

**Dell EMC Ready Architecture for VMware
vCloud NFV 3.2 vCloud Director 9.7**
Architecture and Manual Deployment Guide

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.

© 2019 Dell Inc. or its subsidiaries. All rights reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Overview.....	8
Intended audience.....	8
Acronyms and definitions.....	8
1 Deployment architecture for vCloud NFV.....	10
Architecture design.....	10
Solution bundle network topology.....	10
Solution bundle physical network design and topology.....	11
Solution bundle virtual network design and topology.....	14
Three pod configuration.....	17
Management pod.....	17
Edge pod.....	18
Resource pod.....	18
2 Solution hardware.....	19
Hardware installation and configuration.....	19
Network connectivity and port mapping.....	21
3 Manual deployment.....	24
Solution prerequisites.....	24
Deployment server.....	24
Create standard vSwitch on deployment server.....	27
Create port group on deployment server.....	28
Create datastore on deployment server.....	29
Connectivity overview for deployment VM and server.....	29
Deployment VM.....	29
Create deployment VM using CentOS.....	30
Install CentOS.....	30
Configure deployment VM IP.....	32
Enable automatic connectivity in Network Settings.....	32
Configure deployment settings on deployment VM.....	33
Configure time zone.....	33
Disable DHCP script from adding entries to resolv.conf.....	34
Disable auto mount on CentOS.....	34
Install Google Chrome.....	35
Install OVF tool.....	35
Add network adapters.....	36
Installation of VMRC.....	37
4 ESXi installation and configuration.....	38
Use iDRAC9 to install ESXi on PowerEdge R640, R740, and R740xd servers.....	38
Set VLAN ID for ESXi management network.....	38
Assign ESXi license.....	38
Set SSH policy.....	39

Set firewall rules.....	39
Install DPDK drivers.....	39
5 Auxiliary components.....	41
Install auxiliary components.....	41
Create datastore on first management ESXi server.....	41
Create standard vSwitch on first management ESXi server.....	41
Create port group on first management ESXi server.....	42
NTP server configuration.....	42
Synchronize ESXi clocks using NTP.....	44
Synchronize VM clock using NTP.....	45
Microsoft Windows Server 2016 installation for AD DNS.....	45
Install Microsoft Windows Server 2016.....	45
Active Directory and DNS installation.....	46
Update Windows VM computer name.....	46
Install primary Active Directory and DNS.....	47
Create DNS Reverse Lookup Zone.....	50
Add self-signed certificate to Windows Active Directory.....	51
Configure NTP client in AD VM.....	54
6 VMware vCenter Server deployment and configuration.....	56
VMware vCenter Server Appliance deployment.....	56
Stage 1 - Deploy ISO file for Management vCenter Server Appliance with embedded PSC.....	56
Stage 2 - Set up Management vCenter Server Appliance with embedded PSC.....	57
Change vSAN default storage policy for management vCSA.....	59
Stage 1 - Deploy ISO file for Resource vCenter Server Appliance.....	60
Stage 2 - Set up resource vCenter Server Appliance with embedded PSC.....	61
Change VMware vSAN default storage policy for resource vCSA.....	61
Add AD authentication for vCenter Server.....	62
Assign license to vCSA.....	63
Create data center and cluster on resource vCenter.....	64
Add hosts to vCenter cluster.....	64
Enable VMware enhanced vMotion compatibility.....	65
Enable VMware EVC for management cluster.....	65
Enable VMware EVC for resource and edge cluster.....	66
7 Configure virtual network.....	67
VDS creation and configuration for management pod.....	67
Create VDS for management pod.....	67
VDS configuration settings for management VDS.....	68
Create LAG for management pod.....	68
Create distributed port group for management pod.....	69
Add host to VDS on management pod.....	71
Add hosts to Infra_Network_VDS.....	71
Add hosts to VM_Network_VDS.....	75
vSwitch to VDS Migration on management pod.....	76
Create VDS for resource and edge pods.....	77
VDS configuration settings for resource VDS.....	78
Create LAG for resource and edge pods.....	78

Create distributed port group for resource and edge VDS.....	79
Add hosts to VDS on edge pod.....	80
Add hosts to Edge_Infra_Network_VDS.....	80
Add hosts to Edge_VM_Network_VDS.....	82
Add hosts to VDS on resource pod.....	83
Add hosts to Res_Infra_Network_VDS.....	83
8 Configure VMware vSAN clusters.....	85
Configure vSAN on resource and edge cluster.....	85
Assign vSAN license key to cluster.....	85
Assign a new vSAN license.....	85
Assign vSAN license using an existing license.....	86
Update vSAN HCL database manually.....	86
Enable vSAN performance service.....	86
9 Configure VMware vCenter High Availability.....	88
Management cluster VCSA-HA configuration.....	88
Configure Management vCenter HA.....	88
Resource cluster VCSA-HA configuration.....	90
Configure Resource vCenter HA.....	90
10 NSX-T deployment and configuration.....	97
Install NSX-T Manager Virtual Appliance.....	97
Add license key.....	99
Add Compute Manager for management and resource VCSA.....	99
Deployment of NSX-T node and cluster from NSX-T Manager.....	101
Validate NSX-T node and cluster deployment.....	102
Add Virtual IP.....	103
Configure NSX-T Manager.....	104
Create transport zones.....	104
Create uplink profiles.....	105
Create IP pool for tunnel endpoints.....	108
Create host transport nodes.....	108
Installation of NSX-T edge.....	110
Create edge cluster.....	113
Create logical switches.....	115
Create and configure tier-1 router.....	117
Create router port on tier-1 router.....	118
Configure route advertisement on tier-1 router.....	120
Create and configure NSX-T tier 0 router.....	121
Connect Tier-1 router to NSX-T tier 0 router.....	122
Create logical router port on Tier-0 router.....	123
Redistribution on Tier-0 router.....	124
Configure BGP on NSX-Tier-0 router.....	125
Create and configure VCD-Tier1 router.....	127
Create logical router port on VCD tier-1 router.....	128
11 Configure vCloud Director.....	137
Installation of NFS server.....	137

Installation and configuration of vCloud Director.....	138
Deployment and configuration of vCD Cell 01.....	138
Deployment of vCD Cell 02.....	140
vCD integration with vCenter.....	141
vCD integration with NSX-T.....	142
Creating a session token for vCD.....	143
Retrieve VIM server details.....	144
Update VIM server.....	144
Retrieve the list of available resource pool.....	146
Retrieve NSX-T Manager instance details.....	146
Create a provider VDC.....	147
Create an organization.....	149
Create a new Organization VDC.....	149
Create new catalog.....	150
Create vApp Templates.....	150
Create vApp.....	151
Create virtual machine for vApp template.....	152
Add a network to organization VDC.....	153
Add Network to vApp.....	154
Add Network to Virtual Machine.....	155
Add VM to a vApp.....	156
Move a VM to vApp.....	157
12 VMware vRealize Log Insight deployment and configuration.....	158
Deploy the vRealize Log Insight virtual appliance.....	158
Configure the root SSH password for vRLI virtual appliance.....	159
Master node configuration.....	160
Worker node configuration.....	160
Enable Integrated Load Balancer.....	161
Integrate vRLI with AD.....	161
Integrate vRLI with VMware vCenter.....	161
Configure vRLI to send notifications to vRealize Operations Manager.....	162
Add Log Insight content packs.....	162
Offline update for content pack.....	163
Online update for content pack.....	163
vRLI integration with vCD.....	163
vRLI integration with vRO.....	165
Integrate vRLI with vRO.....	165
13 vRealize Orchestrator.....	166
Installation of vRO.....	166
Configure NTP in vRO.....	166
Configure Orchestrator Server with vSphere Authentication.....	167
Updating the vRO using ISO.....	168
Configure vRO plug-in for vSphere Web Client.....	168
Add a vCenter Server instance to vRO.....	168
Register vRealize Orchestrator as a vCenter Server extension.....	169
Configure vRealize Orchestrator to forward logs to vRLI.....	170

14 VMware vRealize Operations Manager deployment and configuration.....	171
Deployment prerequisites for vRealize Operations Manager.....	171
Deploy vRealize Operations Manager.....	171
Configuration of vRealize Operations Manager.....	172
Add data nodes to scale out vRealize Operations Manager.....	172
Add master replica node.....	172
Enable High Availability for clusters.....	173
Start cluster.....	173
Product license.....	174
vROps integration with other components.....	174
Activate vCenter, vSAN, and vRLI Management packs.....	174
Integrate vROps with VMware vCenter.....	174
vROps integration with AD.....	177
vROps integration with vRLI.....	178
vROps integration with NSX-T.....	178
vROps integration with vSAN.....	179
vROps integration with vCD.....	180
vROps integration with vRealize Orchestrator.....	182
15 vSphere Replication.....	184
Configuring vSphere Replication.....	185
Configure vSphere Replication connection.....	187
16 Set up anti-affinity rules.....	190
Create an anti-affinity rule.....	190
Enable vSphere DRS.....	191
Enabling vSphere availability.....	191
17 Forwarding logs to vRLI.....	192
Forwarding vROps log to vRLI.....	192
Forwarding vSAN logs to vRLI.....	193
Forwarding logs from vCD to vRLI.....	194
Configure syslog server for NSX-T.....	195
Log Message IDs.....	196
A Reference documentation.....	198

Overview

The Dell EMC Ready Solution bundle is designed to consolidate and deliver the networking components that support a fully virtualized infrastructure. The components include virtual servers, storage, and other networks. It uses standard IT virtualization technologies that run on high-volume service, switch, and storage hardware to virtualize network functions.

The Dell EMC Ready Architecture for VMware vCloud NFV 3.2 vCloud Director 9.7 Architecture and Software Deployment Guide provides detailed instructions for the manual deployment of the VMware vCloud NFV 3.2 with VMware vCloud Director 9.7 platform. This guide also provides information about the hardware and software that is recommended for the deployment of the Dell EMC Ready Architecture for VMware NFV 3.2 platform.

The scope of this document is limited to a Greenfield deployment.

Servers:

- Dell EMC PowerEdge R640 or Dell EMC PowerEdge R740 server with the Dell EMC PowerEdge HBA330 disk controller that is based on vSAN Ready Node
- Dell EMC PowerEdge R740xd server with the Dell EMC PowerEdge HBA330 disk controller based on vSAN Ready Node

Networking:

- One Dell EMC Networking S4048T-ON switch as Top of Rack (ToR) switch
- Two Dell EMC Networking S5248-ON, Dell EMC Networking S5232-ON or Dell EMC Networking S6010-ON switches as leaf switches
- Two Dell EMC Networking Z9264F-ON switches as spine switches

This guide consists of three sections:

- Deployment architecture
- Hardware installation and configuration
- Manual deployment

Intended audience

The information in this guide is intended for use by system administrators who are responsible for the installation, configuration, and maintenance of Dell EMC 14G technology along with the suite of VMware applications.

Acronyms and definitions

Dell EMC Ready Solution bundle uses a specific set of acronyms that apply to NFV technology.

Table 1. Acronyms and definitions

Acronyms	Description
CSP	Communication Service Provider
DPDK	Data Plane Development Kit, an Intel led packet processing acceleration technology
iDRAC	integrated Dell Remote Access Controller
NFVI	Network Functions Virtualization Infrastructure
NFV-OI	NFV Operational Intelligence
N-VDS (E)	Enhanced mode when using the NSX-T Data Center N-VDS logical switch that enables DPDK for workload acceleration
N-VDS (S)	Standard mode when using the NSX-T Data Center N-VDS logical switch
ToR	Top of Rack
VIM	Virtualized Infrastructure Manager
VNF	Virtual Network Function running in a virtual machine

Acronyms	Description
VR	vSphere Replication
vRLI	VMware vRealize Log Insight
vRO	vRealize Orchestrator
vROps	VMware vRealize Operations

Deployment architecture for vCloud NFV

This section provides a reference architecture for the design and creation of a Greenfield Network Function Virtualization (NFV) environment using VMware vCloud Director, or VCD with VMware NSX-T and Dell EMC PowerEdge Servers.

This deployment uses the three-pod architecture design as per VMware vCloud NFV 3.0 Reference Architecture Guide to deploy Dell EMC vCloud NFV 3.2 with vCD. By design, the management, resource, and edge pods include a vSphere cluster. You can scale up the clusters by adding ESXi hosts to the clusters.

For more information, see:

- [Architecture design](#)
- [Solution bundle network topology](#)
- [Three-pod configuration](#)

Architecture design

Figure 1 displays the three-pod architecture diagram that is used to deploy the Dell EMC Ready Solution vCloud NFV 3.2.

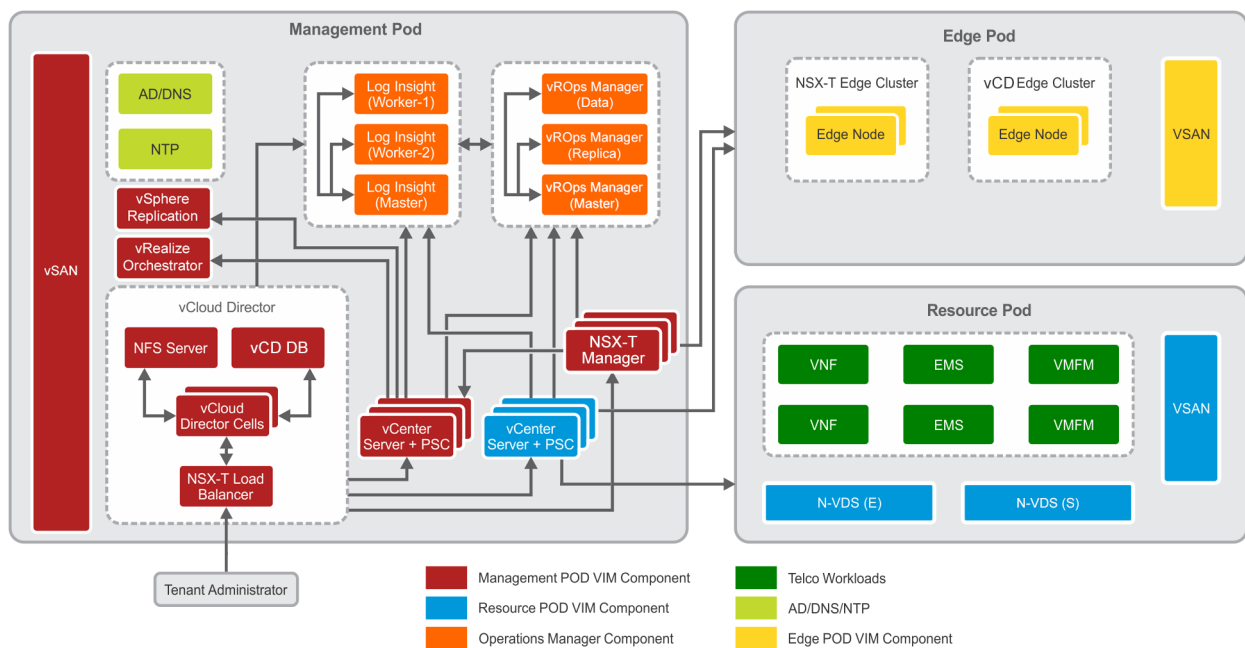


Figure 1. Architecture design

For more information, see the following sections:

- Management pod
- Edge pod
- Resource pod

Solution bundle network topology

This section provides the network information physical network design and virtual network topology design and topology that is used in this deployment. For more information, see:

- [Solution bundle physical network design and topology](#)
- [Solution bundle virtual network design and topology](#)

Solution bundle physical network design and topology

The Dell EMC Ready Solution bundle uses two-tier leaf-and-spine network architecture to build newer data center infrastructure. In this network architecture, leaf switches are connected to spine switches. The leaf switches provide connectivity between the endpoints and data center. The spine switches provide high-speed interconnectivity between the leaf switches. The leaf-and-spine network is connected in a full mesh that provides predictable communication and latency between endpoints. Leaf switches are configured as a Virtual Link Trunking (VLT) pair that enables all connections to be active while providing fault tolerance. The ToR switch provides the external connectivity to the NFV stack.

For this deployment, the Dell EMC Networking S4048T-ON system is used as a ToR switch. Two Dell EMC Networking S5232-ON systems are used as leaf switches, and two Dell EMC Networking Z9264F-ON systems are used as spine switches.

- Physical network topology for the deployment server: Figure 2 displays the deployment server network topology that is used in this deployment:
 - Leaf switches are connected to the VLT using a 100G interface
 - Leaf and spine switches are interconnected using a 100G interface
 - iDRAC is connected to the deployment server and ToR using a 10G interface
 - VM network is connected to ToR and deployment server using a 10G interface - vmnic2
 - ESXi Management Network is connected to the Leaf switches using a 10G interface - vmnic4
 - VM Management Network is connected to the Leaf switches using a 10G interface - vmnic5

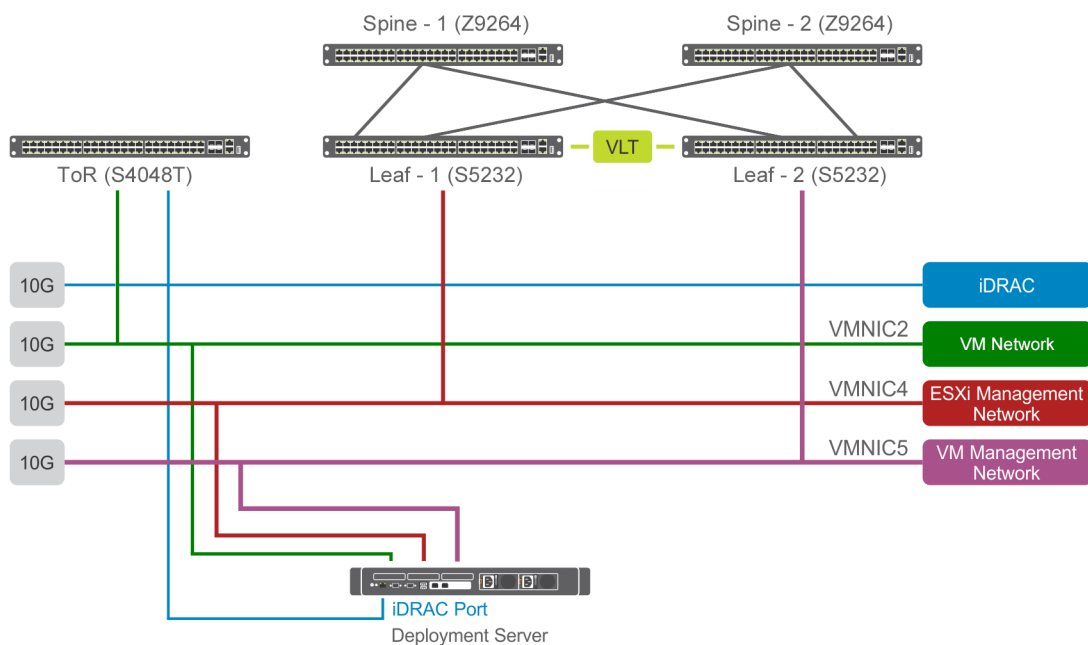


Figure 2. Deployment server networking topology

- Physical network topology for management pod: Figure 3 displays the management pod network topology that is used in this deployment:
 - iDRAC is connected to the ToR switch and management pod using a 10G interface
 - ESXi Management Network is connected to the leaf switches and to the management pods using a 25G interface
 - VM Management Network is connected to the leaf switches and to the management pod using a 25G interface
 - vSAN Network is connected to the leaf switches and to the management pod using a 25G interface
 - vMotion Network is connected to the leaf switches and to the management pod using a 25G interface
 - VCSA HA Network is connected to the leaf switches and to the management pod using a 25G interface
 - Replication Network is connected to the leaf switches and to the management pod using a 25G interface

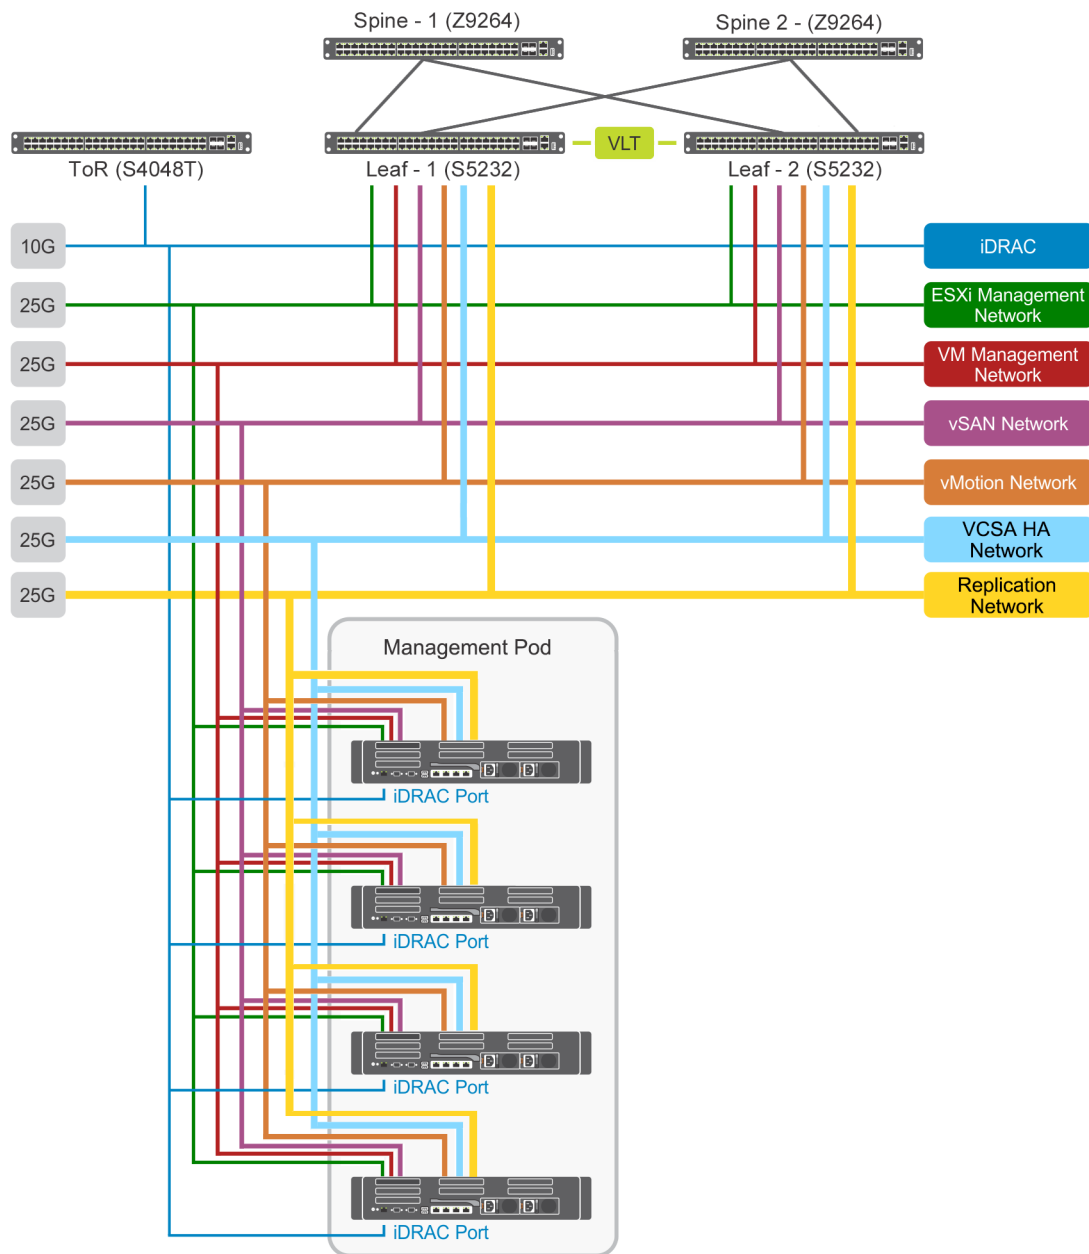


Figure 3. Physical network topology for management pod

- Physical network topology for edge pod: Figure 4 displays the edge pod physical network topology that is used for this deployment:
 - iDRAC is connected to the ToR switch and Edge pod using a 10G interface
 - ESXi Management Network is connected to the leaf switches and to the edge pod using a 25G interface
 - VM Management Network is connected to the leaf switches and to the edge pod using a 25G interface
 - vSAN Network is connected to the leaf switches and to the edge pod using a 25G interface
 - vMotion Network is connected to the leaf switches and to the edge pod using a 25G interface
 - Overlay Network is connected to the leaf switches and to the edge pod using a 25G interface
 - External Network is connected to the leaf switches and to the edge pod using a 25G interface

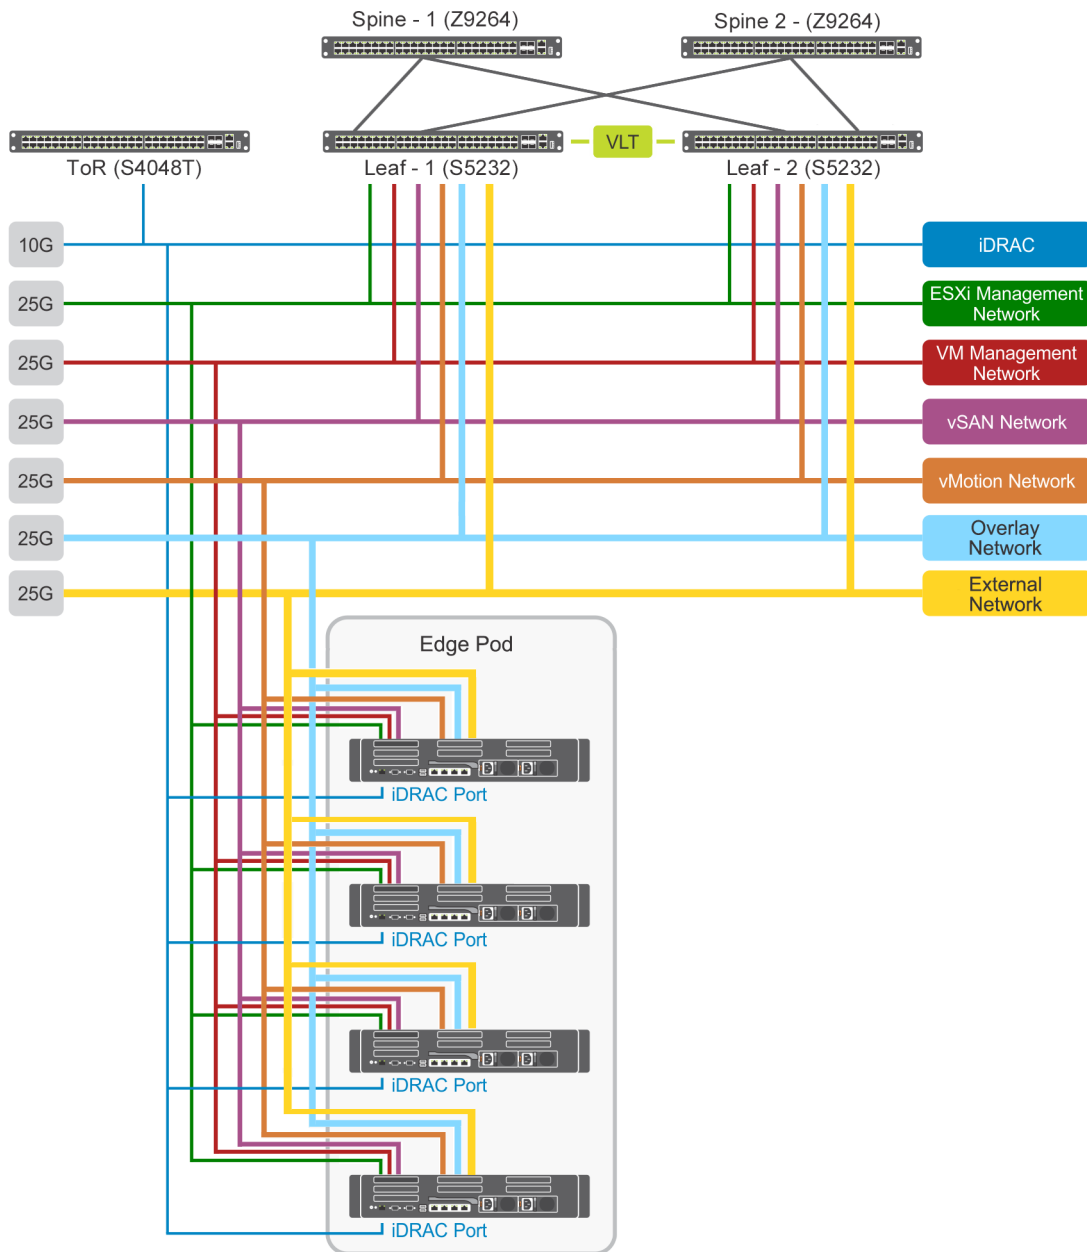


Figure 4. Physical network topology for edge pod

- Physical network topology for resource pod: Figure 5 displays the resource pod physical network topology that is used in this deployment:
 - iDRAC is connected to the ToR switch and resource pod using a 10G interface
 - ESXi Management Network is connected to the leaf switches and to the resource pod using a 25G interface
 - VM Management Network is connected to the leaf switches and to the resource pod using a 25G interface
 - vSAN Network is connected to the leaf switches and to the resource pod using a 25G interface
 - vMotion Network is connected to the leaf switches and to the resource pod using a 25G interface
 - Overlay Network is connected to the leaf switches and to the resource pod using a 25G interface
 - External Network is connected to the leaf switches and to the resource pod using a 25G interface
 - N-VDS (Enhanced mode) Network is connected to the leaf switches and to the resource pod using a 25G interface

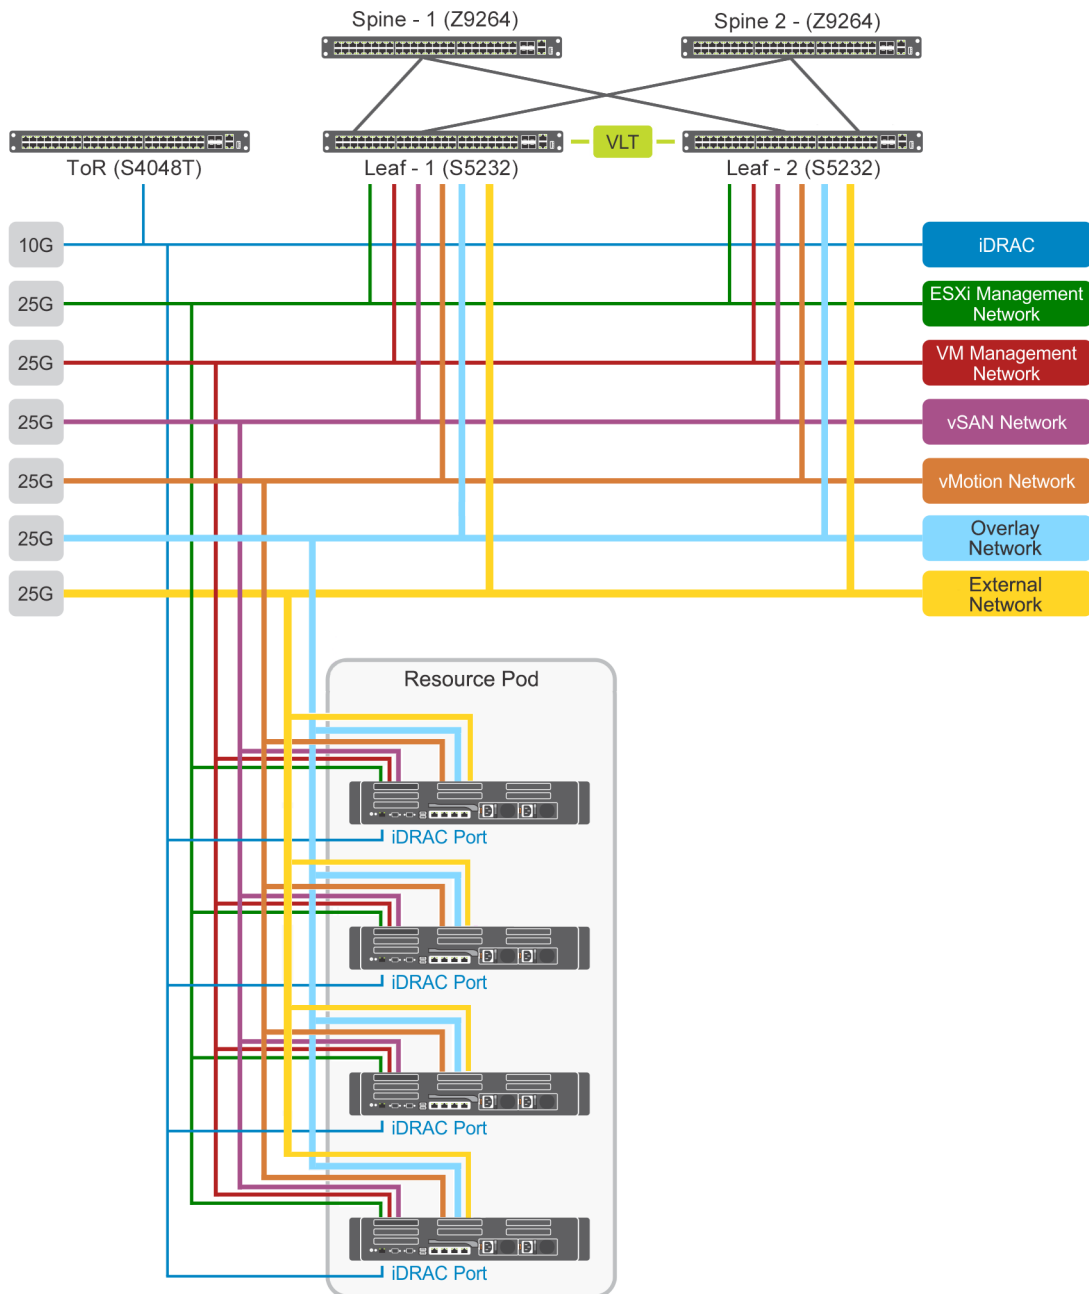


Figure 5. Physical network topology for resource pod

Solution bundle virtual network design and topology

The vCloud NFV platform consist of two networks:

- Infrastructure network
- Virtual Machine (VM) network

Infrastructure networks are host-level networks that are used to connect hypervisors with the physical networks. Each ESXi host has multiple port groups that are configured on each infrastructure network.

The VMware vSphere Distributed Switch (VDS) is configured on the hosts in each pod. This configuration provides a similar network configuration across the multiple hosts. One VDS is used to manage infrastructure network and another is used to manage VM networks. Also, N-VDS is used to manage the traffic between:

- Components running on transport node

- Internal components and physical network

The ESXi hypervisor uses the infrastructure network for Edge overlay, vMotion, and vSAN traffic. The VMs use the VM network to communicate with each other. In this configuration, two distribution switches are used to create a separation. One switch is used for the infrastructure network where the second switch is used for VM network.

Each distribution switch has a separate uplink connection for physical data center network that separates uplink traffic from other network traffic. The uplinks are mapped with a pair of physical NICs on each ESXi host for best performance and resiliency.

