Dell EMC VMware Cloud Foundation 4.0 for PowerEdge MX7000

Deployment Guide

January 2021
Rev. A03

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Contents

1

Audience and scope

This deployment guide includes step-by-step instructions for deployment of VMware Cloud Foundation on Dell EMC PowerEdge MX7000 modular platform. Any deviation from the listed configurations may negatively impact functionality.

This deployment guide makes certain assumptions about the prerequisite knowledge of the deployment personnel. This includes the prerequisite knowledge of:

- Dell EMC products including the location of buttons, cables, and components in the hardware
- Functional knowledge of the items in the Dell EMC owner's manuals for the products being used
- VMware products and the components or features of VMware vSphere
- Data center infrastructure best practices in the areas of server, storage, networking, and environmental considerations such as power and cooling

The scope of this document excludes existing infrastructure components outside of the specific hardware and software that is mentioned in this guide. Dell EMC takes no responsibility for any issues that may be caused to existing infrastructure during deployment.

2

Overview

Deployment of VMware Cloud Foundation on the PowerEdge MX7000 modular platform provides a hyperconverged infrastructure solution incorporating best-in-class hardware from Dell EMC with core VMware products including vSphere, vSAN, NSX, vRealize Log Insight, and SDDC Manager. Virtualization of compute, storage, and networking is delivered in a single package with VMware Cloud Foundation on PowerEdge MX7000.

Dell EMC has determined the compatibility and established certification across hardware and software. The combination of Cloud Foundation software on the Dell EMC PowerEdge MX7000 hardware that is described in this document has been validated in Dell EMC labs and certified by VMware. The PowerEdge MX7000 systems that are described within are certified as vSAN Ready Nodes, as shown in the VMware Compatibility Guide (VCG).

Some of the key benefits of the PowerEdge MX7000 modular platform include:

- Embedded Dell EMC OpenManage Enterprise Modular Edition that provides the features of OpenMange Enterprise systems management within the PowerEdge MX chassis. It includes a unified interface console for managing compute, storage, and networking
- Three I/O networking fabrics—two general purpose, and one storage-specific, each with redundant modules
- Multichassis networking up to 10 chassis
- Single management point for compute, storage, and networking

NOTE: Cloud Foundation builds a strong infrastructure foundation which you can expand with additional products. You can enable a true private cloud consumption model in your environment with optional add-on products from the VMware vRealize suite of software. For more information about these products, see www.vmware.com/products/vrealize-suite.html.

Deploying Cloud Foundation includes several steps. The following figure lists the steps and their sequence, providing a high-level visualization of the overall deployment workflow.

In the figure and throughout the deployment process, the term SDDC bring-up refers to the software-defined data center (SDDC) which is built as the result of the steps documented in this guide.

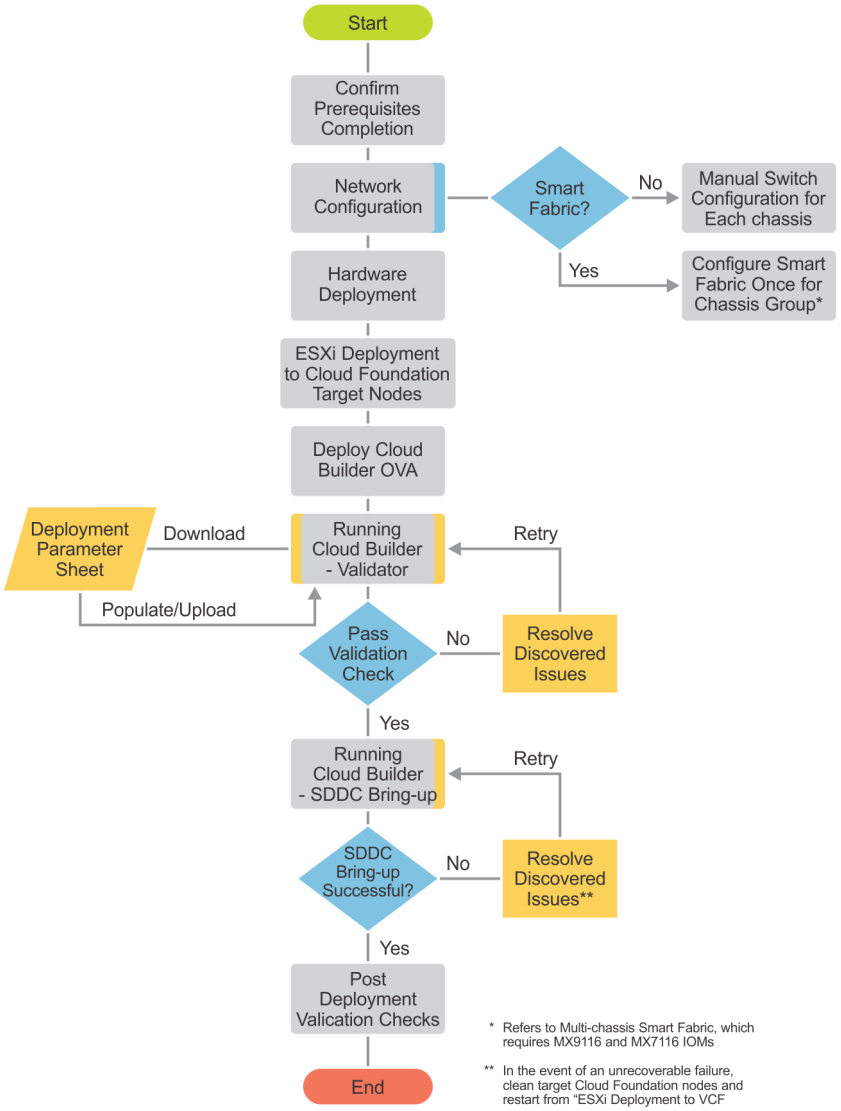

Figure 1. Cloud Foundation deployment workflow

3

Pre-deployment requirements

Topics:

Management host

The deployment of VMware Cloud Foundation is executed by a Cloud Builder VM that is deployed using an Open Virtualization Appliance (OVA). The virtual machine must be deployed on an ESXi host or cluster that is not a part of the Cloud Foundation cluster. If the management network is a private network ensure that the Cloud Builder VM and the Cloud Foundation management hosts have access to the same DNS and NTP services.

In this example, a host in an existing vSphere environment is used to run the services that are required to install Cloud Foundation. The management VLAN and VXLAN VLAN are extended to this management host. An NTP time server is installed on the management VLAN and a DHCP server on the VXLAN VLAN.

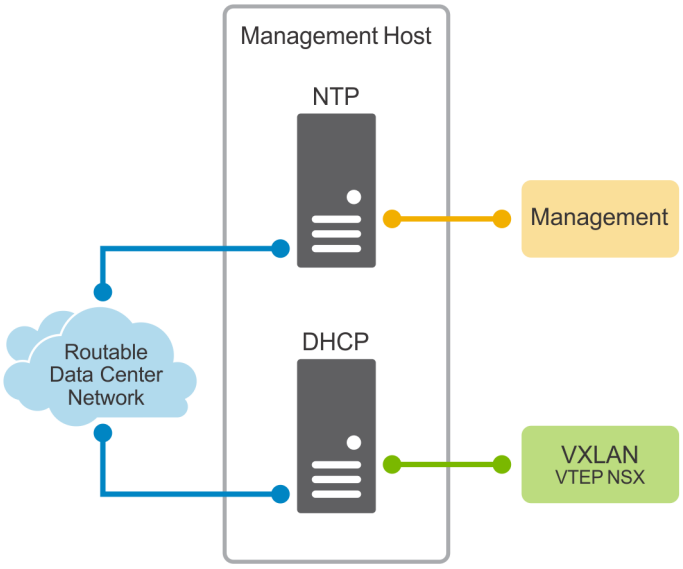

Figure 2. Management host in an existing vSphere environment

Connectivity

The Cloud Builder VM must be able to communicate with the hosts that become the Cloud Foundation management cluster. In this example, the new cluster hosts were given IP addresses on the 172.16.11.0/24 subnet and placed on the management VLAN (1611). A port group on the management host that was tagged with VLAN 1611 was created and connected to a port on a switch that accepted (ingress) traffic. The switchports between that initial ingress port and the new hosts were tagged with VLAN 1611.

Network services

There are three network services that are essential for Cloud Foundation deployment.

NOTE: Misconfiguration or lack of one of these services causes the validation portion of the installation to fail.

The information pertaining to the network services are inserted into the Cloud Builder Deployment Parameter Sheet. The parameter sheet is a spreadsheet that contains the details of the deployment and information specific to these prerequisites.

Domain Name Service

Domain Name Service (DNS) is required to provide both forward and reverse name resolution. The IP addresses of name servers, search domains, and hostnames of all the Cloud Foundation VMs must be inserted into the cloud builder deployment parameter sheet. Forward and reverse DNS entries of any hostname that are indicated in the parameter sheet should be tested and retested for both forward and reverse lookups. Test the DNS entries using their Fully Qualified Domain Name (FQDN) and their short name (hostname).

NOTE: Every DNS hostname and corresponding IP address that is specified in the parameter sheet are tested during the validation phase.

Dynamic Host Configuration Protocol

The Dynamic Host Configuration Protocol (DHCP) provides network addresses to the VXLAN tunnel end points (VTEP). A VTEP is a software-based endpoint of a VXLAN connection that communicates over these IP addresses and are on the same Layer 2 network. An address pool large enough to provide an address to every host connection on the VXLAN network--one address per NIC per host--is required.

NOTE: Software-based VTEPs could transition to hardware-based VTEPs with the launch of new features included in future switch operating systems.

Whenever a new cluster node is added to the Cloud Foundation, a new tunnel endpoint is created which requires DHCP to obtain IP addressing information. The DHCP service is not only a prerequisite for deployment but an ongoing requirement. Place the DHCP server (or VM) on a reliable and well-maintained part of your infrastructure.

The validation process checks to ensure that DHCP is available on the VXLAN network that is specified in the parameter sheet. Validation fails if there is no positive DHCP response on the VXLAN network.

Network Time Protocol

Time synchronization is critical to the Cloud Foundation stack. All hosts and the Cloud Builder VM are synchronized to a reference time source before attempting to run the validation phase of the Cloud Builder process. Network Time Protocol (NTP) traffic is routed from client to source or it can travel over the same L2 network.

4

Validated components

VMware no longer maintains a VMware Compatibility Guide for Cloud Foundation. Since vSAN is an underlying requirement of Cloud Foundation, any hardware specified as a vSAN Ready Node is approved for Cloud Foundation.

Topics:

Hardware components

The following hardware components were used in the validation of this solution.

NOTE: Cloud Foundation automatically configures vSAN disk groups, which requires following a few rules for drive population:

- Identical drive configurations in each target host
- There must be one size for all cache drives, as well as one size for all capacity drives
- The number of capacity drives in a host is cleanly divisible by the number of cache drives (that is, the result is a whole number)

Table 1. Hardware components

Manufacturer	Model	Description	Specifications
Dell EMC	PowerEdge MX7000	Chassis	
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 switch IOM	12 Gbps SAS
Dell EMC	Networking MX9116n	Network fabric switching engine IOM	
Dell EMC	Networking MX7116n	Network fabric expander IOM	
Dell EMC	Networking MX5108n	Network switch IOM	
Dell EMC	HBA330 MX	Disk controller—internal drives	
Dell EMC	HBA330 MMZ	Disk controller—PowerEdge MX5016s drives	
Dell EMC	BOSS MX	Boot/OS device	BOSS Card MX, 2x 256 GB M.2, RAID-1 vDisk
Toshiba	PX05SMB	vSAN cache drive	800 GB, 12 Gbps SAS SSD, 2.5"
Samsung	PM1635a	vSAN capacity drive	1.6 TB, 12 Gbps SAS SSD, 2.5"
QLogic	QL41232HMKR	Network interface card	Slot Mezz 1A, 2 ports x 25 GbE

Software and firmware

NOTE: The VMware Compatibility Guide (VCG) is the system of record for 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 https://www.dell.com/support.

Software

This document is written for Cloud Foundation 4.0 running on VMware ESXi 7.0. The required build and version of VMware ESXi is specified by the version of Cloud Foundation to be installed. It is critical that the versions of both the Cloud Foundation and VMware ESXi correspond.

5

Hardware overview

This section provides additional information about the hardware platform used in the development of this deployment guide.

Topics:

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 or four 4S double-width compute sleds and 12 Gb/s 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.
- Multichassis 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.

Front view of the PowerEdge MX7000 chassis

The front of the PowerEdge MX7000 chassis provides access to compute and storage sleds, fans, KVM, and power supplies. The configuration in the image below includes the following components:

- Four Dell EMC PowerEdge MX740c sleds in slots one through four
- One Dell EMC PowerEdge MX840C sled in slots five and six (not used in this guide)
- Two Dell EMC PowerEdge MX5016s sleds in slots seven and eight

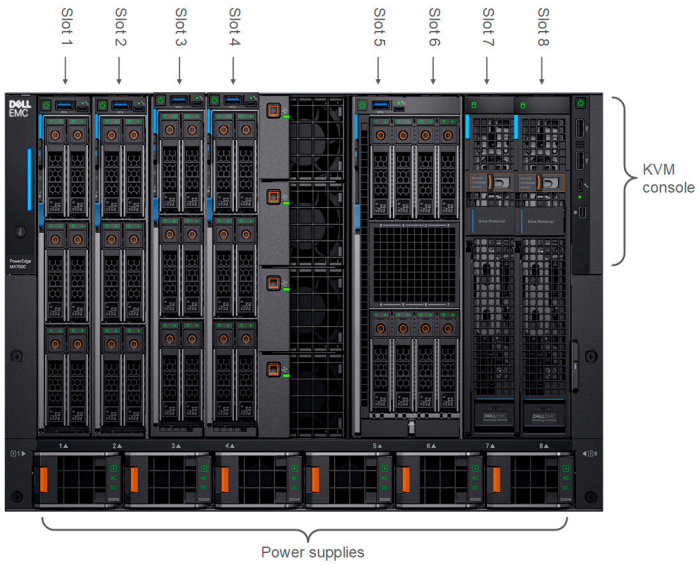

Figure 3. PowerEdge MX7000 chassis—front view

Back view of the PowerEdge MX7000 chassis

The back of the PowerEdge MX7000 chassis provides access to network and storage fabrics, management modules, fans, and power connections. The configuration in the image below includes the following components:

- Two Dell EMC Networking MX5108n Ethernet switches installed in fabric slots A1 and A2
- Two Dell EMC PowerEdge MX9002m management modules that are installed in management slots MM1 and MM2
- Two Dell EMC PowerEdge MX5000s SAS fabric switch modules that are installed in fabric slots C1 and C2
- Empty or available fabric slots B1 and B2

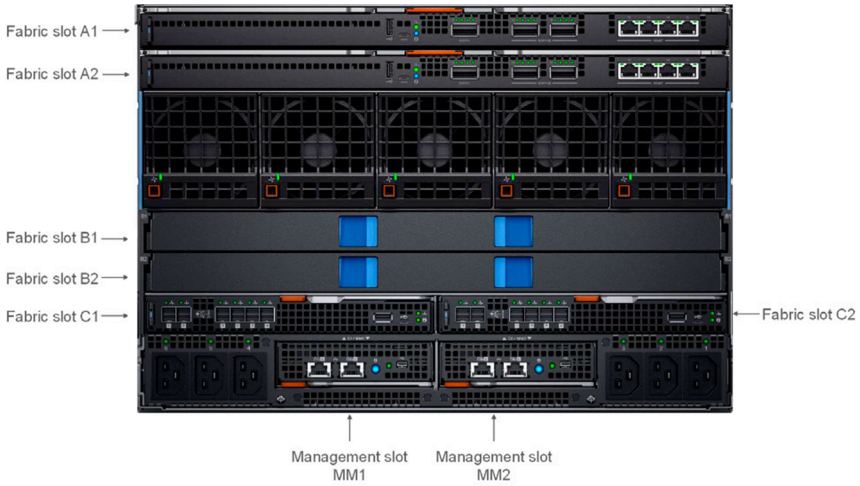

Figure 4. PowerEdge MX7000 chassis—rear view

Logical view of the PowerEdge MX7000 chassis

PowerEdge MX7000 supports three fabrics—two for general use and one for storage only. All three fabrics support redundant modules.

