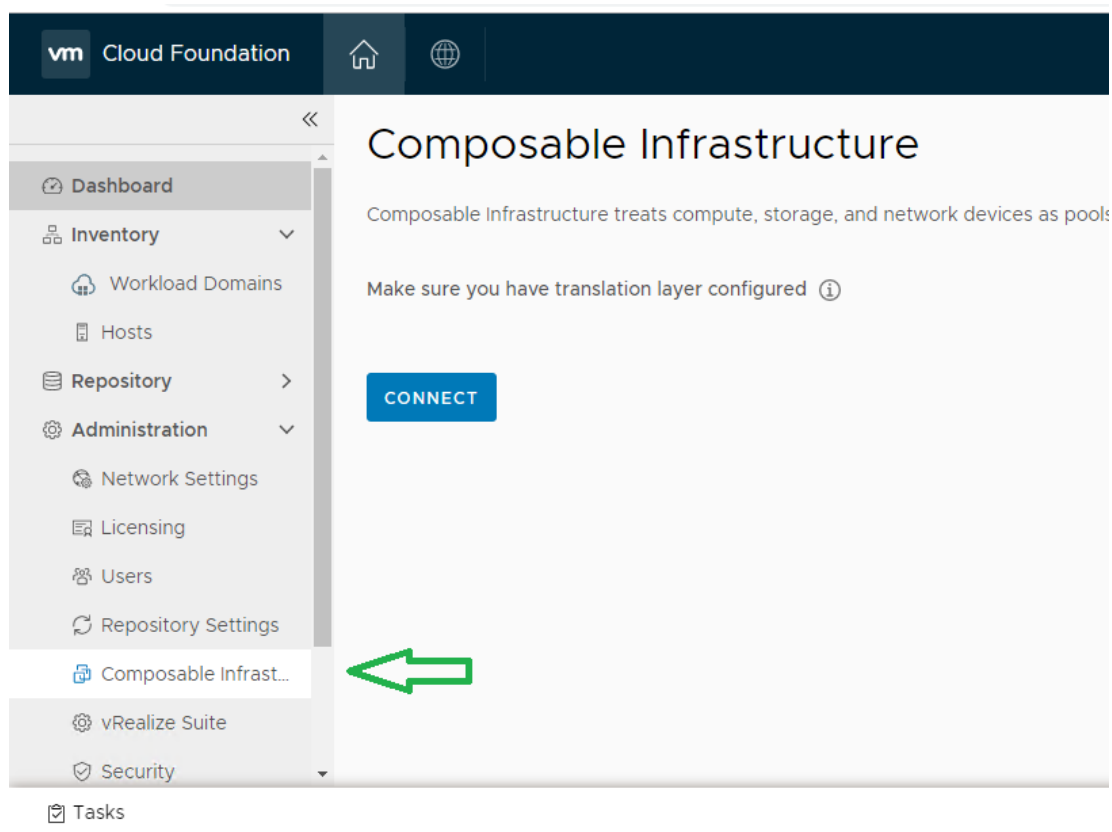


Figure 4 Composable infrastructure

3. Enter the URL of the PowerEdge MX7000 lead chassis of your multi-chassis group. The URL must be suffixed with `/redfish/v1`.
4. Enter the user name and password. It is highly recommended that you use an account other than root.

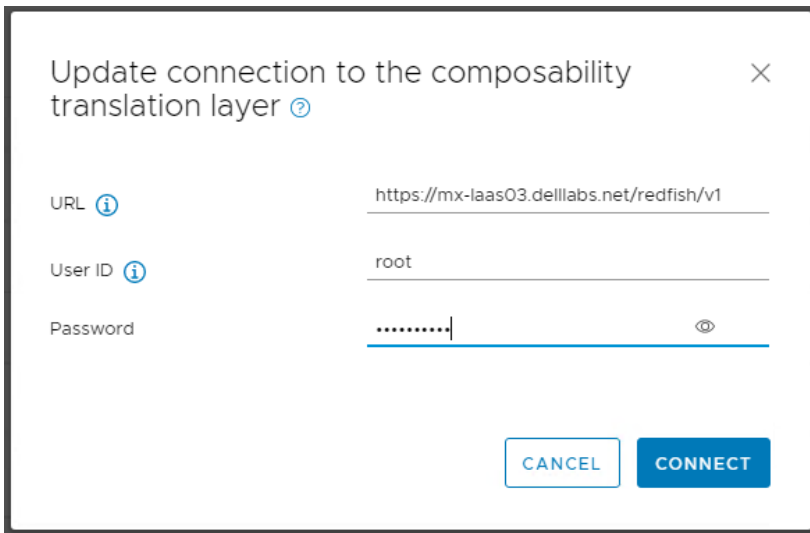


Figure 5 Composability translation layer

5. Click **Connect** to initiate the connection to the translation layer. After establishing the connection, SDDC Manager immediately begins the inventory of your MX7000 group.