NSX-T creates the VLAN-backed logical switches which provide the connectivity to VNF components and VMs. On the ESXi hosts, physical NICs act as uplinks to connect the host virtual switches to the physical switch.

The following infrastructure networks are used in the pods:

- ESXi management network – network for ESXi host management traffic
- vMotion network – network for VMware vSphere vMotion traffic
- vSAN network – network for vSAN shared storage traffic
- Replication network – network used for replication storage traffic

Virtual network topology of management pod

Management pod networking consists of the infrastructure, and VM networks as follows:

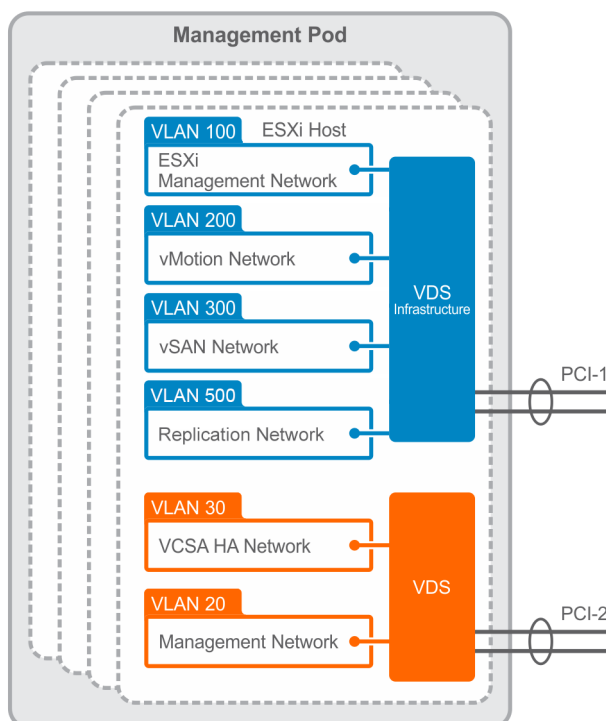


Figure 6. Management Pod virtual network topology

Virtual network topology of edge pod

The virtual network of the edge pod depends on the network topology that is required for VNF workloads. In general, the edge pod has the infrastructure networks, networks for management, and networks for the workloads.

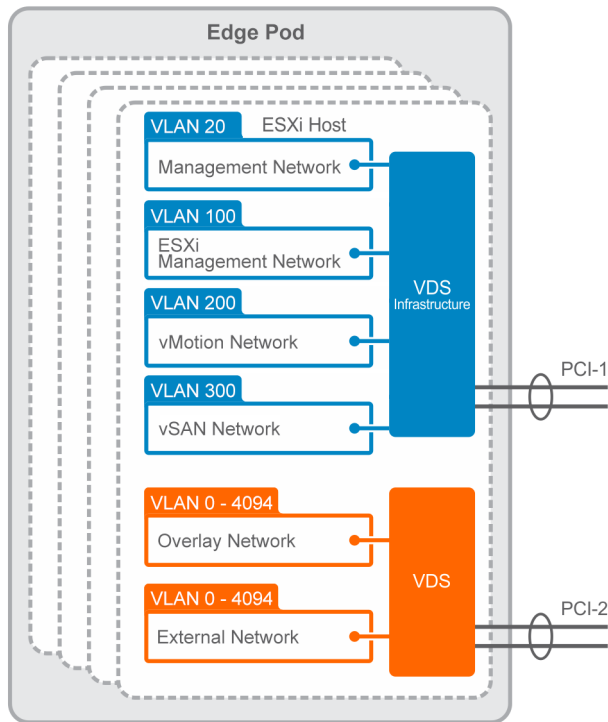


Figure 7. Edge Pod virtual network topology

Virtual Network topology for resource pod

The resource pod virtual network depends on the network topology that is required to deploy tenants. A specific tenant has a certain set of networking requirements.

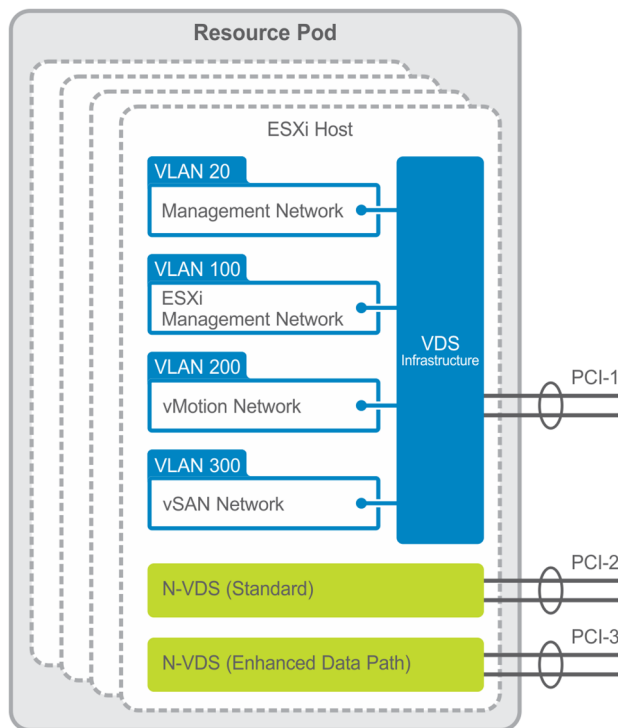


Figure 8. Virtual network topology for resource pod

Three pod configuration

In this deployment, a pod is used to streamline the NFV environment operations and other roles. This deployment architecture illustrates a three-pod configuration where three vSphere clusters are deployed to create the following clusters within the pods:

- Management pod
- Edge pod
- Resource pod

Clusters are the vSphere objects that are used to access the virtual domain resources and manage the resource allocation.

During the initial deployment, Dell EMC recommends:

- Minimum of four servers that consist of either Dell EMC PowerEdge R640 or R740 servers in the management pod
- Minimum of four servers that consist of Dell EMC PowerEdge R740xd servers in the edge pod
- Minimum of four servers that consist of Dell EMC PowerEdge R740xd servers in the resource pod

NOTE: A maximum of 64 server can be added to each pod to scale up the deployment.

Management pod

The management pod hosts and manages all NFV management components:

- vCenter Server Appliance
- NSX-T Manager
- NSX-Controller
- VMware vCloud Director
- AD-DNS
- Network Time Protocol (NTP)

Analytics components such as vRealize Operations (vROps) Manager and vRealize Log Insight (vRLI) are also deployed in the management pod.

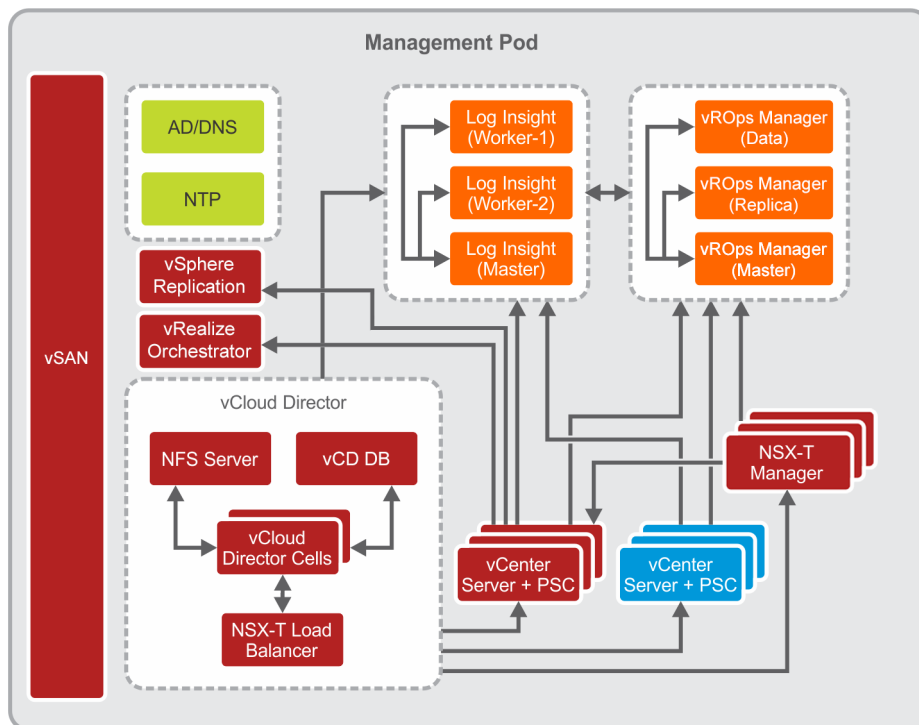


Figure 9. Management pod

Edge pod

Edge pod hosts the NSX-T Edge as a Virtual Machine (VM) and manages all the connectivity to the physical domain within the architecture. The Edge pod also creates different logical networks between VNFs and external networks. The Edge pod host the NSX-T Edge nodes which work as NSX-T data center network components. The NSX-T edge node:

- Participates in east-west connection
- Provides connectivity to the physical infrastructure for north-south traffic management and capabilities

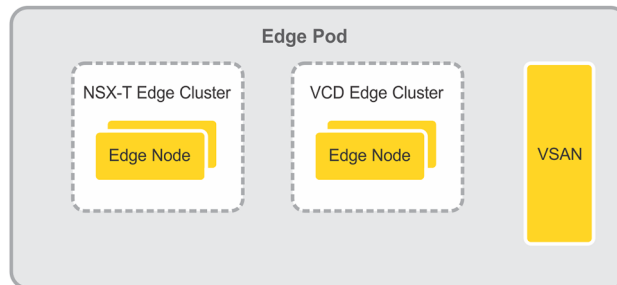


Figure 10. Edge pod

Resource pod

The resource pod provides the virtualized runtime environment, namely compute, network, and storage environments, to fulfill workloads.

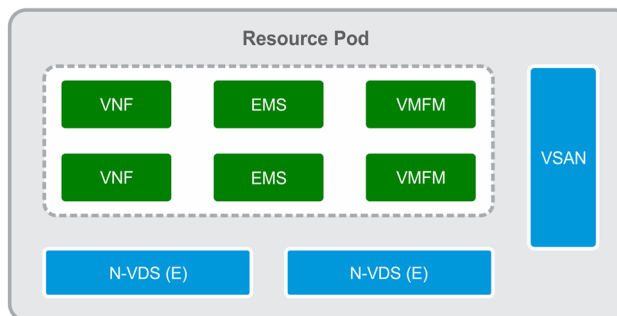


Figure 11. Resource pod

Solution hardware

Hardware installation and configuration

The servers, storage, and other networking component are required to install and configure to deploy Dell EMC Ready Solution bundle.

The following server solution support is used in this deployment:

- Dell EMC PowerEdge R640 or Dell EMC PowerEdge R740 servers
- Dell EMC PowerEdge R740xd servers

This configuration uses the following switches:

- One Dell EMC Networking S4048T-ON switch that serves as a ToR system
- Two Dell EMC Networking S5248-ON, Dell EMC Networking S5232-ON, or Dell EMC Networking S6010-ON switches as leaf switches
- Two Dell EMC Networking Z9264F-ON switches as spine switches

This deployment also uses the Dell Remote Access Controller 9, or iDRAC9, to improve the overall availability of Dell systems.

Unpack and install equipment

After performing all standard industry safety precautions, proceed with the following steps:

1. Unpack and install the racks.
2. Unpack and install the server hardware.
3. Unpack and install the switch hardware.
4. Unpack and install the network cabling.
5. Connect each individual machine to both power bus installations.
6. Apply power to the racks.

 **NOTE:** The Dell EMC EDT team usually performs these steps.

Power on equipment

 **NOTE:** The Dell EMC EDT team usually performs these steps.

To test the installation of the equipment, perform the following steps:

1. Power on each server node individually.
2. Wait for the internal system diagnostic procedures to complete.
3. Power up the network switches.
4. Wait for the internal system diagnostic procedures to complete on each of the switches.

Tested BIOS and firmware

 **CAUTION:** Ensure that the firmware on all servers, storage devices, and switches is up-to-date as outdated firmware may cause unexpected results to occur.

The server BIOS and firmware versions that are tested for the Dell EMC Ready Bundle for NFV platform are as follows:

Table 2. Dell EMC PowerEdge R640/R740 and R740xd tested BIOS and firmware versions

Product	Firmware version
BIOS	2.2.11
iDRAC with Lifecycle Controller	3.34.34.34
rNDC - Intel® 4P X550-t	18.8.9
PCIe - Intel 25G 2P XXC710/ 10G X710	19.00.12
QLLogic QL41262 25GB NIC	15.00.14
HBA330 ADP/Mini storage Controller	16.17.00.03
BP14G PowerEdge R640/PowerEdge R740	4.27
BP14G PowerEdge R740xd	2.41
PowerEdge CPLD firmware for the PowerEdge R640	1.0.2
PowerEdge CPLD firmware for the PowerEdge R740xd or PowerEdge R740	1.1.3

The firmware switch versions that are tested for the Dell EMC Ready Bundle for NFV platform are as follows:

Table 3. Dell Networking tested BIOS and firmware versions

Product	Version
S4048T-ON firmware (1) ToR switch	OS 10.5.0
S5232-ON firmware (2) leaf switch	OS 10.5.0
Z9264F-ON firmware (2) spine-switch	OS 10.5.0

Supported configuration

Table 4 provides the list of VMware component and their supported version that is used and verified for this deployment.

Table 4. VMware vCloud NFV product inventory list

Product	Version
ESXi	6.7 U2
VMware vCenter Server	6.7 U2
VMware NSX-T	2.4.1
VMware vSAN	6.7 U2
VMware vRealize Log Insight	4.8
VMware vRealize Operations Manager	7.5
VMware vCloud Director	9.7
vSphere Replication	8.2
vRealize Orchestrator	7.6

List of components

Various software's are used to create the NFVI environment. Table 5 displays the list of components and their instances that are deployed in this deployment.

Table 5. NFVI components

Product	Instances (count)
ESXi	12 nodes

Product	Instances (count)
AD-DNS	1 VM
NTP	1 VM
VMware vCenter Server	6 VM
VMware vSAN	NA
VMware NSX-T Manager	3 VM
VMware NSX-T Edge	4 VM
VMware vRealize Log Insight	3 VM
VMware vRealize Operations Manager	3 VM
VMware vCloud Director	4 VM
vSphere Replication	1 VM
vRealize Orchestrator	1 VM

Network connectivity and port mapping

Verify network connectivity to server ports

To ensure the network connectivity to server ports, use the information that is provided in the [Network connectivity configuration for Management ESXi host](#), [Network connectivity configuration for Edge ESXi host](#), [Network connectivity configuration for Resource ESXi host](#), and [Network connectivity configuration table for deployment server](#) tables. The information that is provided ensures that the network cables are connected correctly to the servers. The tables also provide the port-mapping information for the VMware VMNIC port references. The installation process requires that the PCIe expansion card slot (riser 1) is used for network connectivity.

NOTE: For more information, see Appendix A to download the appropriate *Dell EMC PowerEdge Owner's Manual* and reference the *Expansion card installation* section.

The configuration process requires that the NIC ports that are integrated on the network adapter card, or NDC, are connected as outlined in the [Network connectivity configuration for Management ESXi host](#), [Network connectivity configuration for Edge ESXi host](#), [Network connectivity configuration for Resource ESXi host](#), and [Network connectivity configuration table for deployment server](#) tables.

NOTE: For more information, see Appendix A to download the appropriate *Dell EMC PowerEdge Owner's Manual* and reference the *Technical specifications* section.

Table 6. Network connectivity configuration for Management ESXi host

	LOM/NDC port				NIC slot 1		NIC slot 2	
	1	2	3	4	1	2	1	2
Port number	1	2	3	4	1	2	1	2
VMware VMNIC port reference	vmnic0	vmnic1	vmnic2	vmnic3	vmnic4	vmnic5	vmnic6	vmnic7
PowerEdge R640/ PowerEdge R740	-	-	-	-	25G	25G	25G	25G

Table 7. Network connectivity configuration for Edge ESXi host

	LOM/NDC port				NIC slot 1		NIC slot 3	
	1	2	3	4	1	2	1	2
Port number	1	2	3	4	1	2	1	2
VMware VMNIC port reference	vmnic0	vmnic1	vmnic2	vmnic3	vmnic4	vmnic5	vmnic6	vmnic7

	LOM/NDC port				NIC slot 1		NIC slot 3	
PowerEdge R740xd	-	-	-	-	25G	25G	25G	25G

Table 8. Network connectivity configuration for Resource ESXi host

	LOM/NDC port				NIC slot 1		NIC slot 3		NIC slot 4	
Port number	1	2	3	4	1	2	1	2	1	2
VMware VMNIC port reference	vmnic0	vmnic1	vmnic2	vmnic3	vmnic4	vmnic5	vmnic6	vmnic7	vmnic8	vmnic9
PowerEdge R740xd	-	-	-	-	25G	25G	25G	25G	25G	25G

Table 9. Network connectivity configuration table for deployment server

	LOM/NDC port				NIC slot 1			
Port number	1	2	3	4	1	2	3	4
VMware VMNIC port reference	vmnic0	vmnic1	vmnic2	vmnic3	vmnic4	vmnic5	vmnic6	vmnic7
PowerEdge R640/ PowerEdge R740	-	-	10G	-	10G	10G	10G	10G

VDS DvPort group mapping with VLAN ID and related ESXi VMNIC

The mapping list provides details about all VSS, VDS, VDS DvPort groups, VLAN ID, and ESXi VMNICs. These details are created and configured under management and resource pod networking.

The [Management pod](#), [Resource pod](#), [Edge pod](#), and [Deployment server](#) tables show the VDS-DvPort group/VSS port group mappings with VLAN ID and corresponding VMNIC present on ESXi, which are assigned as uplinks to the VDS/VSS.

For example, the VDS named Infrastructure Management VDS with the DvPort Group ESXi_Mgmt_Network is configured with VLAN ID 100 and uses a pair of VMNIC which is vmnic4 and vmnic6 as uplinks for the VDS.

Table 10. Management pod

VDS type	VDS name	Port groups	VLAN ID	Uplink NICs	Switch
VDS (Infrastructure)	Infrastructure Management VDS	ESXi_Mgmt_Network	100	vmnic4 vmnic6	Leaf1+Leaf2
		vSAN_Network	300		
		vMotion_Network	200		
		Replication_Network	500		
Virtual Machine Network (VDS)	Management Network VDS	VM_Mgmt_Network	20	vmnic5	Leaf1+Leaf2
		VCSA_HA_Network	30	vmnic7	

Table 11. Resource pod

VDS type	VDS name	Port groups	VLAN ID	Uplink NICs	Switch
VDS (Infrastructure)	Infrastructure Management VDS	ESXi_Mgmt_Network	100	vmnic4 vmnic6	Leaf1+Leaf2
		VM_Mgmt_Network	20		

VDS type	VDS name	Port groups	VLAN ID	Uplink NICs	Switch
VDS (Virtual Machine Network)	N-VDS (S)	vSAN_Network	200	vmnic5 vmnic7	Leaf1+Leaf2
		vMotion_Network	300		
		Overlay_Network	70		
VDS (Virtual Machine Network)	N-VDS (Enhanced Data Path)	Vlan_DPDK_Network	40	vmnic8 vmnic9	Leaf1+Leaf2

Table 12. Edge pod

VDS type	VDS name	Port groups	VLAN ID	Uplink NICs	Switch
VDS (Infrastructure)	Infrastructure Management VDS	ESXi_Mgmt_Network_Edge	100	vmnic4 vmnic6	Leaf1+Leaf2
		VM_Mgmt_Network_Edge	20		
		vSAN_Network_Edge	200		
		vMotion_Network_Edge	300		
VDS (Virtual Machine Network)	Edge VDS	Overlay_Network External_Network	VLAN 0-4094	vmnic5 vmnic7	Leaf1+Leaf2

Table 13. Deployment server

VSS name	Port groups	VLAN ID	Uplink NICs	Switch
vSwitch0	VM Network	0	vmnic2	ToR
vSwitch1	PG-100-ESXI	100	vmnic4	Leaf
	PG-20-VM-Mgmt	20	vmnic5	Leaf

Manual deployment

Solution prerequisites

The following requirements must be satisfied before beginning the Dell EMC VMware vCloud NFV 3.2 platform manual deployment:

NOTE: All compute nodes must have identical hard drive, RAM, and NIC configurations.

- The required hardware must be installed and configured as indicated in the *Hardware installation and configuration* section
- Once the systems are configured as described in the *Hardware installation and configuration* section, power on the systems
- Ensure that there is Internet access, including but not limited to the deployment server
- Verify that the deployment server that is used to deploy the solution, can hold the required VMware Software Appliance files

Deployment server

See the [Solution bundle physical network design and topology](#) section for the deployment server physical network topology that is used in this deployment.

NOTE: ESXi 6.7 U2 or above must be installed on a bare-metal deployment server.

ESXi installation on deployment server

Prerequisites

- iDRAC is configured and accessible
- ESXi 6.7 U2 ISO file is available on the local machine
- ToR switch is configured
- Dell PowerEdge server controllers are set to HBA mode

About this task

This task provides the steps to install ESXi on the deployment server.

Steps

1. Log in to the iDRAC 9 web GUI.
2. From the **Dashboard** screen, click **Launch Virtual Console** within the **Virtual Console** section. The **iDRAC Virtual Console** window displays.
3. On the navigation bar, click **Connect Virtual Media**. The **Virtual Media** screen displays.
4. In the **Map CD/DVD** section, click **Choose File**, select the ESXi image file from your local machine, and click **Map Device**.
5. Click **Close**.

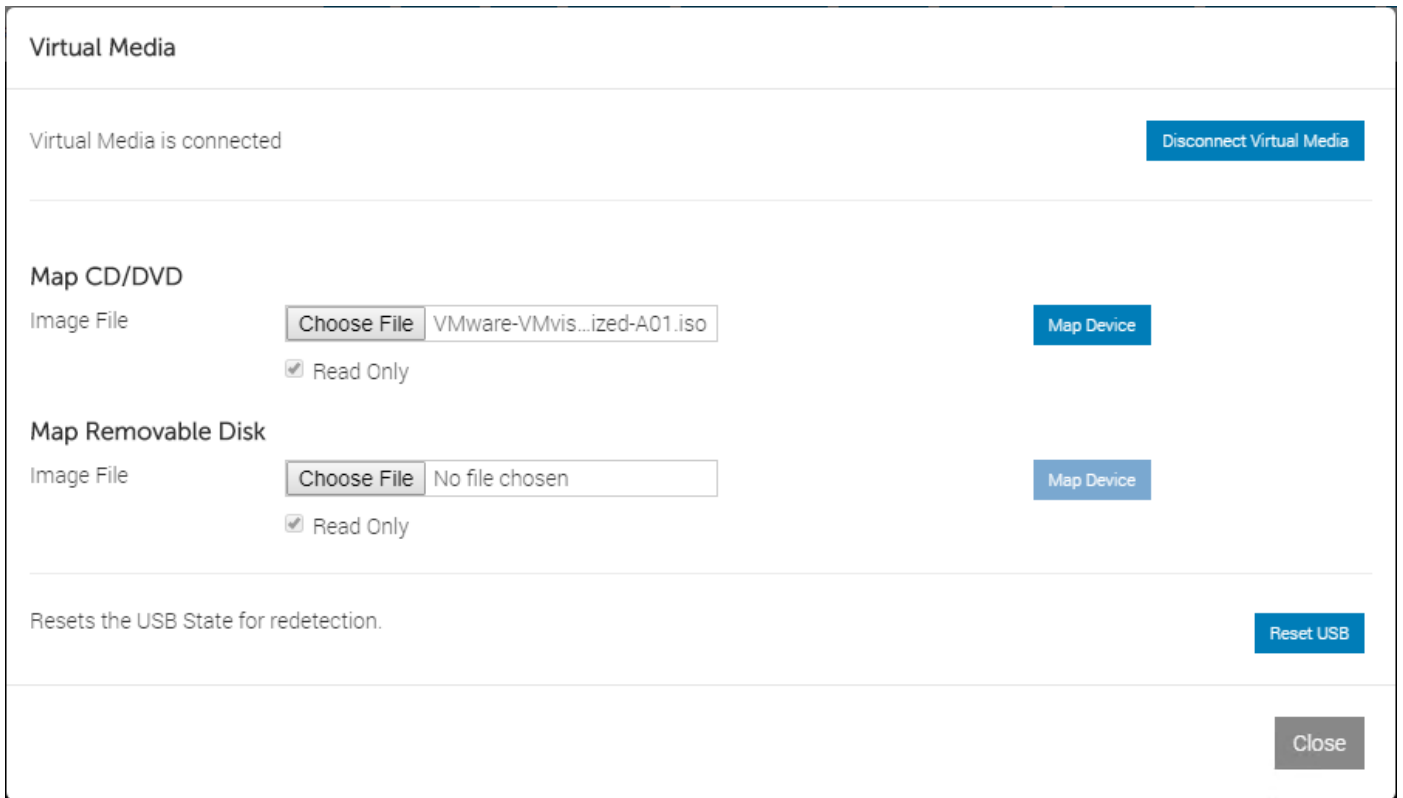


Figure 12. Virtual Media screen

6. In the navigation bar, click **Next Boot**, select **Virtual CD/DVD/ISO**, then click **Save**.
7. In the navigation bar, click **Power**, then select **Power Cycle System (cold boot)**.
The ESXi installation process begins. When complete, the **Welcome to the VMware ESXi 6.7 U2 Installation** window displays.
8. From the **Welcome to the VMware ESXi 6.7 U2 Installation** screen, press **Enter** to continue.
9. Review the contents of the **End User License Agreement (EULA)** and if you agree to the terms, press **F11**.
The **Disk installation** screen displays.
10. Use the arrow keys to select the **Dell Internal Dual SD storage device**, then press **Enter**.
The **Confirm Disk Selection** screen displays.
11. Press **Enter** to confirm the disk selection.
12. From the **Keyboard layout** screen, verify that **US Default** is the option that is selected, then press **Enter**.
The **Enter root password** screen displays.
13. In the fields provided, enter the root password, enter it again to confirm, then press **Enter**.
The **Scanning system** screen displays.
14. Press **F11** to confirm the installation.
After the installation process is done, the **Installation Complete** screen displays.
15. Press **Enter** to reboot the system.
After the system reboots, the **Direct Console User Interface (DCUI)** window displays.

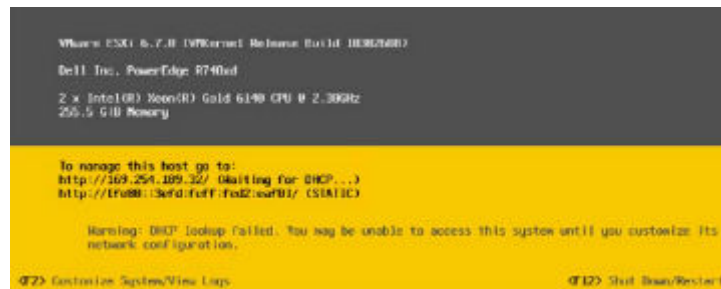


Figure 13. Direct Console User Interface (DCUI) window

Customize ESXi

About this task

The **System Customization** option enables users to customize various ESXi system settings such as:

- Passwords
- Management configuration
- Restart options
- Keyboard settings
- Troubleshooting options
- System reset configurations

To access the **System Customization** screen:

Steps

1. From the **DCUI** screen, press **F2**.
2. Enter the required user credentials in the fields that are provided and then press **Enter**.
The **System Customization** screen displays.
3. Use the arrows to select the option to customize, then press **Enter**.

Management network configuration

About this task

This section provides the steps to configure the management network in the deployment server.

Steps

1. Navigate to the **System Customization** screen and using the arrow keys, select **Configure Management Network**, then press **Enter**.
2. To update the network adapters, use the arrow keys to select the **Network Adapter** option, and then press **Enter**.

Change IPv4 configuration

About this task

This section provides the steps to add the static IPv4 address in the deployment server.

Steps

1. From the **System Customization** screen, select **Configure Management Network, IPv4 Configuration**, and then press **Enter**.
2. Select **Set static IPv4 and network configuration**.
3. In the fields provided, enter the required **IPv4 address**, **Subnet mask**, and **Default gateway** and then press **Enter**.
The changes are saved.

Change DNS configuration

About this task

This section provides the steps to add DNS information in the deployment server.

Steps

1. From the **Configure Management Network** screen, use arrow keys to select **DNS Configuration**, and then press **Enter**.
2. From the **DNS Configuration** screen, use arrow keys to select **Use the following DNS server addresses and hostname** option.
3. In the fields provided, enter the required **Primary DNS Server**, **Alternate DNS Server IP**, and **Hostname information** and then press **Enter**.
4. In the **Suffixes** field, enter the domain name and then press **Enter** to save the settings.

5. Press **Enter** to restart the management network.
6. After the network restarts, select **Test Management Network** and press **Enter**.
The test pings the configured default gateway, primary and alternate DNS servers, and resolves the configured hostname.

Troubleshoot ESXi

About this task

This section provides the steps to access the Troubleshoot option.

Steps

1. Using the arrow keys, select **Troubleshooting Options** and then press **Enter**.
2. From the options provided, use the arrow keys to select the wanted troubleshooting option, then press **Enter**.

Create standard vSwitch on deployment server

About this task

By default, vSwitch0 is available on the deployment server. Create the virtual switches on the deployment using the information provided in the [vSwitch details](#) table.

Table 14. vSwitch details

vSwitch name	Uplink	MTU (bytes)	Link discovery	Security
vSwitch0	vmnic2	1500 Bytes	Listen/CDP	For promiscuous mode and forged transmits, select the Reject radio button
vSwitch1	vmnic4, vmnic5	9000 Bytes	Listen/CDP	For promiscuous mode and forged transmits, select the Accept radio button

Steps

1. Using a web browser, go to the deployment server IP address and log in to it using the necessary credentials.
2. From the navigation panel, click **Networking** then click the **Virtual switches** tab.
3. On the **Virtual switches** tab, select **Add standard virtual switch**.
The **Add standard virtual switch** window displays.
4. In the **vSwitch Name** field, enter the vSwitch name.
5. Using the vSwitch details provided in the [vSwitch details](#) table in this section, select the **MTU**, **Uplink**, **Mode**, **Protocol**, and **Security details** options, and then click **Add**.
6. Once vSwitch1 is created, select the option to **Edit vSwitch1**.
7. In the **NIC Teaming** section, locate the **Load-balancing** drop-down list and select **Route based on IP hash**.

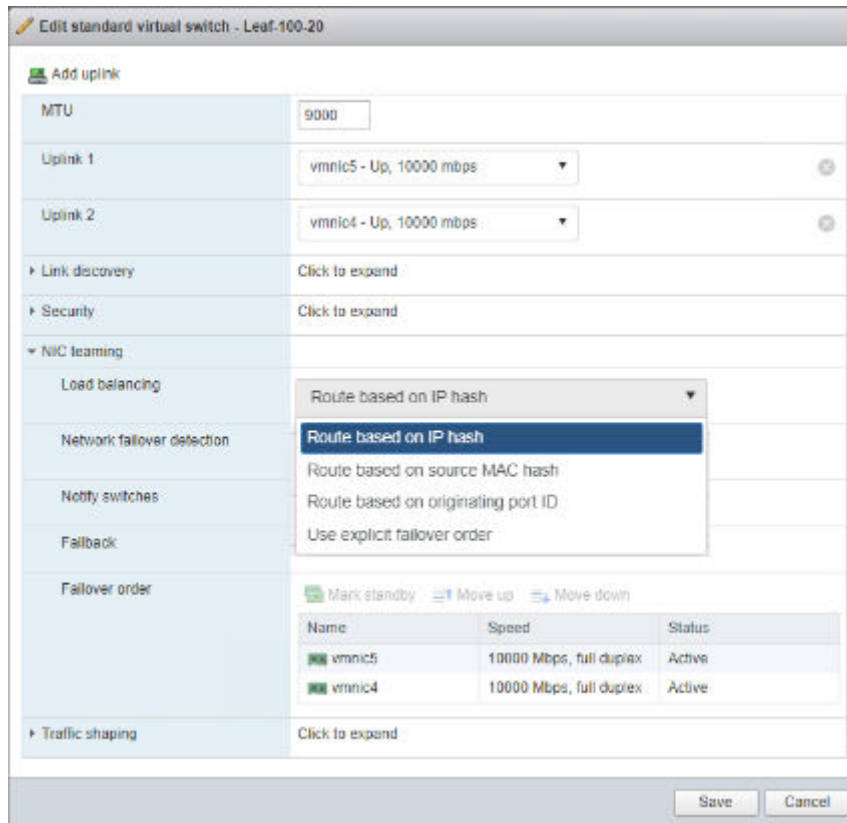


Figure 14. Edit standard virtual switch

Create port group on deployment server

About this task

By default, a VM network port group with VLAN ID 0 is created on the deployment server. When creating an extra port group, assign the VLAN and vSwitch to the port group as specified in the following table:

Table 15. Port group details

Port group name	Description	VLAN ID	Virtual switch	Security
VM network	For iDRAC (OOB management network)	0	vSwitch0	For promiscuous mode and forged transmits, select the Inherit from vSwitch radio button
PG-100-ESXi	For rack ESXi servers (ESXi management network)	100	vSwitch1	For promiscuous mode and forged transmits, select the Inherit from vSwitch radio button
PG-20-VM-Mgmt	For rack server VMs	20	vSwitch1	For promiscuous mode and forged transmits, select the Inherit from vSwitch radio button

Steps

1. Log in to the deployment server web GUI.
2. From the navigation panel, click **Networking**.
3. On the **Port groups** tab, then select **Add port group**. The **Add port group** window displays.

4. Use the information from the [Port group details](#) table in this section to update the required information in the **Add port group** window, then click **Add**.
5. Repeat the steps in this section to create more port groups on the deployment server as specified in the [Port group details](#) table.

Create datastore on deployment server

About this task

Create a datastore on the deployment server. This section provides the steps to create datastore on the deployment server.

Steps

1. Log in into ESXi host using the VMware vSphere web client.
2. From the **Home** screen, click **Storage**, and then click **New datastore**.
3. On the **Select Creation Type** screen, select **Create new VMFS datastore** then click **Next**.
4. In the **Name** field, enter the name of the datastore, select a non-SSD device, and then click **Next**.
5. From the **Select partitioning options** screen, select how you would like to partition the device, then click **Next**. The **Ready to complete** screen displays.
6. Review the options that you selected and if no changes are required, click **Finish**.


Connectivity overview for deployment VM and server

About this task

To deploy the CentOS VM on a deployment server, the VM network must be configured within the network mapping for management purposes.

After the VM is deployed, perform the following steps:

Steps

1. From **Edit settings**, add a **Stamp adapter** to access the VMs.
 2. Add an **ESXi management adapter** to access the ESXi server from the deployment server.
 3. From the console of that VM, assign the static IP addresses to the deployment server for each adapter that is connected to it.
 4. Once the IP address is assigned, open the command console and ping the gateway. If the ping is successful, deploy NFV components.
-  **NOTE: Once the deployment VM and server connectivity are established, install Google Chrome on the access rack servers.**


Deployment VM

About this task

In this document, a deployment VM is used to deploy the solution which can be a virtual machine or a physical server. The deployment VM contains the licenses and required VMware software, ISO, and other required software and licenses necessary for the deployment.