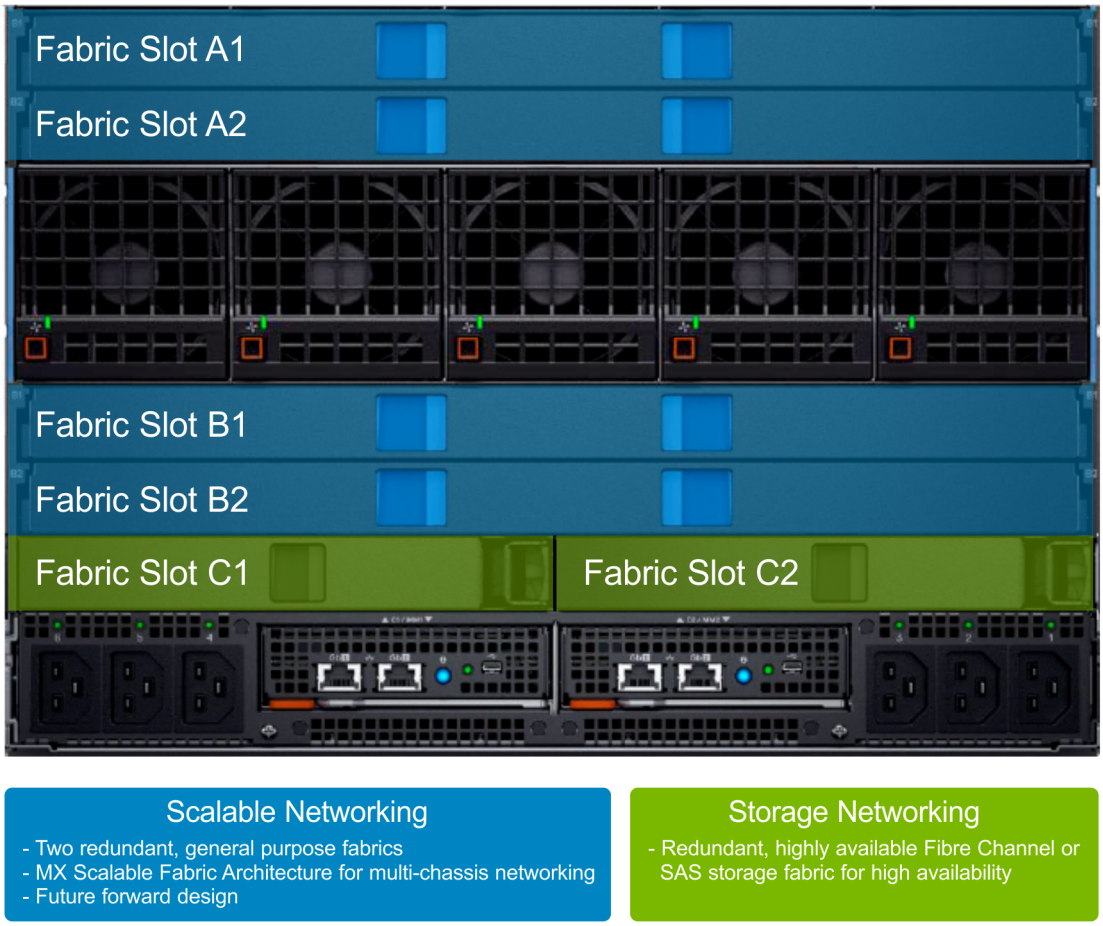

Figure 5. Logical view of the PowerEdge MX7000 chassis

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 can serve 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

Figure 6. Dell EMC PowerEdge MX740c compute sled

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 may 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: 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.

Figure 7. Dell EMC PowerEdge MX5016s storage sled

Dell EMC PowerEdge MX9002m management module

The Dell EMC PowerEdge MX9002m management module controls the overall 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.

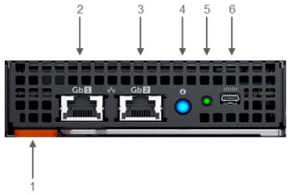

Figure 8. PowerEdge MX9002 management module

1. Handle release
2. Gigabit Ethernet port 1
3. Gigabit Ethernet port 2
4. ID button and health status LED
5. Power status LED
6. Micro-B USB port

Dell EMC Networking MX9116n Fabric Switching Engine

The Dell EMC Networking MX9116n Fabric Switching Engine (FSE) is a scalable, high-performance, low latency 25 GbE switch purpose-built for the PowerEdge MX platform. The MX9116n FSE provides enhanced capabilities and cost-effectiveness for enterprise, mid-market, Tier 2 cloud, and Network Functions Virtualization (NFV) service providers with demanding compute and storage traffic environments.

In addition to 16 internal 25 GbE ports, the MX9116n FSE provides:

- Two 100 GbE QSFP28 ports
- Two 100 GbE QSFP28 unified ports
- Twelve 2x100 GbE QSFP28-Double Density (DD) ports

The QSFP28 ports can be used for Ethernet uplink connectivity. For more information, see Management host section. The unified ports can be used for Ethernet uplink connectivity and supporting eight 32 Gb Fibre Channel (FC) ports for SAN connectivity supporting both NPIV Proxy Gateway (NPG) and direct attach FC capabilities.

The QSFP28-DD ports provide fabric expansion connections for up to nine more PowerEdge MX7000 chassis using the MX7116n Fabric Expander Module. The QSFP28-DD ports also provide capacity for extra uplinks, VLTi links, and connections to rack servers at 10 GbE or 25 GbE using breakout cables. The PowerEdge MX7000 chassis supports up to four MX9116n FSEs in Fabric A or B, or both.

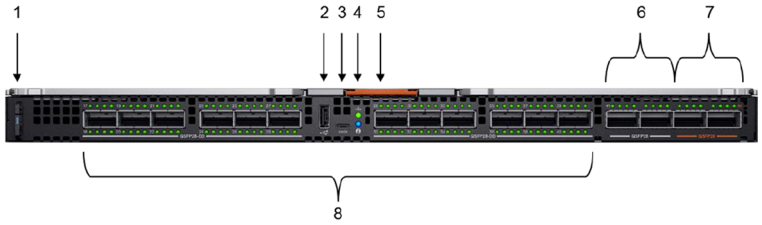

Figure 9. MX9116n FSE

1. Express service tag
2. Storage USB port
3. Micro-B USB console port
4. Power and indicator LEDs
5. Handle release
6. Two QSFP28 ports
7. Two QSFP28 unified ports
8. 12 QSFP28-DD ports

Dell EMC Networking MX7116n Fabric Expander Module

The Dell EMC Networking MX7116n Fabric Expander Module (FEM) acts as an Ethernet repeater, taking signals from an attached compute sled and repeating them to the associated lane on the external QSFP28-DD connector. The MX7116n FEM provides two QSFP28-DD interfaces, each providing up to eight 25 GbE connections to the chassis.

There is no operating system or switching ASIC on the MX7116n FEM, so it never requires an upgrade. There is also no management or user interface, making the MX7116n FEM maintenance-free. The PowerEdge MX7000 chassis supports up to four MX7116n FEMs in Fabric A or Fabric B, or both.

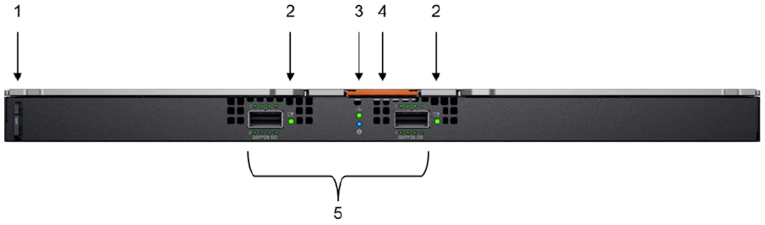

Figure 10. MX7116n FEM

1. Express service tag
2. Supported optic LED
3. Power and indicator LEDs
4. Module insertion or removal latch

5. Two QSFP28-DD fabric expander ports

NOTE: The MX7116n FEM cannot act as a stand-alone switch and must be connected to the MX9116n FSE to function.

Dell EMC Networking MX5108n Ethernet switch

The Dell EMC Networking MX5108n Ethernet switch is targeted at small PowerEdge MX7000 deployments of one or two chassis. Although not a scalable switch, it still provides high-performance and low latency with a non-blocking switching architecture. The MX5108n switch provides line-rate 25 Gbps Layer 2 and Layer 3 forwarding capacity to all connected servers with no oversubscription.

In addition to eight internal 25 GbE ports, the MX5108n switch includes:

- One 40 GbE QSFP+ port
- Two 100 GbE QSFP28 ports
- Four 10 GbE RJ45 BASE-T ports

The ports can be used to provide a combination of network uplinks, VLT interconnects (VLTi), or for FCoE connectivity. The MX5108n switch supports FCoE Initialization Protocol (FIP) Snooping Bridge (FSB) mode but does not support NPG or direct attach FC capabilities. The PowerEdge MX7000 chassis supports up to four MX5106n Ethernet switches in Fabric A or Fabric B, or both.

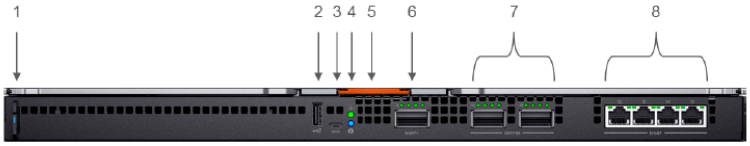

Figure 11. Dell EMC Networking MX5108n Ethernet switch

1. Express service tag
2. Storage USB port
3. Micro-B USB console port
4. Power and indicator LEDs
5. Module insertion or removal latch
6. One QSFP+ port
7. Two QSFP28 ports
8. Four 10GBASE-T ports

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.

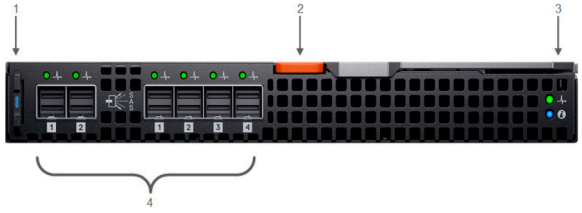

Figure 12. Dell EMC PowerEdge MX5000s storage sled

1. Express service tag
2. Module insertion or removal latch
3. Power and indicator LEDs
4. Six SAS ports

6

Physical layout

There are multiple configurations of Cloud Foundation on PowerEdge MX7000 chassis that are described in this document. The Cloud Foundation software addresses the host servers using their IP Address. Deploying compute sleds across multiple PowerEdge MX7000 chassis has no impact on the software as long as the networking is configured properly on the Networking IO modules and the switches to which the PowerEdge MX7000 chassis connects. The physical layout and resulting cabling is impacted by the number of PowerEdge MX7000 chassis in use but no other changes are made in the environment.

Topics:

Configuration options

Option 1—single PowerEdge MX7000 enclosure

- One Dell EMC PowerEdge MX7000 enclosure
- Four Dell EMC PowerEdge MX740c compute sleds
- Two Dell EMC Networking MX5108n Ethernet switches

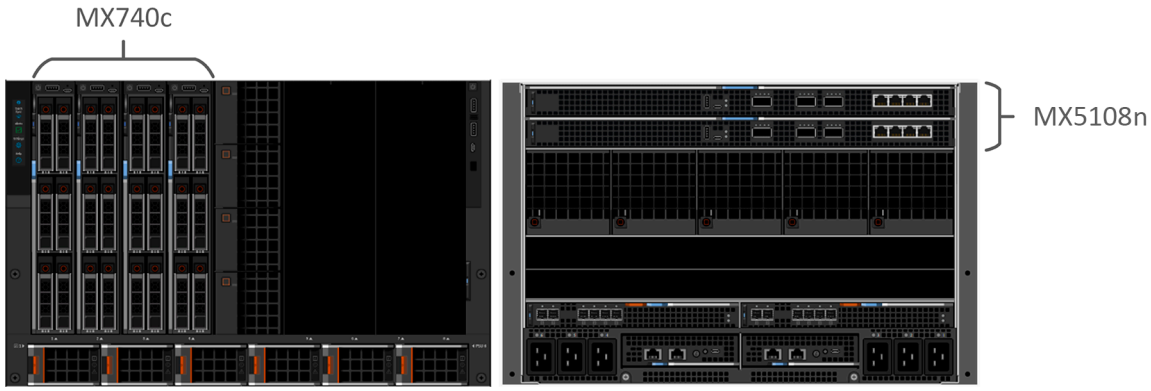

Figure 13. Single PowerEdge MX7000 enclosure

Option 2—single PowerEdge MX7000 with MX5016s storage sled

- One Dell EMC PowerEdge MX7000 enclosure
- Four Dell EMC PowerEdge MX740c compute sleds
- Two Dell EMC Networking MX5108n Ethernet switches
- Two Dell EMC PowerEdge MX5016s storage sleds
- Two Dell EMC PowerEdge MX5000s SAS IO Modules

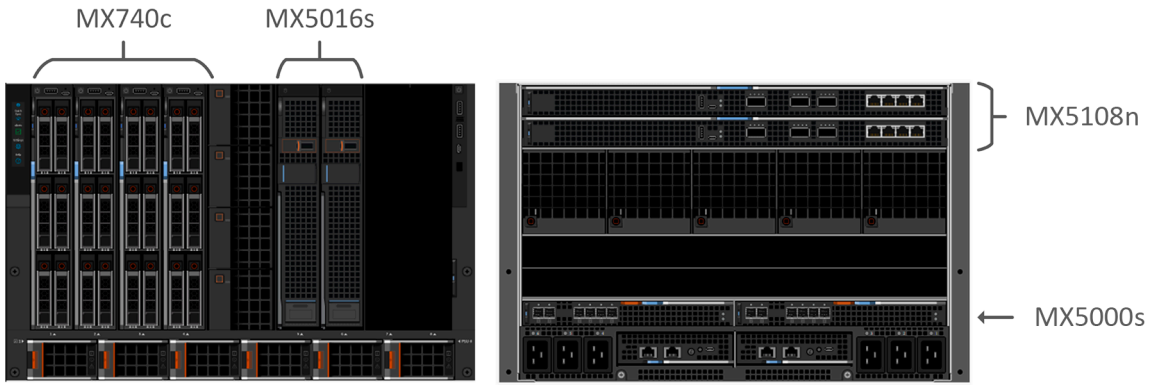

Figure 14. Single PowerEdge MX7000 with MX5016s storage sled

Option 3—two PowerEdge MX7000 enclosures

- Two Dell EMC PowerEdge MX7000 enclosures
- Four Dell EMC PowerEdge MX740c compute sleds
- Four Dell EMC Networking MX5108n Ethernet switches

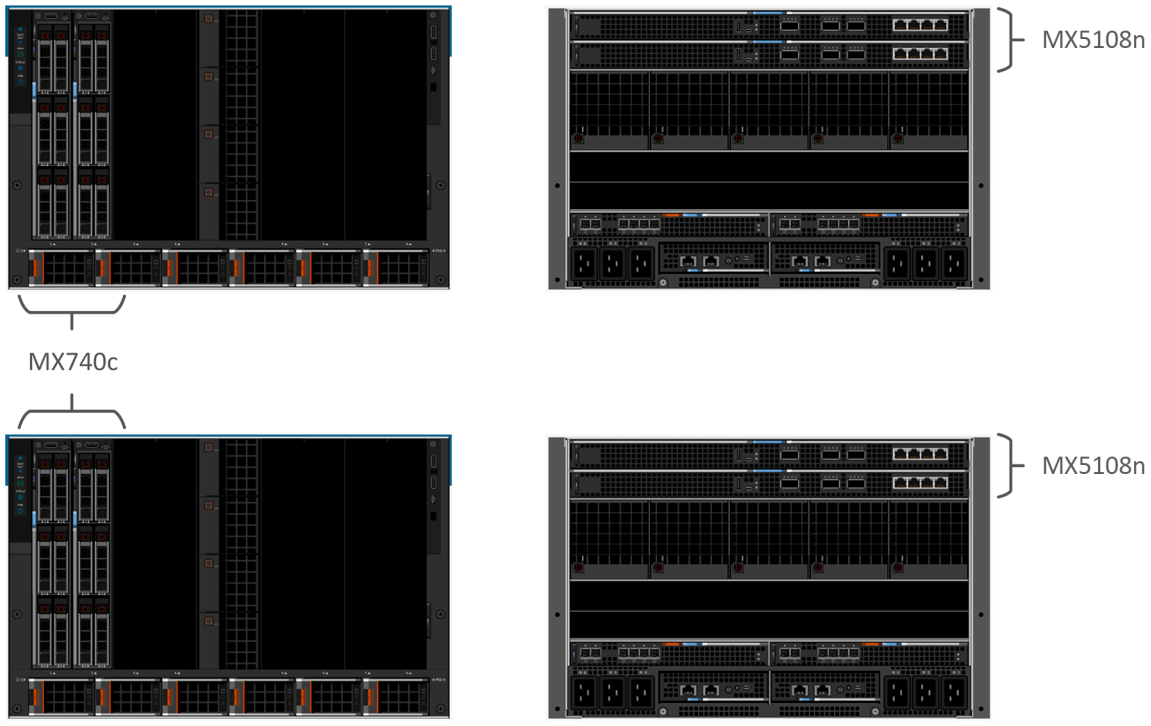

Figure 15. Two PowerEdge MX7000 enclosures

Option 4—two PowerEdge MX7000 with MX5016s storage sled

- Two Dell EMC PowerEdge MX7000 enclosures
- Four Dell EMC PowerEdge MX740c compute sleds
- Two Dell EMC Networking MX5108n Ethernet switches
- Two Dell EMC PowerEdge MX5016s storage sleds
- Four Dell EMC PowerEdge MX5000s SAS IO Modules