6.2 Viewing Available Resources

Once SDDC Manager has inventoried the chassis or chassis group, a summary of each discovered resource is displayed in the **Available Resource** section. Each compute sled is referred to as a Resource Zone.

The chassis that is used for validating the solution contains four PowerEdge MX740c compute sleds and three PowerEdge MX5016s storage sleds. Observe that the **Available Resources** table only displays the four compute zone names for the four compute sleds. All composition uses the compute Zone as the basis of the configuration. Storage resources, disks from the MX5016s storage sleds, will be provisioned later to a compute Resource Zone.

As resource zones are composed, they will be removed from the **Available Resources** table. The **Available Resources** table only displays compute resources that are available.

Available Resources

COMPOSE

	Zone Name	Compute Systems	Storage - SSD	Storage - HDD	Network Interfaces
<input type="radio"/>	> Resource Zone - MXLAAS3 - PFT1267 - Slot 1	1 System(s) 12 Core Max, 4 GHZ Max	0 SSD(s)	20 HDD(s) 43.60 TB Total	4 Interfaces 25.6 Gbps Max
<input type="radio"/>	> Resource Zone - MXLAAS3 - PM740R2 - Slot 2	1 System(s) 14 Core Max, 4 GHZ Max	0 SSD(s)	20 HDD(s) 43.60 TB Total	4 Interfaces 25.6 Gbps Max
<input type="radio"/>	> Resource Zone - MXLAAS3 - PMX7404 - Slot 4	1 System(s) 14 Core Max, 4 GHZ Max	0 SSD(s)	20 HDD(s) 43.60 TB Total	4 Interfaces 25.6 Gbps Max
<input type="radio"/>	> Resource Zone - MXLAAS3 - PMX740C - Slot 3	1 System(s) 12 Core Max, 4 GHZ Max	0 SSD(s)	20 HDD(s) 43.60 TB Total	4 Interfaces 25.6 Gbps Max

Figure 6 Available resources

6.2.1 Resource Zone

To view the detailed information about the resources available in each Resource Zone, click the arrow next to the zone name. The available **Compute Systems**, **Storage-SSD**, **Storage-HDD**, and **Network Interfaces** are displayed. Note that the available **Network Interface** section is not shown in the following figure.

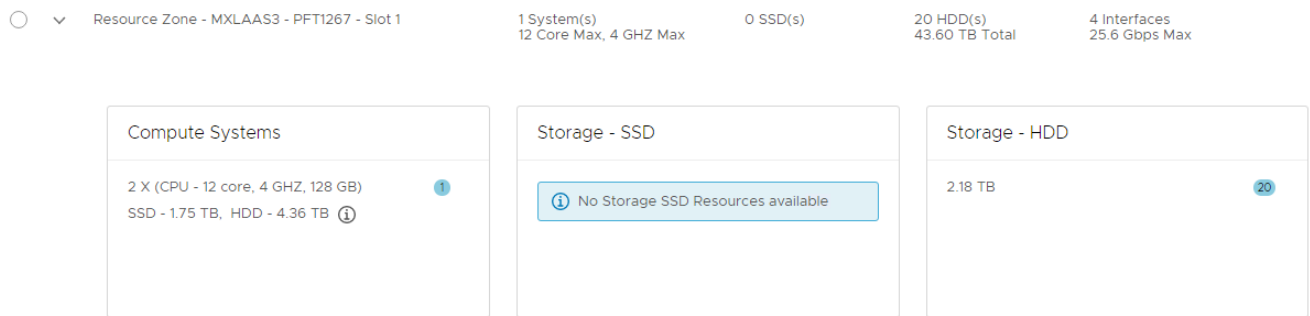


Figure 7 Resource zone

6.2.2 Compute Systems

The Compute Systems pane displays the following information:

- The number of CPUs and the number of cores in each CPU.
- The amount of RAM installed.
- The size of internal SSDs and internal HDDs. Note that only the type and size of SSDs and HDDs are displayed. This indicates that there are certain amount of 1.75 TB SSDs and 4.35 TB of HDDs memory is installed in the compute sleds (as opposed to SSDs or HDDs available on the MX5016s storage sleds).