 **NOTE: To deploy the VM, ensure that the Dell 14G servers and network are accessible.**

The CentOS deployment VM is used in this guide as a base operating system platform for the deployment of the NFV Infrastructure (NFVI). The deployment VM performs all the steps involving installation, configuration, and verification of the VMware software stack.

 **NOTE: Before initiating the deployment, ensure that the necessary software firmware is copied or downloaded in the Deployment VM.**

This document provides the steps necessary to install the following applications:

- VMware-VMvisor-Installer-6.7.0.update02
- Microsoft Windows Server 2016 ISO for AD-DNS
- CentOS 7.7 ISO for NTP

- VMware-VCSA-all-6.7U2
- VMware-vRealize-Log-Insight-4.8
- vRealize-Operations-Manager-Appliance-7.5
- NSX-T Manager 2.4
- VMware-vCloud-Director-9.7

Create deployment VM using CentOS

Prerequisites

- VMware ESXi Server 6.7U2
- CentOS 7.6 (or above) ISO file
- Availability of three network adapters:
 - vnic1: For management network
 - vnic2: For stamp-related network
 - vnic3: For management ESXi network
- Available disk storage is greater than 150 GB

Install CentOS

About this task

Follow the steps provided in this section to install CentOS.

Steps

1. Using a browser, open the ESXi hosts.
2. In the navigation pane, right-click the host and then select **Create/Register VM**. The **Select creation type** screen displays.
3. Click **Create a new Virtual Machine**, then click **Next** to continue.
4. From the **Select a name and guest OS** screen, select the following options:
 - a) In the **Name** field, enter the VM name.
 - b) From the **Compatibility** drop-down list, select **ESXi 6.7 virtual machine**.
 - c) Within the **Guest OS family** drop-down, select **Linux** as the operating system family.
 - d) From the **Guest OS version** drop-down list, select **CentOS 7 (64)** and then click **Next**. The **Select storage screen** displays.
5. Select **Datastore** and then click **Next**. The **Customize settings** screen displays.
6. Select the following options:
 - a) In the **CPU** field, set the number to **8**.
 - b) Expand the **CPU** listing.
 - c) Set the number of **Virtual Sockets** to **2**, then set the number of **Check Sockets** to **4**.
 - d) In the fields provided, set the **Memory** to **16 GB**.
 - e) From the **Hard Disk 1** field, set the size to **150 GB**.
 - f) Expand the **Hard Disk 1** listing and set the **Disk Provisioning** option to **Thin Provisioned**.
 - g) Set **SCSI controller 0** to **LSI Logic Parallel**.
 - h) In the **Network Adapter 1** field, select **VM network**.
 - i) In the **CD/DVD Drive 1** field, select **Datastore ISO file**, then click **Next**.

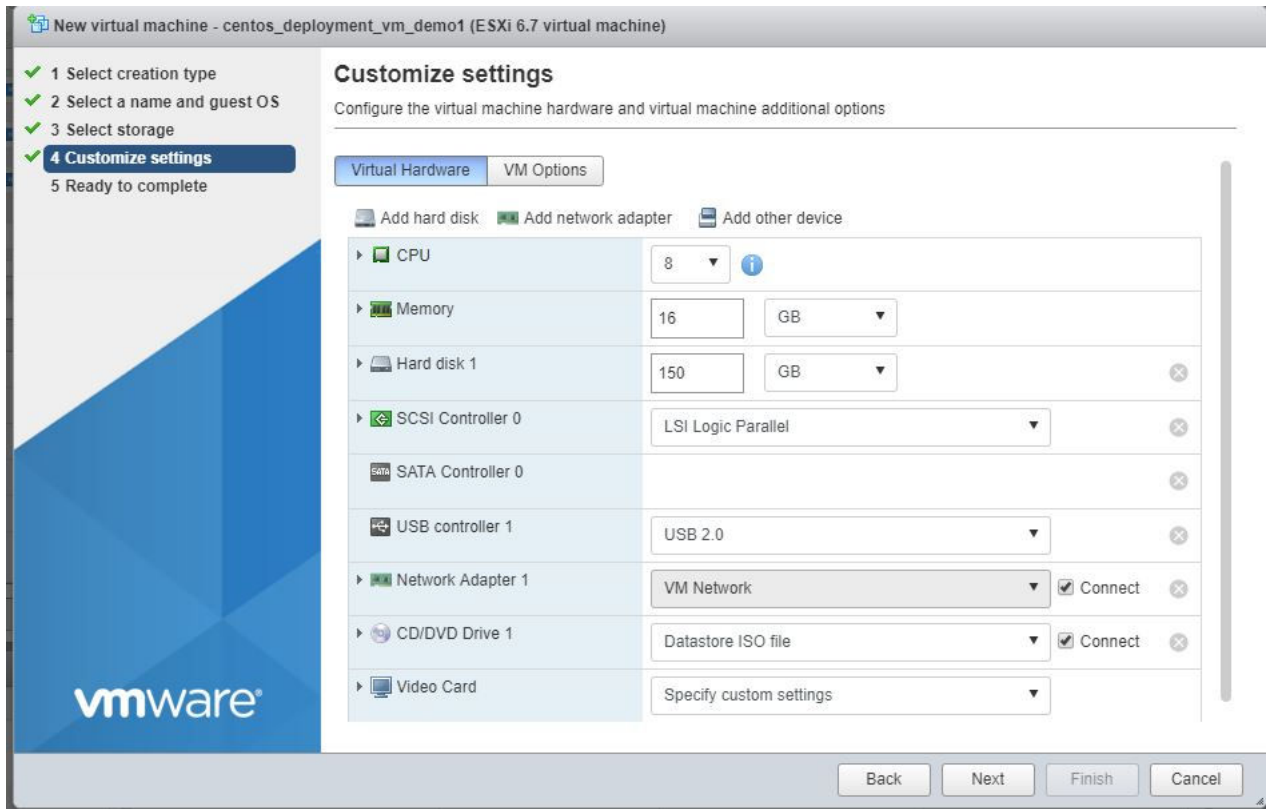


Figure 15. Customize settings screen

The **Ready to complete screen** displays.

7. Review the settings that are selected and then click **Finish**.

8. Power on the **Virtual Machine** and select **Install CentOS 7**.

NOTE: If you do not select the **Install CentOS 7** option, the CentOS installation automatically begins after 60 seconds.

The **Welcome to CENTOS 7** screen displays.

9. Select the wanted language, then click **Continue**.

The **Installation Summary** screen displays.

10. Select the **Software Selection** option, then Select the **GNOME Desktop** radio button, then click **Done**.

The **Installation Summary** screen displays.

11. Click **Installation Destination**.

The **Installation Destination** screen displays.

12. In the Other storage options field, select **Automatically configure partitioning** radio button, then click **Done**.

The **Installation summary** screen displays.

13. Click **Begin Installation**.

NOTE: After the installation is complete, the root password is requested.

14. Set the root password, click **Finish Configuration**, then click **Reboot**.

15. After the system reboots, review the **End User License Agreement (EULA)** and if you agree to the terms, click **Accept**.

16. Review the information provided within the **Privacy** section, then click **Next**.

The **Time Zone** screen displays.

17. Using the selector, choose the appropriate time zone then click **Next**.

The **Online Account** screen displays.

18. Click **Skip**.

19. From the **About You** screen, enter the required username information in the fields that are provided and then click **Next**.

The **Password** screen displays.

20. In the field provided, enter a password, repeat the entry to confirm, and then click **Next**.

21. From the **Ready to go** screen, click **Start using CentOS Linux**.

Configure deployment VM IP

About this task

Configure the deployment VM IP address using the steps provided in this section.

Steps

1. Open the CentOS VM and click **Settings**, then **Network**.
2. From the **Network** screen, click the **Gear** icon.
3. Click the **IPv4** tab and click to select the **Manual** radio button.
4. In the **Addresses** section, enter the **IP address**, **Netmask IP**, and **Gateway IP** for deployment VM in the fields that are provided and then click **Apply**.
The IP is assigned.
5. Restart the network.

Enable automatic connectivity in Network Settings

About this task

Enable the automatic connectivity, to automatically connect with available networks.

Steps

1. From the **Settings** screen, go to **Network**, and then click the **Gear** icon.
2. Click the **Details** tab.
3. Click to place a check in the **Connect Automatically** box, then click **Apply**.

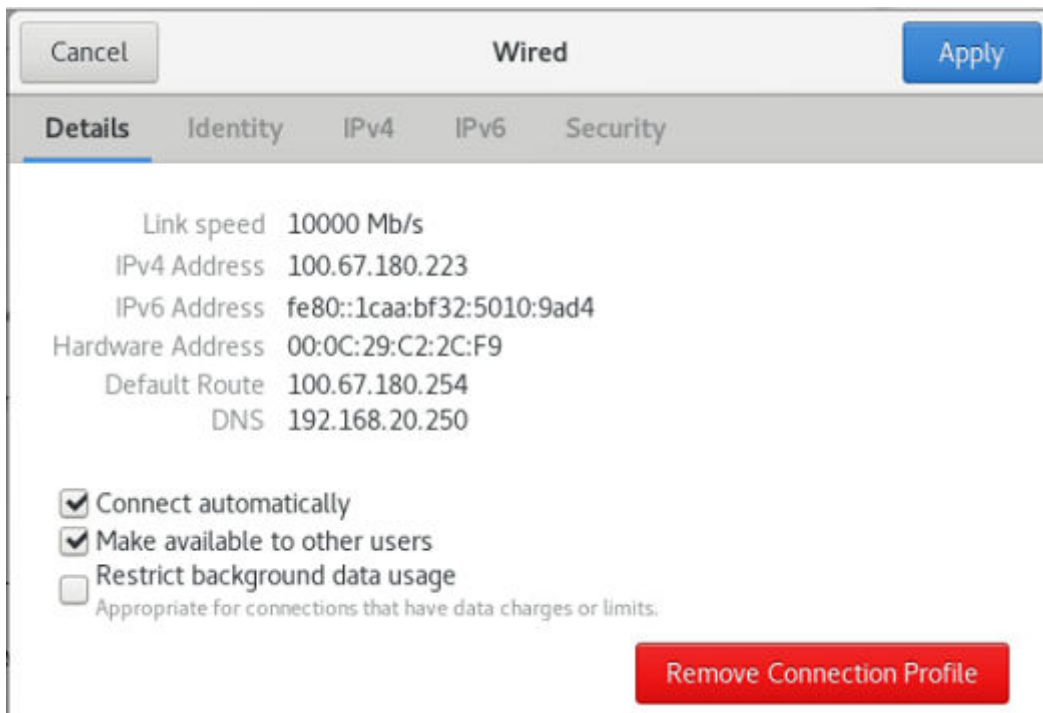


Figure 16. Details tab

Configure deployment settings on deployment VM

Prerequisites

- Deployment VM must have Internet access

About this task

Configure the NTP setting in the deployment VM.

Steps

1. From the CentOS VM, open the terminal.
2. Run the following command to replace the chrony with NTPD:
`# yum remove chrony`
3. Run the following command to disable the firewall:
`# systemctl stop firewalld`
4. Run the following command to install NTP service:
`# yum install ntp`
5. Run the following comment to check NTPD status:
`# systemctl status ntpd.`

NOTE: If the NTPD status shows as running, enter the following command to stop the NTPD service:

```
# systemctl stop ntpd.service
```

6. Run the following commands to restart and enable the NTPD service:

```
# systemctl restart ntpd  
# systemctl enable ntpd
```

Configure time zone

Prerequisites

- Deployment VM must have Internet access

About this task

Set the CentOS timezone to UTC.

Steps

1. From the deployment VM, open the terminal.
2. To verify the time zone being used, run the following command:
`# ls -l /etc/localtime`
3. To search for and change the time zone, run the following command:
`# timedatectl list-timezones | grep UTC`

NOTE: This command displays the list of available time zones. If the UTC time zone is present, the `grep UTC` command displays.

```
[root@localhost ~]# timedatectl list-timezones | grep UTC  
UTC
```

Figure 17. Time zones

4. After verifying that the UTC time zone is present, set the time zone using the following command:
`# timedatectl set-timezone UTC`
5. Run the following command to verify that the time zone is set:
`# ls -l /etc/localtime`

```
[root@localhost ~]# timedatectl set-timezone UTC
[root@localhost ~]# ls -l /etc/localtime
lrwxrwxrwx. 1 root root 25 Jun 12 04:36 /etc/localtime -> ../usr/share/zoneinfo/UTC
[root@localhost ~]#
```

Figure 18. Set time zone confirmation

Disable DHCP script from adding entries to resolv.conf

About this task

Disable the DHCP script from adding entries to the `resolv.conf` during the boot process.

Steps

1. Open the `resolv.conf` file to edit.
2. In the **[main]** section, add the following line:

```
dns=none
```

The DHCP script stops and the DNS entries can be added into the `resolv.conf` file.

```
[main]
#plugins=ifcfg-rh,ibft
dns=none
```

Figure 19. Disable DHCP script

Disable auto mount on CentOS

About this task

By default, the auto mount option is enabled on the CentOS. Disable this file option during the development process to avoid multiple mounts of the ISO files.

Steps

1. From the CentOS VM, open the terminal.
2. Create the `/etc/dconf/db/local.d/00-media-automount` file using the following commands:

```
[org/gnome/desktop/media-handling]
automount=false
automount-open=false
```

3. To check the file, run the following command:
`# cat /etc/dconf/db/local.d/00-media-automount`

NOTE: Confirm that the output displays as:

```
[org/gnome/desktop/media-handling]
automount=false
automount-open=false
```

4. After the file is created, run the following command to save the changes:
`# dconf update`

Install Google Chrome

Prerequisites

- Deployment VM must have Internet access

About this task

This section provides the steps to install Google Chrome on the deployment VM.

Steps

1. From the CentOS VM, open a terminal.
2. Enable the Google YUM repository:
 - a) Create a file, name it `/etc/yum.repos.d/google-chrome.repo`, then enter the following lines of code:

```
[google-chrome]
name=google-chrome
baseurl=http://dl.google.com/linux/chrome/rpm/stable/$basearch
enabled=1
gpgcheck=1
gpgkey=https://dl-ssl.google.com/linux/linux_signing_key.pub
```

3. To install Google Chrome, run the following command:
`# yum install google-chrome-stable.x86_64`
NOTE: During the installation process, confirm each of the installation prompts when presented.
4. Edit the `/usr/bin/google-chrome` file and move the `--no-sandbox -test-type` line of code, to the last line as shown in the following image:

```
# Note: exec -a below is a bashism.
exec -a "$@" "$HERE/chrome" "$@" --no-sandbox -test-type
```

Figure 20. CLI for Google Chrome

NOTE: This line of code removes the need to open Google Chrome using a command line and opens it directly while disabling pop-ups.

5. Run the following command to delete the log in pop-up:
`rm ~/.local/share/keyrings/*`
6. Restart the Google Chrome browser.
7. After the installation of Google Chrome is complete, delete the `google-chrome.repo` file using the following command:
`# rm /etc/yum.repos.d/google-chrome.repo`

Install OVF tool

Prerequisites

- Deployment VM must have Internet access

Steps

1. Download the OVF Tool from following URL: <https://my.vmware.com/group/vmware/details?downloadGroup=OVFTOOL430&productId=742#>

NOTE: You can download the OVF Tool v4.3 from the VMware site using your VMware credentials. Locate the corresponding Linux 64-bit setup file and download it.

The `VMware-ovftool-4.3.0-7948156-lin.x86_64.bundle` downloads.

2. Go to the location where the OVF Tool is downloaded to and open it.

3. From the CentOS VM, open the terminal.
4. Change permissions of the downloaded file: `chmod +x VMware-ovftool-4.3.0-7948156-lin.x86_64.bundle`
5. Run the following command to install the OFV Tool:
`# ./VMware-ovftool-4.3.0-7948156-lin.x86_64.bundle`The **End User License Agreement** (EULA) displays.
6. Review the information that is provided within the EULA and if you accept the terms of the license agreement, click **Next**, and then click **Install**.
The **Installation Complete** screen displays.
7. Click **Finish**.

Add network adapters

Prerequisites


- The respective network adapter must be created.

 **NOTE:** See [Create port group on deployment server](#) for instructions on creating the network adapter.

About this task

In the Deployment VM, the addition of two more network adapters is necessary to access ESXis and VMs. For this deployment, PG-100-ESXi for ESXi management and PG-20-VM-Mgmt for VM Management are added.

Steps

1. Right-click the deployment VM, and select **Edit Settings**.
The **Edit settings** screen displays.
2. Click **Add network adapter** and add two network adapters:
 - For **ESXi management**, select **PG-100-ESXi**
 - For **VM Management**, select **PG-20-VM-Mgmt**
3. Click **Save**.
4. Configure the IP address for each of the network adapters.
 **NOTE:** See [Configure deployment VM IP](#) instructions on assigning the required IP.

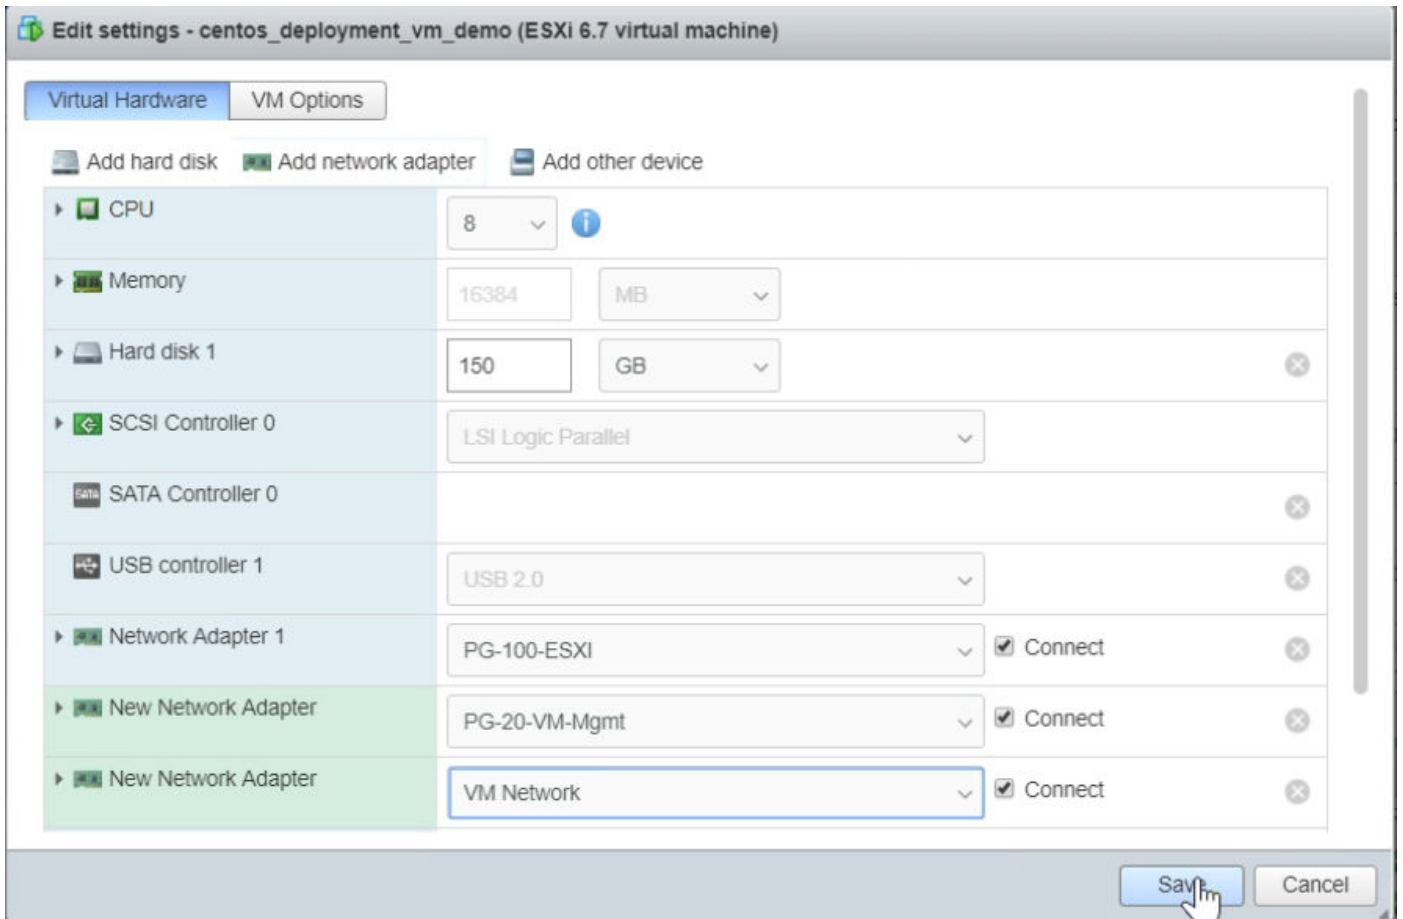


Figure 21. Virtual hardware settings screen

Installation of VMRC

About this task

The VMware Remote Console (VMRC) application is used to open a VM console on a remote host.

Steps

1. From a web browser, download the VMRC from the following location: <https://my.vmware.com/web/vmware/details?downloadGroup=VMRC1006&productId=742>
2. Copy the downloaded `.bundle` file to the local CentOS VM.
3. To make the file an executable, open the terminal and run the following command:
`chmod 777 <VMRC file name>`
4. To install the VMRC, run the following command:
`./VMware-Remote-Console-xx_xx.bundle --console`

NOTE: `xx_xx` is a series of numbers representing the version and build numbers.

5. Follow the installation prompts until the installation is complete.

ESXi installation and configuration

The creation of an NFV infrastructure requires the installation of ESXi on the PowerEdge R640/PowerEdge R740, and PowerEdge R740xd servers based on the vSAN Ready Node.

Use iDRAC9 to install ESXi on PowerEdge R640, R740, and R740xd servers

Prerequisites

- Verify that the minimum required hardware firmware versions are installed on the servers, as described in [Table 2](#)
- ESXi Installer 6.7 U2 or later ISO file
- iDRAC with at least 16 GB SD card enabled

About this task

See the [ESXi installation on deployment server](#) section to install the ESXi on Dell EMC PowerEdge R640/R740 and R720xd servers.

Once the ESXi is installed on the PowerEdge R640, PowerEdge R740, and PowerEdge R740xd servers, see [Customize ESXi](#) and its subsection to configure the ESXi password, update management network configuration, and to change the IPv4 and DNS configuration.

To:

- Add a VLAN ID to the ESXi management network, see [Set VLAN ID for ESXi management network](#)
- Assign licenses to ESXi, see [Assign ESXi license](#)
- Create SSH policies, see [Set SSH policy](#)
- Create firewall rules, see [Set Firewall rules](#)
- Install the DPDK drivers, see [Install DPDK drivers](#)

Set VLAN ID for ESXi management network

About this task

You can set VLAN ID for the ESXi management network.

Steps

1. From the **System Customization** window, select **Configure Management Network**, select **VLAN**, and press **Enter**.
2. In the field provided, enter the configured **VLAN ID** and then press **Enter** to save the change.

Assign ESXi license

About this task

Assign license to ESXi hosts.

Steps

1. From a web browser, open the ESXi, click **Manage**, and then select **Licensing**.
2. In the **License key** field, enter the required license key then click **Check license**.
3. Click **Assign license** then click **Close**.

Set SSH policy

About this task

Set the SSH policy for each ESXi host.

Steps

1. Use the IP address or domain name to go to the ESXi embedded host client.
2. From the left navigation panel, select **Manage** to access the settings for your host.
3. Select the **Services** tab and then select the **TSM-SSH service (SSH)**.
4. Right-click the service name or click to select the **Actions** menu item to set the **Policy** to **Start and stop with host**.

Set firewall rules

About this task

Set the firewall rules for ESXi hosts.

Steps

1. SSH to the ESXi host.
2. When prompted, enter the required credentials.
3. Run the following command to disable the firewall rule:
`esxcli network firewall set --enabled false`
4. To get the status, run the following command:
`esxcli network firewall get`

Install DPDK drivers

Prerequisites

- Download the updated NIC driver from the VMware Compatibility Guide in either VIB or offline bundle format
- Verify that the resource ESXi hosts are in maintenance mode


About this task

On the resource pod, install the necessary drivers for the Intel 25G 2P XXC740/10G X710 NIC cards. The installation of the drivers is required on each of the ESXi hosts that use N-VDS Enhanced mode.

 **NOTE: QLogic drivers do not support the N-VDS Enhanced mode feature.**

Steps

1. Copy the downloaded VIB or offline bundle file to the `/tmp/` directory in the ESXi server.
2. Using the SSH terminal, run the following command:
 - If you are using the VIB file to install the driver, run: `esxcli software vib install -v {VIBFILE path}`
 - If you are using an offline bundle to install the driver, run: `esxcli software vib install -d {OFFLINE_BUNDLE path}`

 **NOTE: Enter the complete path of VIB or offline bundle path of the ESXi server in place of {VIBFILE path} and {OFFLINE_BUNDLE path}. For example, `esxcli software vib install -v /tmp/VMware_bootbank_net-driver.1.1.0-1vmw.0.0.372183.vib`**

Installation of the driver begins.

3. After the installation of the driver is complete, reboot the host.
4. Verify the VMware installation bundle (VIB) versions that are installed on the resource pod hosts.

```
[root@esxi6:~] esxcli software vib list|grep ens
i40en-ens          1.0.4-1OEM.670.0.0.7535516      INT  VMwareCertified  2019-06-10
vmxnet3-ens       2.0.0.21-1vMW.670.0.0.8169922    VMW   VMwareCertified  2019-06-07
```

Figure 22. Installed DPDK drivers

NOTE: Repeat the steps provided in the [ESXi installation and configuration](#) section to install and configure the remaining ESXi servers.

5. Update the VMware NIC drivers. Refer the Operations Guide to update.

NOTE: If you are using Qlogic NICs, add and set the ESXi hosts/Physical NICs to Auto-negotiations.

Next steps

NOTE: Repeat the steps provided in the [ESXi installation and configuration](#) section to install and configure the remaining ESXi servers.

Auxiliary components

Auxiliary components are required for installation on the Dell EMC Ready Solution bundle. Network Time Protocol (NTP), Active Directory (AD), and Domain Name System (DNS) serve as the auxiliary components.

- AD provides a centralized authentication source for management components
- DNS provides forward and reverse lookup services to all platform components
- NTP provides a time synchronization source to all components

Deploy the AD-DNS and NTP virtual machines on the first management ESXi server.

Install auxiliary components

Prerequisites

To install and configure the auxiliary components, create a datastore, standard vSwitch, and a port group. The information in the following sections aid in the installation process:

- [Create datastore on first management ESXi server](#)
- [Create standard vSwitch on first management ESXi server](#)
- [Create port group on first management ESXi server](#)

Create datastore on first management ESXi server

About this task

Create the datastore on the first management ESXi server:

Steps

1. Using a web browser, open the ESXi and log in using the required credentials.
2. From the **Home** screen, click **Storage**, and then click **New datastore**. The **Select Creation Type** screen displays.
3. Select **Create new VMFS datastore** then click **Next**.
4. In the **Name** field, enter the datastore name, select a non-SSD disk device, and then click **Next**. The **Select partitioning options** screen displays.
5. Using the options provided, select how you would like to partition the device, then click **Next**. The **Ready to complete** screen displays.
6. Review the options that you selected and if no other changes are required, click **Finish**.

Create standard vSwitch on first management ESXi server

About this task

By default, vSwitch0 exists on the first ESXi server. Use the settings in the following table to create more virtual switches.

Table 16. vSwitch details

vSwitch name	Uplink	MTU (bytes)	Link discovery	Security
vSwitch1	Vmnic5	9000 bytes	Listen/CDP	• For Promiscuous mode and Forged

vSwitch name	Uplink	MTU (bytes)	Link discovery	Security
				<p>transmits, select the Reject radio button.</p> <ul style="list-style-type: none"> For MAC address changes, select the Accept radio button.

Steps

1. Log in into first management ESXi server.
2. From the navigation panel, click **Networking**.
3. Click the **Virtual switches** tab and select **Add standard virtual switch**. The **Add standard virtual switch** screen displays.
4. Using the information provided in the [vSwitch details table](#) above, update the required information in the **Add standard virtual switch** screen, then click **Add** to create vSwitch.

Create port group on first management ESXi server

About this task

By default, the VM network port group, VLAN ID 0, exists on the ESXi server. To add more port groups, you must create them and assign the VLAN and vSwitch information to the port group.

Table 17. Port group details

Port group name	Description	VLAN ID	Virtual switch	Security
Appliance_Network	For rack ESXi servers (VM management network)	20	vSwitch1	<ul style="list-style-type: none"> For Promiscuous mode and Forged transmits, select the Reject radio button. For MAC address changes, select the Accept radio button.

To create a port group on the first management ESXi server:

Steps

1. Log in into first management ESXi server.
2. From the navigation panel, click **Networking**.
3. On the **Port groups** tab, select **Add port group**. The **Add port group** window displays.
4. Use the information provided in the [Port group details table](#) above, update the required information in the **Add port group** window, then click **Add**.

NTP server configuration

Prerequisites

- Linux CentOS 7.6 or higher VM is installed, running, and configured for network use
- Verify presence of Internet connectivity on the VM

 **NOTE: The Linux VM acts as the NTP server.**

About this task

This section provides the steps to configure the NTP server.

Steps

1. Install the NTP daemon that is provided by default, within the CentOS repository.
2. Run the following command: `# yum install ntp`
NOTE: The `# yum install ntp` command runs when an Internet connection is available.
3. Run the following command to set the time zone to UTC:
`# timedatectl set-timezone UTC`
NOTE: Set the NTP client and NTP server to have the same time zone.
4. After setting the time zone, verify that the system and hardware clocks are synchronized.
NOTE: If the clocks are not synchronized, enter the following command:

```
[root@localhost ~]# date
Wed Jul 25 15:25:44 UTC 2018
[root@localhost ~]# hwclock
Wednesday 26 July 2018 03:25:54 PM UTC -0.646933 seconds
[root@localhost ~]# hwclock --systohc
[root@localhost ~]# hwclock
Wednesday 26 July 2018 03:37:39 PM UTC -0.771275 seconds
[root@localhost ~]# date
Wed Jul 25 15:27:45 UTC 2018
```

NOTE: The `#date` command shows the system clock. The `#hwclock` command shows the hardware clock. The `#hwclock --systohc` command synchronizes the hardware and system clocks.

5. Within the `/etc/ntp.conf` file, enter the following command to enable clients from your network to synchronize the time with this server:

```
restrict ::1
# Hosts on local network are less restricted.
restrict 192.168.20.0 netmask 255.255.255.0 nomodify notrap
# Use public servers from the pool.ntp.org project.
# Please consider joining the pool (http://www.pool.ntp.org/join.html).
#server 0.centos.pool.ntp.org iburst
#server 1.centos.pool.ntp.org iburst
#server 2.centos.pool.ntp.org iburst
#server 3.centos.pool.ntp.org iburst
server 127.127.1.0
```

NOTE: The `restrict 192.168.20.0 netmask 255.255.255.0 nomodify notrap` command allows you to restrict access to the network. The IP range sees the network address in the production environment. Multiple IP ranges can be added.

6. Enter the following IP command for the NTP server configuration: `127.127.1.0`
NOTE: The `server 127.127.1.0` line in the `ntp.conf` file configures itself as an NTP server.
7. Enter the following commands to add firewall rules:

```
# firewall-cmd --add-service=ntp --permanent
# firewall-cmd --reload
```

8. Activate the NTPD service by entering the following commands:

```
# systemctl restart ntpd
# systemctl enable ntpd
# systemctl status ntpd
```

NOTE: When using the commands to restart the VMs, services are required to run again.

9. Enter the following command to ensure that NTP is running properly:
`# ntpq -p`

```

[root@localhost ~]# ntpq -p
      remote           refid      st t when poll reach  delay  offset  jitter
-----
*LOCAL(0)          .LOCL.    5 l  43  64   1   0.000   0.000   0.000
[root@localhost ~]#

```

Figure 23. Confirmation of NTP status

Synchronize ESXi clocks using NTP

About this task

Before you install vCenter Server or deploy the vCenter Server Appliance, ensure that the machines on the vSphere network have their clocks that are synchronized.

Steps

1. Using a web browser, connect to the ESXi host.
2. Select **Manage** then click **System, Time & Date**, and then **Edit Settings**. The **Edit time configuration** screen displays.
3. Click to select **Use Network Time Protocol (Enable NTP client)**.

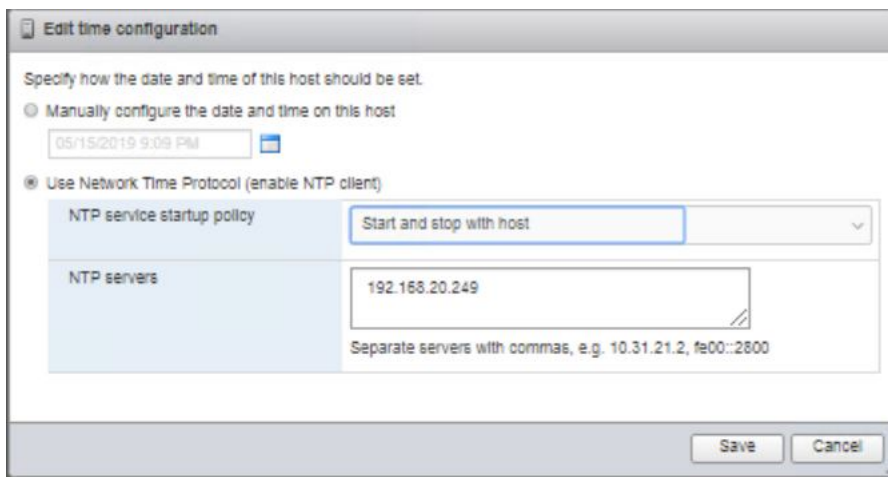


Figure 24. Edit time configuration window

4. In the **NTP Server** field, enter the IP address or fully qualified domain name of one or more NTP servers to synchronize.
5. Set the startup policy and service status as **Start and stop with host** then click **OK**.
6. To save the changes and service, select **Time and date**, click **Actions, NTP service**, and then click **Start**.