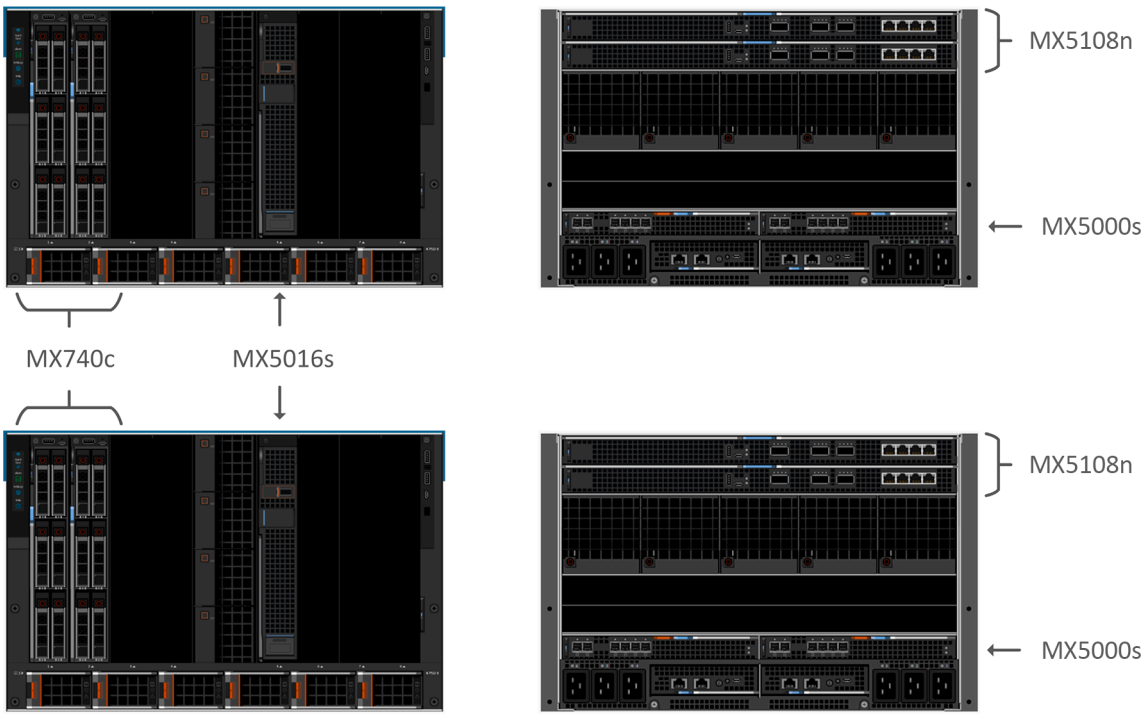

Figure 16. Two PowerEdge MX7000 with MX5016s storage

Option 5—two PowerEdge MX7000 enclosures using Fabric Switching Engine

- Two or more (up to a maximum of 10) Dell EMC PowerEdge MX7000 enclosures
- Four Dell EMC PowerEdge MX740c compute sleds
- Two Dell EMC Networking MX9116n Fabric Switching Engines
- Two Dell EMC Networking MX7116n Fabric Expansion Modules
  - Plus two MX7116n modules for each additional chassis
- Four Dell EMC PowerEdge MX5000s SAS IO Modules (only if using the MX5016s storage sleds)

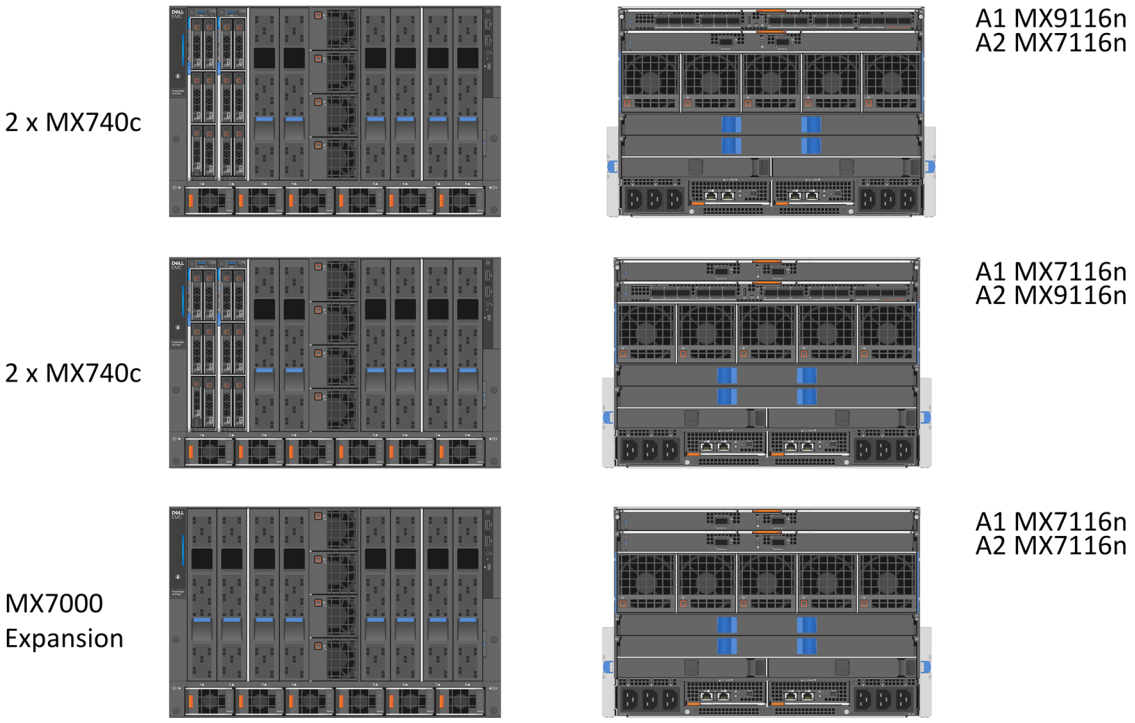

Figure 17. Two PowerEdge MX7000 enclosures using Fabric Switching Engine

Cabling

PowerEdge MX5016s storage sleds are internally cabled and the PowerEdge MX5000s SAS IOM has no impact on external cabling.

Cabling for a dual PowerEdge MX7000 enclosure configuration using Fabric Switching Engines

The following figures show the external cabling for a multiple PowerEdge MX7000 enclosure configuration when the MX9116n Fabric Switching Engines and MX7116n Fabric Expansion Modules are used. The Customer Network Link Aggregation is shown as an example as the upper layer connection is not specified except that it must use an LACP-enabled link aggregation group. You can add more enclosures (up to a maximum of 10) that connect back to the upper level devices in the infrastructure. Additional PowerEdge MX7000 enclosures require only two MX7116n Fabric Expansion Modules whose ports appear as additional ports on the MX9116n Fabric Switching Engines on the first two PowerEdge MX7000 enclosures.

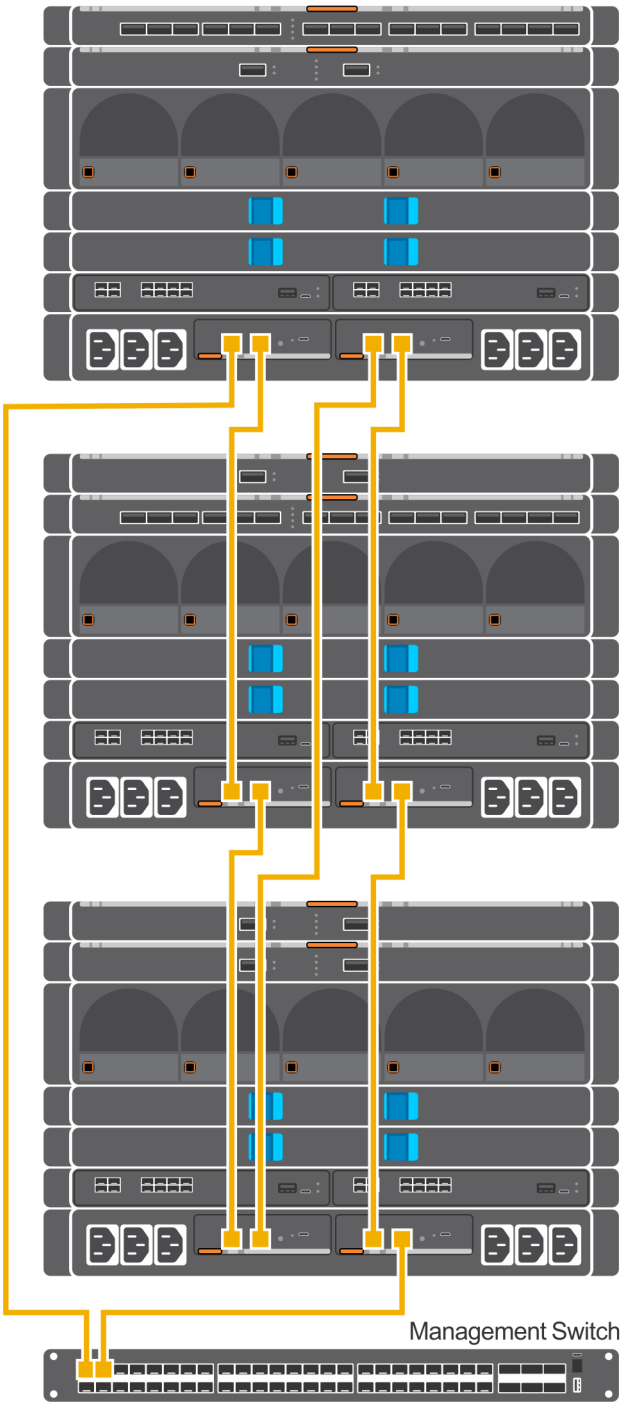

Figure 18. MX9002m Management module cabling

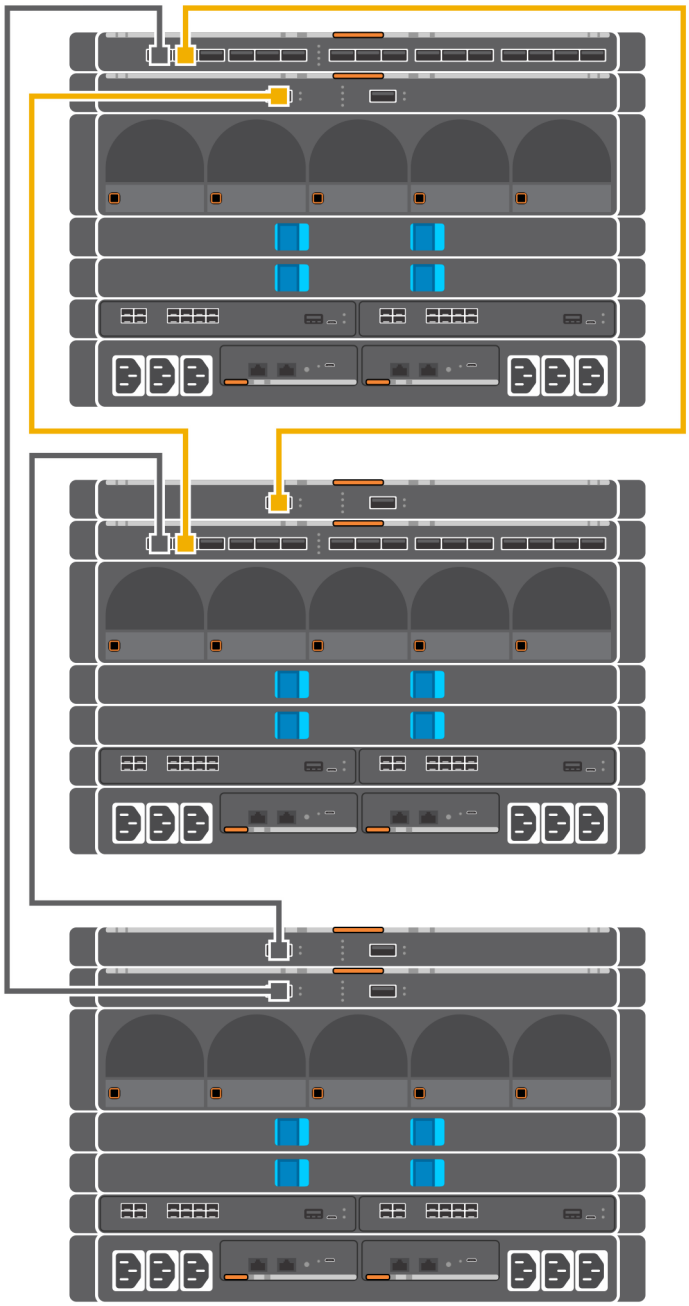

Figure 19. Connectivity between FSE modules and FEM modules

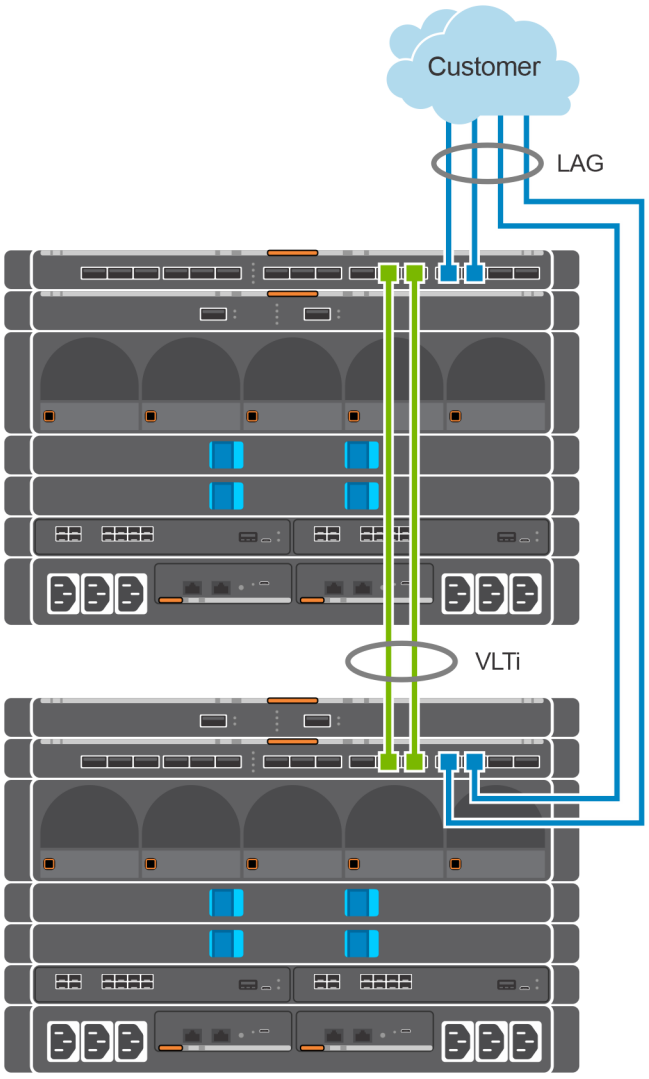

Figure 20. Uplinks to customer network environment

7

Cloud Foundation and SDDC design considerations

VMware Cloud Foundation relies on a set of key infrastructure services to be made available externally. You must configure these external services before you begin deployment.

NOTE: This section is universal for Cloud Foundation deployments regardless of hardware platform. The content in this section is also available in the VMware Cloud Foundation Planning and Preparation Guide, and is included here for reference. The original content in the VMware website includes additional sections which are not in the scope of this document.

Topics:

External services overview

Many external services are required for the initial deployment of Cloud Foundation and for the deployment of other optional components such as vRealize Operations or vRealize Automation. The following table lists the required and optional external services and dependencies:

Table 2. Required and optional external services and dependencies

Service	Purpose
Active Directory (AD)	(Optional) Provides authentication and authorization. **NOTE:** AD is required if you are deploying vRealize Automation.
Dynamic Host Configuration Protocol (DHCP)	Provides automated IP address allocation for VXLAN Tunnel Endpoints (TEPs).
Domain Name Service (DNS)	Provides name resolution for the various components in the solution.
Network Time Protocol (NTP)	Synchronizes time between the various components.
Simple Message Transfer Protocol (SMTP)	(Optional) Provides method for email alerts.
Certificate Authority (CA)	(Optional) Allows replacement of the initial self-signed certificates that are used by Cloud Foundation. **NOTE:** A CA is required if you are deploying vRealize Automation.

Active Directory

Cloud Foundation uses Active Directory (AD) for authentication and authorization to resources. The Active Directory services must be reachable by the components that are connected to the management and vRealize networks.