6.2.3 Storage – SSD

The details provided in the **Storage - SSD** pane indicates the combined number of SSDs available in all MX5016s storage sleds in the chassis that the compute sled exists. In [figure 7](#), the combined storage sleds chassis contain no SSDs. There may be additional SSD storage in one or more of the storage sleds, but they have already been assigned to other compute resources. Note that a compute sled can only be assigned drives (SSD or HDD) from storage sleds in the same chassis.

6.2.4 Storage – HDD

The details provided in the **Storage - HDD** pane indicates the combined number of HDDs available in all MX5016s storage sleds in the chassis that this compute sled exists. In [figure 7](#), the combined storage sleds chassis contain 20 HDDs. There may be additional HDD storage in one or more of the storage sleds, but they have already been assigned to other compute resources. Note that a compute sled can only be assigned drives (SSD or HDD) from storage sleds in the same chassis. Only 2.18 TB HDDs are listed. This indicates that, all the available 20 HDDs are 2.18 TB drives.

6.2.5 Network Interfaces

The Network Interfaces pane displays the number and speed of the NICs physically present in the compute sled.

6.3 Identifying composable resources

Composing resources in a single chassis is a relatively simple process. Composing resources in an eight chassis multi-chassis group is a complex process. Eight chassis can contain up to a total of 64 compute or storage resources. The maximum number of chassis in a multi-chassis group is 20 that can contain up to 160 compute or storage resources. Consider that all compute sleds may or may not be configured similarly. The storage sleds can contain different quantities, sizes and types of drives, and a compute sled can only access storage resources contained in the same chassis as the target storage sled.

When all these variables are considered, you must be able to identify the compute resources based on configuration and capability. SDDC Manager can filter the **Available Resources** based on:

- Number of CPU cores
- Speed of CPU cores
- Size range of SSD drives available in a storage sled located in the same chassis
- Size range of HDD drives available in a storage sled located in the same chassis
- Network interfaces based on range of speeds

6.3.1 Resource selection

Resources can be identified by filtering the compute and available storage resources based on a capacity required for a workload. For example, if your workload requires 16 cores per CPU and minimum clock speed of 3.7 GHz, you can use the **Compute Systems** filter to view only the computes resources with this specification. Click the **Filter** button next to the **Compute Systems** column header.

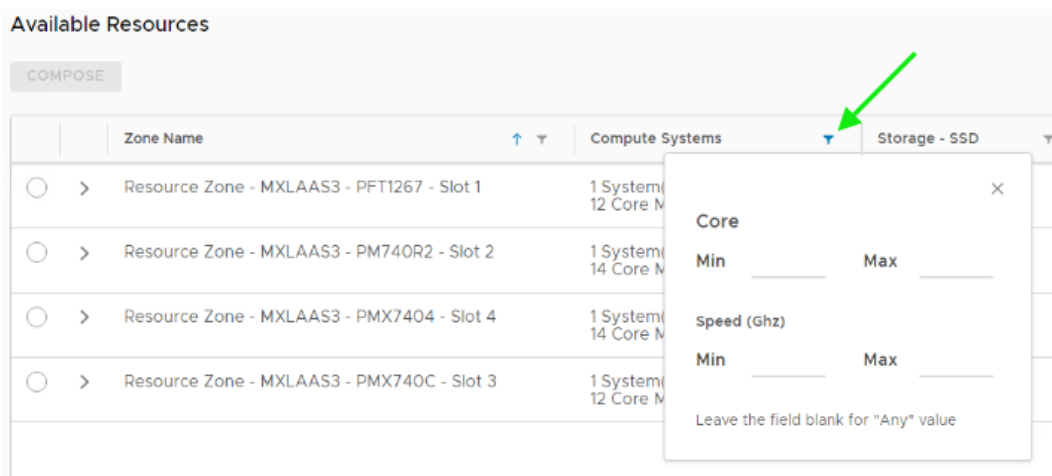


Figure 8 Available resources

Each time a filtering criterion is added, the number of available compute resources will reduce accordingly. The **Available Resources** lists the resources that meet the specified criteria. The filtering criteria can be adjusted to refine the filtering until the required resources are displayed.