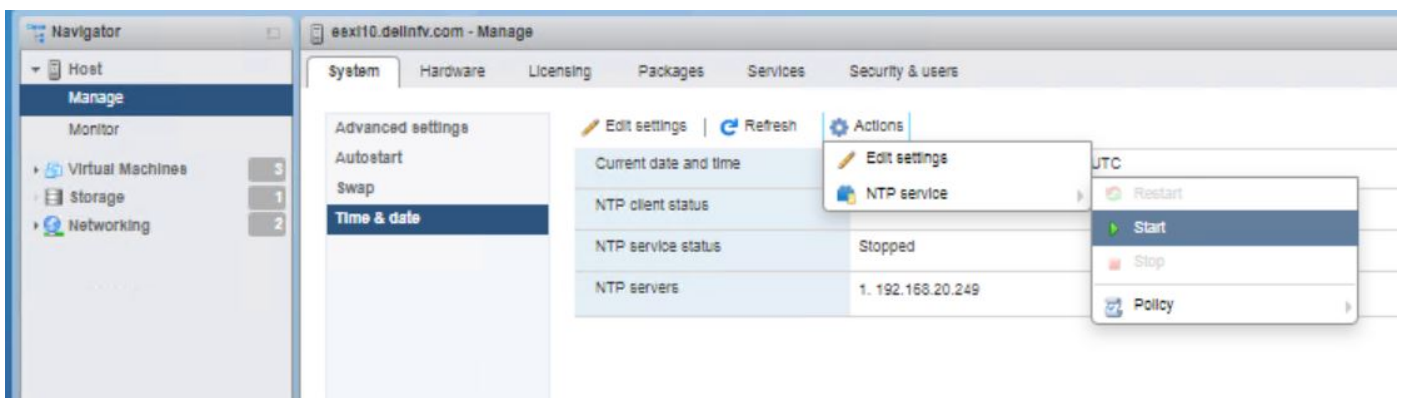


Figure 25. Time and date start action

7. To synchronize the ESXi clock with NTP, repeat the steps in this section on each of the ESXi servers.

Synchronize VM clock using NTP

About this task

Synchronize the Microsoft Windows machine time with the NTP server if the system running Microsoft Windows is not part of a domain controller.

Steps

1. From the task bar, right-click the time that is listed and select **Adjust date/time settings**.
The **Date and Time** screen displays.
2. Select the **Internet Time** tab, then click **Change Settings**.
The **Internet Time Settings** screen displays.
3. Click to place a check in the **Synchronize with an Internet time server** option.
4. Enter your NTP server IP or FQDN in the **Server** field that is provided, and then click **Update Now**.
5. Click **OK** to save the changes.

Microsoft Windows Server 2016 installation for AD DNS

Install Microsoft Windows Server 2016

Microsoft Windows Server 2016 uses AD, and DNS auxiliary services. This section describes the steps that are used to install Microsoft Windows Server 2016.

NOTE: The server using AD, DNS, and NTP auxiliary services are deployed on the ESXi datastore and not on the vSAN datastore. It is recommended that you create a datastore on the ESXi server for AD, DNS, and NTP and that is part of the management cluster. Use the ESXi datastores for deployment of the AD, DNS, and NTP server.


When the cluster is in place, you must migrate the AD, DNS, NTP server, and their VMs from the ESXi datastore to the vSAN datastore. Once the AD, DNS, and NTP server VMs are migrated, delete the ESXi datastore.

Prerequisites

- Minimum 1.4 GHz 64-bit processor
- Minimum 4 GB memory
- 40 GB of disk space or greater
- ESXi datastore
- Verify that the VM has connectivity to the Internet

Steps

1. Create a VM and select the Windows Server 2016 ISO, or later, to deploy the AD-DNS services.
NOTE: Setup of the necessary files may take several minutes.
2. From the **Start** screen, select **Next**.
3. When prompted, select the required language, time, and keyboard selections, then click **Next**.
4. Click **Install now**.
The **Setup is starting** window displays until the application setup is complete.
5. On the **Active windows** screen, enter the required **Microsoft Windows license key for Windows Server 2016** and then click **Next**.
6. From the **Select the operating system** screen, select **Windows Server 2016 Datacenter/Desktop Experience** and click **Next**.
7. Review the **End User License Agreement** and if you agree to the terms, click to select the **I accept the license terms** box and then click **Next**.
The **Windows Setup Installation Type Selection** screen displays.
8. Select **Custom: Install Windows only (advanced)**.
The **Where do you want to install Windows?** screen displays.

9. Verify the location of the drive or partition, then click **Next**.
10. From the **Settings** window, enter the desired administrator password, reenter the password to confirm it, then click **Finish**. The setup process finalizes the settings.
 **NOTE: This process may take several minutes to complete.**
11. After the setup is complete, log in with the Administrator credentials. The **Server Manager** displays.
12. After the VM is created on ESXi, launch the VM using the VMware Remote Console.

Configure network for Microsoft Windows Server 2016 and VMware tool installation

About this task

Configure the network for Microsoft Windows Server 2016.

Steps

1. From the **Network and Sharing Center**, click right-click the Ethernet to configure and then select **Properties**. The **Internet Protocol Version 4 (TCP/IPv4)** screen displays.
2. Configure the IP address for the AD/DNS server and then click **OK**.

Install VMware tools

About this task

To install VMware tools:

Steps

1. From the Windows VM, go to the location where the ESXi is installed.
2. Right-click the **VM preview** option and select **Guest OS**, and then **Install VMware Tools**.
3. From the **Windows VM console (UI)**, locate the drive.
4. Double-click the **VMware Tools installation** option.
5. For the **Setup type**, select **Typical**, then click **Install**.

Active Directory and DNS installation

Update Windows VM computer name

About this task

Before creating the AD DNS server, you must change the Windows VM computer name.

Steps

1. From the **Server Manager** window, select **Local Server** in the left-navigation pane. The **Properties** window displays.
2. Right-click the **Computer Name** field, and select **System properties**. The **System Properties** screen displays.
3. Click the **Computer Name** tab, then select the **Change** button. The **Computer Name/Domain Changes** screen displays.
4. Locate the **Computer Name** field and enter a computer name for the **AD DNS server**.
5. Click **Restart Now**. The virtual machine restarts.

Install primary Active Directory and DNS

About this task

This section provides the steps to install the primary AD-DNS.

Steps

1. From the **Server Manager** dashboard, click **Add roles and features**.
2. Review the information that is provided on the **Before you Begin** screen, then click **Next**.
3. Select the **Role-based or feature-based installation** option then click **Next**.
The **Server Selection** screen displays.
4. Select the server role to install from the options that are provided and click **Next**.
5. In the **Server Roles** window, check the **Active Directory Domain Services** box then click **Next**.

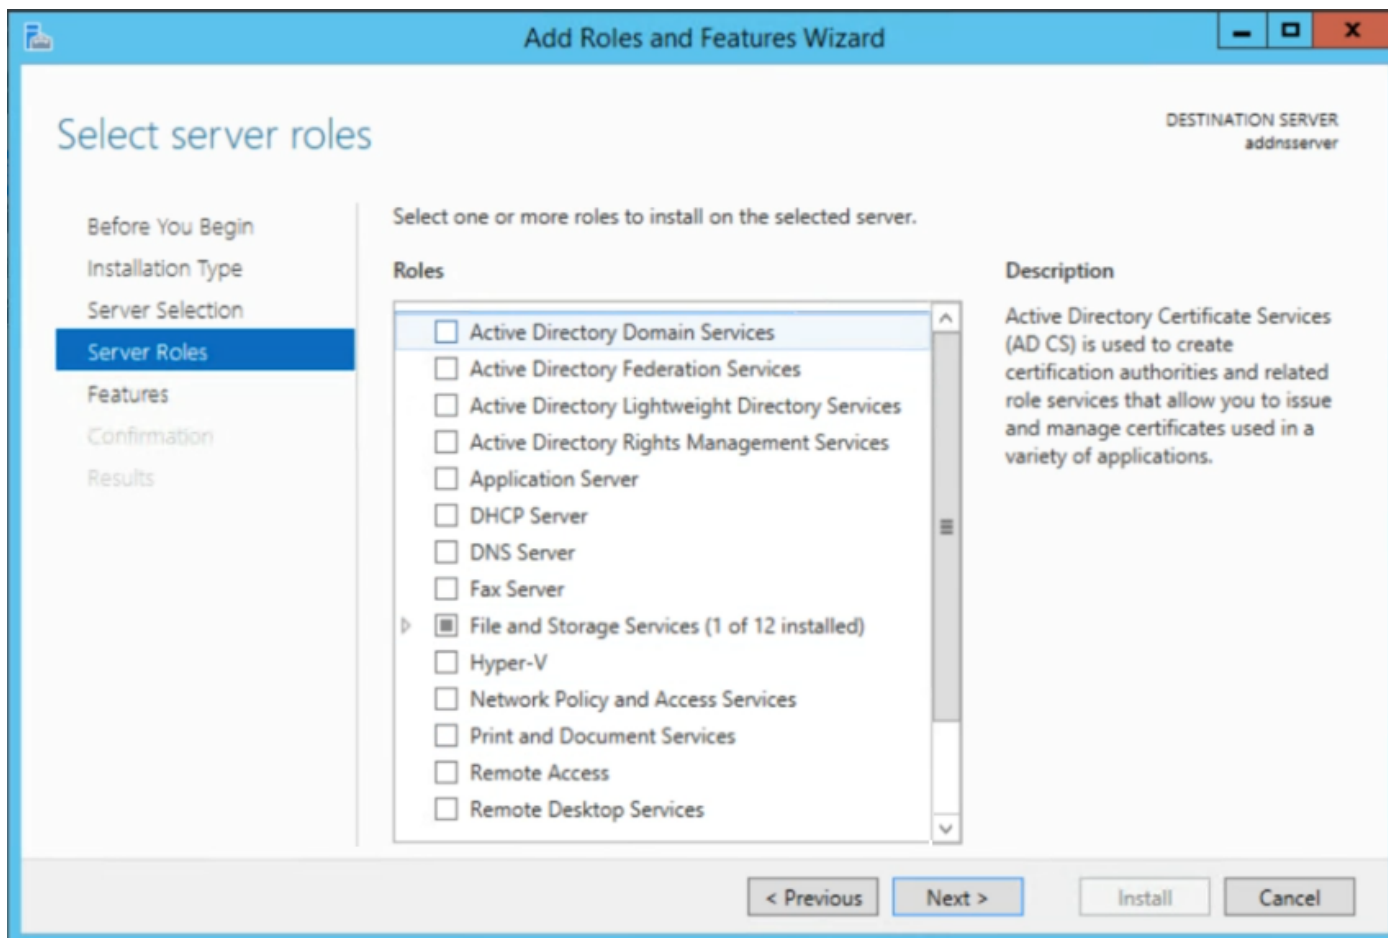


Figure 26. Server roles selection screen

6. From the **Add Roles and Features Wizard**, select the **Add Features** button and then click the **DNS Server** listing.
7. From the **Add features that are required for DNS Server** window, click the **Add Features** button.

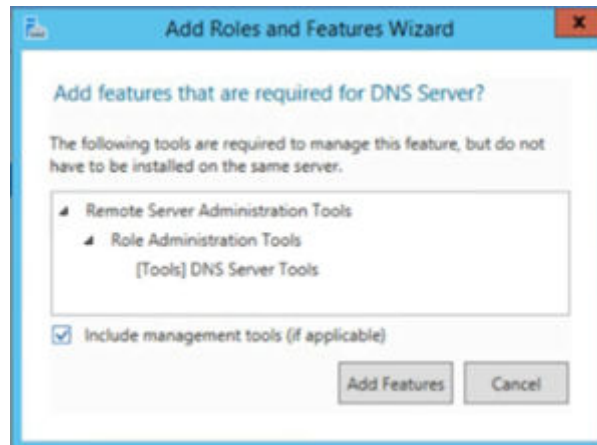


Figure 27. DNS Server required features window

8. Click **Next**.
9. From the **Features** screen, review the options that are selected and verify that the default selections are kept, then click **Next**.
10. Review the information provided on the **AD DS** screen, then click **Next**.
11. Review the information that is provided on the **DNS Server** screen, then click **Next**.
12. From the **Confirm installation selections** screen, carefully review each of the selections, and then click **Install**.

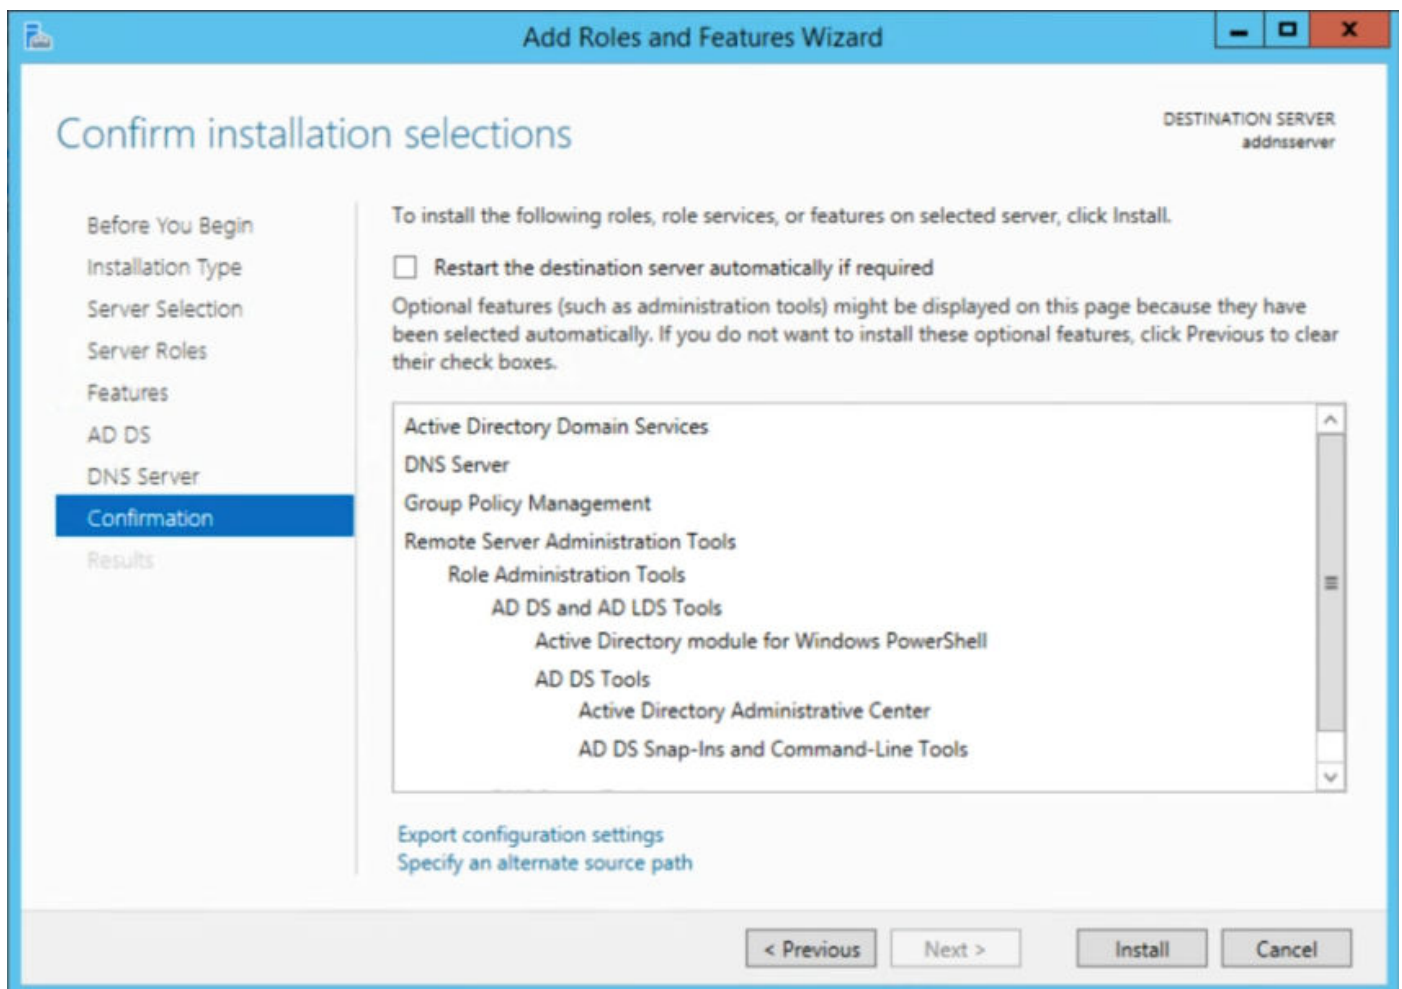


Figure 28. Confirm installation selections window

13. Once the installation is complete, click **Close** to exit the wizard.
14. Reboot the server to complete the installation.
15. Once the server reboots, access the **Server Manager** and from the **Dashboard**, click the **Task flag**, that is located in the navigation bar.

16. From the options provided, click the **Promote this server to a domain controller** hyperlink.

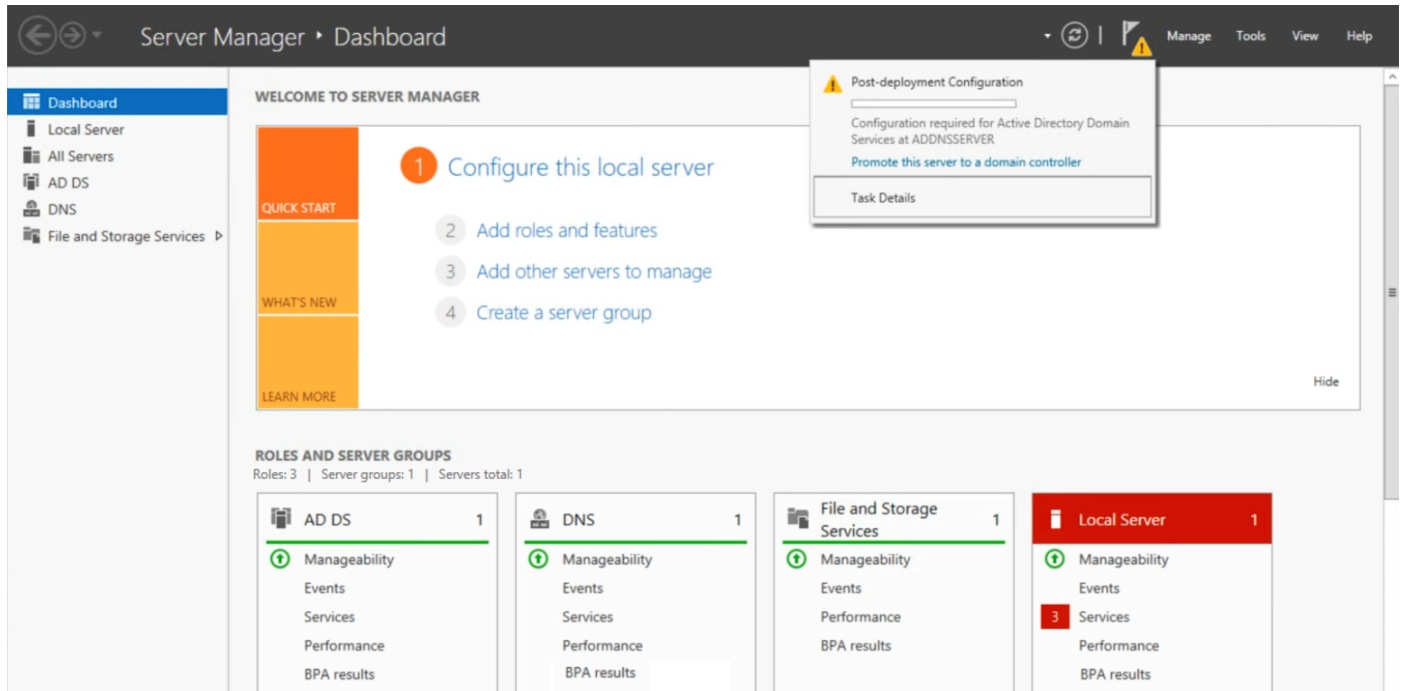


Figure 29. Server Manager Dashboard task flag advisory

The **Deployment Configuration Introduction** screen displays.

17. Click **Add a new forest**.

18. In the **Root domain name** field, enter the wanted name then click **Next**.

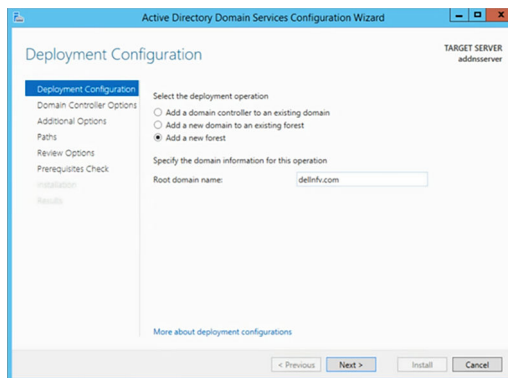


Figure 30. Deployment configuration window

19. Review the selections within the **Domain Controller Options** screen.

- a. In the **Password** field, enter the password for the **Directory Services RestoreMode (DSRM)**, reenter the new password in the **Confirm password** field, and then click **Next**.

NOTE: If a delegation error displays within the **DNS Options** window, disregard the error and click **Next**.

The **Additional Options** screen displays.

20. Review the NetBIOS name in the field that is provided then click **Next** to continue.

The **Paths** screen displays.

21. The **Paths** screen allows you to adjust the assigned locations of the AD DS database, log files, and SYSVOL folders. Perform the necessary changes, and then click **Next**.

The **Review Options** screen displays.

22. Review the selections, and click **Next**.

The system check begins to verify the compatibility of the system with the options selected. When complete, review the results of the prerequisites check.

NOTE: A successful prerequisites check displays a green check at the top of the window. Any critical errors that are found must be addressed before the option to begin the installation is provided.

23. Click **Install**. The server automatically reboots after the installation process is complete.
24. Log in to the server using domain administrator credentials.

Create DNS Reverse Lookup Zone

About this task

The Reverse Lookup Zone is not created by default if the DNS server is newly configured. The steps provided in this section to help with the creation of the Reverse Lookup Zone for the IPv4 address range. In this deployment, two DNS reverse lookup zones are created: one for VM management network and one for the ESXi management network.

Steps

1. From the **Server Manager Dashboard**, click **Tools** in the top navigation panel, then click **DNS**. The **DNS Manager** window opens.
2. Right-click the **Reverse Lookup Zone** in the left-navigation panel, and then click **New Zone**. The **Welcome to the New Zone Wizard** window opens.
3. Click **Next**.
4. From the **Zone Type** screen, select the zone to create then click **Next**.

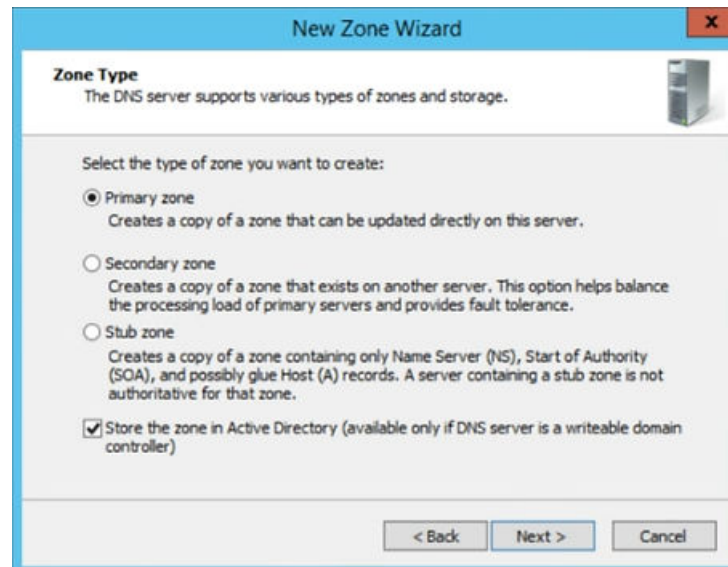


Figure 31. Zone Type selection window

5. From the **Active Directory Zone Replication Scope** screen, select how the DNS data is replicated, then click **Next**. The **Reverse Lookup Zone Name** screen displays.
6. Select **IPv4 Reverse Lookup Zone** and click **Next**.
7. Verify that the **Network ID** option is selected and enter the **Network ID** in the field that is provided, then click **Next**.
8. Select **Allow both Secure dynamic updates (recommended for Active Directory)**, then click **Next**.
9. Click **Finish** to complete the configuration.
10. Repeat the steps provided in this section to set up the ESXi Management network.

NOTE: For ESXi Management setup, enter **192.168.100.x**, where **x** is the remainder of the ID, in the Network ID field.

Add DNS server host

About this task

This section provides the required steps to add DNS server host.

Steps

1. From the **Server Manager Dashboard**, click **Tools**, and then **DNS**.
2. In the left-navigation panel, click to expand the **Forward lookup zone** listing.
3. Right-click your domain, and select **New Host**.
The **New Host** window opens.
4. In the fields provided, enter the **Hostname**, **IP address**, click to place a check in the **Create associated pointer (PTR) record** box, and then click **Add Host**.

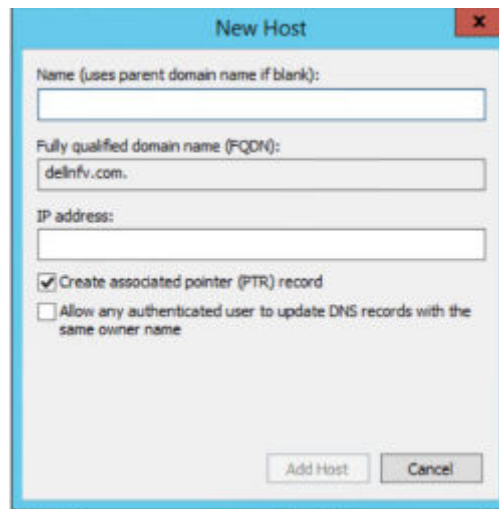


Figure 32. New Host configuration screen

Disable firewall

About this task

This section provides the steps to disable the firewall.

Steps

1. From the **Control Panel**, select **System and Security**, **Windows Firewall**, and then **Customize Settings**.
2. Set the **Windows firewall for the network settings**, to **Off** and then click **OK**.

Add self-signed certificate to Windows Active Directory

Prerequisites

- WinSCP installed on a Microsoft Windows VM to copy files from Windows to Linux
- Installation of PuTTY to run commands from a Microsoft Windows VM (optional)

About this task

This section provides the steps to add a self-signed certificate in AD.

Steps

1. On a Linux machine with open-SSL installed, run the following command to generate the `ca.key` file:

```
$ openssl genrsa -des3 -out ca.key 4096
```
2. When prompted, enter the wanted password.
3. On the Linux machine, run the following command to generate the `ca.crt` file:

```
$ openssl req -new -x509 -days 3650 -key ca.key -out ca.crt
```
4. Using WinSCP or a similar tool, copy the `ca.crt` file to the **Active Directory Windows VM** from the Linux machine.
5. From the **Windows Active Directory VM**, click **Run** and enter `certlm.msc`.
6. Once the location is found, import the copied certificate.

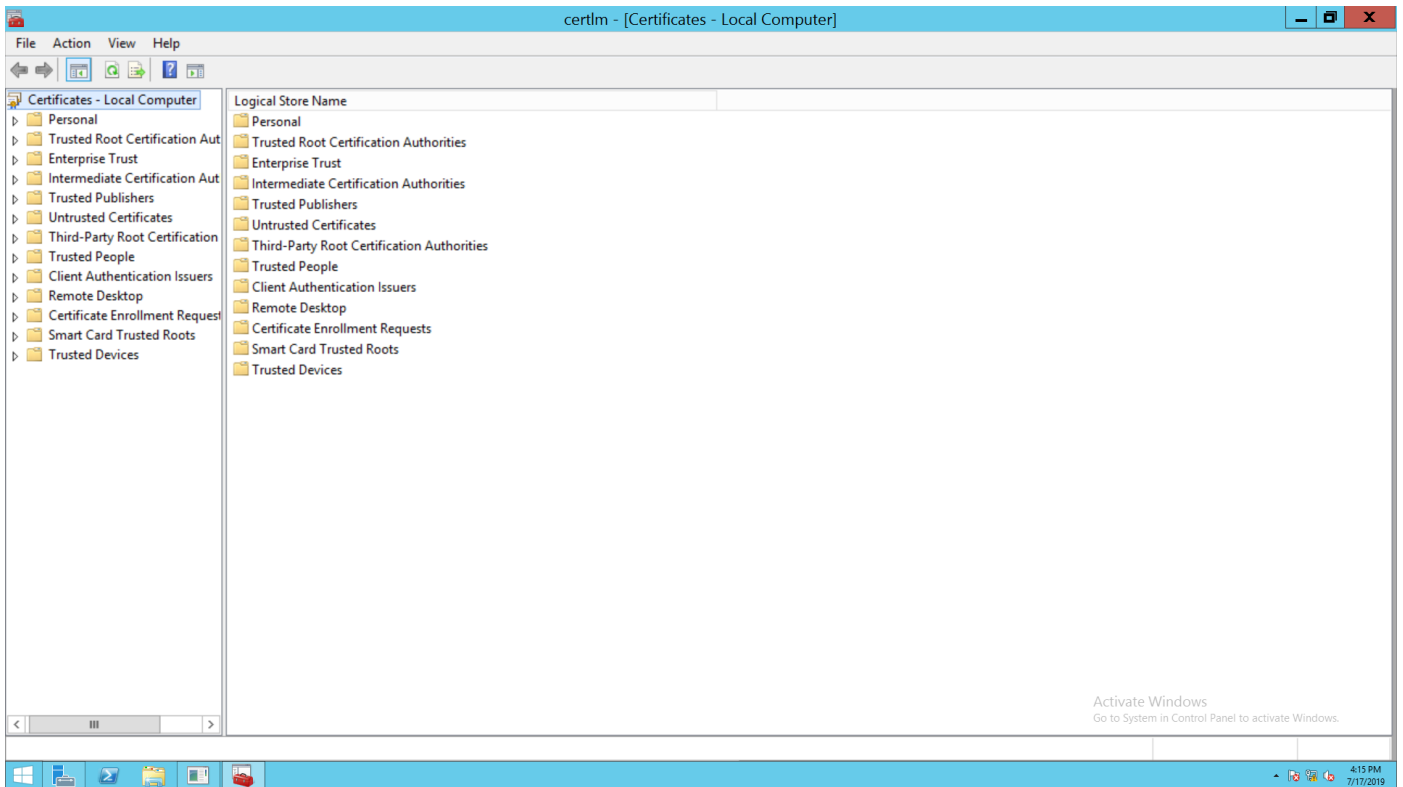


Figure 33. Manage computer certificates window

7. Right-click the **Trusted Root Certification Authority** listing within the **Logical Store Name** section, select **All Tasks**, then select **Import**.
The **Welcome to the Certificate Import Wizard** window opens.
8. Click **Next** to continue.
9. Click the **Browse** button and select the `ca.crt` file and then click **Next**.
10. Verify that the **Certificate store** listed is correct, then click **Next**.
11. Click **Finish** to import the certificate.
The **Import was successful** message displays.
12. Click **OK**.
13. Verify that the imported certificate displays in the **Trusted Root Certification Authority** section.




Figure 34. Trusted Root Certification Authority window

- From the **Windows AD VM**, create a `request.inf` file.
- Copy and paste the following text into the file:

```
[Version]
Signature="$Windows NT$"

[NewRequest]
Subject = "CN=ACTIVE_DIRECTORY_FQDN"
KeySpec = 1
KeyLength = 1024
Exportable = TRUE
MachineKeySet = TRUE
SMIME = FALSE
PrivateKeyArchive = FALSE
UserProtected = FALSE
UseExistingKeySet = FALSE
ProviderName = "Microsoft RSA SChannel Cryptographic Provider"
ProviderType = 12
RequestType = PKCS10
KeyUsage = 0xa0

[EnhancedKeyUsageExtension]
OID = 1.3.6.1.5.5.7.3.1 ; Server Authentication
```

- Replace the `ACTIVE_DIRECTORY_FQDN` text in quotes, with the FQDN of the Windows AD VM, and save the changes.

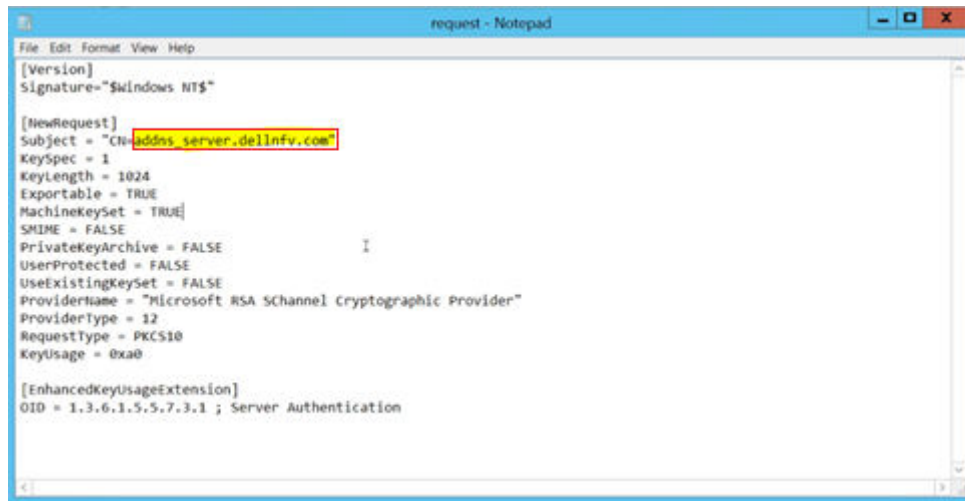


Figure 35. Example of request.inf file

17. Open the PowerShell and 'cd' to the directory where the request.inf file is saved.
18. Run the following command:
certreq -new request.inf client.csr The client.csr file is created.
19. Using WinScp or a similar tool, copy the client.csr file to the Linux system.
20. Create a v3ext.txt file on the Linux system.
21. Paste the following text into the v3ext.txt file:

```

keyUsage=digitalSignature,keyEncipherment
extendedKeyUsage=serverAuth
subjectKeyIdentifier=hash

```

22. Using PuTTY or by pasting it directly, enter the following command within the Linux system:
\$ openssl x509 -req -days 3650 -in client.csr -CA ca.crt -CAkey ca.key -extfile v3ext.txt -set_serial 01 -out client.crt
23. When prompted, enter the password for the ca.key.
The client.crt file is created on the Linux system.
24. Use WinSCP to copy the client.crt file to the Windows AD VM.
25. Run the following command in the Windows AD VM PowerShell, within in the same directory as the client.crt file:
certreq -accept client.crt
26. In the **Windows AD VM**, create a file that is named ldap-renewservercert.txt.
27. Paste the following text within the file:

```

dn:
changetype: modify
add: renewServerCertificate
renewServerCertificate: 1
-

```

28. Save the changes that you made to the file.
29. From the PowerShell, enter the following command:
ldifde -i -f ldap-renewservercert.txt

NOTE: Ensure that the PowerShell is in the same directory as the ldap-renewservercert.txt file.

Configure NTP client in AD VM

Prerequisites

- Verify that the NTP server is up and running before the adding NTP on a Microsoft Windows VM
- Ensure that you can ping the NTP server from the Windows VM

About this task

This section provides the steps to configure NTP client on windows machine.