You must configure user and group accounts in AD before adding them to the SDDC manager and assigning privileges.

NOTE: If you plan to deploy vRealize Automation, Active Directory services must be available. For more information on AD configuration, see the vRealize Automation documentation.

Dynamic Host Configuration Protocol

Cloud Foundation uses Dynamic Host Configuration Protocol (DHCP) to automatically configure each VM kernel port of an ESXi host that is used as a TEP with an IPv4 address. One DHCP scope must be defined and made available for this purpose.

The DHCP scope that is defined must be large enough to accommodate all the initial and future servers that are used in the Cloud Foundation solution. Each host requires two IP addresses, one for each TEP configured.

Domain Name System

During deployment, you must provide the DNS domain information to be used to configure the various components. The root DNS domain information is required and, optionally, you can also specify subdomain information.

DNS resolution must be available for all the components that are contained within the Cloud Foundation solution, which includes servers, virtual machines, and any virtual IPs that are used. For more information on the components that are required for DNS resolution before starting a Cloud Foundation deployment, see Host names and IP addresses.

Ensure that both forward and reverse DNS resolutions are functional for each component before deploying Cloud Foundation or creating any workload domains.

Network Time Protocol

All components must be synchronized against a common time provider by using the Network Time Protocol (NTP) on all nodes. Important components of Cloud Foundation, such as vCenter Single Sign-On (SSO), are sensitive to a time drift between distributed components. Synchronized time between the various components also assists with troubleshooting.

Requirements for the NTP sources include the following:

- The IP addresses of two NTP sources are provided during the initial deployment.
- The NTP sources must be reachable by all the components in the Cloud Foundation solution.
- Time skew is less than 5 minutes between NTP sources.

Simple Mail Transfer Protocol mail relay (optional)

Certain components of the SDDC, such as vCenter, Log Insight, and vRealize Automation, can send status messages to users by email. To enable this functionality, a mail relay that does not require user authentication must be available through SMTP. As a best practice, limit the relay function to the networks allocated for use by Cloud Foundation.

Certificate Authority (optional)

The components of the SDDC require SSL certificates for secure operation. During deployment, self-signed certificates are used for each of the deployed components. These certificates can be replaced with certificates that are signed by an internal enterprise CA or by a third-party commercial CA.

If you plan to replace the self-signed certificates, the CA must be able to sign a Certificate Signing Request (CSR) and return the signed certificate. All endpoints within the enterprise must also trust the root CA of the CA.

If you plan to deploy vRealize Automation, a Certificate Authority is required, and the certificates are required during installation.

Physical network requirements

Before deploying Cloud Foundation, configure the physical network to enable the following features:

- VLAN Tagging (802.1Q)
- Jumbo frames
  - A minimum MTU value of 1600 is required, however it is recommended that you set the MTU to 9000.

Network pools

Cloud Foundation uses a construct that is called a network pool to automatically configure VM kernel ports for vSAN, NFS, and vMotion.

Cloud Foundation uses an Internet Protocol Address Management (IPAM) solution to automate the IP configuration of VM kernel ports for vMotion, vSAN, and NFS (depending on the storage type being used).

When a server is added to the inventory of Cloud Foundation, it goes through a process called host commissioning. During this process, the hosts are associated with an existing network pool. When the host is provisioned during the create VI workload domain, add cluster, or add host workflow, it automatically configures the VMkernel ports and allocates IP addresses for vMotion, vSAN, and NFS from the network pool the host was associated with.

You can expand the included IP address range of a network pool at any time, however you cannot modify the other network information. Ensure that you have defined each subnet in the network pool to account for current and future growth in your environment.

VLANs and IP subnets

Network traffic types within Cloud Foundation are isolated from each other by using VLANs. Before deploying your SDDC, you must allocate VLAN IDs and IP subnets for each required traffic type. Configure the VLAN IDs and IP subnets in your network to pass traffic through your network devices. Before you start the Cloud Foundation deployment, verify that the allocated network information is configured and does not conflict with pre-existing services before starting your Cloud Foundation deployment.

The number and size of the subnets that are required for a deployment depends on:

- The number of workload domains that are created
- The number of clusters defined
- The optional components that are installed

The following table lists the basic allocation of VLANs and IP subnets for a sample deployment. Use this sample to define the VLANs and IP subnets in your environment.

Table 3. VLANs and IP subnets for a sample deployment

Workload Domain	Cluster	VLAN Function	VLAN ID	Subnet	Gateway
Management	Cluster-01	Management	1711	172.17.11.0/24	172.17.11.253
		vMotion	1712	172.17.12.0/24	172.17.12.253
		vSAN	1713	172.17.13.0/24	172.17.13.253
		VXLAN (NSX VTEP)	1714	172.17.14.0/24	172.17.14.253
		vRealize Suite (optional)	1715	172.17.15.0/24	172.17.15.253
		Uplink 1	2711	172.27.11.0/24	172.27.11.253
		Uplink 2	2711	172.27.12.0/24	172.27.12.253

NOTE: Cloud Foundation deploys vRealize Suite products to a dedicated VLAN-backed vSphere Distributed Port Group. The IP subnet must be routable to the Cloud Foundation management network and the firewall. Also, the networks should be disabled or configured as prescribed in the Cloud Foundation documentation.

Host names and IP addresses

Before deploying a Cloud Foundation, or creating or expanding a workload domain, you must define the hostnames and IP addresses for various system components.

Most of the defined hostnames and IP addresses need to exist in DNS and be resolvable, through forward and reverse lookups.

The required hostnames and IP addresses are categorized as follows:

- External services—services that are external to the Cloud Foundation solution and are required for proper operation.

- Virtual infrastructure layer—components that provide for the basic foundation of the Cloud Foundation solution.
- Operations management layer—components used for day-to-day management of the environment, for example, vRealize Operations.
- Cloud management layer—services that use the infrastructure layer resources, for example, vRealize Automation.

Host names and IP addresses for external services

External services such as Active Directory (AD) and NTP must be accessible and resolvable by IP Address and Fully Qualified Domain Name (FQDN). Acquire the hostnames and IP addresses for AD and NTP before deploying Cloud Foundation.

Allocate hostnames and IP addresses to the following components:

- NTP
- AD
- DNS
- Certificate Authority (CA)

The following table provides sample information for the external services. This example uses a DNS domain called `osevcf17.local` for illustration purposes. Modify the sample information to conform to the configuration of your site.

Table 4. Configuration for external services

Component Group	Hostname	DNS	IP Address	Description
DNS	dc01sfo	sfo01.osevcf17.local	172.18.11.5	AD and DNS server for the sfo01 subdomain
NTP	Ntp	sfo01.osevcf17.local		Round-robin DNS pool containing the NTP servers
	0.ntp	sfo01.osevcf17.local	172.18.11.251	First NTP server
	1.ntp	sfo01.osevcf17.local	172.18.11.252	Second NTP server
AD, DNS or CA	dc01rpl	sfo01.osevcf17.local	172.18.11.4	Windows host that contains the AD configuration, the DNS

Host names and IP addresses for the virtual infrastructure layer

Most of the virtual infrastructure components that are installed by Cloud Foundation require their hostnames and IP addresses to be defined before deployment.

During the initial deployment of Cloud Foundation, the management domain is created and you must define components specific to the management domain before installation.

After the initial deployment, additional workload domains are created and you must define components specific to each additional workload domain.

Planning ahead for the initial deployment and the workload domains to be created avoids delays in deployment.

The following table provides an example of the information for the virtual infrastructure layer using the standard deployment model with a single workload domain. This example uses a DNS domain called `osevcf17.local` for illustration purposes. Modify the sample information to conform to the configuration of your site.

Table 5. Configuration for the virtual infrastructure layer

Workload Domain	Hostname	DNS Zone	IP Address	Description
Management	sddc-manager	osevcf17.local	100.71.101.107	SDDC Manager VM
	vcenter-m01-17	osevcf17.local	100.71.101.120	vCenter VM
	sddc17-m01-nsx01	osevcf17.local	100.71.101.131	NSX-T Management Cluster
	sddc17-m01-nsx02	osevcf17.local	100.71.101.132	NSX-T Virtual Appliance 1
	sddc17-m01-nsx03	osevcf17.local	100.71.101.133	NSX-T Virtual Appliance 2

Table 5. Configuration for the virtual infrastructure layer (continued)

Workload Domain	Hostname	DNS Zone	IP Address	Description
	sddc17-m01-nsx04	osevcf17.local	100.71.101.134	NSX-T Virtual Appliance 3
	vcfmgmthost01	osevcf17.local	100.71.101.171	Management Host 1
	vcfmgmthost02	osevcf17.local	100.71.101.172	Management Host 2
	vcfmgmthost03	osevcf17.local	100.71.101.173	Management Host 3
	vcfmgmthost04	osevcf17.local	100.71.101.174	Management Host 4

8

Networking requirements

This section covers the networking requirements from both the Cloud Foundation software perspective and from a networking hardware connectivity perspective. This section also briefly describes the configuration options for configuring networks on a Dell EMC PowerEdge MX7000 chassis. The actual networking configuration procedures are described in the later sections.

Topics:

VMware Cloud Foundation networking

A successful VMware Cloud Foundation deployment relies heavily on networks that are constructed and allocated to Cloud Foundation. The networks are used by Cloud Builder during the installation and configuration process and then used by Cloud Foundation to carry out various activities. The different networks are allocated to specific purposes and have different requirements.

VMware Cloud Foundation requires six networks and at least one connection to a customer network (for external access to your Cloud Foundation stack). In the following example, a private IP address range is used for all connectivity within the management stack. There is also an IP network that connects back to an external network.

Each of these networks is propagated to the Cloud Foundation stack using tagged VLANs. Using tagged VLANs enables mapping of port groups to VLANs allowing access to resources as needed. All these networks are routable to and from each other. The routing task is executed at some layer above the access level switched fabric that is deployed here.

The networks required to deploy Cloud Foundation are listed in the following table:

Table 6. Networks required to deploy Cloud Foundation

Network	Description
Management	Dedicated to communication between all the deployed resources and services. When the SDDC Manager Utility needs to communicate to any other service or resource, it uses the management network.
vSAN	Used to communicate and synchronize vSAN storage traffic across multiple hosts to ensure data integrity and resiliency.
vMotion	Used to quickly redistribute virtual machine state and or storage between hosts.
Host Overlay	The host overlay network is used by NSX-T for control plane communication between the hosts of the VMware Cloud Foundation cluster.
Uplink 1 and Uplink 2	The uplink networks are used by NSX-T for data traffic into and out of the cluster.
Edge Overlay	This network is used by the Edge Nodes in an NSX-T environment to allow the transport nodes to access the capabilities of the NSX-T Data Center.

Network configuration options

There are two different approaches to configuring the network switches used in this deployment guide. The first approach is the manually configured approach where a startup-configuration is built up by manually configuring the ports, VLANs, aggregations

assigning tagging rules, QoS, and other settings typical to this type of deployment. The other approach is to use Dell EMC SmartFabric which uses a simplified, reusable approach to configure the switches on the PowerEdge MX7000 chassis and assign those configurations to the compute resources in the chassis.

The advantage of using the manual configuration method is that every aspect of the switch configuration is available. The switch startup-configuration reflects every change that is made by the network administrator but this method is slower to deploy and more prone to human error.

The advantage of SmartFabric is in the time it takes to deploy a configuration. With a relatively small number of configuration steps, a fabric and profile that can be assigned to the compute sleds are created. When a change is made to the fabric or profile, it can be easily pushed out to the switches and compute sleds in the PowerEdge MX7000 chassis. The `startup-config` and `running-config` commands do not reflect the actual configuration of the chassis switches.

Review the sections in both Manual switch configuration and SmartFabric network configuration before choosing a path.

Networking and NSX-T

Cloud Foundation version 4.0 has made a change to NSX-T from NSX, also seen as NSX-V in Cloud Foundation 3.9. NSX-T introduces the concept of an Edge Node Cluster. An Edge Node Cluster is a resilient cluster of NSX-T edge nodes that are used to connect to upstream networks. All NSX traffic leaving the cluster passes through these edge nodes and the edge nodes rely on physical NICs.

Each member of an Edge Node Cluster contains identical configurations of physical NICs so that the NSX network constructs can fail over if there is a node failure. This diagram shows a conceptual view of the different uplinks from the top of rack switches to the spine. The MLAG is a VLT port channel that includes connections from both top of rack switches to multiple devices in the spine layer.

The connections that are shown are for conceptual purposes only and are not accurate to any specific model of switch. In the case of a MX7000 modular deployment, the MX9116n IOMs are the Top of Rack or leaf switches.

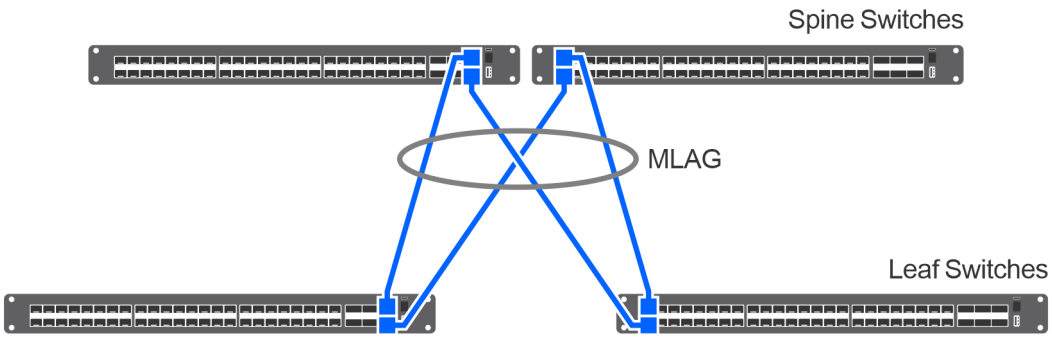

Figure 21. Top of Rack to Spine

Physical Hardware

Installing physical hardware to provide additional network interfaces in an MX7000 environment requires the installation of devices into the B fabric of the MX7000 chassis. More MX9116n IOMs would be installed into slots B1 and B2 on the rear of the chassis. Extra network mezzanine cards would be installed into the B connector of each of MX740c compute sleds to interface with the B fabric IOMs.

Installing dedicated hardware provides better throughput, lower latency, and simplified monitoring of network performance and metrics. Using NIC Partitioning (NPAR) to create virtualized network interfaces is not supported.

Network connectivity

When deploying Dell EMC Networking MX5108n switch modules, the switches are installed in Fabric A of the PowerEdge MX7000 enclosure. Each modular switch has eight internal ports with one port being connected to each compute sled. Two modules provide a redundant (A1 and A2) connection to each of the PowerEdge MX740c compute sleds. The connection between the compute sleds and the MX modular switches do not use any kind of link aggregation protocol. The connections are

separate network connections that are managed by the Cloud Foundation stack. Due to the limited ports available for uplink the MX5108n should not be used where NSX-T edge node capabilities are desired.

Deploying Dell EMC Networking MX9116n Fabric Switching Engines and Dell EMC Networking MX7116n Fabric Expansion Modules is a different process. The FSEs are installed into the A1 fabric of the first chassis and the A2 fabric of the second chassis. The FEMs are distributed across both PowerEdge MX7000 chassis in the remaining A fabric slots. The FEMs connect back to the FSEs using a double data rate, 200 Gbps cable. The connection between the compute sleds and the MX modular switches do not use any kind of link aggregation protocol. Additional PowerEdge MX7000 chassis (up to 8 more) can be added and require only the MX7116n FEMs in the A1 and A2 fabric slots. The additional ports available for uplink make the MX9116n an excellent choice for the deployment of NSX-T edge nodes. Since NSX-T is now a part of Cloud Foundation 4.0 we will place the emphasis of this document on the MX9116n Scalable Fabric Architecture.

The connections from the modular switches to the external network are implemented using Virtual Link Trunking (VLT) link aggregation. VLT allows you to create a single LACP-managed link aggregation from the two modular switches to an LACP-managed aggregation in the external network. Use the link aggregation only on the link between the modular switches and the customer network.

VLAN and subnets for networking configuration

The following table shows the VLAN and networking data that are used for the Cloud Foundation deployment. In our example, private addresses are used for Management, vSAN, vMotion, and VXLAN networks. However, this is not mandatory.