6.4 Compose resources

Resources are composed by selecting and clicking **COMPOSE**. As each server is composed, drives from the available Storage-SSD or Storage-HDD resources are allocated to the server being composed. After composing, a server is moved from **Available Resources** to **Composed Servers**. Composed servers can be later commissioned into SDDC Manager to use as a Workload Domain.

The compose resources workflow is as follows:

- Identify servers that fit the intended workload.
- Select a server from the list that was identified.
- Assign storage from the PowerEdge MX5016s storage sleds as required.
- Compose the server(s).

Using the information provided in the [Resource selection](#) section, use the available filters to identify servers that meet the requirements of the intended workflow. If multiple similar servers are available, then they can be composed as a group.

To compose the resources, do the following:

1. From the filtered list of available resources, select a server or group of servers (if multiples have been found) and click **COMPOSE**.
2. In the **Allocate Resources** window, select the server. From the **Allocate** drop-down menu, select the number of drives to be assigned.

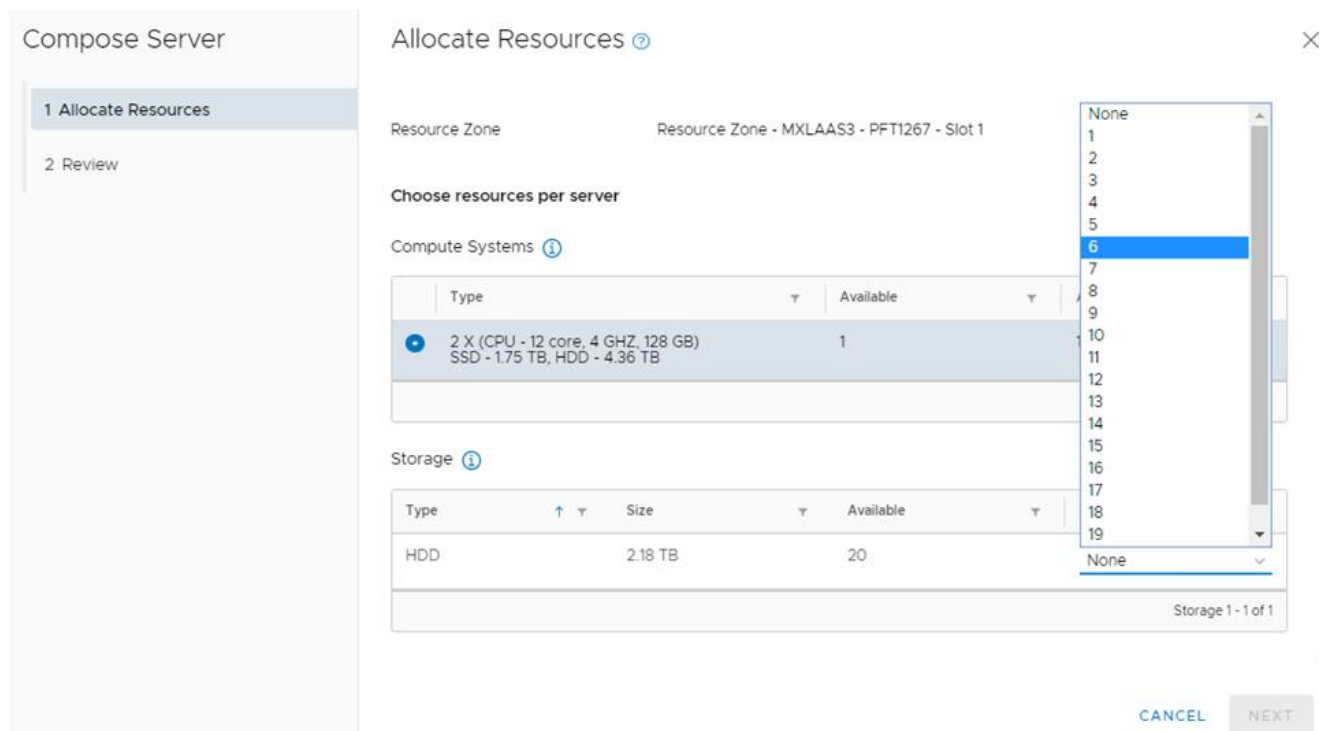


Figure 9 Allocate Resources

3. Scroll down the page and select the number of servers to be allocated (if there are more than one).
4. Click **NEXT** and review the server(s) to be composed.
5. Click **COMPOSE**.