Steps

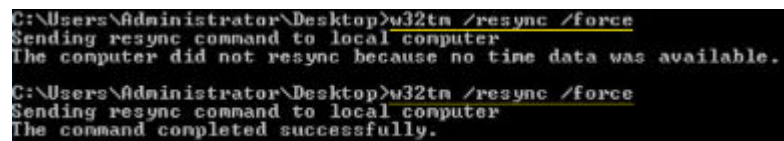
1. Enter the following commands in the order provided:

```
C:\ net stop w32time
C:\ w32tm /config /syncfromflags:manual /manualpeerlist:"<enter ntp server ip here>"
C:\ w32tm /config /reliable:yes
C:\ net start w32time
```

2. After the commands have been entered, enter the following command:

```
C:\ w32tm /resync /force
```

i **NOTE:** If *The computer did not resync because no time data was available.* message displays, reenter the command.



```
C:\Users\Administrator\Desktop>w32tm /resync /force
Sending resync command to local computer
The computer did not resync because no time data was available.

C:\Users\Administrator\Desktop>w32tm /resync /force
Sending resync command to local computer
The command completed successfully.
```

Figure 36. Command to resync windows with NTP server window

VMware vCenter Server deployment and configuration

In the VMware vCloud NFV 3.0 architecture, two instances of the VMware vCenter Server Appliance (VCSA) with embedded PSC are deployed. One VCSA instance manages the management cluster, while the second instance manages the resource cluster and edge cluster. Management pod hosts both of these instances. Each VCSA instance required to configure in HA mode and consist of active, passive, and witness vCenter Server instance.

VMware vCenter Server Appliance deployment

Prerequisites

- ESXi 6.7 U2 is configured and running
- AD-DNS and NTP are running
- Manual creation of forward and reverse lookup entries for all VCSA instances on the DNS server before deployment

About this task

In the vCloud NFV 3.0 architecture, the deployment of the VCSA is done in two stages. The first deployment is for the VCSA instance that manages the management cluster, and the second VCSA deployment manages the Edge and resource cluster.

Stage 1 - Deploy ISO file for Management vCenter Server Appliance with embedded PSC

About this task

Stage 1 of the deployment process involves the deployment of the ISO file in the VCSA installer, as a vCenter Server Appliance with an embedded PSC. To deploy the ISO file:

Steps

1. Mount the VMware-VCSA-all-6.7.x ISO on Windows/Linux VM Deployment server.
2. Go to the path where VCSA installer is mounted and go to the `vcsa-ui-installer` directory, then to the subdirectory for your operating system, and run the installer executable file.
 - For **Microsoft Windows**, go to the **win32** subdirectory, and run the `installer.exe` file
 - For **Linux**, go to the `lin64` subdirectory, and run the installer file.
3. From the **vCenter Server Appliance 6.7 Installation** window, click **Install**. The **Introduction** screen displays.
4. Review the installation overview, and click **Next**.
5. Review the information provided within the **End User License Agreement (EULA)** and if you agree to the terms check the **I accept the terms of the license agreement** box and, click **Next**. The **Select deployment type** screen displays.
6. From the **Embedded Platform Services Controller** section, select **vCenter Server with an Embedded Platform Services Controller** then click **Next**. The **Appliance deployment target** screen displays.
7. Enter the **ESXi host**, **HTTPS port**, **User name**, and **Password** in the fields that are provided, then click **Next**.
8. When prompted, review the contents of the **Certificate Warning**, then click **Yes** to accept the certificate.
9. On the **Set up appliance VM** window, enter the **VM name**, set the root password in the fields that are provided, and then click **Next**.

NOTE: The appliance name can contain upper and lower-case letters, however special characters such as %, \, or / are not permitted. The appliance name must not exceed 80 characters in length.

- From the **Select deployment size** screen, use the drop-down within the **Deployment size** and **Storage size** sections to select the sizes necessary for the vCenter Server Appliance for your vSphere inventory, then click **Next**.
- From the **Select datastore** screen, select **Install on new vSAN cluster containing the target host option**, enter the **Datacenter** and **Cluster Name**, and then click **Next**.
The **Claim disks for vSAN** screen displays/
- Select all of the disks, select the **Enable Thin Disk Mode** box then click **Next**.
- In the **Configure Network Settings** window, enter the appropriate network settings details, then click **Next**.
- NOTE:** Dell EMC recommends the use of an FQDN. However, if an IP address is used instead, use a static IP address allocation for the appliance as IP addresses allocated by DHCP may change.
- From the **Ready to complete stage 1** screen, review the deployment settings for the vCenter Server Appliance and click **Finish**.
The deployment process starts.

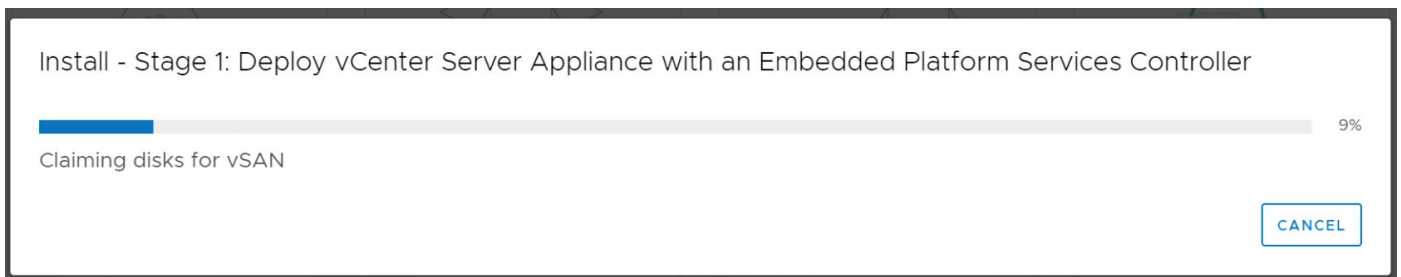


Figure 37. Stage 1 status window

- Once the deployment is complete, click **Continue** to proceed with the stage 2 deployment process.

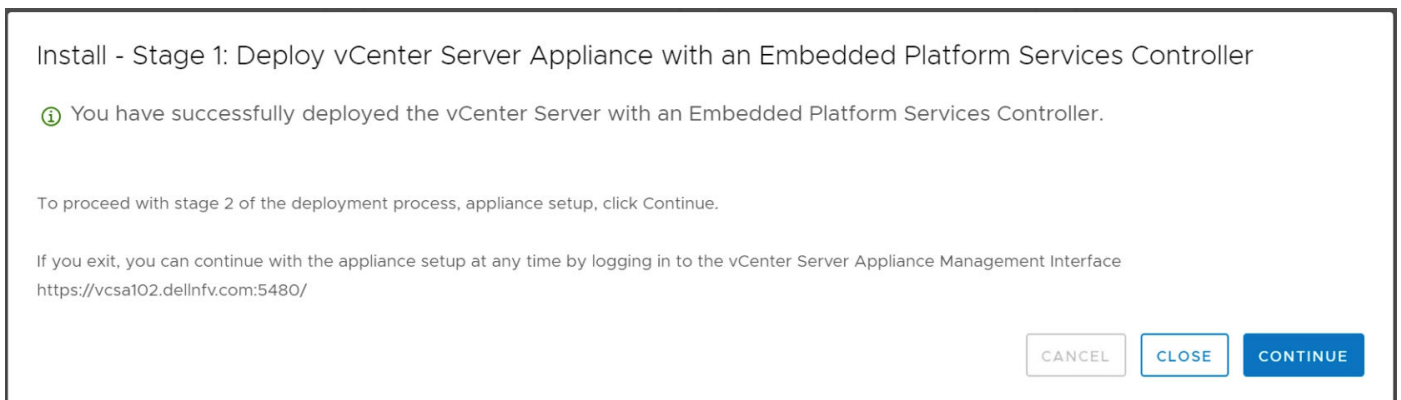


Figure 38. Stage 1 completion and continuation window

Stage 2 - Set up Management vCenter Server Appliance with embedded PSC

About this task

After the stage 1 deployment is complete, you are redirected to stage 2 of the deployment process to set up and start the services of the newly deployed vCenter Server Appliance with an embedded Platform Services Controller.

Steps

- Review the information within the **Introduction to stage 2** page, then click **Next**.
- From the **Appliance configuration** screen:
 - Locate the **Time synchronization mode** drop-down list and select **Synchronize time with NTP servers**.
 - In the **NTP Servers** field, enter the NTP server IP address.
 - From the **SSH access** drop-down list, select **Enabled**, then click **Next**.

NOTE: From this screen, the option to enable remote SSH access to the appliance, is provided.

The SSO configuration screen displays.

- In the field provided, enter the **vCenter Single Sign-On domain name**, and **administrator password**, then click **Next**.
 - Optionally, you can opt to participate in the **Customer Experience Improvement Program** by selecting the **Join the VMware Customer Experience Improvement Program (CEIP)** box, then click **Next**.
- NOTE:** The Customer Experience Improvement Program (CEIP) provides VMware with information that enables VMware to improve its products and services and to fix problems. By choosing to participate in CEIP, you agree that VMware may collect technical information about your use of VMware products and services regularly. This option is enabled by default.
- From the **Ready to complete** page, review the configuration settings for the **vCenter Server Appliance**, click **Finish**, and then click **OK**.

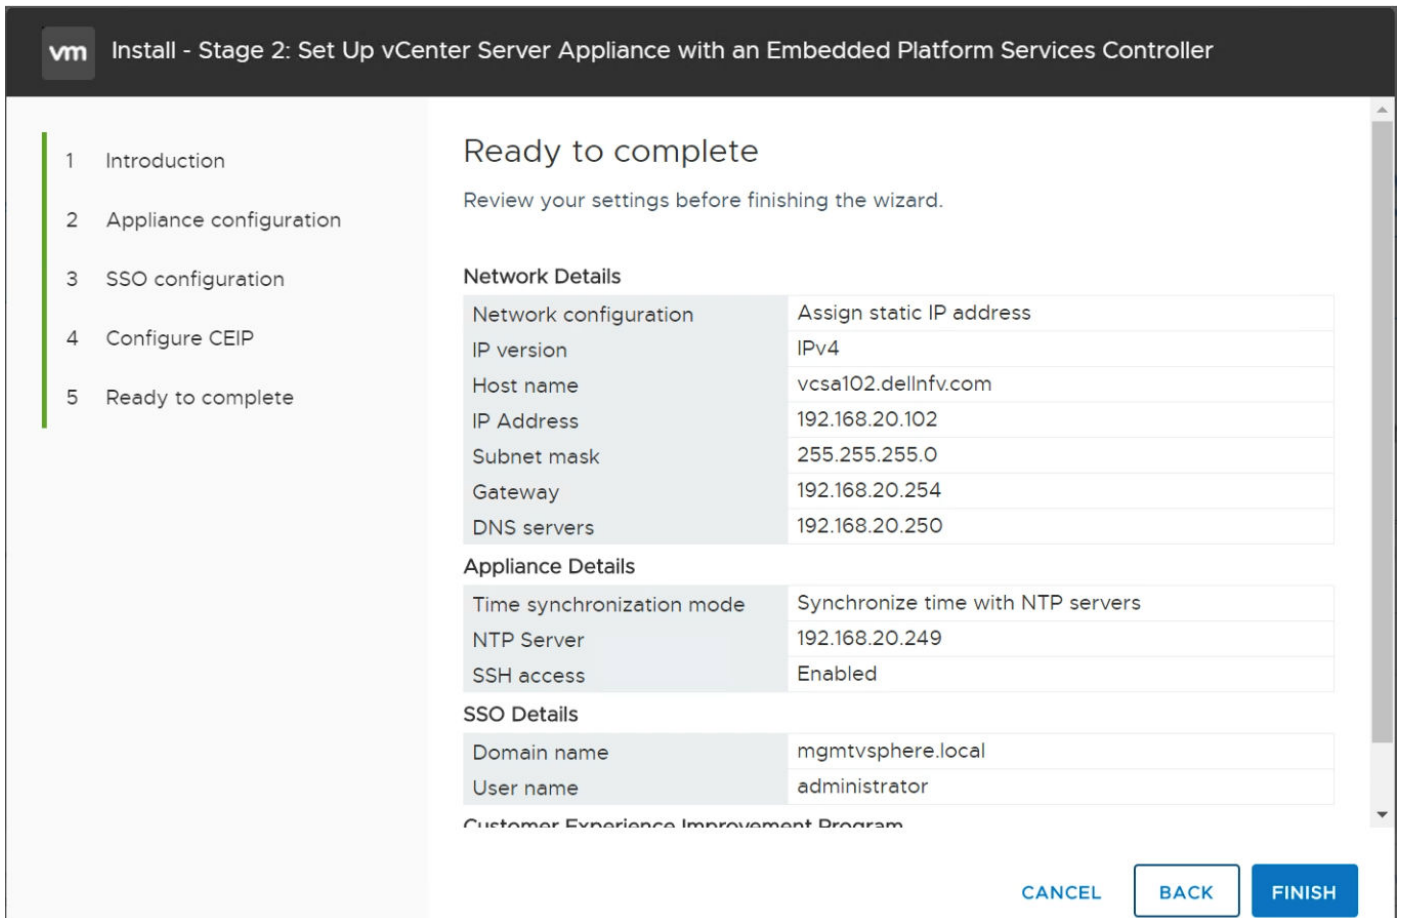


Figure 39. Ready to complete window

Stage 2 of the deployment process and set up of the appliance is complete.

- Optionally, once the initial setup is complete, open the browser and go to the following URL: https://<vcenter_server_appliance_fqdn>:443. Optionally, from the **Appliance Getting Started Page** click the https://<vcenter_server_appliance_fqdn>:443 link to access the **vCenter Server Appliance Getting Started** page. Otherwise, click **Close** to exit the wizard.

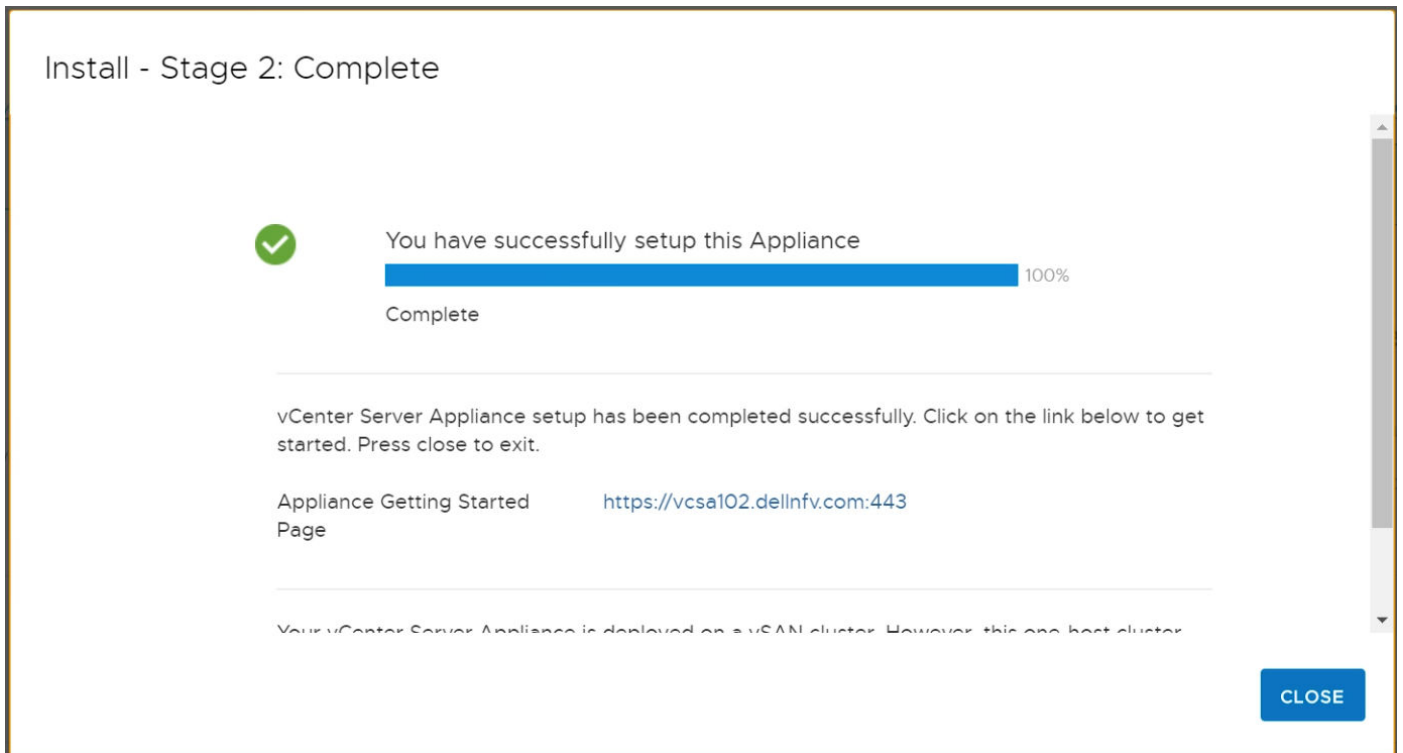


Figure 40. Stage 2 completion window

Change vSAN default storage policy for management vCSA

About this task

Before deploying extra VMs to the system, you must first update the VMware vSAN storage policy.

Steps

1. From a web browser, log in to the vCSA web client.
2. Click the **Home** icon and select **Policies and Profiles**.
3. In the left navigation pane, click **VM Storage Policy**, and select **vSAN Default Storage Policy**.
4. On the **Manage** tab, select the **Rule-set 1: VSAN** option.
5. Click the **Edit** button.
The **Rule-set 1** window displays.
6. Locate the **Primary level of failures to tolerate** option and enter **1** in the field provided.
7. From the **Force provisioning** drop-down, click to select **Yes**.

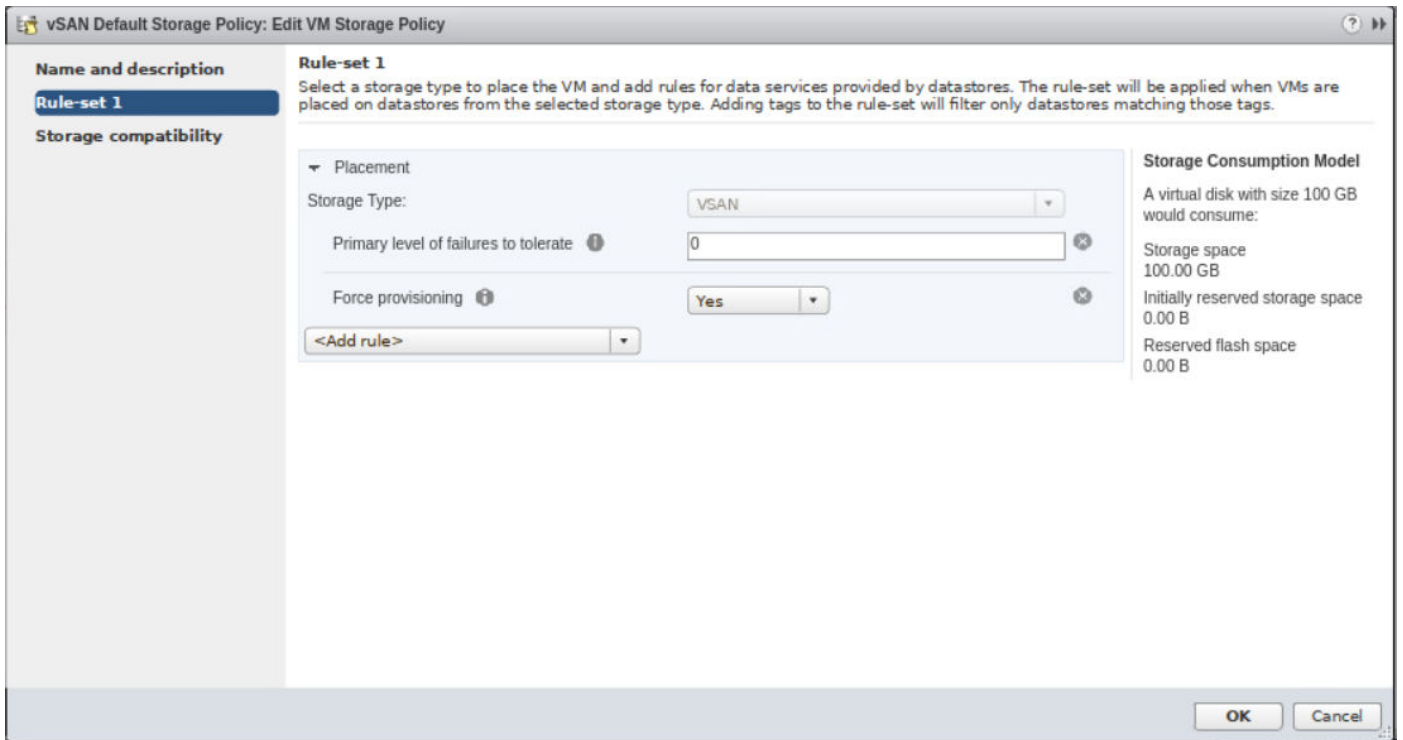


Figure 41. Rule-set 1 settings window

8. Click **OK** to save the changes.



Stage 1 - Deploy ISO file for Resource vCenter Server Appliance

About this task

Stage 1 of the deployment process involves the deployment of the ISO file. The ISO file is part of the vCenter Server Appliance installer as a vCenter Server Appliance with an embedded PSC on the management vCenter.

Steps

1. Mount the VMware-VCSA-all-6.7.x ISO on Windows/Linux VM Deployment server.
2. Go to the path where vCenter Server Appliance installer is mounted, and go to the **vcsa-ui-installer** directory, then to the subdirectory for your operating system, and run the installer executable file.
 - For **Microsoft Windows**, go to the **win32** subdirectory, and run the `installer.exe` file.
 - For **Linux**, go to the **lin64** subdirectory, and run the installer file.
3. From the **vCenter Server Appliance 6.7 Installation** window, click **Install**. The **Introduction** screen displays.
4. Review the installation overview then click **Next**. The **End User License Agreement (EULA)** screen displays.
5. Review the information that is provided within the EULA and if you agree to the terms, select the **I accept the terms of the license agreement** box and, click **Next**. The **Select deployment type** screen displays.
6. Select **vCenter Server with an Embedded Platform Services Controller** option in the **Embedded Platform Services Controller** section, and click **Next**. The **Appliance deployment target** screen displays.
7. In the fields provided, enter the **vCenter Server name**, **HTTPS port**, **User name**, and **Password** and then click **Next**. The **Certificate Warning** window display.
8. Review the contents of the **Certificate Warning**, if agreed then click **Yes** to accept the certificate. The **Select folder** window opens.


9. Select the **Management Datacenter** option, and then click **Next**.
10. From the **Select compute resource** screen, select **Require ESXi to deploy**, then click **Next**. The **Set up appliance VM** screen displays.
11. In the fields provided, enter the **VM name** and set the **Root password** and then click **Next**.
 **NOTE: The appliance name can contain upper and lower case letters, however special characters such as %, \, or / are not permitted. The appliance name must not exceed 80 characters in length.**
- The **Select deployment size** screen displays.
12. Select the **Deployment size** and **Storage size**, then click **Next**. The **Select datastore** screen displays.
13. Select the **vsanDatastore** option, and check the **Enable Thin Disk Mode** box, then click **Next**. The **Configure Network Settings** screen displays.
14. Enter the appropriate network settings details, and then click **Next**.
 **NOTE: Dell EMC recommends the use of an FQDN. If an IP address is used, the use of static IP address allocation for the appliance is recommended, as the IP addresses that the DHCP allocates may change.**
15. From the **Ready to complete Stage 1** screen, review the deployment settings for the VCSA and click **Finish**. The deployment process starts.
16. Once the Stage 1 deployment is complete, click **Continue** to proceed with the **Stage 2** deployment process.

Stage 2 - Set up resource vCenter Server Appliance with embedded PSC

About this task

After the deployment of Stage 1 is complete, you are redirected to Stage 2. The Stage 2 deployment process sets up and starts the services of the newly deployed VCSA with an embedded Platform Services Controller.

Steps

1. Review the information within the **Introduction to Stage 2** screen, then click **Next**.
2. From the **Appliance configuration** screen, perform the following steps:
 - a. From the **Time synchronization mode** drop-down list, select **Synchronize time with NTP servers**.
 - b. In the **NTP Servers** field, enter the NTP server IP address.
 - c. From the **SSH access** drop-down list, select **Enabled**, then click **Next**.
 **NOTE: From this window, the option to enable remote SSH access to the appliance is provided.**
3. From the **SSO configuration** screen, enter the **vCenter Single Sign-On domain name**, **User name**, and **Administrator password** in the fields that are provided, then click **Next**.
4. On the **Ready to complete** window, review the configuration settings, click **Finish**, then click **OK**. Stage 2 of the deployment process and set up of the appliance is complete.
 Once the initial setup is complete, open a web browser and go to https://<vcenter_server_appliance_fqdn>:443. Alternately, from the appliance, go to the **Getting Started** page option and click the https://<vcenter_server_appliance_fqdn>:443 link to access the **vCenter Server Appliance Getting Started** page.
5. Click **Close** to exit the wizard.

Change VMware vSAN default storage policy for resource vCSA

About this task

Before deploying more VMs to the system, you must first update the VMware vSAN storage policy.

Steps

1. From a web browser, log in to the VCSA web client.

2. Click the **Home** icon, and select **Policies and Profiles**.
3. In the left navigation pane, click **VM Storage Policy**, and select **vSAN Default Storage Policy**.
4. From the **Manage** tab, select the **Rule-set 1: VSAN** option.
5. Click the **Edit** button.
The **Rule-set 1** window opens.
6. Locate the **Primary level of failures to tolerate** option and enter **1** in the field provided.
7. From the **Force provisioning** drop-down, click to select **Yes**.
8. Click **OK** to save the changes.

Add AD authentication for vCenter Server

Prerequisites

- Deployment of the Management VCSA must be complete
- Deployment of the Resource VCSA must be complete

About this task

This section provides the steps to add AD authentication for vCenter Server.

Steps

1. Open the VCSA web client and log in as single sign-on administrator.
 ⓘ **NOTE: To access the vCSA web client, go to https://VCSA_FQDN_OR_IP.**
2. From the **Home** screen, click **Administration, System Configuration**, select the appropriate VCSA, and then click the **Manage** tab. The **Manage** screen displays.
3. Locate the **Active Directory** section and click the **Join** button.
4. In the fields provide, enter the **Domain name**, **User name**, and **Password**, then click **OK**.
 ⓘ **NOTE: The Organizational unit information is optional.**
5. Reboot the node to apply the changes. Once the node reboots, log in to the **VMware vCenter Server**.
6. Click **Home, Administration, System Configuration**, select the appropriate VCSA, and then select the **Manage** tab.
7. Locate the **Active Directory** section and confirm that the domain is listed in the **Domain** field.

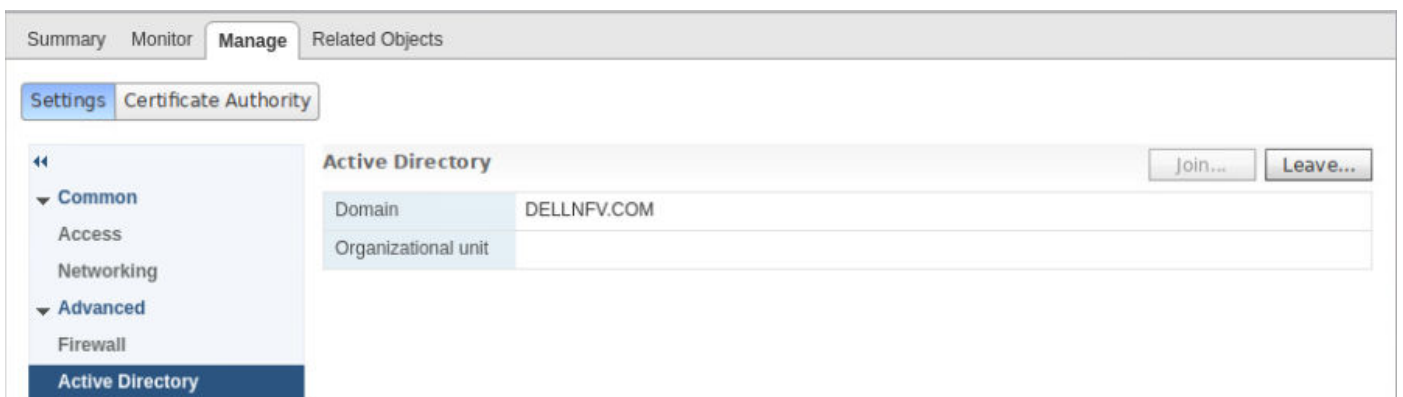


Figure 42. Active Directory window

8. Click the **Home** icon, and then select **Administration, Single Sign-On**, and then click **Configuration**. The **Configuration** screen displays.
9. Click the **Identity Sources** tab, and then click the **Add (+)** icon to add an identity source. The **Add Identity source** screen displays.
10. From the **Add identity source** screen, select the **Active Directory (Integrated Windows Authentication)** source type from the options listed.
 ⓘ **NOTE: The underlying system must be a member of the Active Directory domain.**
11. Enter the **Domain name** in the field provided, and then click **Next**.

- Review the domain name and then click **Finish**.
The identity source displays in the **Identity Sources** tab.



Figure 43. Identity Sources tab with new identity listed

NOTE: Repeat the steps in this section for resource vCenter server.

Assign license to vCSA

About this task

This section provides the steps to add license to VCSA.

Steps

- From a web browser, log in to the **vCenter** through **vSphere Web Client**.
- Click the **Home** icon, and select **Administration**.
- From the left navigation pane, click **Licenses**.
- Click the **Add (+)** icon and in the field that is provided, enter the license key, and then **Next**.
- In the field provided, add a name for the license and then click **Next**.
The license displays in the **Licenses** tab and is added to vCenter.

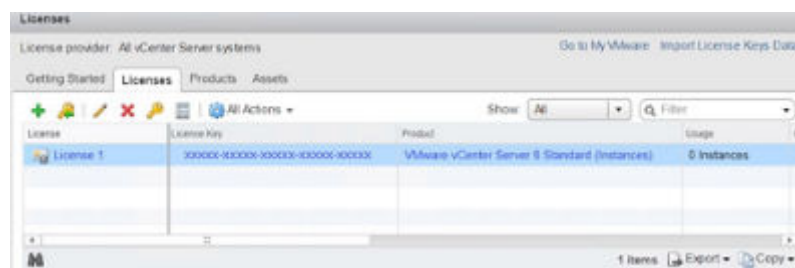


Figure 44. Listing of licenses within Licenses tab

- Click the **Assets** tab, and then the **Solutions** tab.
- Select the **vCenter Server** from the **Solutions** list.
- From the **All Actions** drop-down menu, select **Assign** license.
- Select the license that you want to apply and then click **OK** to assign the license.

Create data center and cluster on resource vCenter

About this task

After the installation of the management and resource vCenter, create the data center, resource cluster, and edge cluster to enable vSAN, DRS, and vSphere HA.

 **NOTE:** The management data center and cluster are created while the management vCSA is deployed.

Steps

1. Using the required credentials, log in to the **VMware vCenter Server Web Client** using the required credentials.
2. From the left-navigation panel, right-click the desired vCenter server.
3. Click **New Datacenter** from the options provided.
4. To create resource cluster, right-click the newly created data center and select **New** cluster.
5. Enter a name for the new cluster in the **Name** field, then click **OK**.
6. To create the Edge cluster, repeat steps 5 and 6.

Add hosts to vCenter cluster

About this task

When adding hosts to the vCenter cluster, a minimum four hosts must be added to the management, resource, and edge clusters. Once the clusters are created, add the hosts to the cluster.

Steps

1. To add a host to the vCenter cluster, right-click the cluster and select **Add host**.
The **Add Host** window opens.
2. Enter the name or IP address of the host and click **Next**.
The **Connections settings** screen displays.
3. Enter the required host **User name** and **Password** for the connection, and click **Next**.
The **Security Alert** screen displays.
4. Click **Yes** to replace the host certificate with a new certificate that is signed by the **VMware Certificate Server**.
The **Host Summary** screen displays.
5. Review the information, and click **Next** to continue.
6. From the **Assign license** screen, click to select the license that is listed, then click **Next**.
7. Within the **Lockdown mode** screen, select **Disabled** and click **Next**.
8. Review the selected configurations within the **Ready to complete** screen and if no other changes are required, click **Finish**.
Ensure that all the hosts are added to the management, edge, and resource cluster:

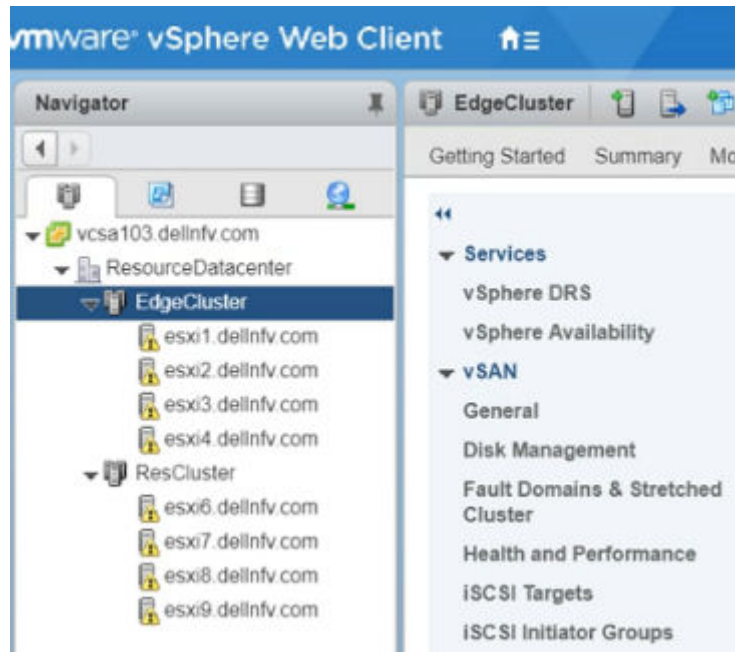


Figure 45. Sample Edge and resource cluster listing

9. Repeat the steps within this section to add more ESXi hosts to the vCenter Cluster. Each cluster must have a minimum of four hosts added to it.

Enable VMware enhanced vMotion compatibility

Prerequisites

Table 18. VMware vMotion compatibility requirements

Requirements	Description
ESXi version	ESXi 6.7 U2
vCenter Server	The host must be connected to a vCenter Server system
CPUs	A single vendor, either AMD or Intel
Advanced CPU features enabled	Enable these CPU features in the BIOS if they are available: <ul style="list-style-type: none"> • Hardware virtualization support (AMD-V or Intel VT) • AMD No eXecute(NX) • Intel eXecute Disable (XD)

- Power off all virtual machines in the cluster that are running on hosts with a feature set greater than the EVC mode that you intend to enable
- All hosts in the cluster must meet the following requirements

About this task

To improve CPU compatibility between hosts that have varying CPU feature sets, you can hide some host CPU features from the virtual machine by placing the host in an Enhanced vMotion Compatibility (EVC) cluster. Hosts in an EVC cluster and hosts that you add to an existing EVC cluster must meet EVC requirements.

Enable VMware EVC for management cluster

About this task

This section provides the steps to enable the VMware EVC for management cluster.

Steps

1. Select the **Management Cluster** in the **vSphere Web Client**.
2. On the **Configure** tab, go to **VMware EVC** and click **Edit**.
3. Select the **Enable EVC for Intel Hosts** radio button from the **Select EVC Mode** option and from the **VMware EVC Mode** drop-down list select appropriate processor for the hosts to add to the cluster, then click **OK**.

Enable VMware EVC for resource and edge cluster

About this task

This section provides the steps to enable the VMware EVC for resource cluster and edge cluster.