Table 7. VLAN and networking data used for the Cloud Foundation deployment

Name	VLAN ID	Subnet	Mask	Default Gateway
Management	1711	172.17.11.0	255.255.255.0	172.17.11.253
vSAN	1712	172.17.12.0	255.255.255.0	172.17.12.253
vMotion	1713	172.17.13.0	255.255.255.0	172.17.13.253
Host Overlay	1714	172.17.14.0	255.255.255.0	172.17.14.253
Uplink1	2711	172.27.11.0	255.255.255.0	172.27.11.1
Uplink2	2712	172.27.12.0	255.255.255.0	172.27.12.1
Edge Overlay	2713	172.27.13.0	255.255.255.0	172.17.13.253

MTU Settings

Configuring jumbo frames is a best practice for both vMotion and vSAN networks, both of which are core components of Cloud Foundation. All the switch ports on the modular switches and up to the aggregation switches used to connect multiple PowerEdge MX7000 enclosures together must be configured for jumbo frames. It is also recommended to configure jumbo frames on the VXLAN network. The validation phase run as part of the Cloud Foundation installation process tests end-to-end connectivity of all specified devices, but fails if the jumbo frames are not correctly configured.

9

Manual switch configuration

This section describes the configuration of the MX9116n FSEs (Fabric Switching Engines) switches. Each PowerEdge MX7000 has one MX9116n and one MX7116n in the A fabric. The MX9116n in chassis 1 should be placed in the A1 slot and the MX9116n in chassis 2 should be placed in the A2 slot. This distributes the fabric's switching engines across both chassis. In the event of a loss of one of the MX9116n, only one half of the fabric is impacted. The two MX7116n expanders are installed in the remaining A fabric slots. The IP address of each switch module is assigned using the chassis management interface (MSM).

Topics:

Switch operating mode

The MX9116n switches operate either in Full Switch or SmartFabric mode. The switch should be operating in Full Switch mode for this example.

```
MX7K-IOM-A2# show switch-operating-mode
MX7K-IOM-A2# Switch-Operating-Mode : Full Switch Mode
```

VLANs and subnets for manual switch configuration

Here are the VLANs created on the MX9116N switch that is required for Cloud Foundation. These VLANs and subnets are created on both switches.

```
interface vlan1711
description 1711-Mgmt
no shutdown
mtu 9216

interface vlan1712
description 1712-VMotion
no shutdown
mtu 9216

interface vlan1713
description 1713-VSAN
no shutdown
mtu 9216

interface vlan1714
description 1714-host-overlay
no shutdown
mtu 9216

interface vlan2711
description 2711-Uplink1
no shutdown
mtu 9216
```

```
interface vlan2712
description 2712-Uplink2
no shutdown
mtu 9216

interface vlan2713
description Edge-Overlay
no shutdown
mtu 9216
```

Uplink and VLTi ports

VLT synchronizes Layer 2 table information between two switches and enables them to display as a single logical unit from outside the VLT domain. The VLT interconnect (VLTi) between two Dell EMC Networking MX9116N switches is a port group that is generated by configuring a VLT domain and specifying the discovery interfaces.

VLTi ports are specified during the configuration of the VLT domain by specifying the interconnect ports as discovery interfaces. The resulting VLTi port group cannot be managed manually. The VLT domain must be created on both VLT peer switches.

Configure the ports for VLTi

The VLTi ports for this example are ethernet 1/1/37 through 1/1/40. These ports are placed into a link aggregation by the VLT configuration process. Do not create the link aggregation manually.

To configure the ethernet1/1/37 through 1/1/40 ports, ensure that the:

- Ports are preconfigured correctly.
- Ports are set to the correct frame size.
- Ports are not set in any switchport (VLAN) mode.
- Ports are not shut down.

```
interface ethernet1/1/40
 no shutdown
 no switchport
 mtu 9216
 flowcontrol receive off
```

Configure VLT domain

The next step is to configure a VLT domain on each switch. It contains the exact same VLT domain on each switch. In this case, create a VLT domain '10' on each switch. The backup destination for the first switch is the management IPv4 address of the second switch. Similarly, the backup destination for the second switch is the management IPv4 address of the first switch.

```
MX9116-A1(config)#
MX9116-A1(config)# vlt-domain 10
MX9116-A1(conf-vlt-100)# discovery-interface ethernet1/1/37,1/1/38,1/1/39,1/1/40
MX9116-A1(conf-vlt-10)# backup destination 100.71.242.220
MX9116-A1(conf-vlt-10)# exit
```

Verify VLT settings

Verify the VLT settings by running the following command:

```
MX9116-A1# show vlt 10
Domain ID                  : 10
Unit ID                    : 2
Role                       : secondary
Version                    : 2.3
Local System MAC address   : 3c:2c:30:80:a8:80
```

```
Role priority               : 32768
VLT MAC address             : aa:bb:cc:11:11:11
IP address                  : fda5:74c8:b79e:1::2
Delay-Restore timer         : 90 seconds
Peer-Routing                : Disabled
Peer-Routing-Timeout timer : 0 seconds
VLTi Link Status
    port-channel1000        : up

VLT Peer Unit ID    System MAC Address    Status    IP Address              Version
-------------------------------------------------------------------------------
  1                 3c:2c:30:80:aa:80      up       fda5:74c8:b79e:1::1     2.3
MX9116-A1# show port-channel summary

Flags:  D - Down    I - member up but inactive    P - member up and active
        U - Up (port-channel)    F - Fallback Activated
--------------------------------------------------------------------------------
Group Port-Channel           Type     Protocol  Member Ports
--------------------------------------------------------------------------------
1    port-channel1    (D)     Eth      DYNAMIC   1/1/41(P) 1/1/42(P)
1000 port-channel1000 (U)     Eth      STATIC    1/1/37(P) 1/1/38(P) 1/1/39(P) 1/1/40(P)
```

Verify the VLTi (port-channel)

To view the port channel summary, enter the following command:

```
MX9116-A1#
MX9116-A1# show port-channel summary

Flags:  D - Down    I - member up but inactive    P - member up and active
        U - Up (port-channel)    F - Fallback Activated
--------------------------------------------------------------------------------
Group Port-Channel           Type     Protocol  Member Ports
--------------------------------------------------------------------------------
1000 port-channel1000 (U)     Eth      STATIC    1/1/37(P) 1/1/38(P) 1/1/39(P) 1/1/40(P)
MX9116-A1#
```

Configure the Link Aggregation Control Protocol

After the VLT is enabled, create the uplinks to the network layer above the MX9116N switches. The connections are a Link Aggregation Control Protocol (LACP) active link aggregation of two or more ports.

A VLT link aggregation is created by creating a VLT port channel on each of the MX9116N switches. First create the uplink VLT port channel on both switches and assign the appropriate VLANs.

Here are the port channels that are created on switch one:

```
MX9116-A1# show running-configuration interface port-channel 1
!
interface port-channel1
 description "Uplink to DataCenter"
 no shutdown
 switchport mode trunk
 switchport trunk allowed vlan 96, 1711-1714, 2711-2713
 mtu 9216
 vlt-port-channel 1

MX9116-A1# show port-channel summary

Flags:  D - Down    I - member up but inactive    P - member up and active
        U - Up (port-channel)    F - Fallback Activated
--------------------------------------------------------------------------------
Group Port-Channel           Type     Protocol  Member Ports
--------------------------------------------------------------------------------
1    port-channel1    (D)     Eth      DYNAMIC   1/1/41(P) 1/1/42(P)
1000 port-channel1000 (U)     Eth      STATIC    1/1/37(P) 1/1/38(P) 1/1/39(P) 1/1/40(P)
MX9116-A1#
```

Here are the port channels that are created on switch two:

```
MX9116-A1# show running-configuration interface port-channel 1
!
interface port-channel1
 description "Uplink to DataCenter"
 no shutdown
 switchport mode trunk
 switchport access vlan 1
 switchport trunk allowed vlan 96, 1711-1714, 2711-2713
 mtu 9216
 vlt-port-channel 1
```

Configure the host facing ports

To support multiple VLANs, you must place the server facing ports in trunk mode. All the VLANs assigned to the ports are tagged to allow the port groups to identify and direct traffic appropriately.

On switch one:

```
MX9116-A2(config)#
MX9116-A2(config)# interface range ethernet 1/1/1-1/1/16
MX9116-A2(conf-range-eth1/1/1-1/1/16)# switchport mode trunk
MX9116-A2(conf-range-eth1/1/1-1/1/16)# switchport trunk allowed vlan 96, 1711-1714,
2711-2713
MX9116-A2(conf-range-eth1/1/1-1/1/16)# mtu 9216
MX9116-A2(conf-range-eth1/1/1-1/1/16)# no shutdown
MX9116-A2(conf-range-eth1/1/1-1/1/16)# exit
MX9116-A2(config)#
```

On switch two:

```
MX9116-A2(config)#
MX9116-A2(config)# interface range ethernet 1/1/1-1/1/16
MX9116-A2(conf-range-eth1/1/1-1/1/16)# switchport mode trunk
MX9116-A2(conf-range-eth1/1/1-1/1/16)# switchport trunk allowed vlan 96, 1711-1714,
2711-2713
MX9116-A2(conf-range-eth1/1/1-1/1/16)# mtu 9216
MX9116-A2(conf-range-eth1/1/1-1/1/16)# no shutdown
MX9116-A2(conf-range-eth1/1/1-1/1/16)# exit
```

Save your switch configuration on each MX9116N switch using the following command:

```
MX9116-A2(config)# do write mem
```

Verify switch configuration

The MX9116N switches are now configured with the minimum required settings to support the deployment of Cloud Foundation. Each of the following should now be configured:

- All required VLANs and subnets
- VLT Domain
- VLTi
- VLT Link Aggregations
- Link Aggregations
- Host facing ports

10

SmartFabric network configuration

The PowerEdge MX9002m management module hosts the OpenManage Enterprise Modular (OME-M) console. Creation and deployment of SmartFabric topologies is facilitated using the OME-Modular console in conjunction with the MX9116n switch operating system. SmartFabric is a web-based mechanism to create a reusable networking template that can be applied to a PowerEdge MX7000 chassis, the IO modules (switches) and the compute sleds.

SmartFabric creates and configures the switches based on networking best practices. By selecting the topology, SmartFabric creates the VLT domain and VLTi connections and creates the uplink LACP enabled link aggregations. It assigns VLANs to ports as either tagged or untagged based on the networks that are created by the administrator through the SmartFabric interface.

NOTE: For detailed instructions to deploy a SmartFabric, see www.dell.com/networking.

To deploy a SmartFabric, complete the following steps using the OME-Modular console:

1. Create chassis groups.
2. Create the networks to be used in the fabric.
3. Select IOMs and create the fabric that are based on the required physical topology.
4. Create uplinks from the fabric to the existing network and assign networks to those uplinks.
5. Deploy the appropriate server templates to the compute sleds.

Topics:

Create chassis groups

About this task

CAUTION: The firmware version of the PowerEdge MX9002m management module is critical to the creation of a chassis group. Ensure that the latest version of the Management Module firmware is used. It is available on www.dell.com/support. All MX9002m modules in a chassis group must run the same firmware.

Chassis groups are defined in the PowerEdge MX9002m management modules in the PowerEdge MX7000 chassis. Before a chassis group can be created, the chassis PowerEdge MX9002m modules must be cabled together as shown in the following figure.

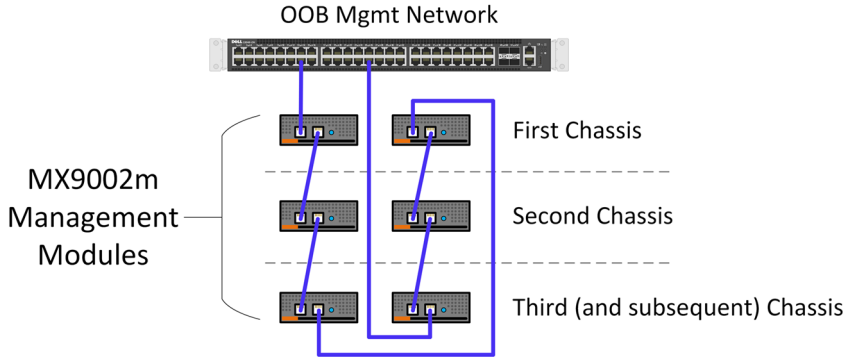

Figure 22. Chassis PowerEdge MX9002m cabling

Steps

1. After the PowerEdge MX9002m modules have been cabled together, log in to the OME-Modular web interface of the chassis that will be the lead chassis of the new chassis group.
2. From the **Chassis Overview** menu, click **Configure**, and then select **Create Chassis Group**.
3. Enter the group name and group description.

   NOTE: The group name must be one word without any spaces.

4. Select chassis onboarding permissions that propagate to each chassis that is added to the group, and then click **Next**.
   The **Add Members** page displays the chassis that can be added as members to the group. The lead chassis is not listed here.
5. Select the chassis to be added, and then click **Add Chassis**.
   The selected chassis moves to the **Current Members** list.
6. Click **Finish**.
   A chassis group is created and the selected members are added.
7. To view the list of jobs that is started by the group creation, click **Monitor**, and then click **Jobs**.
8. To view the group members, click **OME-M home**.

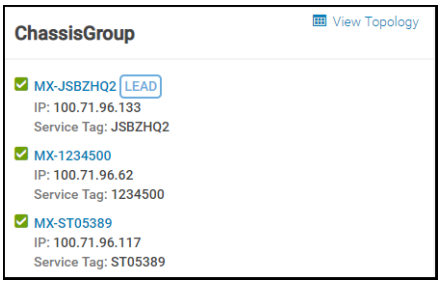

Figure 23. Chassis group members

9. To view the topology, select **View Topology** in the **ChassisGroup** pane.

Define networks

Prerequisites

Networks or subnets should be defined to meet the Cloud Foundation requirements. The prerequisites to define networks are as follows:

1. Create chassis groups
2. Cloud Foundation management network
3. vMotion network
4. vSAN network
5. VXLAN network
6. Uplink1 network
7. Uplink2 network

To define networks using the OME-Modular console, perform the following steps:

Steps

1. Open the **OME-M** console.
2. From the **Configuration** menu, click **Networks**.
3. In the **Network** pane, click **Define**.
4. In the **Define Network** window, complete the following:
   a. Enter the name of the network.
   b. Optionally, enter the description in the **DESCRIPTION** box.
   c. Enter the value **1711** in the **VLAN ID** box.
   d. From the **Network Type** list, select **General Purpose (Bronze)**.

e. Click **Finish**.

5. Repeat steps 1-4 to create the remaining six VLANs and any other VLANs required.
   A sample completed configuration is shown in the following figure:

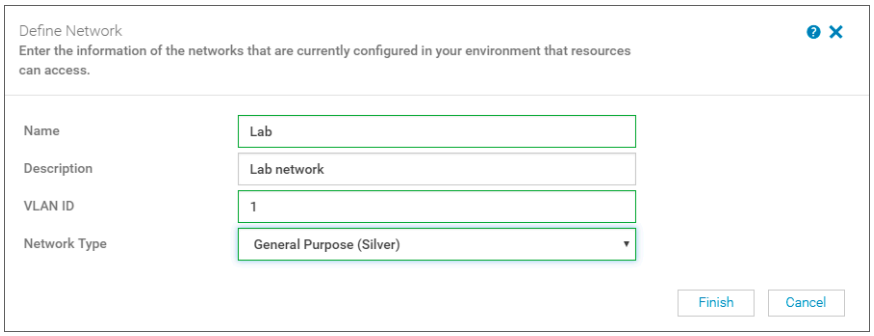

Figure 24. VLAN configuration

Create SmartFabric

Creation of the SmartFabric depends on the IOM selected and the number of PowerEdge MX7000 chassis to be installed. The devices eligible for SmartFabric deployment are:

- MX5108n Ethernet switch
- MX9116n Fabric Switching Engine
- MX7116n Fabric Expansion Module

When deploying the MX5108n switch, create chassis groups for improved management but the IOMs in each chassis function independently of the IOMs in other chassis. Networks and templates that are created for the chassis groups can be applied to multiple MX5108n switches for fast, reliable deployment of up to 20 PowerEdge MX7000 chassis containing up to 40 MX5108n Ethernet switches.

When deploying the MX9116n Fabric Switching Engine (FSE) and the MX7116n Fabric Expansion Module (FEM), the configured FSE can manage FEMs across up to 10 PowerEdge MX7000 chassis. The FSE is single point of management and configuration and the FEMs extend the configured ports to the local PowerEdge MX7000 chassis.