The compose task will begin and you can monitor the task in the **Tasks** pane. After completing the task, the server is listed in the **Composed Servers** pane. Repeat this process to select required number of servers to complete the intended workload.

6.5 Commission servers

Before a Workload Domain can be created in SDDC Manager, you must have servers that are commissioned (commissioned as opposed to composed). Servers are commissioned from the SDDC Manager Dashboard.

Ensure the following prerequisites are met before you commission the servers:

- Host for vSAN workload domain must be vSAN compliant and certified as per the [VMware Hardware Compatibility Guide](#). BIOS, HBA, SSD, and HDD must be compatible as mentioned in the [VMware Compatibility Guide](#).
- Host has a standard switch with two NIC ports with a minimum 10 Gbps speed.
- Host has the drivers and firmware versions specified in the [VMware Compatibility Guide](#).
- Host has ESXi installed on it. The host must be preinstalled with supported versions (6.7.0-14320388)
- SSH and syslog are enabled on the host.
- Host is configured with DNS server for forward and reverse lookup and FQDN.
- Hostname should be same as the FQDN.
- Management IP is configured to first NIC port.
- Ensure that the host has a standard switch and the default uplinks with 10 GbE or greater speed are configured starting with traditional numbering (for example, vmnic0) and increasing sequentially.
- Host hardware health status is healthy without any errors.
- All disk partitions on HDD/SSD are deleted.
- Ensure required network pool is created and available before host commissioning.
- Ensure hosts to be used for vSAN workload domain are associated with vSAN enabled network pool.
- Ensure hosts to be used for NFS workload domain are associated with NFS enabled network pool.
- Ensure hosts to be used for VMFS on FC workload domain are associated only with NFS or VMOTION enabled network pool.

Servers can be commissioned one at a time by manually entering data or they can be imported by means of a JSON template. A sample JSON template is available in the Commission Hosts workflow in SDDC Manager. Either approach requires the following information for each host to be commissioned:

- Host FQDN
- Storage type
- Network Pool Name
- Username
- Password

After all your servers have met the requirements listed above, follow the COMMISSION HOSTS workflow from the SDDC Manager Dashboard.

1. Click **COMMISSION HOSTS**, and then select **Add new**.
2. Enter the **FQDN** of the host.
3. Select the required **Storage Type**.
4. Select the **Network Pool Name** from the list of Network Pools already created.
5. Enter the host username and password.

Figure 10 Host Addition and Validation

After commissioning the required servers to meet the requirements of your workflow, follow the normal process to deploy a Workload Domain. The Composable Infrastructure creates and configures the resources to be commissioned, and then deployed as a Workload Domain.

6.6 Decompose resources

When compute resources are no longer required, they can be returned to the list of **Available Resources** through the Decompose workflow. Before you decompose the servers, remove the workload domain and decommission the servers. After decomposing, any drive allocation to storage sled drives will be lost permanently.

1. In SDDC Manager, after removal of the Workload Domain and the decommissioning of the servers return to the **Composable Infrastructure** page. This page displays the list of composed servers at the bottom.
2. Select the check box next to the servers that you want to decompose and click **DECOMPOSE**. After selecting all the servers, click **DECOMPOSE**.

	Server Name	Status	Mac Address	IP Address
<input type="checkbox"/>	VCF-Compose-Server-1	Composed	34:80:0d:18:67:8c	Unassigned
<input checked="" type="checkbox"/>	VCF-Compose-Server-2	Composed	34:80:0d:15:e5:51	Unassigned
<input checked="" type="checkbox"/>	VCF-Compose-Server-3	Composed	34:80:0d:15:cb:05	Unassigned

Figure 11 Decompose resources

A **Decompose Servers** warning message is displayed.

3. Click **DECOMPOSE** again to decompose the selected servers. After decomposing the servers, the servers will be listed in the **Available Resources** table.

One of the reasons to decompose compute resources is to return any storage resources back to the MX5016s storage blades. Another reason to decompose compute resources is that only Available Resources can be sorted using the filters available in the **Available Resources** table.