Steps

1. From the **VMware vSphere Web Client**, go to the **Resource Cluster**.
2. Click the **Configure** tab.
3. From the **Services** tab, go to **VMware EVC** and click **Edit**.
4. Select the **Enable EVC for Intel Hosts** radio button from the **Select EVC Mode** option, and from the **VMware EVC Mode** drop-down list select the appropriate processor for the hosts to add to the cluster, then click **OK**.
5. To enable EVC for the edge cluster, repeat the steps in this section on the edge cluster.

Configure virtual network

Figure 6, Figure 7, and Figure 8 display the underlying virtual distributed switch, or VDS, for the management, edge, and resource clusters. Different VLAN IDs can be used in the physical environment.

VDS creation and configuration for management pod

See the following sections to create and configure VDS on management pod:

- [Create VDS for management pod](#)
- [VDS configuration settings for management VDS](#)
- [Create LAG for management pod](#)
- [Create distributed port group for management pod](#)
- [Add host to VDS on management pod](#)

Create VDS for management pod

About this task

This section provides the steps to create vSphere Distributed Switch (VDS) for the deployment. The [VDS-management settings](#) table in this section displays the VDS configuration settings that are used for VDS in the management cluster.

Table 19. VDS-management settings

Distributed switch name	Version	Number of uplinks	Network I/O control	Discovery protocol type/operation	MTU (bytes)
Infra_Network_VDS	6.6.0	2	Enabled	CDP/Both	9000 Bytes
VM_Network_VDS	6.6.0	2	Enabled	CDP/Both	9000 Bytes

Steps

1. In the **VMware vSphere Web Client**, open the **Networking View** tab.
2. Right-click the **Management Datacenter** and select **Distributed switch**, and then **New Distributed Switch**.
3. In the **New Distributed Switch** window, see the [VDS-management settings](#) table in this section and enter the **Switch name**, then click **Next**.
4. From the **Select version** screen, select **Distributed switch: 6.6.0** and click **Next**.
5. From the **Edit settings** screen, see the [VDS-management settings](#) table and select the number of required uplinks.

NOTE: The number of uplinks that are used depends on the number of physical NICs associated to the VDS.
6. Verify that the **Network I/O Control** option is set to **Enabled**.
7. From the **Default port group** drop-down, verify that the **Create a default port group** box is not selected, and click **Next**.
8. Review the selected settings on the **Ready to complete** screen, and if no changes are required, click **Finish**.

NOTE: Using the information provided within the [VDS-management settings](#) table in this section, create the **VM_Network_VDS** distributed switch.

Once the VDS' are created, the VDS' display in the **Networking** tab.




Figure 46. VMware vSphere Web Client Networking tab

VDS configuration settings for management VDS

Steps

1. After the distributed switches are created, select each distributed switch and click the **Configure** tab, **Properties**, then click **Edit**.
2. From the **Advanced** tab, set the **MTU** to **9000 bytes**.
3. Locate the **Discovery protocol** section and from the **Type** drop-down list select **Cisco Discovery Protocol**.
4. Use the drop-down arrow next to the **Operation** field and select **Both** and then click **OK**.
5. Repeat the steps in this section for each of the remaining management VDS switches.

Create LAG for management pod

About this task

This section provides the steps to create LAGs on VDS for management pod. The following table displays the required LAG configuration settings to create LAG.

Table 20. VDS-LAG settings

Distributed switch	Name	Number of ports	Mode	Load-balancing mode
Infra_Network_VDS	MgmtLag1	2	Active	Source and destination IP address and TCP/UDP port
VM_Network_VDS	MgmtLag1	2	Active	Source and destination IP address and TCP/UDP port

Steps

1. From the left navigation panel, click to select the **Infra_Network_VDS** and then click the **Configure** tab.
2. In the **Settings** section, select **LACP**, click **Add (+)** to create the **Migrating network traffic to LAGs**.
3. Using the information from the [VDS-LAG settings](#) table, enter the **Lag name**, **Number of ports**, **Mode**, and **Load-balancing mode** and then click **OK**.

Figure 47. New Link Aggregation Group screen

NOTE: Using the information provided within the [VDS-LAG settings](#) table in this section, create the LAG on the `VM_Network_VDS` distributed switches.

Create distributed port group for management pod

About this task

The information in this section assists with the creation of more distributed port groups for the distributed switch. The port group settings that are used for VDS-management cluster are shown in the following table:

Table 21. VDS-Management port group settings

Port group	VLAN type	VLAN ID	Teaming and failover settings				
			Load balancing	Network failure detection	Notify switches	Failback	Active uplinks
ESXi_Mgmt_Network (under Infra_Network_VDS)	VLAN	100	Route based on IP hash	Link status only	Yes	Yes	MgmtLag1

Port group	VLAN type	VLAN ID	Teaming and failover settings				
			Load balancing	Network failure detection	Notify switches	Failback	Active uplinks
vSAN_Network (under Infra_Network_VDS)	VLAN	300	Route based on IP hash	Link status only	Yes	Yes	MgmtLag1
vMotion_Network (under Infra_Network_VDS)	VLAN	200	Route based on IP hash	Link status only	Yes	Yes	MgmtLag1
Replication_Network (under Infra_Network_VDS)	VLAN	500	Route based on IP hash	Link status only	Yes	Yes	MgmtLag1
VM_Mgmt_Network (under VM_Network_VDS)	VLAN	20	Route based on IP hash	Link status only	Yes	Yes	MgmtLag1
VCSA_HA_Network (under VM_Network_VDS)	VLAN	30	Route based on IP hash	Link status only	Yes	Yes	MgmtLag1

Steps

1. Right-click the newly created distributed switch, and select **Distributed Port Group**, and then **New Distributed port group**.
2. In the **New Distributed Port Group** window, use the information in the table above to enter the **Port group name** in the fields provided, then click **Next**.
3. From the **Configure settings** screen, set the general properties of the new port group as follows:
 - **Port binding:** Static binding
 - **Port allocation:** Elastic
 - **Number of ports:** 8
 - **Network resource pool:** (default)
 - **VLAN type:** VLAN
 - **VLAN ID:** See the [VDS-Management port group settings](#) table to add VLAN ID
4. Select the **Customize default policies** configuration box, then click **Next**. The **Security and Traffic Shaping** screen displays.
5. Click **Next** as defaults settings are used. The **Teaming and failover** screen displays.
6. Using the information from the [VDS-Management port group settings](#) table, select the **Load balancing**, **Network failure detection**, **Notify switches**, **Failback**, and **Active uplinks** options and then click **Next**.
7. Verify that the default settings remain in the **Monitoring**, **Miscellaneous**, and **Edit additional settings** fields and then click **Next**.
8. Review the selected settings in the **Ready to complete** screen, and if no changes are required, click **Finish**. The distributed port group displays within the VDS:

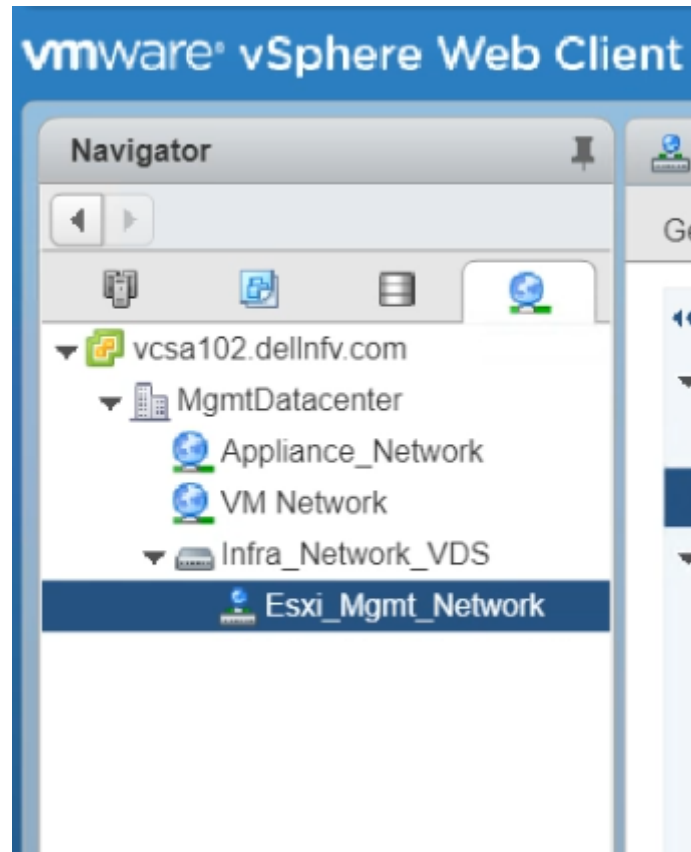


Figure 48. VMware vSphere Web Client window

NOTE: Repeat the steps in this section to create the port groups described in the [VDS-Management port group settings](#) table in this section.

Add host to VDS on management pod

About this task

Add hosts to each VDS switch that you create for the management pod in the [VDS creation and configuration](#) section.

Add hosts to Infra_Network_VDS

About this task

This section provides the steps to add hosts to Infra_Network_VDS.

Steps

1. Right-click **Distributed switch Infra_Network_VDS** and select **Add and Manage Hosts**. The **Add and Manage Hosts** screen displays.
2. Select the **Add hosts** radio button then click **Next**. The **Select hosts** screen displays.
3. Click the **(+) New hosts** icon.
4. Select each of the hosts for the **management cluster**, then click **OK**.
5. Check the **Configure identical network settings on multiple hosts (template mode)** box then click **Next**.
6. From the **Select a template host** screen, select the network configuration host to apply to the other hosts, then clicks **Next**. The **Select network adapter tasks** screen displays.

7. Check the **Manage physical adapter (template mode)** box to associate the necessary uplinks to the DVS, select the **Manage VMkernel adapters (template mode)** box then click **Next**.
8. From the **Manage Physical network and adapters (template mode)** screen, add **Physical network adapters** to the distributed switch:
 - a. Select **vmnic4**, then click the **Assign uplink** icon and assign it to **Mgmtlag1-0**.
 - b. Select **vmnic6**, then click the **Assign uplink** icon and assign it to **Mgmtlag1-1**.
 - c. After assigning the uplink on the template host, click **Apply to all** to apply the physical network adapter assignments on the switch to each of the hosts, and click **Next**.
9. From the **Manage VMkernel network adapters (template mode)** screen:
 - a. Select **vmk0** and click the **Assign port group** icon.
 - b. Select the **destination port group**, for example, **Esxi_Mgmt_Network** to migrate the VMkernel adapters from the source port group, then click **OK**.
 - c. Click **New adapter** to create network adapters for the **vSAN port group**:
 1. On the **Select target device** screen, click the **Browse** button and select the **vSAN_Network** then click **OK**, then click **Next**.
 2. On the **Port properties** screen, specify the VMkernel port settings:
 - **Network label:** vSAN_Network
 - **IP settings:** IPv4
 - **TCP/IP stack:** Default
 - **Enabled services:** vSAN
 3. Click **Next**.
 4. From the **IPv4 settings** screen, select the **Use static IPv4 settings** option and enter the **IPv4 address** and **Subnet mask** **IP** in the provided fields, then click **Next**. For this deployment, *192.168.3.XX* is used as IPv4 address for the **vSAN_Network**.
 5. From the **Ready to complete** screen, review the settings and selections then click **Finish** to create network adapter.
 - d. Click **New adapter** to create network adapters for the **vMotion port group**:
 1. On the **Select target device** screen, click the **Browse** button, select the **vMotion_Network**, click **OK**, then click **Next**.
 2. From the **Port properties** window, specify the VMkernel port settings:
 - **Network label:** vMotion_Network
 - **IP settings:** IPv4
 - **TCP/IP stack:** Default
 - **Enabled services:** vMotion
 3. Click **Next**.
 4. From the **IPv4 settings** screen, click the **Use static IPv4 settings** option, enter the **IPv4 address** and **Subnet mask** in the fields that are provided, and then click **Next**. For this deployment, *192.168.2.XX* is used as IPv4 address for **vMotion_Network**.
 5. From the **Ready to complete** screen, review the settings and selections, and then click **Finish** to create the network adapter.
 - e. Click **New adapter** to create network adapters for the **Replication network port group**:
 1. From the **Select target device** screen click the **Browse** button, select the **Replication_Network**, click **OK**, and then click **Next**.
 2. On the **Port properties** screen, specify the VMkernel port settings:
 - **Network label:** Replication_Network
 - **IP settings:** IPv4
 - **TCP/IP stack:** Default
 - **Enabled services:** vSphere Replication
 3. Click **Next**.
 4. From the **IPv4 settings** screen, click the **Use static IPv4 settings** option, enter the **IPv4 address** and **Subnet mask** in the provided fields, and then click **Next**. For this deployment, *192.168.5.XX* is used as IPv4 address for **Replication_Network**.
 5. From the **Ready to complete** screen, review the settings and selections, and then click **Finish** to create network adapter.
 - f. Once all the port groups are assigned to network adapters, click **Apply to all**:

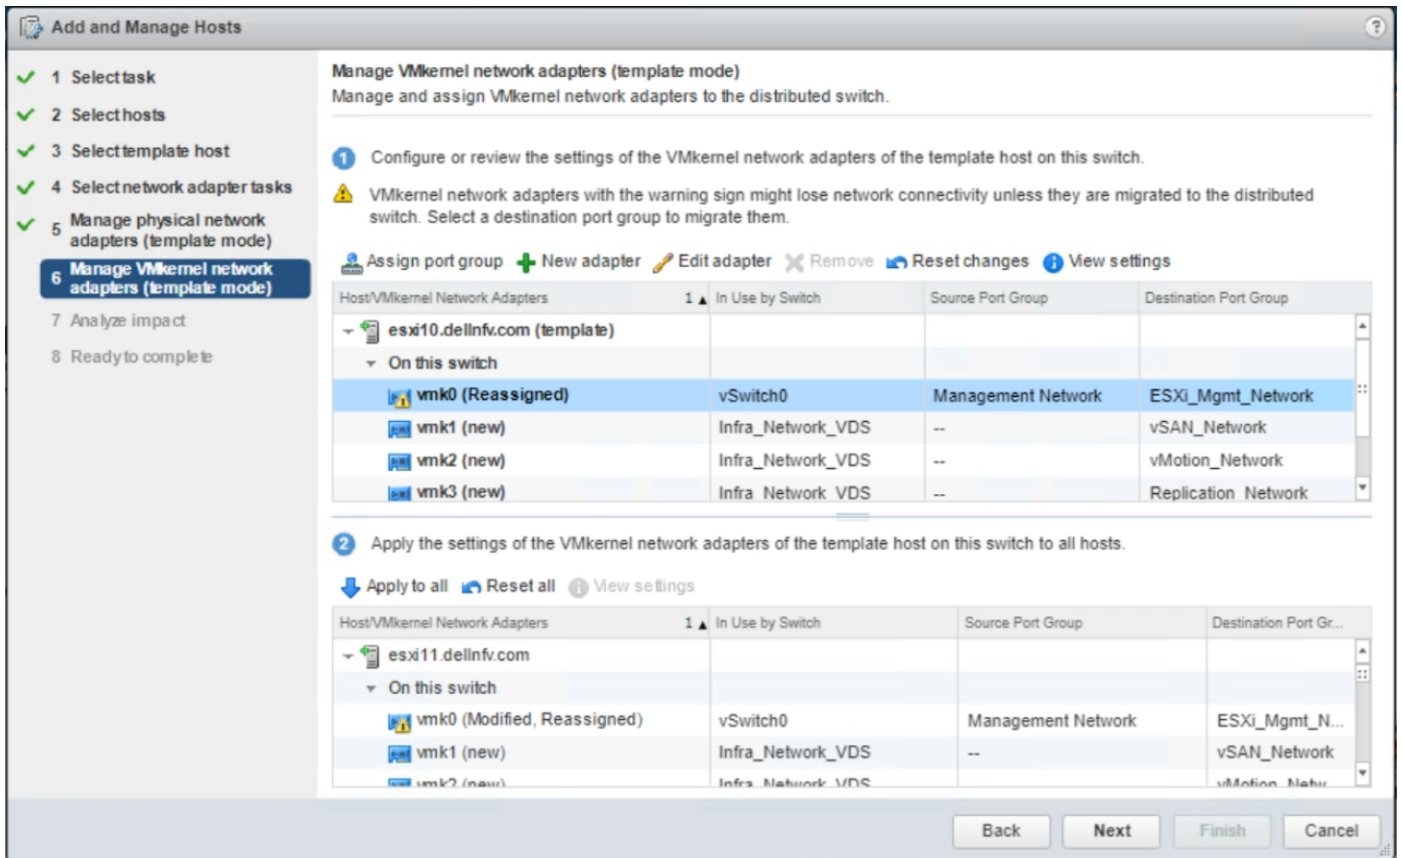


Figure 49. Manage VMkernel network adapters - template mode

10. From the **Apply VMkernel network adapter configuration to other hosts** screen, add the IPv4 addresses of the other hosts to apply the settings of the VMkernel network adapters of the template host on the switch to all hosts, click **OK** and click **Next**.
11. From the **Analyze impact** screen, review the impact of the configuration change might have on some network-dependent services, then click **Next**.
12. From the **Ready to complete** screen, review the settings and selection and click **Finish** to add the host.

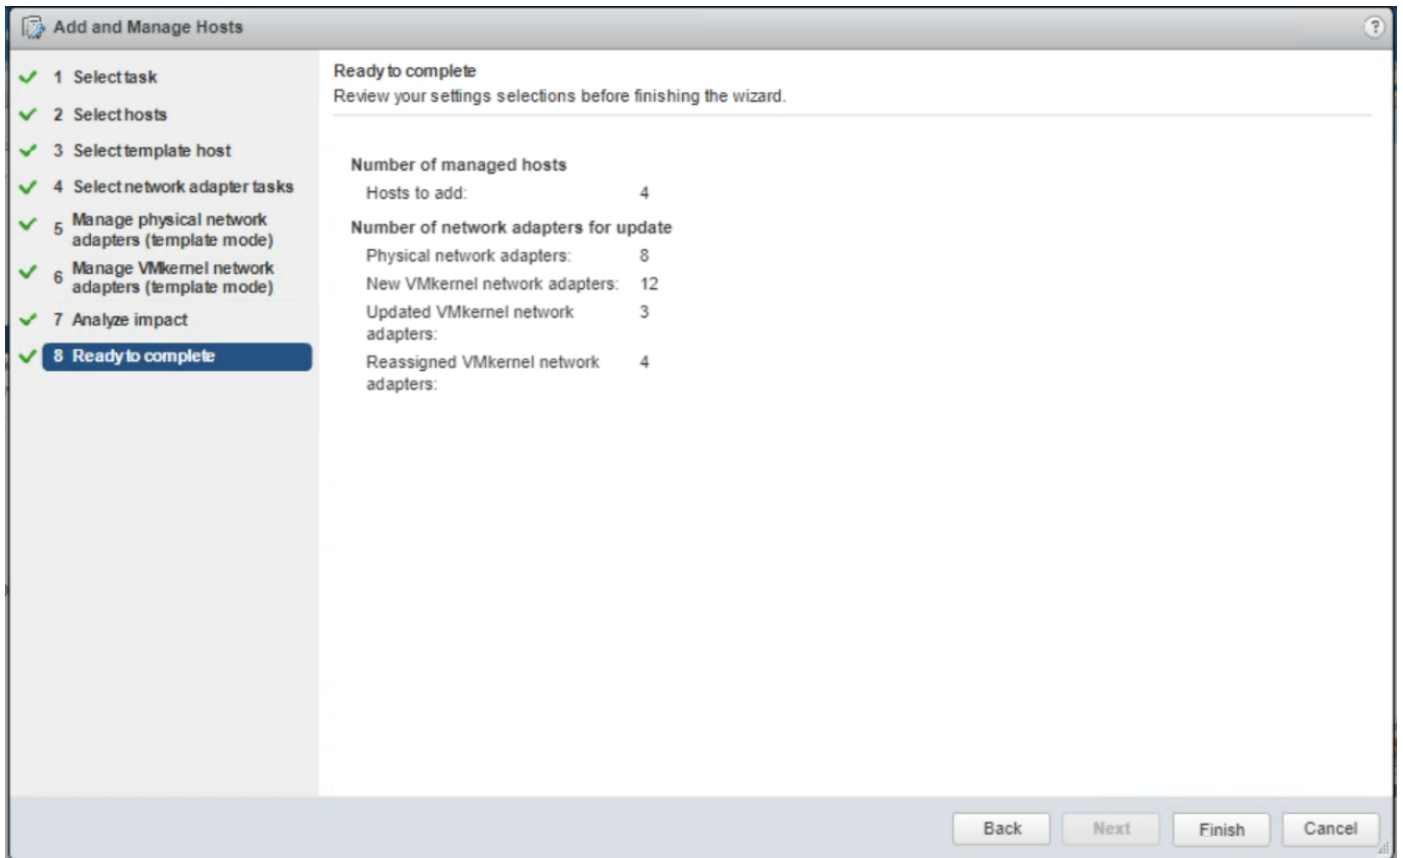


Figure 50. Ready to complete window

Configure vSAN on management cluster

About this task

To migrate the VM to VDS, configure the vSAN first. This section provides the steps to configure vSAN on management cluster.

Steps

1. Using a web browser, open the VMware vSphere Web Client for the management cluster. The **Management Cluster** window opens.
2. Select the **Configure** tab.
3. In the left navigation panel, click **Disk Management**, expand the **vSAN** section, click the **Claim disks** icon in **Disk Groups**.

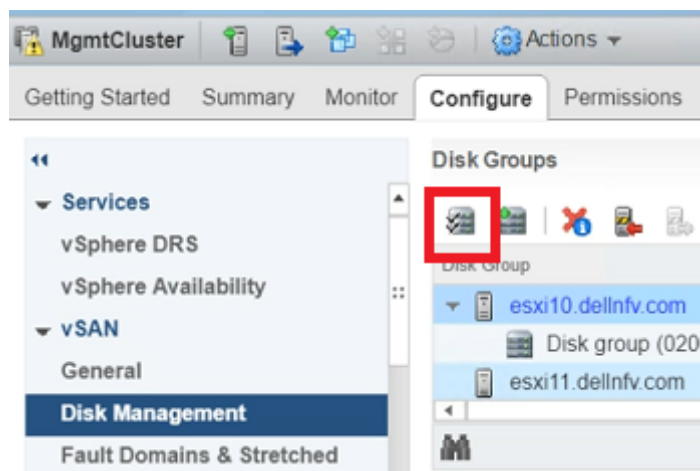


Figure 51. Disk management window

- In the **Claim Disk for vSAN use** screen, ensure that the **Capacity tiers** option is selected for HDD, and that the **Cache tier** has **Flash** selected, then click **OK**.

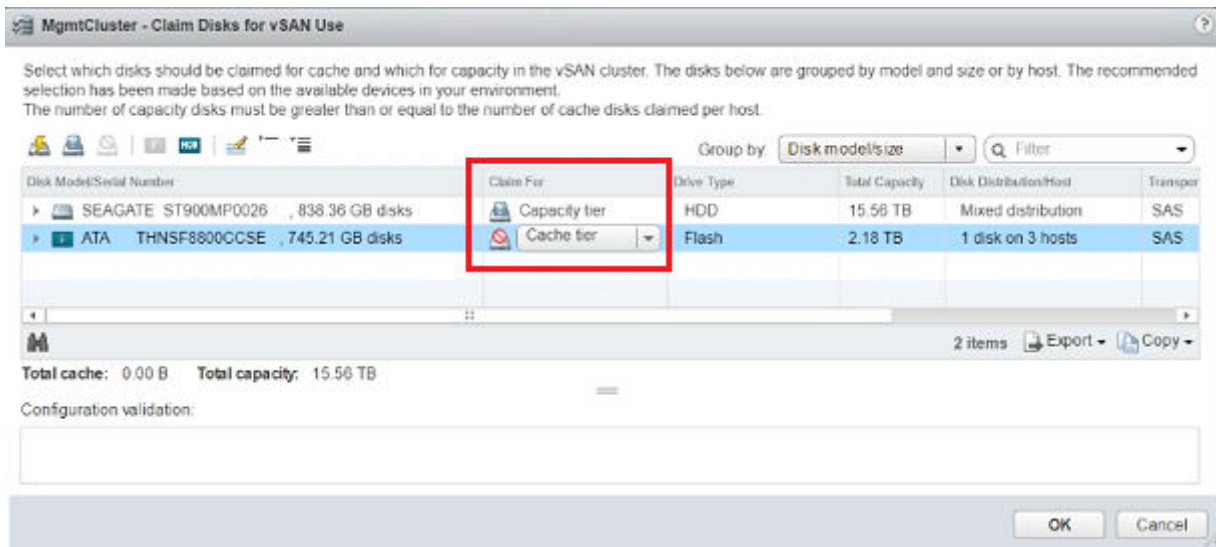


Figure 52. Claim disks window

- From the **Configure** tab, verify that the disk groups are created on each ESXi and that the **vSAN Datastore** is created.

Add hosts to VM_Network_VDS

Prerequisites

- All ESXi hosts have an extra uplink to the management network

About this task

This section provides steps to add hosts to VM_Network_VDS.

Steps

- Right-click the **VM_Network_VDS** distributed switch and select **Add and Manage Hosts**. The **Add and Manage Hosts** screen displays.
- Select the **Add hosts** radio button then click **Next**. The **Add hosts** screen displays.
- Click the **(+) New hosts** icon.
- On the **New hosts** screen, select all of the management cluster hosts, then click **OK**.

NOTE: Do not select the hosts that are associated with any VM.
- Select the **Configure identical network settings on multiple hosts (template mode)** box then click **Next**. The **Select a template host** screen displays.
- Select the network configuration host to apply to the other hosts, then click **Next**. The **Select network adapter tasks screen** displays.
- Select the **Manage physical adapter (template mode)** box to associate the necessary uplinks to the DVS, then select the **Manage VMkernel adapters (template mode)** box then click **Next**.
- From the **Manage Physical network and adapters (template mode)** window, add physical network adapters to the distributed switch.
 - Select **vmnic5**, then click the **Assign uplink** icon and assign it to **Mgmtlag1-0**.
 - Select **vmnic7**, then click the **Assign uplink** icon and assign it to **Mgmtlag1-1**.
 - After assigning the uplink on the template host, click **Apply to all** to apply the physical network adapter assignments on this switch for all the hosts, and then click **Next**.
- Review the information within the **Manage VMkernel network adapter (template mode)** window, then click **Next**.
- From the **Analyze impact** screen, review the impact of the configuration change might have on some network-dependent services, then click **Next**.

11. From the **Ready to complete** window, review the settings and selection and click **Finish** to add the host.

vSwitch to VDS Migration on management pod

About this task

vSwitch is migrated to VDS for both management and resource PODs. The VSS to VDS migration for both PODs follows similar steps. Virtual machines residing on individual PODs are migrated to the VDS during this operation.

Migrate Windows AD-DNS VM

About this task

Migrate the Windows AD-DNS VM to a specific host or cluster and to a specific datastore.

Steps

1. Open the **VMware vCenter Server Virtual Appliance**.
2. From the datastore, right-click the **AD-DNS VM** and select **Migrate**.
The **Select migration type** screen displays.
3. Select the **Change both compute resource and storage** radio button and then click **Next**.
The **Select a compute resource** screen displays.
4. Select the compute resource that you added to the VM_Network_VDS for VM migration, then click **Next**.
The **Select storage** screen displays.
5. Select **vSAN Default Storage Policy** from the **VM storage policy** drop-down, then select **vsanDatastore** from the datastore listing, and click **Next**.
The **Select networks** screen displays.
6. Select the destination network, for example, VM_Mgmt_Network, from the drop-down list, and click **Next**.
7. From the **Select vMotion priority** screen options, verify that the **Schedule vMotion with high priority** radio button is selected, and then click **Next**.
8. From the **Ready to complete** screen, review the selections and then click **Finish**.

Migrate NTP VM

About this task

Follow the steps described in [Migrate Windows AD-DNS VM](#) to migrate the NTP VM that is deployed on the management cluster.

NOTE: Once the VMs are migrated from the ESXi host to VDS, it is required to add the remaining ESXi host to VDS. Follow the steps that are described in [Add hosts to VM_Network_VDS](#) to add the host.

Migrate management VCSA active VM

About this task

Migrate the management VCSA active VMs to a specific host or cluster.

Steps

1. Open the VMware vCenter Server Virtual Appliance.
2. From the datastore, right-click the management VCSA active VM and select **Migrate**.
3. On the **Select the migration type** screen, select the **Change compute resource only** radio button then click **Next**.
4. On the **Select a compute resource** screen, select the compute resource for VM migration, then click **Next**.
5. On the **Select networks** screen, select the destination network from the drop-down list, and click **Next**.
6. On the **Select vMotion priority window options**, make sure that **Schedule vMotion with high priority** radio button is selected, then click **Next**.
7. From the **Ready to complete** screen, review the chosen selections then click **Finish** to complete the operation.

Migrate resource VCSA active VM

About this task

Follow the steps described in the [Migrate Management VCSA active VM](#) section to migrate the NTP VM that is deployed on the management cluster.

NOTE: Once the VMs are migrated from the ESXi host to VDS, it is required to add the remaining ESXi host to VDS. Follow the steps that are described in [Add hosts to VM_Network_VDS](#) to add the host.

Create VDS for resource and edge pods

About this task

In the resource cluster, create a VDS for Infrastructure network and for VM networks. The [VDS-Resource settings](#) table in this section provides the VDS information and the recommended settings for this deployment.

Table 22. VDS-Resource settings

VDS name	Version	Number of uplinks	Network I/O control	Discovery protocol type/operation	MTU (bytes)
Edge_Infra_Network_VDS	6.6.0	2	Enabled	CDP/Both	9000 bytes
Edge_VM_Network_VDS	6.6.0	2	Enabled	CDP/Both	9000 bytes
Res_Infra_Network_VDS	6.6.0	2	Enabled	CDP/Both	9000 bytes

To create and configure the virtual distributed switch, perform the following steps:

Steps

1. From the **VMware vSphere Web Client**, select **Home/Networking** to switch to the **Networking** view.
2. Right-click the **ResourceDatacenter**, and select **Distributed switch**, and then **New Distributed Switch**. The **New Distributed Switch wizard** opens.
3. Using the information from the [VDS-Resource settings](#) table, enter the **VDS name** in the field that is provided, and then click **Next**. The **Select version** screen displays.
4. Select **Distributed switch: 6.6.0** then click **Next**.
5. From the **Edit settings** window, use the information from the [VDS-Resource settings](#) table to select the **Number of uplinks**, **Network I/O control**.
6. Click to clear the **Create a default port group** box.
7. Review the selected settings in the **Ready to complete** window, and if no changes are required, click **Finish**.

NOTE: Repeat the steps provided in this section to create each of the VDS switches described in the [VDS-Resource settings](#) table.

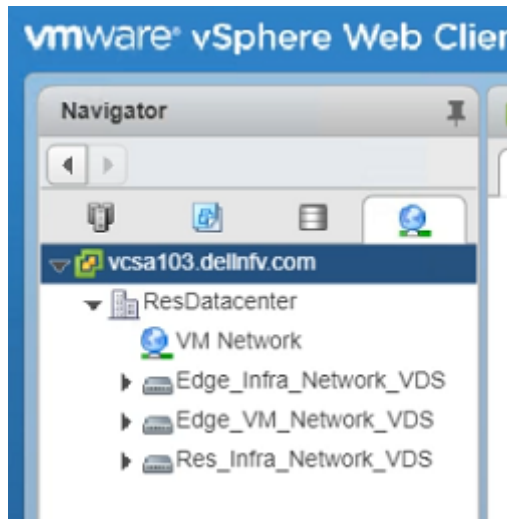


Figure 53. ESXi_Mgmt_VDS listing

VDS configuration settings for resource VDS

About this task

Once all the VDS switches are created as described in the [VDS-Resource settings](#) table, update the configuration settings of each resource VDS switch. To update the VDS configuration settings:

Steps

1. Select the distributed switch then select the **Configure** tab, **Properties**, then click the **Edit** button.
2. Click the **Advanced** option in the left-navigation panel.
3. Set the **MTU** to **9000 bytes**.
4. From the **Discovery protocol** section, select the **Type** to **Cisco Discovery Protocol**, and the **Operation** option to **Both**.
5. Click **OK**.

NOTE: Repeat the steps in this section to configure the settings for the remaining VDS switches.

Create LAG for resource and edge pods

About this task

The VDS-LAG settings table in this section displays the configuration settings that are used to create LAG for VDS.

Table 23. VDS-LAG settings

VDS	Name	Number of ports	Mode	Load-balancing mode
Edge_Infra_Network_VDS	EdgeLag1	2	Active	Source and destination IP address and TCP/UDP port
Edge_VM_Network_VDS	EdgeLag1	2	Active	Source and destination IP address and TCP/UDP port
Res_Infra_Network_VDS	ResourceLag1	2	Active	Source and destination IP address and TCP/UDP port

Steps

1. From the left navigation panel, click to select the **Edge_Infra_Network_VDS**, and then click the **Configure** tab.

2. Select the **Settings** listing, **NES**, and then click **(+) Add** to create the Migrating network traffic to the LAGs.
 3. Using the information in the [VDS-LAG settings](#) table, enter the **LAG name**, **Number of ports**, **Mode**, and **Load-balancing mode** options and then click **OK**.
- NOTE:** Using the information provided in the [VDS-LAG settings](#) table, create the LAG on **Edge_VM_Network_VDS** and **Res_Infra_Network_VDS** distributed switches.

Create distributed port group for resource and edge VDS

About this task

The [VDS-port group settings](#) table displays the port group settings that are used for the resource and edge pod VDS.