Create SmartFabric using MX9116n Fabric Switching Engine IOMs

About this task

To create a SmartFabric using the OME-M console, perform the following steps:

Steps

1. Open the OME-M console.
2. From the **Devices** menu, click **Fabric**.
3. In the **Fabric** pane, click **Add Fabric**.
4. In the **Create Fabric** window, complete the following:
   a. Enter SmartFabric in the **Fabric Name** box.
   b. Optionally, enter the description in the **Description** box.
   c. Click **Next**.
5. Based on your IOMs and number of chassis select from the **Design Type** list:
   a. 2x MX5108n Ethernet Switches in same chassis
   b. 2x MX9116n Fabric Switching Engines in same chassis
   c. 2x MX9116n Fabric Switching Engines in different chassis

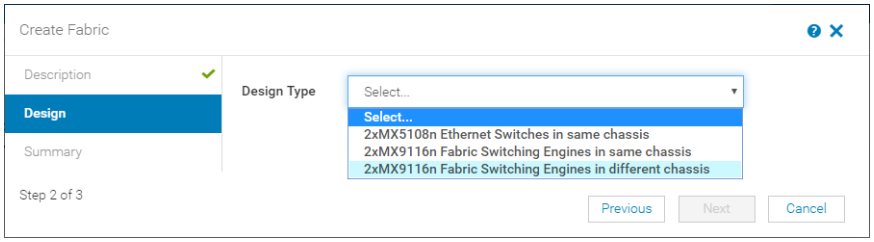

Figure 25. Create SmartFabric using MX9116n Fabric Switching Engine IOMs

6. From the **Chassis-X** list, select the first PowerEdge MX7000 chassis containing an MX9116n FSE.
7. From the **Switch-A** list, select **Slot-IOM-A1**.
8. From the **Switch-B** list, select **Slot-IOM-A2**.
9. Click **Next**.
10. On the **Summary** page, verify the proposed configuration, and then click **Finish**.
    The fabric displays a health error which is resolved in the next section by adding uplinks to your fabric.

Configure uplinks

About this task

The newly created fabric requires uplinks to connect to the rest of the network. These uplinks are created as a single logical link to the upstream network using the same Virtual Link Trunking (VLT).

Perform the following steps to configure the uplinks:

Steps

1. From the **Devices** menu, click **Fabric**.
2. Click **SmartFabric**.
3. In the **Fabric Details** pane, click **Uplinks**.
4. Click **Add Uplinks**.
5. In the **Add Uplink** window, complete the following steps:
   a. Enter the name in the **Name** box.
   b. Enter the description in the **Description** box.
   c. From the **Uplink Type** list, select **Ethernet**.
   d. Click **Next**.
   e. From the **Switch Ports** list, select the appropriate Ethernet ports:
      i. Ethernet 1/1/11 from both MX5108n IOMs
      ii. Ethernet 1/1/41 and Ethernet 1/1/42 from both MX9116n IOMs
   f. From the **Tagged Networks** list, select all tagged **Cloud Foundation VLANs**.
   g. From the **Untagged Network** list, select any required untagged VLAN.
   h. Click **Finish**.
6. To verify the switch topology, from the **Devices** menu, click **Fabric**, and then perform the following steps:
   a. Click the fabric that was created.
   b. Click **Topology**.
   Make sure the displayed topology matches the desired configuration.:

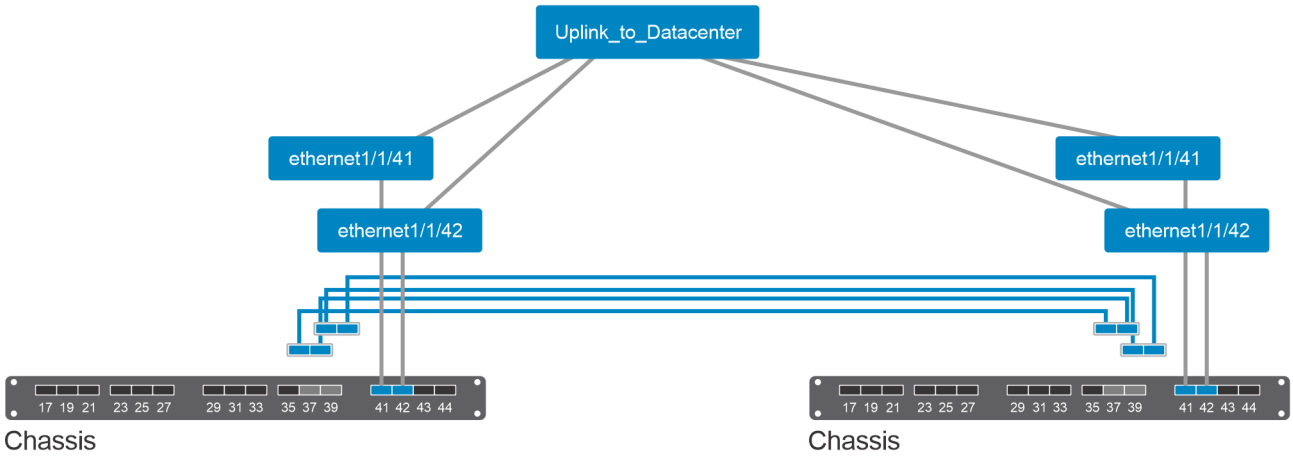

Figure 26. Topology of the SmartFabric using MX9116n Fabric Switching engine IOMs

Configure jumbo frames

About this task

Cloud Foundation requires jumbo frames on all links.

NOTE: By default, SmartFabric does not configure the jumbo MTU (frame size) on switch ports.

To configure jumbo frames, set the MTU (frame size) using the following procedure:

Steps

1. From the **Devices** menu, click **I/O Modules**.
2. Select the **IO Module**.
3. From the **IOM banner** menu, click **Hardware**.
4. Click **Port Information**.
5. Select ports Ethernet 1/1/1-Ethernet 1/1/16 and the uplink port channels.
6. Click **Configure MTU** and set MTU to 9216.
7. Click **Finish**.

Results

Jumbo frames that are required for the Cloud Foundation are configured.

Server templates

A server template contains the parameters that are extracted from a reference server and allows the parameters to be quickly applied to multiple compute sleds. First, create an empty template by copying one from any existing compute sled. Then, modify the template and then apply it to new compute sleds as needed.

Create a server template

About this task

To create a server template, perform the following steps:

Steps

1. Open the OME-M console.
2. From the **Configuration** menu, click **Deploy**.
3. In the **Center pane**, click **Create Template**, and then click **From Reference Device**.
4. In the **Create Template** window, complete the following steps:
   a. In the Template Name box, enter **`MX740c with Intel mezzanine`**.
   b. Optionally, enter the description in the Description box.
   c. Click **Next**
   d. In the **Device Selection** pane, click **Select Device**.
   e. In the **Select Devices** window, select **Sled-1** from **Chassis-1**
   f. Click **Finish**.
   g. From the **Elements to Clone** list, select the **iDRAC**, **System** and **NIC** options
   h. Click **Finish**.

Associate server template with a VLAN

Prerequisites

Before the server template is associated with VLAN, a server template must be created. For more information about creating the server template, see Create a server template.

About this task

To associate the server template with VLAN, perform the following steps:

Steps

1. In the **Deploy** pane, select the **MX740c with Intel mezzanine** server template and then click **Edit Network**.
2. In the **Edit Network** window, perform the following steps:
   a. Optionally, from the **Identity Pool** list, select **Ethernet ID Pool**.
   b. From the **Untagged Network** list, select the VLAN previously created to be the untagged VLAN for both ports.
   c. From the T**agged Network** list, select all the tagged VLANs for both ports.
3. Click **Finish.**

Results

The server template is associated with the VLAN network.

Deploy the server template

Prerequisites

Before deploying the server template, you must associate the server template with a VLAN.

About this task

To deploy the server template, perform the following steps:

Steps

1. In the **Deploy** pane, select the **MX740c with Intel mezzanine** server template, and then click **Deploy Template**.
2. In the **Deploy Template** window, complete the following steps:
   a. Click **Select** to choose the slots or compute sleds to deploy the template to.
   b. Select **Do not forcefully reboot the host OS**
   c. Click **Next.**.
   d. Select **Run Now.**.
   e. Click **Finish.**

Results

The interfaces on the switches are updated automatically. SmartFabric configures each interface with an untagged VLAN and tagged VLANs.

11

Deploy ESXi to cluster nodes

Only perform the steps listed in this section if the compute sleds were not pre-installed with ESXi 7.0. If the compute sleds have been pre-installed with ESXi 7 jump ahead to Configure ESXi settings—using DCUI. Below are the steps to install VMware ESXi on each of the PowerEdge MX740c hosts that are part of the management cluster. This guide covers the steps to install VMware ESXi remotely using iDRAC Virtual Console with Virtual Media. In this example, a static IP address is assigned to the management interface of the ESXi hosts, which is required for Cloud Foundation.

NOTE: This guide assumes that the steps in this document are being followed comprehensively and sequentially. The tasks of previous sections are prerequisites for this section.

Topics:

Prerequisites

The following items are required to complete this section of the deployment guide:

- OME-Modular and iDRAC IP addresses or FQDNs
- OME-Modular and iDRAC credentials
- iDRAC Enterprise license that is applied on all nodes
- Dell EMC customized ESXi image
- Host names, Management VLAN ID, IP address information
- Credentials for vSphere
- Hostnames added to DNS server

Installation of ESXi

Complete the following steps on each physical compute node targeted for Cloud Foundation deployment before moving on to the next section.

Connect to iDRAC and boot installation media

Prerequisites

- The virtual console should be in HTML5 mode, which is the default setting.
- The location of the Dell EMC customized ESXi image (ISO image) file should not be changed during the installation process.

Steps

1. Using a web browser, go to Open Manage Enterprise Modular (embedded chassis management) web interface at **https://<OME Modular Address>**.
2. Log in with your credentials.

   NOTE: The default user name is **root** and the password is **calvin**.

3. From the **Devices** menu, click **Compute**.
4. Select the required compute node.
5. Click **Launch Virtual Console**, and then enable the support in the pop-up window for each iDRAC device.

NOTE: When the virtual console is launched for the first time, repeating this step may be necessary due to browser pop-up blocker.

6. The mapping screen for the virtual media is displayed on the **Virtual Media** menu.
7. In the **Map CD/DVD** section, click **Choose File**.
8. Browse and select the required Dell EMC customized ESXi image (ISO image) file.
9. Click **Map Device** and then click **Close**.
10. From the **Virtual Console** menu, click **Boot**, and then click **Virtual CD/DVD/ISO**.
11. Click **Yes**.
12. From the **Power** menu, click **Power on System**.
13. If the system is not turned on, click **Power on System**. If the system is ON, click **Power Cycle System (cold boot)**.

Results

The server is connected to the iDRAC devices and boots into the ESXi installer.

Install VMware ESXi

Prerequisites

Before installing VMware ESXi, you must connect to the iDRAC devices and boot into the ESXi installer. For more information, see Connect to iDRAC and boot installation media.

Steps

1. In the **Welcome to ESXi Installation** window, press Enter.
2. Review the End User License Agreement (EULA), and then press F11 to accept and continue.
3. On the **Select a Disk to Install or Upgrade** page, select the **DELLBOSS VD** device, and then press Enter.

   NOTE: If DELLBOSS VD device has been used for previous ESXi installation, in the **ESXi/VMFS Found** dialog box, choose to overwrite the existing VMFS datastore, and then press Enter.

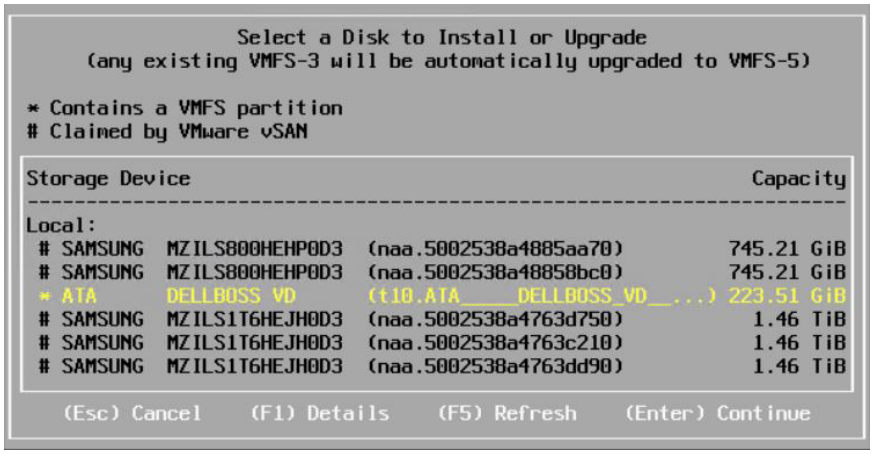

Figure 27. ESXi disk partition page

4. Select the required keyboard layout, and then press Enter.
5. Enter the root password, and then press Enter.
6. In the **Confirm Install** window, press F11 to install the VMware ESXi.
7. In the **Installation Complete** window, press Enter to reboot the server.
   The installation completes and the server boots into ESXi.
8. From the **Virtual Media** menu, click **Disconnect Virtual Media**.

Results

VMware ESXi is installed in the server.

Configure ESXi settings—using DCUI

About this task

The Direct Console User Interface (DCUI) is a menu-based interface that is accessed from the host console and used to configure ESXi running on vSphere hosts.

Steps

1. After the server reboots and fully loads ESXi, press F2 to log in to the DCUI.
2. Enter the credentials that were created during the ESXi installation, and then press Enter.
3. From the **System Customization** menu, select **Configure Management Network**.
4. From the **VLAN (Optional)** menu, press Enter.

   NOTE: Step 4 is mandatory although the name of the menu item includes the word optional.

5. Enter the required **management VLAN ID**, and then press Enter.
6. Select **IPv4 Configuration** and press Enter.
7. Select **Set static IPv4 address** and press the spacebar.

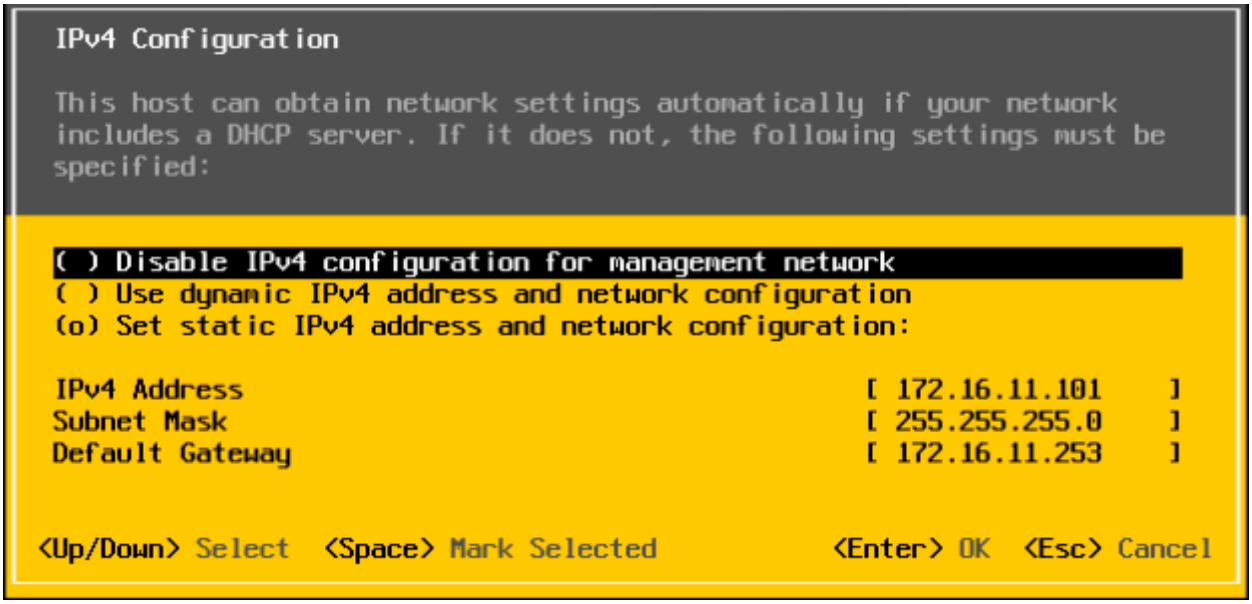

Figure 28. IPv4 configuration page

8. Enter the **IPv4 Address**,**Subnet Mask**, and the **Default Gateway**, and then press Enter to confirm.
9. Select **DNS Configuration**, and then press Enter.
10. Enter the IP addresses of the DNS servers and FQDN of the host.
11. Press Esc to return to the main menu, and then press Y to confirm the changes and restart the management network.
12. From the main menu, click **Test Management Network**.
    The target IP addresses and DNS hostname are pre-populated.
13. Press Enter to perform the network test, and after the test is completed, press Enter to return to the main menu.

    CAUTION: If the network test fails, troubleshoot and resolve the issues before proceeding further.

14. From the main menu, select **Troubleshooting Options** > **Enable ESXi Shell**, and then select **Enable SSH** (required during validation and deployment phases) to enable the ESXi shell.
15. Press Esc to return to the main menu.

Configure ESXi settings using web interface

Prerequisites

Before configuring ESXi settings using web interface, you must configure the ESXi settings using DCUI. For more information, see Configure ESXi settings—using DCUI.

Steps

1. Using a web browser, go to the ESXi host-level management web interface at **https://<ESXi Host Address>/ui**.
2. Enter the credentials that were created during the ESXi installation, and then click **Log in**.
3. In the **Navigator** pane, click **Networking**.
   The current port groups on the host are displayed by default. One is configured with VLAN information that is entered during installation and the other is still at zero.
4. Right-click **VM Network** and then click **Edit settings**.