Table 24. VDS-port group settings

Port group	VLAN type	VLAN ID	Teaming and failover settings				
			Load balancing	Network failure detection	Notify switches	Failback	Failover order
ESXi_Mgmt_Network_Edge (under Edge_Infra_Network_VDS)	VLAN	100	Route based on IP hash	Link status only	Yes	Yes	Active uplinks - EdgeLag1
VM_Mgmt_Network_Edge (under Edge_Infra_Network_VDS)	VLAN	20	Route based on IP hash	Link status only	Yes	Yes	Active uplinks - EdgeLag1
vSAN_Network_Edge (under Edge_Infra_Network_VDS)	VLAN	300	Route based on IP hash	Link status only	Yes	Yes	Active uplinks - EdgeLag1
vMotion_Network_Edge (under Edge_Infra_Network_VDS)	VLAN	200	Route based on IP hash	Link status only	Yes	Yes	Active uplinks - EdgeLag1
Overlay_Network (under Edge_VM_Network_VDS)	VLAN trunking	0-4094	Route based on IP hash	Link status only	Yes	Yes	Active uplinks - EdgeLag1
External_Network (under Edge_VM_Network_VDS)	VLAN trunking	0-4094	Route based on IP hash	Link status only	Yes	Yes	Active uplinks - EdgeLag1
VM_Mgmt_Network_Resources (under Res_Infra_Network_VDS)	VLAN	20	Route based on IP hash	Link status only	Yes	Yes	Active uplinks - ResourceLag1
ESXi_Mgmt_Network_Res (under Res_Infra_Network_VDS)	VLAN	100	Route based on IP hash	Link status only	Yes	Yes	Active uplinks - ResourceLag1
vSAN_Network_Res (under Res_Infra_Network_VDS)	VLAN	300	Route based on IP hash	Link status only	Yes	Yes	Active uplinks - ResourceLag1
vMotion_Network_Res (under Res_Infra_Network_VDS)	VLAN	200	Route based on IP hash	Link status only	Yes	Yes	Active uplinks - ResourceLag1

Port group	VLAN type	VLAN ID	Teaming and failover settings				
			Load balancing	Network failure detection	Notify switches	Failback	Failover order
Res_Infra_Network_VD S)							

Steps

- Right-click the VDS, select **Distributed Port Group**, and then **New Distributed port group**. The **New Distributed port group**
 - Using the information from the [VDS-port group settings](#) table, enter the **Port group** name then click **Next**. The **Configure settings** screen displays.
 - Set the **General properties** of the new port group as follows:
 - Port binding:** Static binding
 - Port allocation:** Elastic
 - Number of ports:** 8
 - Network resource pool:** (default)
 - VLAN type:** Use the information provided in the [VDS-port group settings](#) table
 - VLAN ID:** Use the information provided in the [VDS-port group settings](#) table
 - After setting the properties, click **Next**.
 - Review the selected settings on the **Ready to complete** screen, and if no changes are required, click **Finish**. The distributed port group is created, and displays under the **VDS** in the **Topology** window.
 - In the left navigation panel, locate the newly created port group, right-click the listing, and select **Edit Settings**. The **Edit Settings** window opens.
 - Select the **Teaming and failover** tab.
 - Using the information from the [VDS-port group settings](#) table, select the **Load balancing**, **Network failure detection**, **Notify switches**, **Failback**, and **Active uplinks** options, and then click **OK**.
- NOTE:** Repeat steps from step in this section to create the other port groups as described in the [VDS-port group settings](#) table.

Add hosts to VDS on edge pod

Add hosts to the edge cluster of each VDS. See the following sections to add hosts to the edge pod VDS:

- [Add hosts to Edge_Infra_Network_VDS](#)
- [Add hosts to Edge_VM_Network_VDS](#)

Add hosts to Edge_Infra_Network_VDS

About this task

This section provides the steps to add host to Edge_Infra_Network_VDS.

Steps

- Right-click the **Edge_Infra_Network_VDS** and select **Add and Manage Hosts**. The **Add and Manage Hosts** window opens.
- Select the **Add hosts** radio button then click **Next**.
- From the **Select hosts** screen, click the **(+) New hosts** icon.
- Select the hosts for the edge cluster, then click **OK**.
- Click to select the **Configure identical network settings on multiple hosts (template mode)** box then click **Next**.
- From the **Select a template host** screen, select the network configuration host to apply to the other hosts, then click **Next**. The **Select network adapter tasks** screen displays.
- Click to select the **Manage physical adapter (template mode)** box to associate the necessary uplinks to the VDS, click to select the **Manage VMKernel adapters (template mode)** box, and then click **Next**.

The **Manage Physical network and adapters (template mode)** screen displays.

8. Add the physical network adapters to the distributed switch:
 - a. Select **vmnic4**, then click the **Assign uplink** icon and assign it to **EdgeLag1-0**.
 - b. Select **vmnic6**, then click the **Assign uplink** icon and assign it to **EdgeLag1-1**.
 - c. After assigning the uplink on the template host, click **Apply to all** to apply the physical network adapter assignments on the switch for all the hosts, and then click **Next**.

The **Manage VMkernel network adapters (template mode)** screen displays.

9. Select the following settings:
 - a. Select **vmk0** and click the **Assign port group** icon.
 - b. Select the **Destination port group**, for example, **ESXi_Mgmt_Network_Edge**, to migrate the VMkernel adapters from the source port group, then click **OK**.
 - c. Click **New adapter** to create network adapters for vSAN port group:
 1. On the **Select target device** screen, click the **Browse** button, select the **vSAN_Network_Edge**, click **OK**, and then click **Next**.
 2. From the **Port properties** screen, specify the **VMKernel port settings**:
 - **Network label:** vSAN_Network_Edge
 - **IP settings:** IPv4
 - **TCP/IP stack:** Default
 - **Enabled services:** vSAN
 3. Click **Next**.
 4. From the **IPv4 settings** window, click the **Use static IPv4 settings** option and enter the **IPv4 address** and **Subnet mask IP** in the provided fields, then click **Next**. For this deployment, **192.168.3.XX** is used as the IPv4 address for vSAN_Network.
 5. From the **Ready to complete** screen, review the settings and selections then click **Finish**.
 - d. Click **New adapter** to create network adapters for the vMotion port group:
 1. From the **Select target device** screen click the **Browse** button, select **vMotion_Network_Edge**, click **OK**, and then click **Next**.
 2. In the **Port properties** section, specify the following VMKernel port settings:
 - **Network label:** vMotion_Network_Edge
 - **IP settings:** IPv4
 - **TCP/IP stack:** Default
 - **Enabled services:** vMotion
 3. Click **Next**.
 4. On the **IPv4 settings** screen, select the **Use static IPv4 settings** option and enter the **IPv4 address** and **Subnet mask** in the provided fields, then click **Next**. For this deployment, **192.168.2.XX** is used as the IPv4 address for vMotion_Network.
 5. From **Ready to complete** screen, review the settings and selections then click **Finish**.
 - e. Once the port groups are assigned to network adapters, click **Apply to all**.

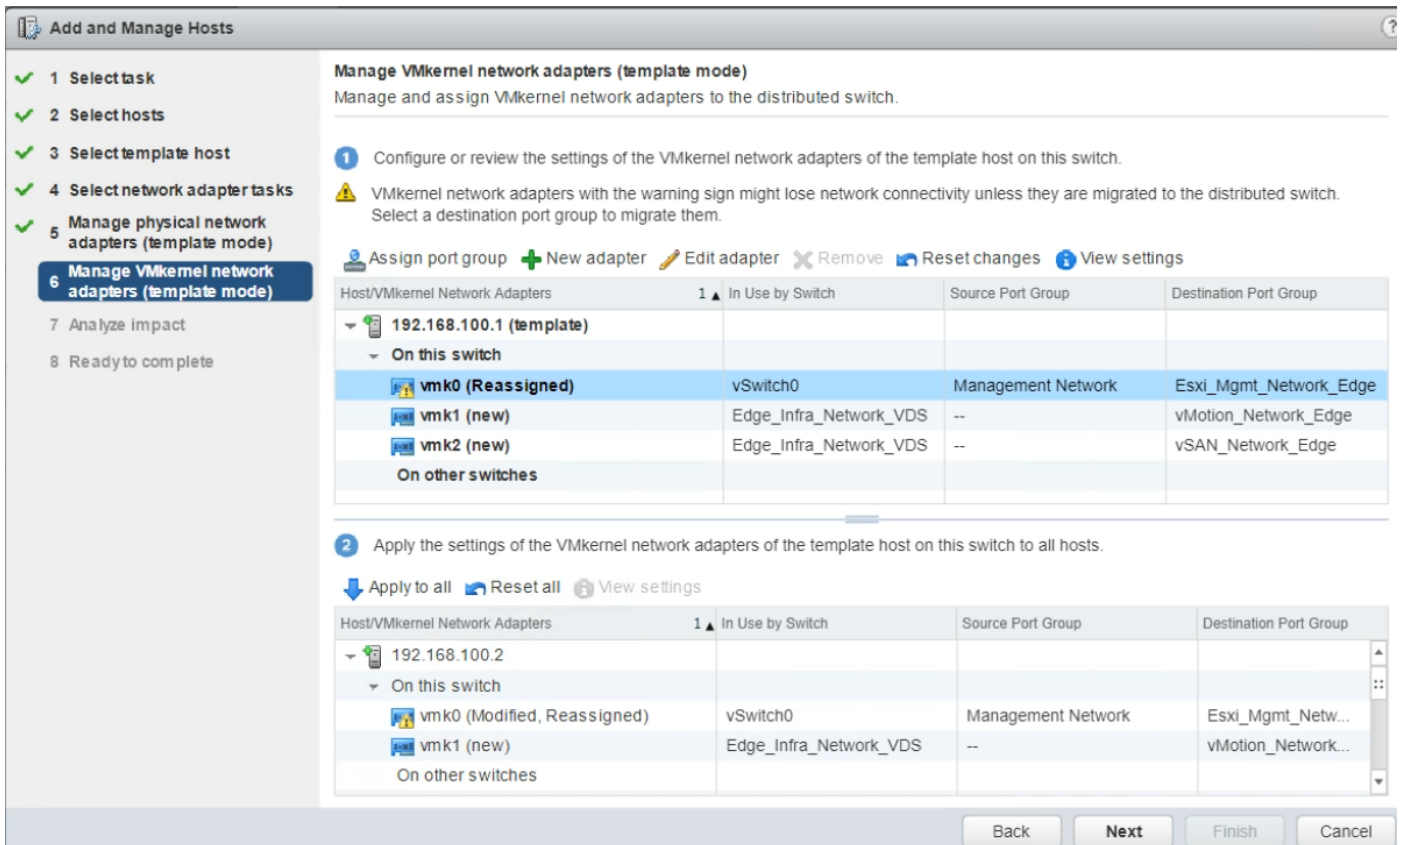


Figure 54. Manage VMkernel network adapters (template mode)

The **Apply VMkernel network adapter configuration to other hosts** screen displays.

10. To apply the settings of the VMkernel network adapters of the template host on the switch to the hosts, enter IPv4 addresses of the other edge hosts.
11. Click **OK**, and then click **Next**.
The **Analyze impact** screen displays.
12. Review the impact that the configuration change may have on some of the network-dependent services, then click **Next**.
13. From the **Ready to complete** screen, review the settings and selections, then click **Finish**.

Add hosts to Edge_VM_Network_VDS

About this task

This section provides steps to add hosts to Edge_VM_Network_VDS.

Steps

1. Right-click the **Edge_VM_Network_VDS** distributed switch and select **Add and Manage Hosts**.
The **Add and Manage Hosts** screen displays.
2. Select the **Add hosts** radio button then click **Next**.
3. From the **Select hosts** screen, click the **(+) New hosts** icon.
4. On the displayed window select all the edge cluster hosts, then click **OK**.
5. Check the **Configure identical network settings on multiple hosts (template mode)** box then click **Next**.
The **Select a template host** screen displays.
6. Select the network configuration host to apply to the other hosts, then clicks **Next**.
7. From the **Select network adapter tasks** screen, select the **Manage physical adapter (template mode)** box to associate the necessary uplinks to the DVS, select the **Manage VMKernel adapters (template mode)** box then click **Next**.
The **Manage Physical network and adapters (template mode)** screen displays.
8. Add the physical network adapters to the distributed switch:

- a. Select **vmnic5**, click the **Assign uplink** icon and assign it to **EdgeLag1-0**.
 - b. Select **vmnic7**, click the **Assign uplink** icon and assign it to **EdgeLag1-1**.
 - c. After assigning the uplink on the template host, click **Apply to all** to apply the physical network adapter assignments on this switch for all the hosts, and then click **Next**.
9. Review the information that is provided in the **Manage VMkernel network adapter (template mode)** screen, and then click **Next**. The **Analyze impact** screen displays.
 10. Review the impact of the configuration change might have on some network-dependent services, then click **Next**.
 11. From the **Ready to complete** window, review the settings and selection and click **Finish**.

Add hosts to VDS on resource pod

The information in this section provides the steps needed to add hosts to the each VDS of the resource cluster.

Add hosts to Res_Infra_Network_VDS

About this task

This section provides the steps to add hosts to Res_Infra_Network_VDS.

Steps

1. Right-click **Res_Infra_Network_VDS** and select **Add and Manage Hosts**. The **Add and Manage Hosts** screen displays.
2. Select the **Add hosts** radio button then click **Next**. The **Select hosts** screen displays.
3. Click the **(+) New hosts** icon.
4. Select the hosts for the resource cluster, then click **OK**.
5. Select the **Configure identical network settings on multiple hosts (template mode)** box then click **Next**.
6. From the **Select a template host** window, select the network configuration host to apply to the other hosts, then click **Next**. The **Select network adapter tasks** screen displays.
7. Select the **Manage physical adapter (template mode)** box to associate the necessary uplinks to the DVS, and select the **Manage VMkernel adapters (template mode)** box then click **Next**.
8. From the **Manage Physical network and adapters (template mode)** screen, add the physical network adapters to the distributed switch:
 - a. Select **vmnic4**, then click the **Assign uplink** icon and assign it to **ResourceLag1-0**.
 - b. Select **vmnic6**, then click the **Assign uplink** icon and assign it to **ResourceLag1-1**.
 - c. After assigning the uplink on the template host, select **Apply to all** to apply the physical network adapter assignments on the switch for all the hosts, and click **Next**.
9. From the **Manage VMkernel network adapters (template mode)** screen:
 - a. Select **vmk0** and click the **Assign port group** icon
 - b. Select the destination port group, for example, **ESXi_Mgmt_Network_Res**, to migrate the VMkernel adapters from the source port group, then click **OK**.
 - c. Click **New adapter** to create network adapters for the vSAN port group:
 1. On the **Select target device** screen click the **Browse** button, select **vSAN_Network_Res**, click **OK**, and then click **Next**.
 2. From the **Port properties** window, specify the VMkernel port settings:
 - **Network label:** vSAN_Network_Res
 - **IP settings:** IPv4
 - **TCP/IP stack:** Default
 - **Enabled services:** vSAN
 3. Click **Next**.
 4. From the **IPv4 settings** screen, select the **Use static IPv4 settings** option and enter the **IPv4 address** and **Subnet mask IP** in the provided fields, and then click **Next**. For this deployment, **192.168.3.XX** is used as the IPv4 address for vSAN_Network.
 5. From the **Ready to complete** screen, review the settings and selections then click **Finish**.
 - d. Click **New adapter** to create network adapters for vMotion port group:

1. From the **Select target device** screen, click **Browse**, select **vMotion_Network_Res**, click **OK**, and then click **Next**.
2. On the **Port properties** window, specify the VMKernel port settings:
 - **Network label:** vMotion_Network_Res
 - **IP settings:** IPv4
 - **TCP/IP stack:** Default
 - **Enabled services:** vMotion
3. Click **Next**.
4. From the **IPv4 settings** screen, select the **Use static IPv4 settings** option, enter the **IPv4 address** and **Subnet mask** in the provided fields, and then click **Next**. For this deployment, **192.168.2.XX** is used as the IPv4 address for vMotion_Network.
5. From the **Ready to complete** screen, review the settings and selections then click **Finish**.
- e. Once the port groups are assigned to network adapters, click **Apply to all**.
10. From the **Apply VMkernel network adapter configuration to other hosts** screen, enter IPv4 addresses of the other edge hosts. This selection applies the settings of the VMkernel network adapters of the template host on the switch, to each of the hosts.
11. Click **OK**, and click **Next**.
12. From the **Analyze impact** screen, review the impact that the configuration change may have on the network-dependent services, then click **Next**.
13. From the **Ready to complete** screen, review the settings and selection and click **Finish**.

Configure VMware vSAN clusters

The VMware Virtual SAN (vSAN) is a distributed layer of software that runs natively as a part of the ESXi hypervisor. vSAN aggregates local or direct-attached capacity devices of a host cluster and creates a single storage pool that is shared across all hosts in the vSAN cluster.

Each server has a hybrid mix of SSD and HDD drives for storage deployment. For vSAN, solid-state disks are required for the cache tier, while the spinning disks make up the capacity tier. Each Dell EMC server is configured for Host Bus Adapter (HBA) in non-RAID or pass-through mode since the vSAN software handles the redundancy and storage cluster information.

Configure vSAN on resource and edge cluster

About this task

Configure vSAN on the resource and edge cluster.

Steps

1. Log in to the **VCSA using vSphere Web Client** and go to the resource cluster.
2. Click the **Configure** tab.
3. From the vSAN listing, select **General**, and then click **Configure** in the upper-right corner to edit the **Virtual SAN configuration**.
4. In the **vSAN Capabilities** window, leave the default setting as-is and click **Next**.

NOTE: It is not possible to delete disks from a disk group after deduplication and compression is enabled on the cluster. Consider the configuration ahead of time and add all the capacity that you need before activating the deduplication.
5. Confirm the information within the **Network Validation** screen, then click **Next**.
The **Claim disks** screen displays.
6. From the **Claim For** column, verify that **HDD** is claimed for the **Capacity tier** and that the **Cache tier** has **Flash** claimed.
7. Verify the settings within the **Ready to complete** screen, and then click **Finish**.
8. From the **Virtual SAN** listing, click the **General listing** in the left navigation panel and click the **Configure** tab.
9. Locate the **vSAN status** on the **Configure** tab and verify that **vSAN** status shows as **Turned ON**.

NOTE: Repeat the steps in this section to configure the vSAN for the edge cluster.

Assign vSAN license key to cluster

NOTE: Add the license to the management, resource, and edge clusters.

To assign a vSAN license key to a cluster, select from the following options:

- [Assign a new vSAN license](#)
- [Assign vSAN license using an existing license](#)

Assign a new vSAN license

About this task

Assign the vSAN license to a vSAN cluster.

Steps

1. From the **VMware vSphere Web Client**, go to a cluster where vSAN is enabled.

2. On the **Configure** tab, locate the **Configuration** section, select **Licensing**, and click **Assign License**.
3. Click the **(+) Add** icon.
4. In the **New Licenses** dialog box, enter the Virtual SAN license key and click **Next**.
5. From the **Edit license names** page, rename the new license as appropriate and click **Next**.
6. Click **Finish**.
7. In the **Assign License** dialog box, select the newly created license, and click **OK**.

Assign vSAN license using an existing license

About this task

Assign the vSAN license using an existing license.

Steps

1. From the **VMware vSphere Web Client**, go to a cluster where you vSAN is enabled.
2. On the **Configure** tab, locate the **Configuration** section, select **Licensing**, and click **Assign License**.
3. Select the licensing option, then select an existing license.
4. Click **OK**.

Update vSAN HCL database manually

About this task

 **NOTE:** Update the vSAN HCL database on the management, resource, and edge clusters.

This section provides the steps to manually update vSAN HCL database.

Steps

1. Log in to the vCenter Server using VMware vSphere Web Client using administrator credentials.
2. From the **vSAN cluster**, click the **Configure** tab.
3. In the left navigation panel, locate the **vSAN** section, select **Health and Performance**, and from the **HCL Database** section:
 - a. If the vCenter Server **can** communicate with the Internet, click **Get latest version online to update the HCL Database**.
 - b. If the vCenter server **does not** have proxy and is unable to communicate with the Internet, download the updated file locally using the following steps:
 1. Log in to a workstation with Internet access.
 2. Open the following link in a browser: <https://partnerweb.vmware.com/service/vsan/all.json>
 3. Save the file as `all.json`.
 4. Copy the file to the vCenter Server for upload.
 5. Under the **Health** tab in the **HCL Database** section, select the **Update from file** option and select the `all.json` file and upload the file.
4. To retest the health, click **Monitor, Virtual SAN, Health**, and then **Retest**.

Enable vSAN performance service

Prerequisites

- Cluster should be configured on both Management and Resource vCenter Server.

About this task

This section provides the steps to enable the vSAN performance services. You must enable the performance services on Management cluster, resource cluster, and edge cluster.

Steps

1. Log in into VCSA using vSphere Web Client and go to the **Cluster**.
2. On the **Configure** tab, in the **vSAN** section, select **Health and Performance**, and click **Edit** to change the performance service settings.
3. Select the **Turn On Virtual SAN performance service** check box.
4. Select a storage policy, and click **OK**.
After the **Performance service** turned on, review the settings.




Figure 55. Health and Performance settings

Configure VMware vCenter High Availability

VMware vCenter High Availability, or vCenter HA, protects the VMware vCenter Server Appliance against host and hardware failures. The active-passive architecture of the solution helps to reduce downtime when you patch vCenter Server Appliance.

The VCSA-HA feature works as a cluster of three VMs. The three-node cluster contains active, passive, and witness nodes. A different configuration path is available but depends on your current configuration.

NOTE: As part of solution deployment, VCSA-HA is configured differently for the management and resource clusters.

Management cluster VCSA-HA configuration

Management Cluster VCSA-HA is configured using the Basic Option. The Basic Option allows the vCenter HA wizard to create and configure a second network adapter on the VCSA. The vCenter HA wizard also clones the active and witness nodes and configures the vCenter HA network.

Configure Management vCenter HA

About this task

Configure the management vCenter HA.

Steps

1. Log in to the Management VMware vSphere Web Client.
2. Right-click the top-level **vCenter Server** in the inventory and select **vCenter HA Settings**.
3. On the top-right corner, click the **Configure** button.
4. From the **Select a configuration option** screen, select **Basic**, and then click **Next**.
5. On the **Add a vCenter HA network adapter for Active node** window:
 - a. In the **IPv4 address** field, enter the IP address for VCSA HA Active node.
 - b. In the **IPv4 subnet mask** field, enter the subnet mask IP address for VCSA HA Active node.
 - c. In the **Select vCenter HA network** field, click **Browse** to select **VCSA HA network**, and then click **Next**.
6. From the **Select IP settings for Passive and Witness nodes** screen, enter the IP addresses for the **Passive Node** and **Witness Node** in the respective fields, and then click **Next**.
7. From the **Select a deployment configuration** screen, click **Edit** to select the deployment configuration for **Passive Node**. The **vCenter HA Passive Node - Edit Deployment Configuration** window opens.
 - a. From the **Select a name and folder** screen, enter the **VM name**, select a data center to deploy the VM, and then click **Next**.
 - b. On the **Select a compute resource** screen, select a target host to run the VM, then click **Next**.

NOTE: Configure the VCSA Active node, Passive node, and Witness node on different compute hosts.
 - c. On the **Select storage** window, locate the **VM storage policy** drop-down list, select **vSAN Default Storage Policy**, and then click **Next**.
 - d. In the **Select networks** screen, select the appropriate networks, then click **Next**.
 - e. From the **Ready to complete** screen, verify the options that are selected, then click **Finish**.
 - f. Repeat the steps in this section to complete the deployment configuration setting for witness nodes, then click **Next**.

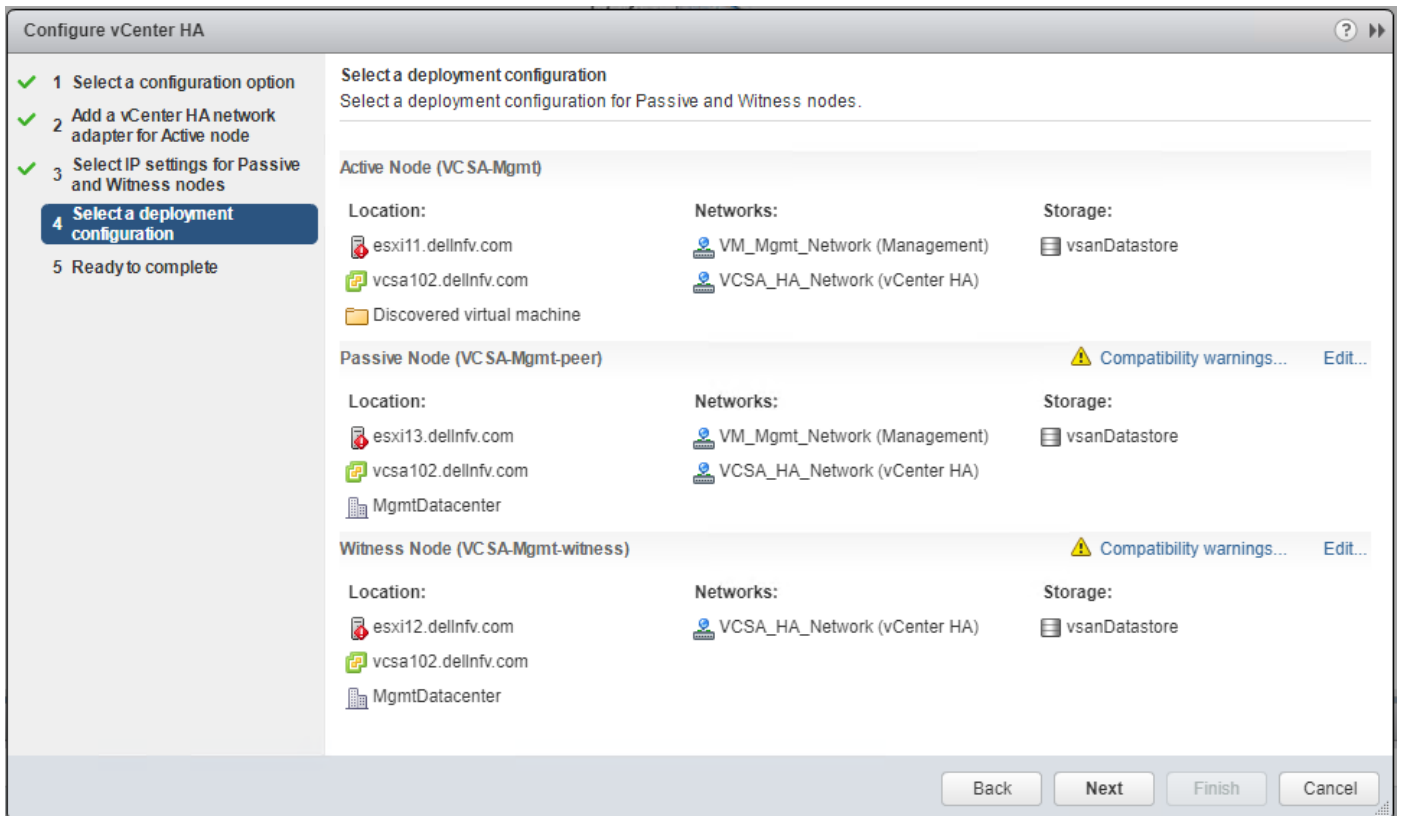


Figure 56. Select a deployment configuration screen

8. From the **Ready to complete** window, verify the options that are selected, then click **Finish**.

NOTE: Monitor the Tasks pane as it may take several minutes to clone and deploy the vCenter HA cluster nodes. When complete, the vCenter HA status shows Enabled and the nodes in the cluster show the Up status.

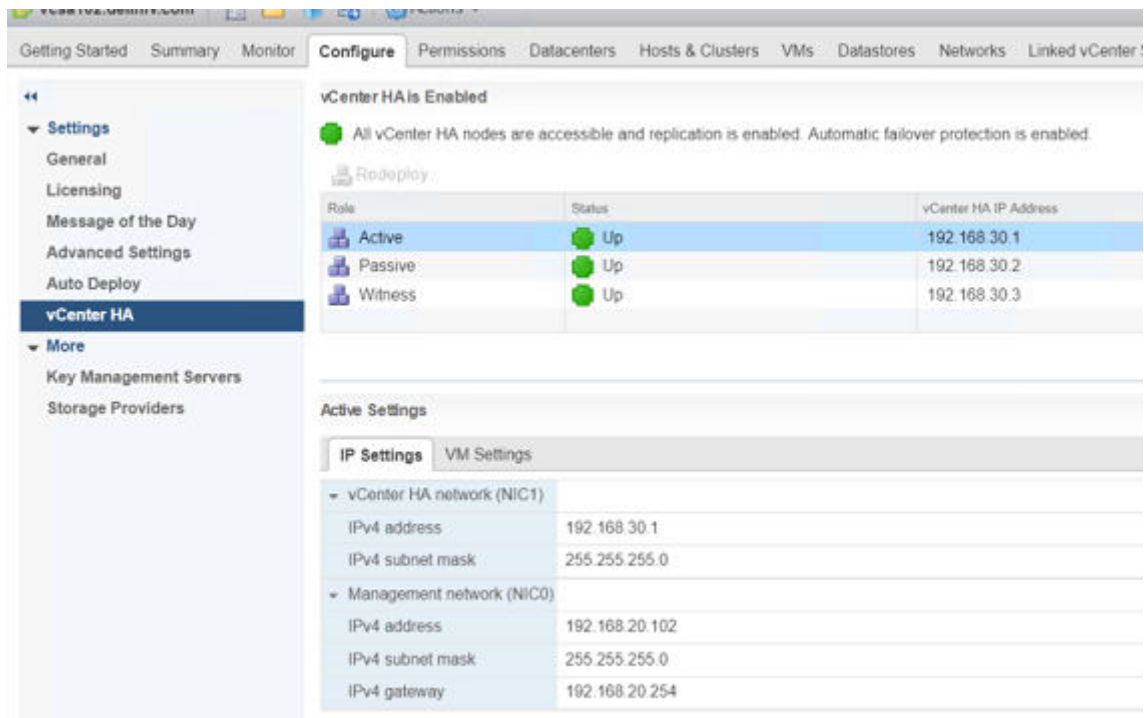


Figure 57. Status screen

NOTE: You can edit the status of vCenter HA at any time by going back into the vCenter HA menu and clicking Edit. For the options that are available, see the following vCenter HA option screen:

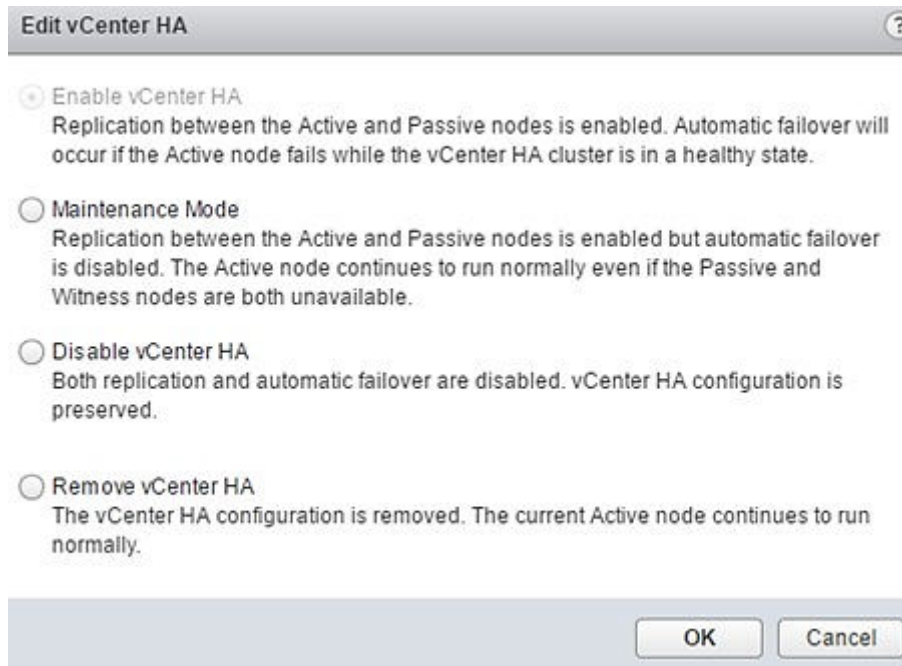


Figure 58. vCenter HA option screen

Resource cluster VCSA-HA configuration

Prerequisites

- `dvPortGroup` is created on VDS for the private HA network.
- HA private network is set to reside on a different subnet than the one used for management.
- One IP address for management, three private IP addresses, and one for each HA node, are set.

Resource Cluster VCSA-HA is configured using the Advanced Option. Using the Advanced Option to configure the vCenter HA cluster, makes you responsible to add a second NIC to VCSA and clone the active node to passive and witness nodes, and configuring the clones.

i **NOTE: Each HA node is configured to reside on a different host. Verify that the IPv4 and IPv6 addressing was not mixed when networking was configured on the nodes. Ensure that a gateway for the HA network was not specified when configuring the nodes.**

Configure Resource vCenter HA

About this task

The Configure Resource vCenter HA requires the manual addition of a second network card to the vCSA and the cloning of the appliance, two times. Also, the creation of passive and witness node clones must be done half way through the Resource vCenter HA configuration process.

i **NOTE: Each VCSA HA node requires its own ESXi host. In the installation process, set the Inventory Layout to identify and select the ESXi host where the vCSA appliance and HA instances are deployed. These steps are different than the steps that are used in the Configuring Resource vCenter HA section as the HA is being configured using the Advanced Option.**

Steps

1. In the VMware vSphere Web Client, add a second network adapter to the vCSA by editing the **Resource VCSA VM** settings and associate it to the VCSA HA Network.