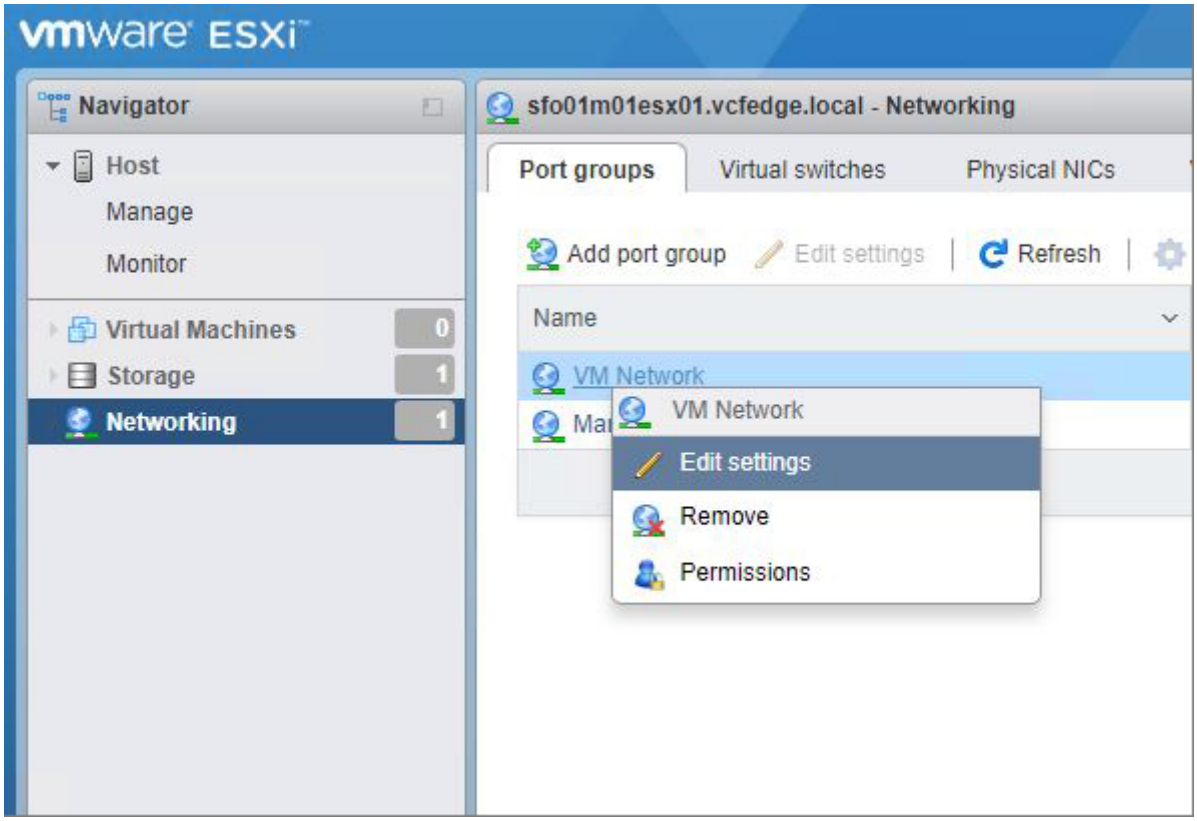

Figure 29. ESXi web interface—Edit settings page

5. In the **Edit Port Group** window, enter the **Management VLAN ID**, and then click **Save**.

   CAUTION: Leaving the VLAN ID at default setting causes pre-deployment validation to fail during a later step.

6. In the **Navigator** pane, click **Manage** to set up the NTP.
7. In the right pane, click **Time & Date**.
8. Click **Edit Settings** and then select **Use Network Time Protocol (enable NTP client)**.
9. In the **NTP Servers** box, enter the NTP server IP addresses.

   NOTE: If multiple IP addresses are provided, separate the IP addresses with commas.

10. Click **Save**.

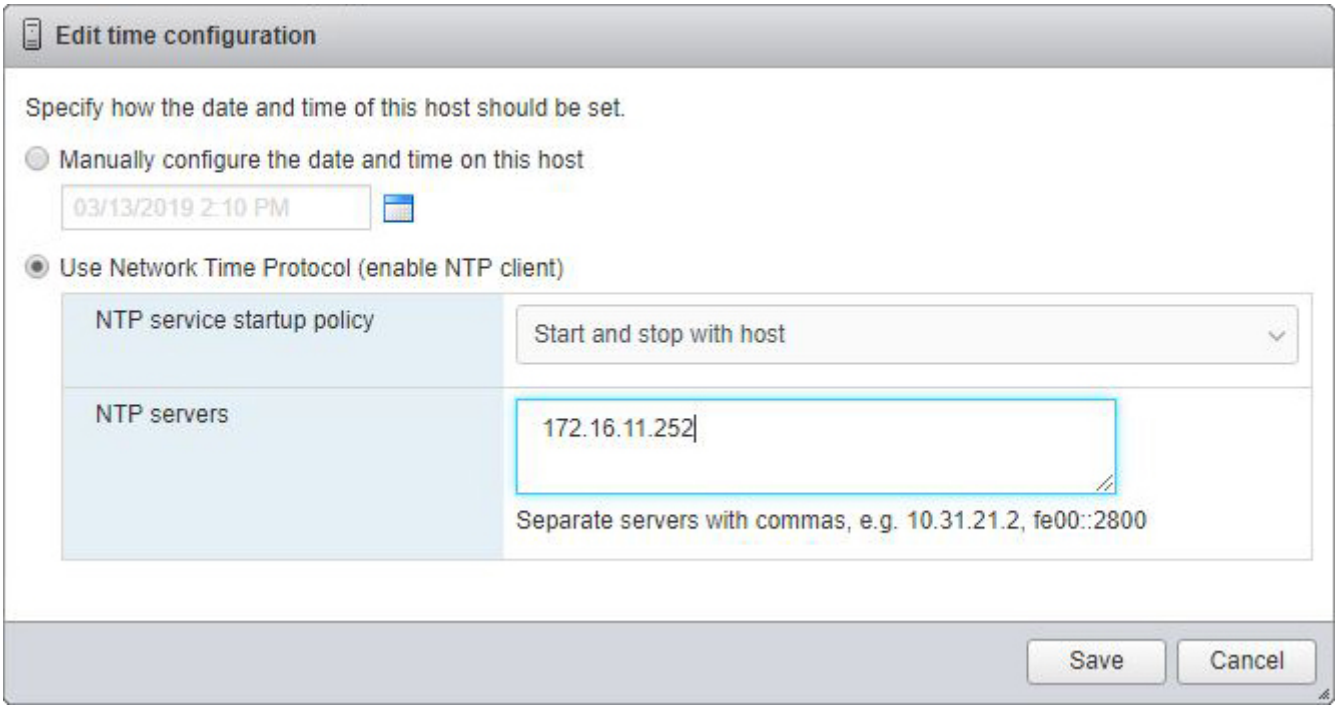

Figure 30. ESXi web interface—Edit time configuration page

11. In the **Manage** pane, select the **Services** tab.
    The resulting page is as shown in the following figure:

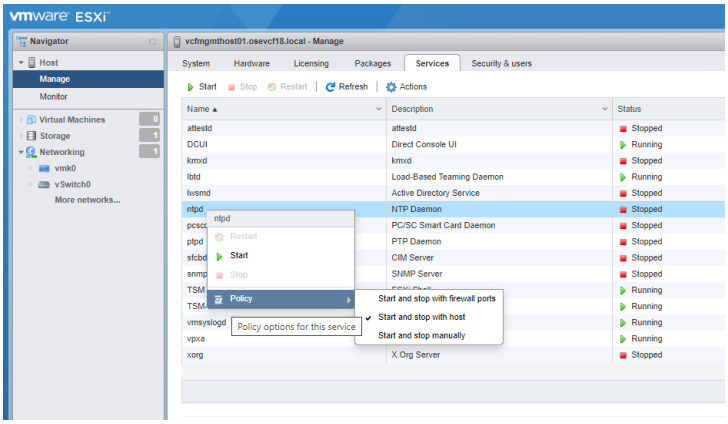

Figure 31. ESXi settings web interface—Manage pane

12. Right-click on the **ntpd** service and set the policy to **Start and stop with the host**.

Next steps

Once the policy is set, start the ntpd service. If the ntpd service is already running, restart the service.

Repeat all the steps for each host targeted for Cloud Foundation management domain deployment. Validate that each ESXi host can access the NTP servers by establishing an SSH connection to each host and executing the `ntpq -p` command. A refid value of **`.INIT.`** indicates that your host is not obtaining time from your specified time source.

12

Cloud Builder and SDDC deployment

The primary software installation tool for Cloud Foundation 4.x is Cloud Builder. It is delivered as a virtual appliance in the standard OVA format. This section describes the steps to deploy the OVA. The Cloud Builder VM is a temporary tool to facilitate deployment of Cloud Foundation. It can be discarded after the deployment.

Topics:

Deploy Cloud Builder

Prerequisites

The prerequisites to deploy Cloud Builder are as follows:

- Cloud Builder OVA file with version and build numbers that are specified in Validated Components section
- The following must be deployed to an ESXi host outside of the Cloud Foundation target nodes and have network access to the Cloud Foundation Management VLAN on which the Cloud Builder VM resides
- All prior sections are reviewed and any listed steps completed

Steps

1. Using a web browser, go to the vCenter web interface at **https://<vCenter Address>**.
2. In the **Navigator** pane, on the **Hosts and Clusters** tab, locate an available ESXi host.
3. Right-click the ESXi host and select **Deploy OVF Template**.
4. Locate the OVA file locally or from the URL, and then click **Next**.
5. Select the required ESXi server to host the Cloud Builder VM, and then click **Next**.

   NOTE: The selected EXSi server should not be targeted for Cloud Foundation deployment.

6. Click **Next**.
7. Review the EULA, and if you agree, click **Accept**, and then click **Next** to continue.
8. Select the required datastore and then click **Next**.
9. Select the required network and then click **Next**.
10. In the **Customize Template** page, enter username, password, network 1 IP address and subnet mask, default gateway, hostname, DNS, and NTP resources in the appropriate fields, and then click **Next**.

    CAUTION: You must adhere to the following password rules that are listed on the screen:

    - **Noncompliant passwords cause an unrecoverable error when attempting to access and use the Cloud Builder VM in the next section of the document.**
    - **The password rules vary for different user accounts. For example, minimum 8-character rule is applicable to some users and minimum 12-character rule is applicable to other users.**
    - **You must not enter dictionary words as passwords.**

    NOTE: IP Address and related fields must align with the IP subnet of the Cloud Foundation Management VLAN.

    CAUTION: You must adhere to the following password rules that are listed on the screen:

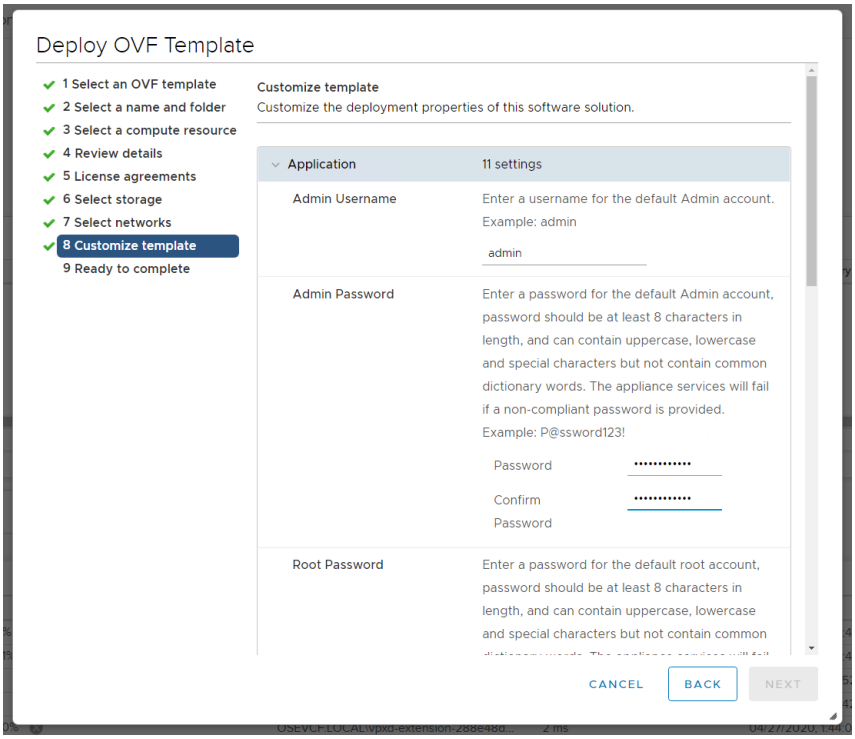

Figure 32. OVF customize template page

11. Review the **Ready to Complete** final configuration page, and then click **Finish**.
12. In the **Recent Tasks** pane, check the OVA deployment status.
    When the OVA deployment is complete start the Cloud Builder VM..

Check Time Synchronization

After the Cloud Builder VM is started, it takes some time to for all the services to start and for time synchronization to complete. It is recommended that you access the command line of the Cloud Builder VM and verify time synchronization status using Linux `ntpq` commands.

```
root@vcloudbuilder# ntpq -p
     remote           refid      st t when poll reach   delay   offset  jitter
==============================================================================
*100.71.100.2    143.166.255.32   2 u  113 1024  377    0.218   -0.755   0.037
+ntp-cent-3.osev 143.166.226.32   2 u  510 1024  377    0.205    0.612   0.443
```

13

VCF Deployment using Cloud Builder

In the previous section, you deployed the Cloud Builder virtual appliance. In this section, the software within the virtual machine is used to validate the target environment and deploy the entire Cloud Foundation stack.

NOTE: Before proceeding with the Cloud Builder validation process, take a snapshot of your Cloud Builder VM.

Topics:

Prerequisites

The prerequisites to deploy the Cloud Foundation stack are as follows:

- Cloud Builder OVA file with version and build numbers specified in the Validated components section must be deployed to an available ESXi host outside of the Cloud Foundation target nodes with network access to the Cloud Foundation Management VLAN.
- Target hardware certified on VMware Compatibility Guide (VCG) with appropriate firmware versions are specified in the Validated components section.
- Clean install of ESXi on each target server node, with SSH/NTP/DNS enabled and configured.
- Physical switches configured with jumbo frames and necessary VLANs.
- All prior sections are reviewed and any listed steps completed.
- A completed parameter sheet (can be downloaded during the CloudBuilder process

Launch Cloud Builder web interface

About this task

Complete the following steps to launch the Cloud Builder web interface:

Steps

1. Turn on **Cloud Builder VM**.
   The VM must finish booting up and load the application stack.
   NOTE: You can monitor the progress of Cloud Builder VM from a VM console session.
2. Using a web browser, go to the Cloud Builder web interface at **https://<Cloud Builder IP Address>**.

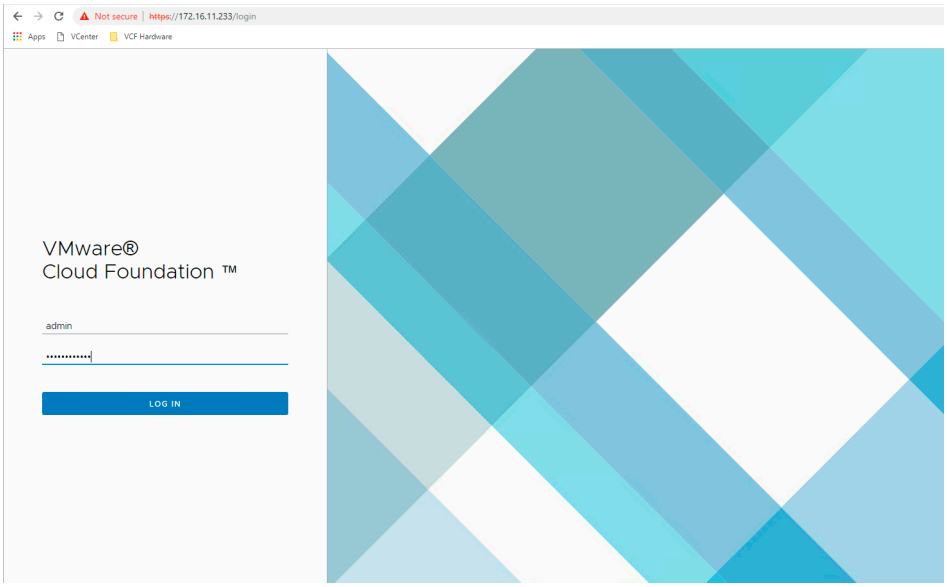

Figure 33. Cloud Builder web interface

3. Log in using the credentials that you specified during OVA deployment.
4. Click **Check All** to review the checklist of pre-bring-up steps and confirm that all the steps that are completed, and then click **Next**.
5. Review the EULA, and if you agree, click **Agree to End User License Agreement**, and then click **Next**.
6. If you have not obtained and completed the Cloud Foundation Information Spreadsheet, click **Download Deployment Parameter Sheet**.

Cloud Builder Deployment Parameter Sheet

All parameters that are needed for Cloud Builder to deploy your Cloud Foundation stack are passed to the tool using the deployment parameter sheet. It is critical that you complete the sheet fully and accurately. The configuration details that you enter into the deployment parameter sheet should be tied to existing services and records before proceeding. For example, the DNS servers you specify in the deployment parameter sheet must correspond to DNS servers running in your infrastructure.

NOTE: This guide covers the process to deploy Cloud Foundation using a spreadsheet, for example, .xlsx format. However, Cloud Builder also supports the use of a .json file for uploading deployment parameters. The latter is outside the scope of this document.

This worksheet is incorporated into the Cloud Builder VM and specific version of Cloud Foundation that is deployed. Save the worksheet to a safe location and edit the file modifying the information to match your deployment environment.

Note the following points before deploying Cloud Foundation using the deployment parameter sheet:

- Do not deploy Cloud Foundation using the deployment parameter sheet intended for a different version of Cloud Foundation.
- Any parameter field with a yellow background indicates that data format validation rules for that field have been specified.
- If when adding information to the parameter sheet your font turns red, you have not entered the right information.
- Do not copy and paste data between fields as different fields may have different requirements.
- Copying fields can change or delete the specified input data validation for that data.

Cloud Builder parameters

Some of the Cloud Builder parameters are:

- Management Workload tab
- Users and Groups tab
- Hosts and Networks tab
- Deploy Parameters tab

Management Workload tab

License keys are required for the following items:

- ESXi hosts
- vSAN
- vCenter
- NSX-T
- SDDC Manager Appliance

Users and Groups tab

In the **Users and Groups** tab, you can set the passwords for your initial Cloud Foundation components.

CAUTION: Do not make a mistake on this page because if any of the passwords do not meet the indicated specifications, you must redeploy your Cloud Builder VM, unless you elected to create a snapshot after you created your VM.

NOTE: Minimum password length should be 12 characters for NSX controllers.

Hosts and Networks tab

In the **Hosts and Networks** tab, perform the following steps:

- Specify the VLAN IDs, subnets, gateways, and MTU for the networks that are used for the four required Cloud Foundation networks.
  - The MTU value should be the MTU size that can be tested during validation. For an actual MTU value of 9000 set this value to 8972
- Specify the management IP addresses of the deployed hosts which are on the Cloud Foundation management VLAN.
- Specify the range of IP addresses for both the vSAN and vMotion networks.
- Host names that you enter in the **Management Domain ESXi Hosts** section are also auto-populated in the **ESXi Host** column of the **ESXi Host Security Thumbprint** section.

By default, use of DSA fingerprints as a more secure method to validate a host's authenticity is optional and disabled. Unique fingerprints are automatically generated on each host during installation. You can view the fingerprints using the **DCUI** under **View Support Information**. Due to the complexity of fingerprints, it is easier to access the hosts using SSH where they can be copied and pasted into the cloud builder deployment parameter sheet. Change the **Validate ESXi Thumbprints** option to **Yes** to enable fingerprint validation.

Deploy Parameters tab

The **Deploy Parameters** tab contains all the IP addresses and hostnames that must be added to DNS.

- All the hostnames must be configured for forward and reverse look-ups.
- Test all the hostnames before attempting validation.
- Ensure that all the values provided are consistent with the values used to deploy the Cloud Builder VM.
- The NTP server would have already been tested as the ESXi hosts have already been configured to this same NTP server(s).
- Change the hostnames that are listed and the domain that meets the customer required specifications.
- For a simplified network installation where a multi-tiered BGP environment is not available make sure that **field K27** is set to **No**.

Run Cloud Builder Deploy SDDC

About this task

Complete the following steps to launch the Cloud Builder web interface:

Steps

1. Turn on Cloud Builder VM.
2. The VM must finish booting up and load the application stack.
3. Using a web browser, go to the Cloud Builder web interface at https://<Cloud Builder IP Address>.
4. Log in using the credentials that you specified during OVA deployment.
5. Review the **EULA**, and if you agree, click **Agree** to End User License Agreement, and then click **Next**.
6. Select **VMware Cloud Foundation**. Be sure not to select VMware Cloud Foundation on VxRAIL.
7. Review the list of **Prerequisites** and ensure that they have all been completed
8. **Place a checkmark** to indicate that the prerequisites have been met and select **Next**
9. If you have not obtained and completed the Cloud Foundation Information Spreadsheet, click **Download Deployment Parameter Sheet**. If you have the parameter sheet click **Next** to move forward
10. If you have not completed your parameter sheet do it now. If the parameter sheet is complete, click **Next** to continue.
11. Click **Select File** to browse to your completed parameter sheet.
12. Locate your completed parameter sheet and click **Open**.
13. Click **Next** to begin the validation phase.
    If validation discovers any issues in your configuration or parameter file it will show an error and provide clear and detailed information on the issue. Read the message provided, correct the issue(s) and run the validation again. Validation can be run as many times as necessary to obtain a successful result.
14. Once Validation has passed, and you have the **green banner** indicating that 'Configuration file validated successfully', click **Next**
15. A pop-up will appear on the screen to Deploy SDDC Manager.
16. Click **DEPLOY SDDC** to continue
    The deployment process will begin. The deployment will take in excess of an hour. A successful deployment will result in the **SDDC Deployment Complete** pop-up message. If there were any issues you can identify them in the **VMware Cloud Foundation** deployment screen.
17. Review any task failures and correct the issue. You can restart the deployment when you have resolved the issue(s).
18. Click **LAUNCH SDDC MANAGER**

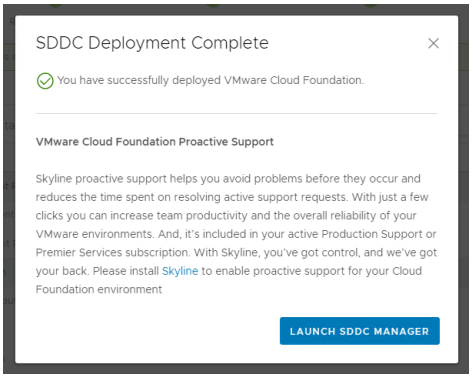

Figure 34. Pop-up window on successful SDDC deployment

Results

VMware Cloud Foundation on Dell EMC servers and networking hardware is now successfully deployed.

Cloud Builder Configuration Validation

About this task

The Cloud Builder Configuration Validation is a critical step in the Cloud Foundation deployment process. It probes your target servers, required services, and network environment to detect potential issues.

NOTE: The validation may fail initially and can be run as many times as necessary to address any issues.

Steps

1. After you have completed the deployment parameter spreadsheet, click **Upload**, select the file, and then click **Open**. A message is displayed to acknowledge successful upload of the parameter sheet.
2. Click **Validate** and monitor progress on the **Configuration File Validation** page.

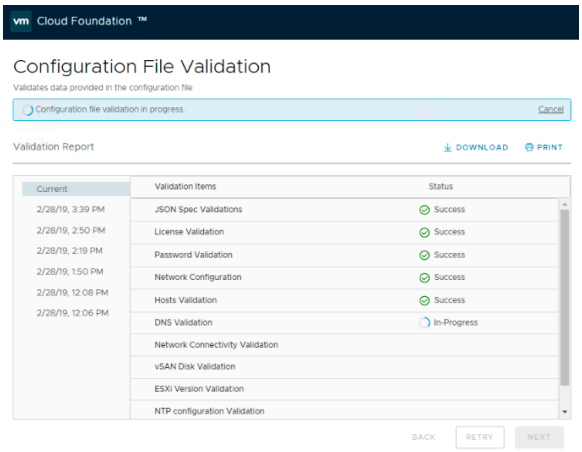

Figure 35. Configure Cloud Builder validation

NOTE: Validation may take 15 minutes or more. However, if there are issues such as the DNS server being down or if you provided a wrong IP address, validation may take longer.

NOTE: On the **Validation Report** page, you can access the information about previous validation attempts. Each validation attempt is tracked with an entry that is designated by the date and time of execution.

3. If the validation fails, expand the failed line item and review the detailed error messages. Some of the possible validation failures are shown in the following example.

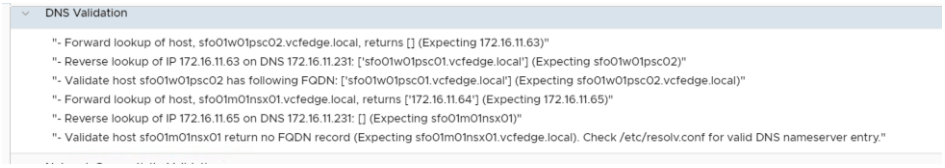

Figure 36. Configuration validation error messages

4. Resolve all the reported failures within the targeted servers, required services, network environment, or parameter spreadsheet, and then re-run validation in Cloud Builder by clicking **Retry**.
5. Repeat Steps 1-4, until the validation is successful.

SDDC bring-up

About this task

Perform the following steps for the SDDC bring-up:

Steps

1. In the **Bringing Up the SDDC** page, click **Next** to start the deployment of Cloud Foundation.

   NOTE: You can monitor the deployment progress on the **Bringing Up the SDDC** page and this deployment process may require two hours or more, depending on your environment.

2. If a bring-up failure occurs, expand the failed line item, review the detailed error messages. Depending on the nature of the error, resolve all the reported failures and click **Retry**.

NOTE: If an unrecoverable failure occurs during SDDC bring-up, it is necessary to resolve the root cause issue, wipe the Cloud Foundation target servers including all disk partitions, and then restart from the Deploy ESXi to cluster nodes section.

3. The SDDC bring-up process is completed when the Cloud Builder reports that the **SDDC has been successfully created**. When the SDDC bring-up process is complete, this indicates that the Cloud Foundation is successfully deployed.
4. Access the recently deployed SDDC Manager by clicking the link on the page.
5. Confirm deployment of vCenter, vSAN, NSX, and vRealize Log Insight.

   NOTE: Hostnames and IP addresses are in the Cloud Foundation Deployment Parameters spreadsheet.

   The Cloud Builder VM can now be discarded.

14

Post-install validation

Topics:

Cloud Foundation Cluster Verification

After installing Cloud Foundation, perform the steps in the following sections to verify that the components are installed and available.

SDDC Manager

Log in to SDDC Manager using a web browser at: **https://<ip address or DNS name>**. The SSO user ID is `administrator@vsphere.local` and the password is the one you specified during installation.

NOTE: Use the domain `vsphere.local`. Do not use the DNS domain that was used for the Cloud Foundation deployment.

Customer Experience Improvement Program

Join Customer Experience Improvement Program (CEIP) when prompted. CEIP provides statistical data to SDDC Manager, allowing for improved performance and troubleshooting.

The following figure shows a sample SDDC Manager dashboard. In this example, the host usage is yellow because the management cluster has consumed all of the hosts currently available to Cloud Foundation.

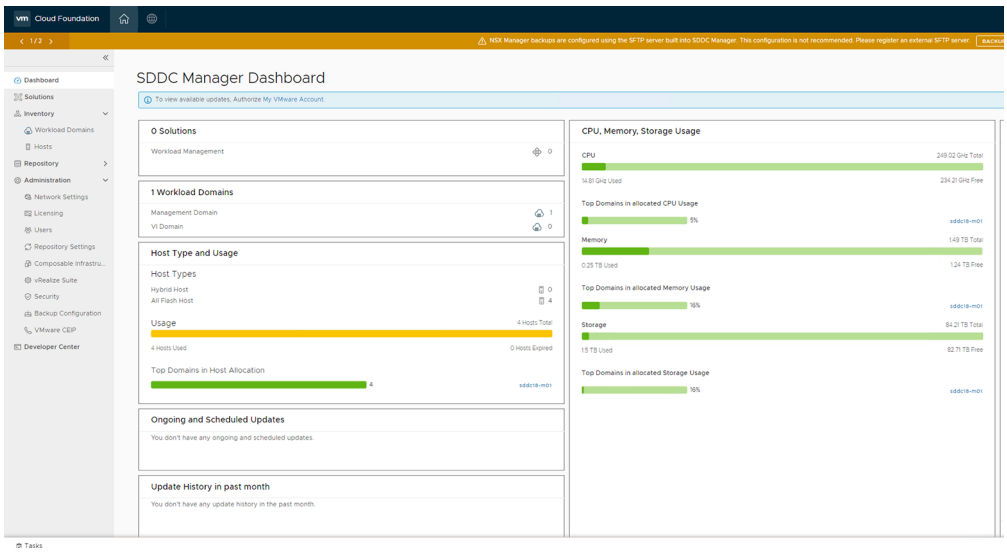

Figure 37. SDDC Manager dashboard

vCenter

Log in to vCenter using the same SSO credentials—`administrator@vsphere.local`. Verify that the cluster appears healthy and that there are no warnings.

Validate that the VMs created during deployment are all up and running. Validate that your VSAN is running and available. Look at the disk groups that are created and ensure they are consistent with the disks available in the hosts. Look at the virtual networking components that are deployed to vCenter.

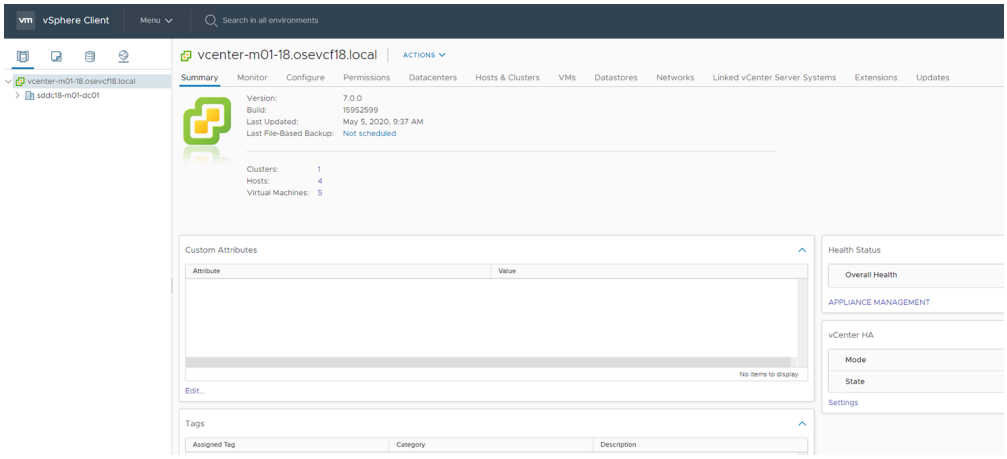

Figure 38. vCenter dashboard

NSX Manager

Log into the NSX Manager through a web browser using the Admin credentials set in your parameter sheet.

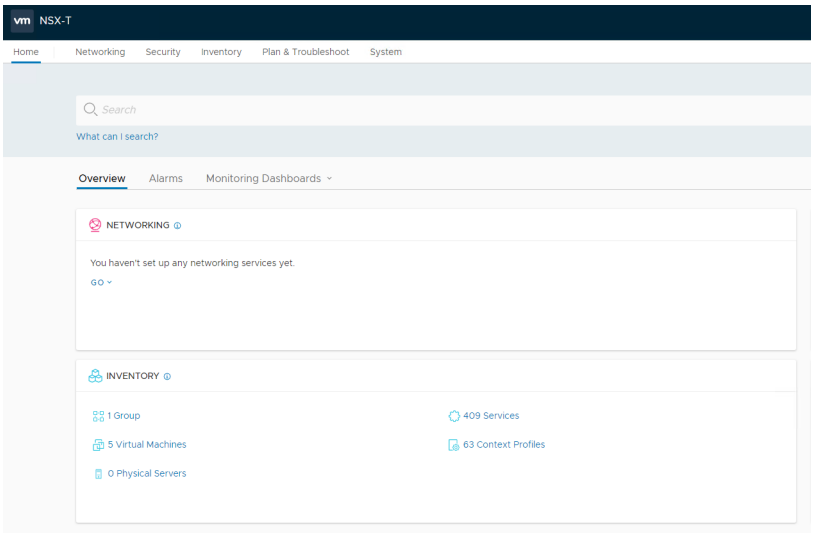

Figure 39. NSX dashboard

VMware Cloud Foundation installation complete

Cloud Foundation has been successfully deployed and is ready for use. Typical tasks from this point would be:

- Configure your VMware account credentials in SDDC manager
- If you chose simplified networking configure your SDDC for VLAN backed Application Virtual Networks
- Install additional products such as:
  - Life Cycle Manager
  - vRealize Suite
  - Other licensed bundles.
- Deploy a Workload Domain