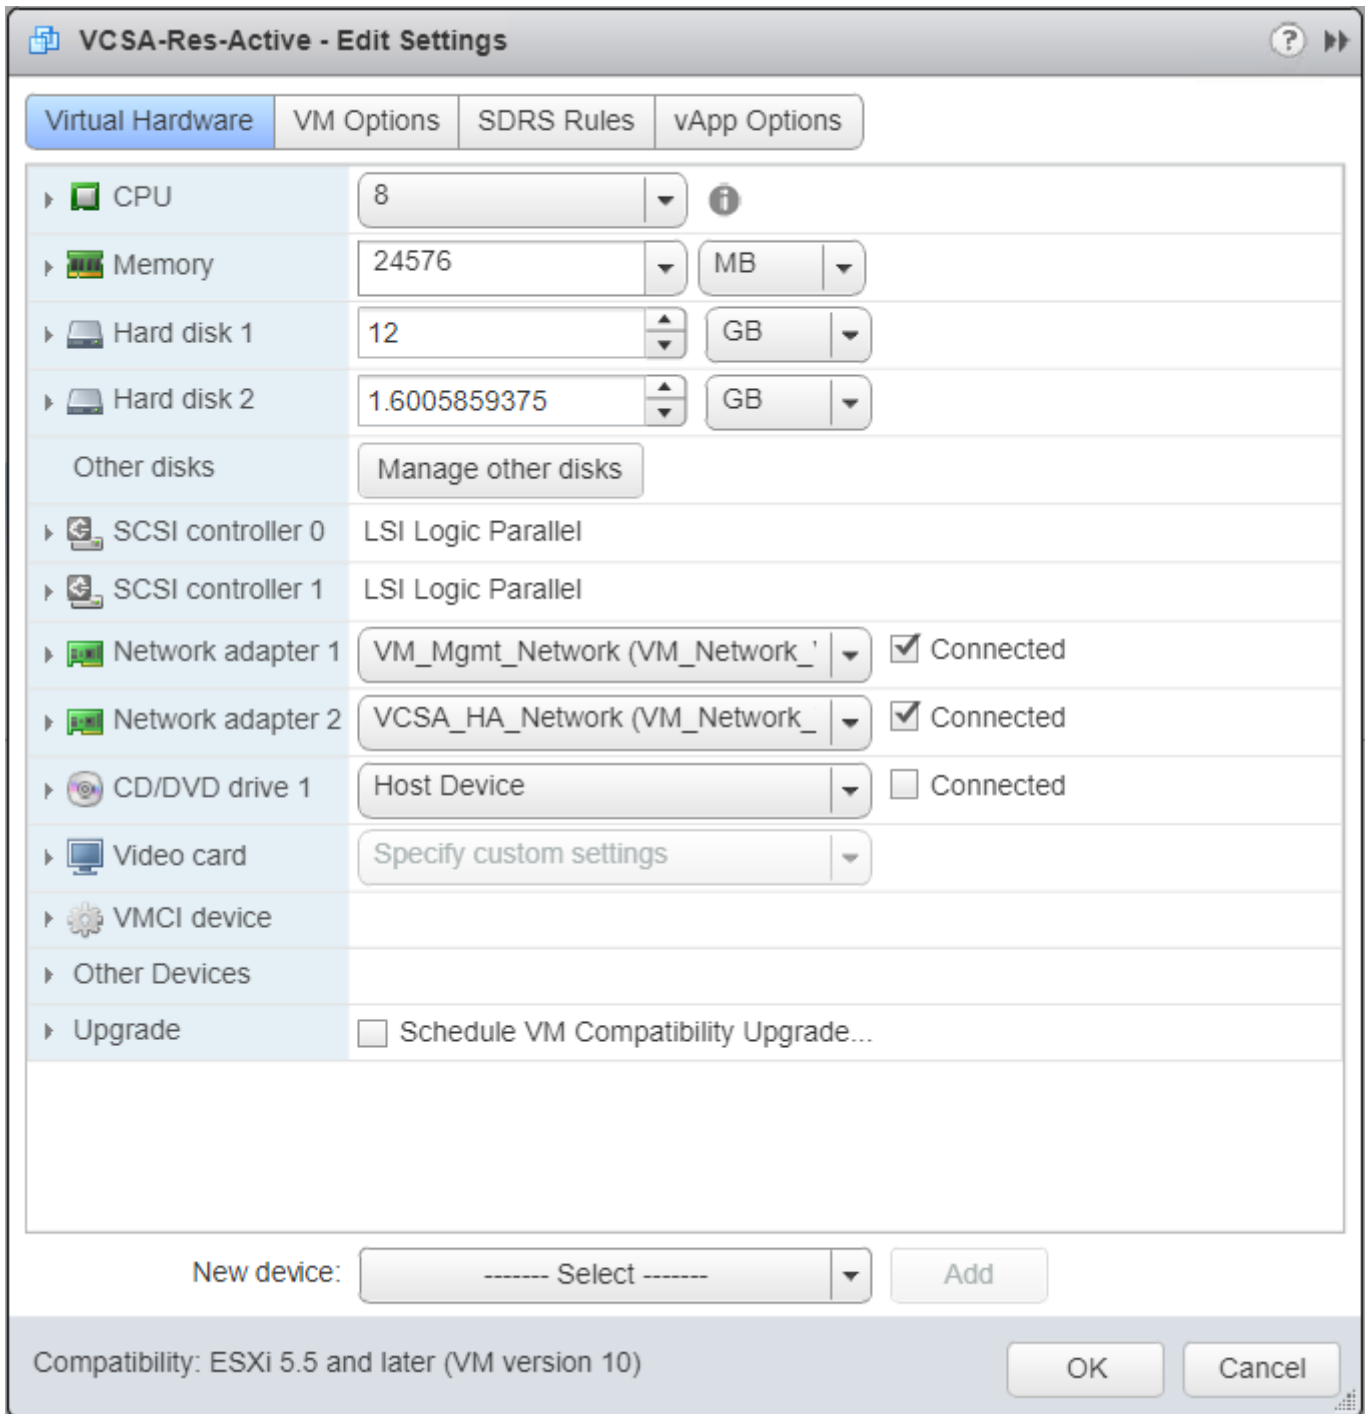


Figure 59. Edit Settings screen

2. To configure the IP settings for the second network adapter:
 - a. Log in to the resource VMware vSphere Web Client.
 - b. Click **Home, Administration, System Configuration** tab, and then select the node in which you must configure the **Network Adapter**.
 - c. From the **Manage** tab, select **Settings**. On the **Edit Setting** window, in the **Networking** tab, click **Edit** to configure the second NIC.
 - d. From the **IP4 settings** section:
 1. Select the **Use the following IPv4 settings** radio button.
 2. In the **IP address** field, enter the IP address.
 3. In the **Subnet prefix length** field, provide the subnet prefix length.
 4. In the **Default gateway** field, provide the default gateway IP address.

5. Click **OK**.

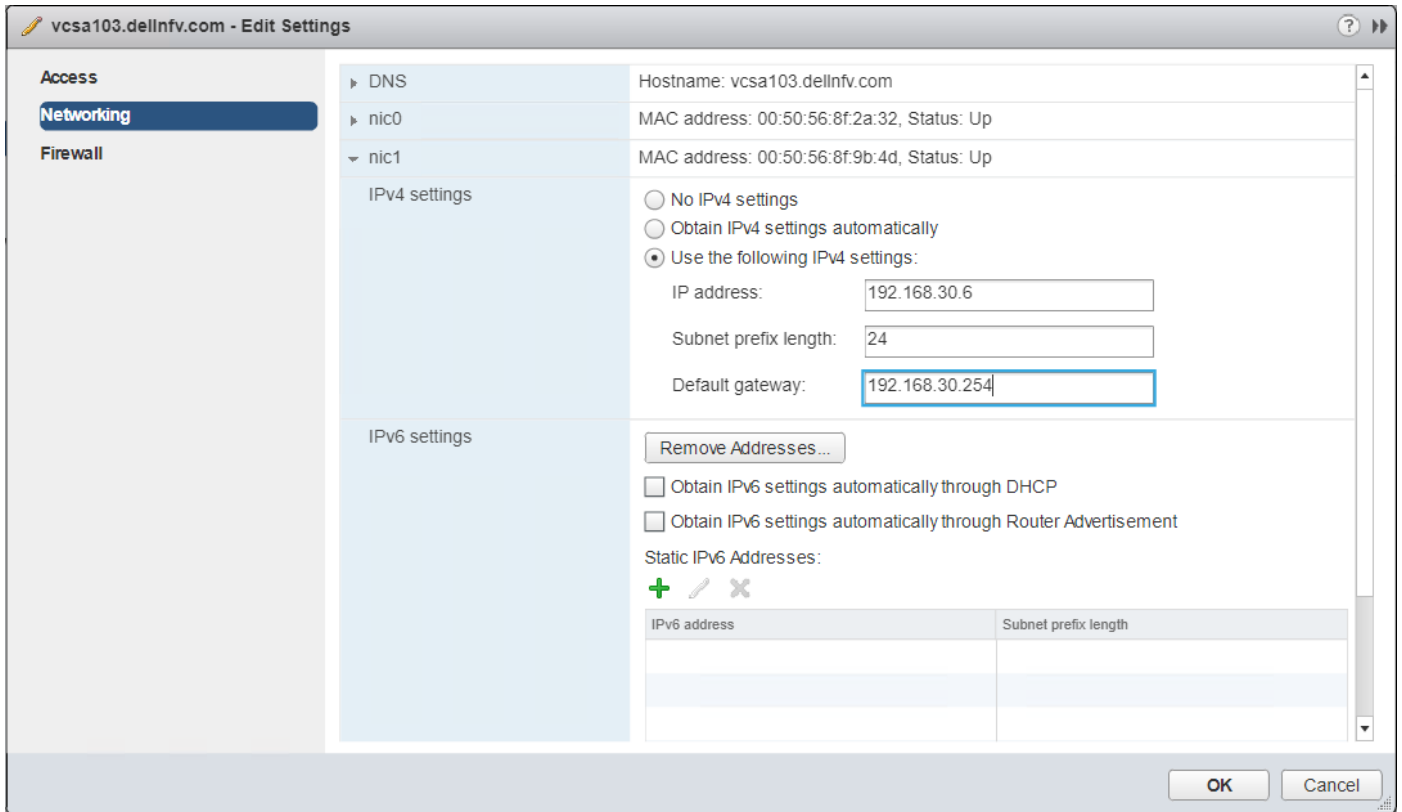


Figure 60. Networking tab

3. Right-click the top-level vCenter Server in the inventory and select **vCenter HA Settings**.
 4. From the top-right corner, click **Configure** to open **Configure vCenter HA** window.
 5. In the **Configure vCenter HA** window, select **Advanced** then click **Next**.
 6. In the fields provided, enter the **Private IP** address and **Subnet mask** for both the **Passive** and **Witness** nodes.
 7. Click **Next**.
- NOTE:** Leave the **Configure vCenter HA** window open and perform the cloning tasks. As part of the **Advanced** configuration process, clone the **Active** node to create the **Passive** and **Witness** nodes. Do not exit from the **Configure vCenter HA** window while you perform the cloning tasks.
8. Direct the **Clone VCSA passive node** and **Clone VCSA witness node** to clone the VCSA Active node and create VCSA passive and witness node. Once the cloning of VCSA passive and witness node is complete, go to the **Configure vCenter HA** window, and click **Finish**.
For additional information, see the [Clone VCSA passive node](#) and [Clone VCSA witness node](#) sections.

Clone VCSA passive node

About this task

The section provide the steps to clone vCSA passive node.

Steps

1. Log in to the Management vCenter Server.
2. Right-click the **Resource vCenter Server Appliance virtual machine (Active node)**, select **Clone**, and then **Clone to Virtual Machine**.
3. On the select a name and folder window, enter the **VM name**, select the **VM location**, and click **Next**.
4. On the **Select a compute resource** screen, select target host to run the VM, then click **Next**.
5. On the **Select storage** screen, from the **VM storage policy** drop-down list, select **vSAN Default Storage Policy**, then click **Next**.
6. On the **Select clone options** screen, check the **Customize the operating system** and **Power on virtual machine after creation** boxes, then click **Next**.

7. On the **Customize guest OS** window, click the **New Customization Spec** icon.
8. In the **New Customization Specification** window, enter a name in the **Customization Spec Name** field, and click **Next**.
9. On the **Set Computer Name** screen, enter the VCSA active node hostname in the **Enter a name** field, enter the domain name in the **Domain name** field, and then click **Next**.
10. From the **Time Zone** screen:
 - a. From the **Area** drop-down list, select **Etc**.
 - b. From the **Location** drop-down list, select **UTC**.
 - c. From **Hardware Clock Set To** drop-down list, select **UTC**.
11. In the **Configure Network** screen, select the **NIC1**, and click the **Edit** icon.
12. From the **NIC1 – Edit Network** screen, select the **Use the following IP setting** radio button, then enter the **IP address**, **Subnet mask**, and **Default gateway IP for NIC1**, and then click **OK**.
13. On the **Configure Network** screen, select the **NIC2**, and click the **Edit** icon.
14. From the **NIC2 – Edit Network** screen, select the **Use the following IP setting** radio button, then enter the **IP address** and **Subnet mask for NIC2**, and then click **OK**.

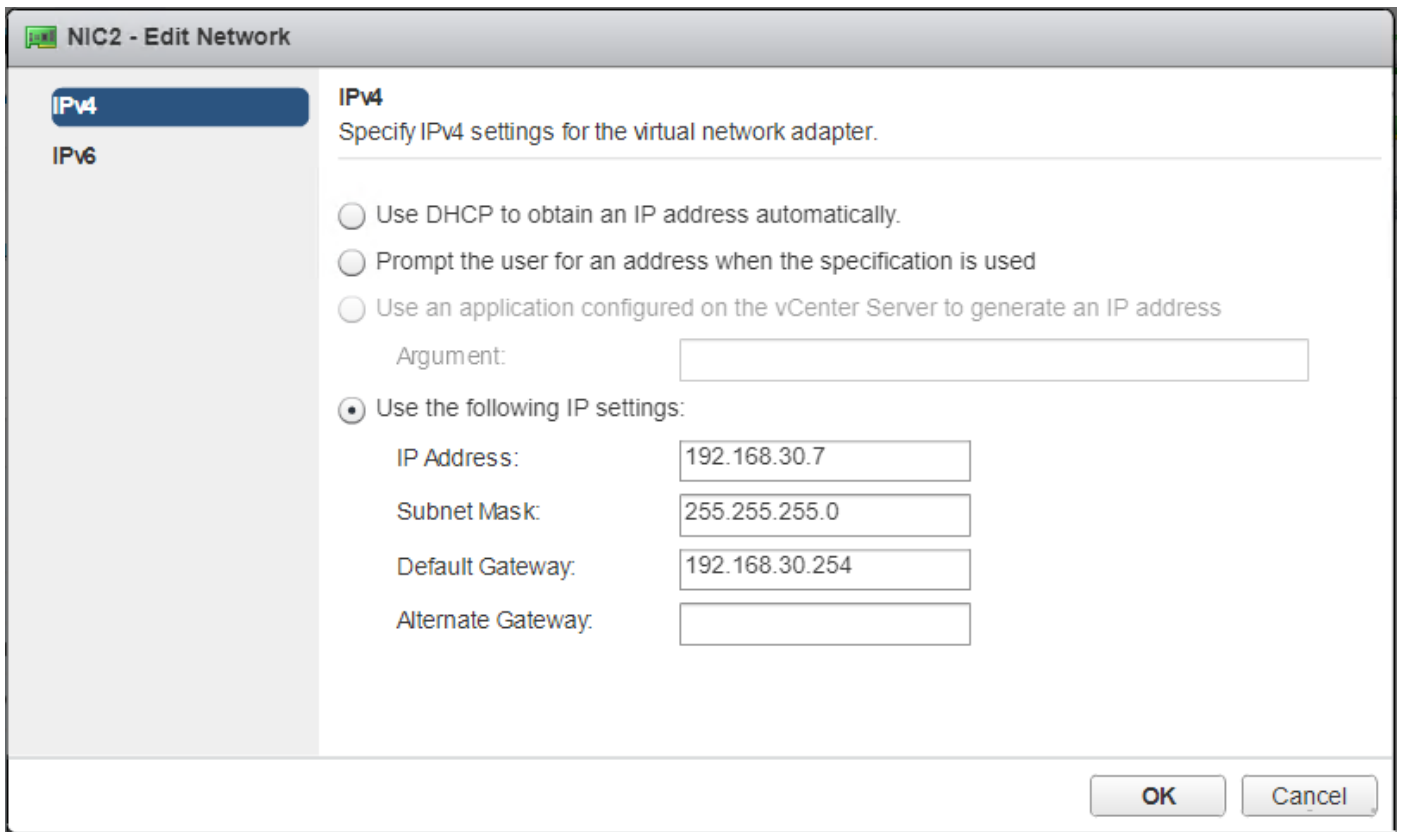


Figure 61. NIC2 – Edit Settings screen

15. On the **Configure Network** screen, keep the default setting, and then click **Next**.
16. On the **DNS and Domain Settings** window:
 - a. In the **Primary DNS** field, enter the **Primary DNS name**.
 - b. In the **DNS Search path** field, enter the domain name, then click **Add**.
 - c. Click **Next**.
17. From the **Ready to complete** screen, review the options and then click **Finish** to create guest VM.

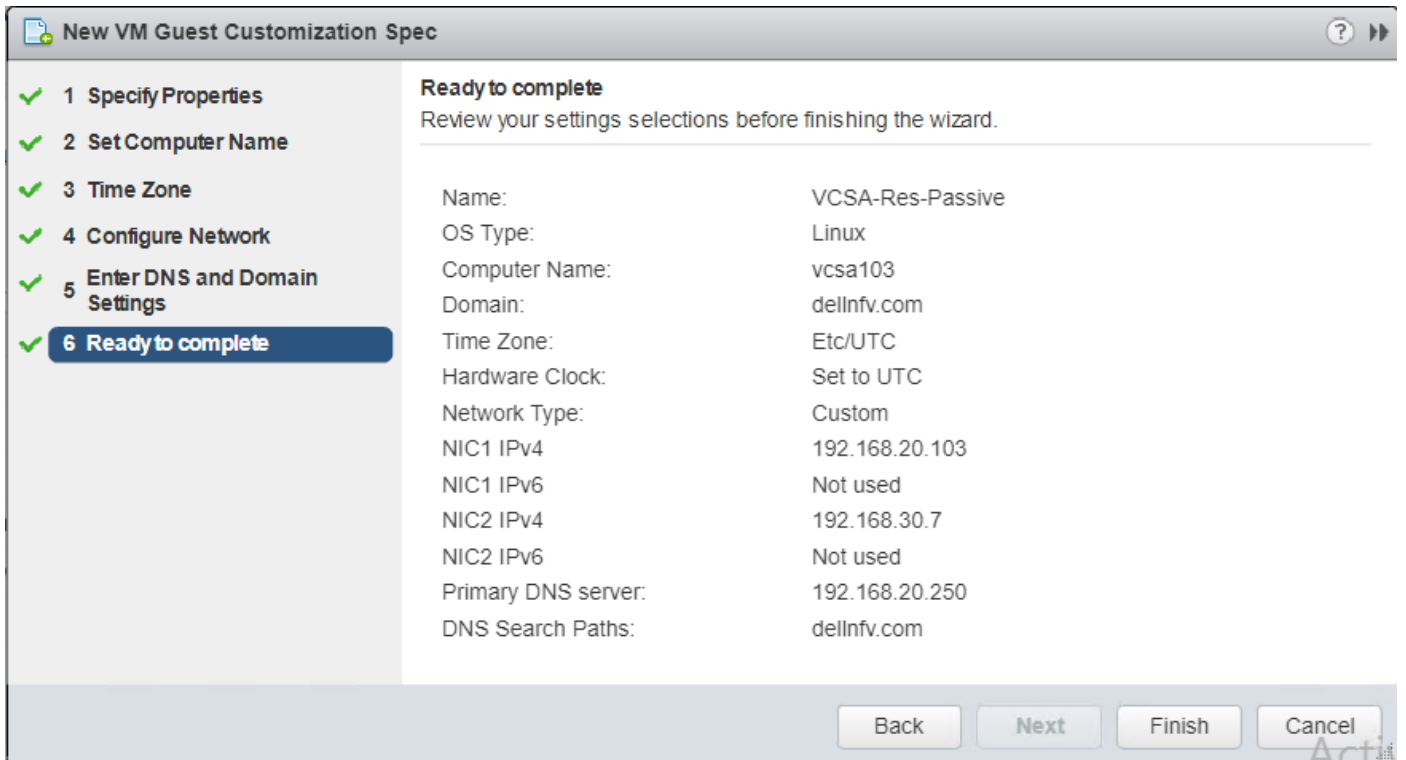


Figure 62. Ready to complete screen

18. After the VM is created, go to the **Customize guest OS** screen, select the VM, and then click **Next**.
19. On the **Customize vApp properties** screen, click **Next**.
20. From the **Ready to complete** screen, review the options that are selected then click **Finish**.

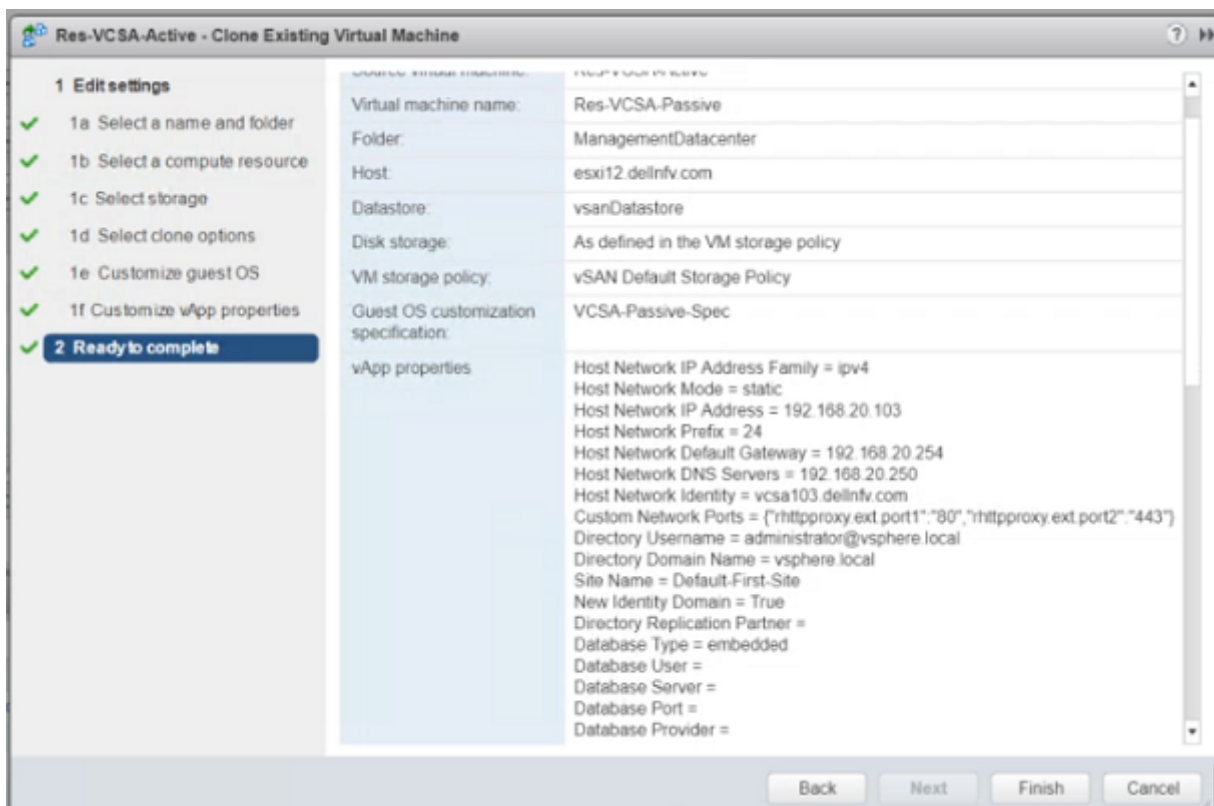


Figure 63. Ready to complete screen

Clone VCSA witness node

About this task

After cloning the VCSA passive node, clone the active node again for the witness node. The cloning of the witness node uses the same steps as the cloning of the passive node, with only a few exceptions in the process.

Steps

1. Log in to the Management vCenter Server.
2. Right-click the **Resource VCSA VM (Active node)**, select **Clone**, and then **Clone to Virtual Machine**.
3. From the **Select a name and folder** window, enter the **VM name**, **VM location**, and click **Next**.
4. On the **Select a compute resource** screen, select the target host to run the VM, then click **Next**.
5. From the **Select storage** window, select **vSAN Default Storage Policy**, and then click **Next**.
6. On the **Select clone options** screen, check the **Customize the operating system** and **Power on virtual machine after creation** boxes, then click **Next**.
7. From the **Customize guest OS** window, click the **New Customization Spec** icon.
8. In the **New Customization Specification** screen, enter a name in the **Customization Spec Name** field, and click **Next**.
9. On the **Set Computer Name** screen, enter the **VCSA active node hostname** and the **Domain name** in the fields that are provided, and then click **Next**.
10. From the **Time Zone** screen:
 - a. From the **Area** drop-down list, select **Etc**.
 - b. From the **Location** drop-down list, select **UTC**.
 - c. From **Hardware Clock Set To** drop-down list, select **UTC**.
11. On the **Configure Network** screen, select the **NIC1**, and click the **Edit** icon.
12. On the **NIC1 – Edit Network** screen, select the **Use the following IP setting** radio button, then enter the **IP address**, **Subnet mask**, and **Default gateway IP for NIC1**, and click **OK**.
13. From the **Configure Network** screen, select the **NIC2**, and click the **Edit** icon.
14. On the **NIC2 – Edit Network** window, select the **Use the following IP setting** radio button, then enter the **IP address** and **Subnet mask for NIC2**, then click **OK**.
15. From the **Configure Network** screen, click **Next**.
16. On the **DNS and Domain Settings** window:
 - a. In the **Primary DNS** field, enter the primary DNS name.
 - b. In the **DNS Search path** field, enter the domain name, then click **Add**.
 - c. Click **Next**.
17. On the **Ready to complete** screen, review the options and click **Finish** to create guest VM.
18. Once the VM is created, on the **Customize guest OS** screen, select the **VM**, and click **Next**.
19. From the **Customize vApp properties** window, locate the **SSO configuration** section and complete the following fields:
 - a. In the **Directory Username** field, enter the username for resource directory.
 - b. In the **Directory Password** field, set the password for resource directory.
 - c. In the **Directory Domain Name** field, enter the domain name for resource directory.
 - d. Click **Next**.

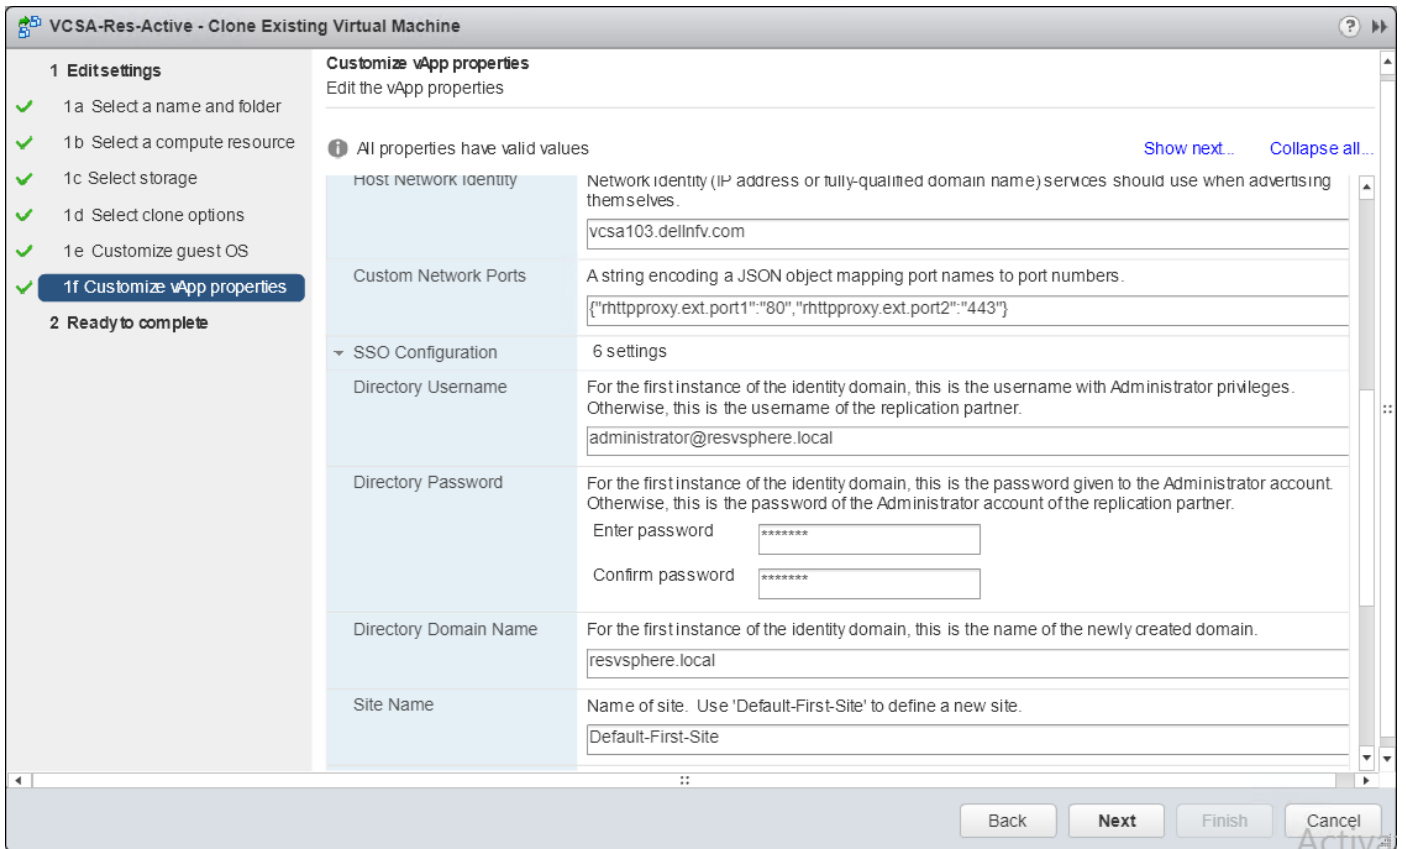


Figure 64. Customize vApp properties screen

20. On the **Ready to complete** window, review the options that are selected then click **Finish**.

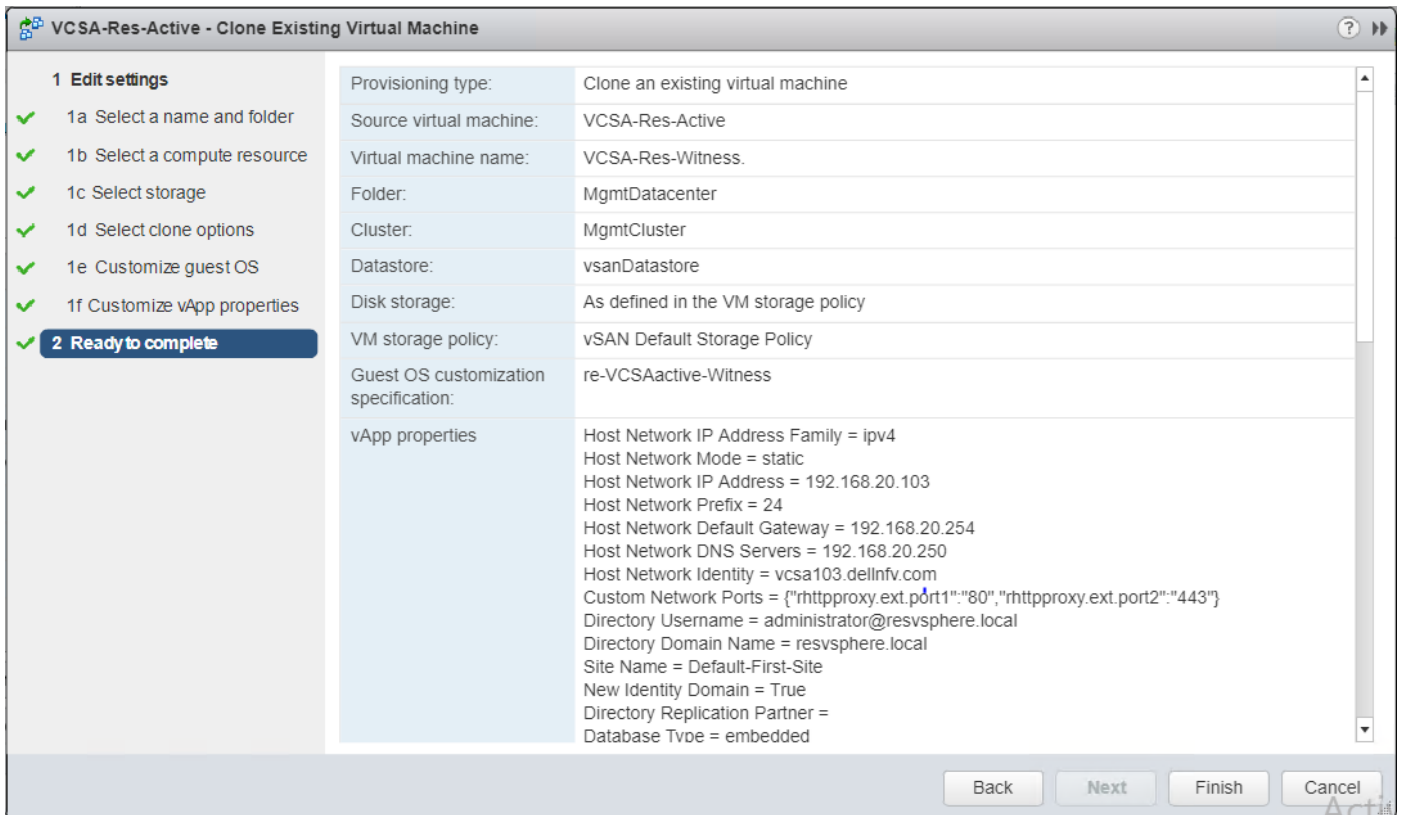


Figure 65. Ready to complete screen

NSX-T deployment and configuration

Prerequisites

- Review the necessary hardware requirements for NSX-T as specified in the System Requirements for NSX-T section of the NSX-T Installation Guide
- NSX-T OVA should be present in the deployment VM
- Verify that the following VMware products are installed:
 - VMware vCenter Server 6.7U2
 - VMware ESXi 6.7U2
- DNS entries must be added in the DNS server for all the NSX-T instances.
- For client and user access, consider the following:
 - For ESXi hosts added to the vSphere inventory by name, ensure that forward and reverse name resolution is working, otherwise, the NSX-T Manager cannot resolve the IP addresses
 - Permissions are provided to add and power on virtual machines
 - The VMware Client Integration plug-in must be installed
 - A web browser that is supported for the version of vSphere Web Client you are using
 - IPv4 IP addresses are used as IPv6 is not supported in the previously mentioned version of NSX-T

About this task

NSX-T Data Center is the software defined networking component for the vCloud NFV platform. It allows you to create, delete, and manage software-based virtual networks. In this deployment, one NSX-T manager VM and two NSX-T manager nodes are deployed.

Install NSX-T Manager Virtual Appliance

About this task

The NSX-T Manager provides a Graphical User Interface (GUI) and REST API for the creation, configuration, and monitoring of NSX-T components such as logical switches, logical routers, and firewalls. The NSX-T Manager provides an aggregated system view and is the centralized network management component of NSX. NSX-T Manager is installed as a virtual appliance on any ESXi host in the vCenter environment.

The NSX-T Manager virtual machine is packaged as an OVA file, which allows for the use of the vSphere Web Client to import the NSX-T Manager into the datastore and virtual machine inventory.

Only one instance of NSX-T Manager can be installed in an environment. When NSX-T Manager is deployed on an ESXi host, the vSphere high availability (HA) feature can be used to ensure the availability of NSX-T Manager.

NOTE: The NSX-T Manager virtual machine installation includes VMware Tools. Do not attempt to upgrade or delete VMware Tools on the NSX-T Manager.

Steps

1. In a web browser, open the vCenter Server using vSphere Web Client.
2. Select **VMs and Templates**, right-click **vCenter Server**, and select **Deploy OVF Template**.
3. Enter the download URL or click **Browse** to select the `.ova` file on your computer.

NOTE: Deploy the NSX-T manager OVA file within the Management Cluster vCenter server.
4. If required, edit the **NSX-T Manager name**, then select the folder or data center location for the deployed NSX-T Manager and click **Next**.

NOTE: The name entered displays in the vCenter inventory. The folder that is selected is used to apply permissions to the NSX-T Manager.
5. Within the **Select a resource** screen, select a **Host**, **Cluster**, **Resource pool**, or **vApp** to deploy the NSX-T Manager appliance and click **Next**.

NOTE: NSX-T Manager should be placed in a cluster that provides network management utilities.

6. In the **Review details** section, review the details of the OVA template then click **Next**.
7. From the **Select configuration** window, select the **Configuration** from the drop-down menu and click **Next**.
8. On the **Select storage** window:
 - a. From the **Select virtual disk format** drop-down list, select **Thin provision**.
 - b. From the **VM storage policy** drop-down list, and then select **vSAN Default Storage Policy**.
 - c. Select the **datastore** and click **Next**.
9. From the **Select networks** screen, select the port group or destination network for the NSX-T Manager and click **Next**.
10. On the **Customize template** screen, specify the **Root**, **Admin**, and **Audit user passwords**.
11. On the **Customize template** screen, locate the **Network properties** section, and fill the following fields:
 - a. **Hostname**: Enter the hostname.
 - b. **Default IPv4 Gateway**: Enter the IP address of default gateway for NSX-T Manager
 - c. **Management Network IPv4 Address**: The IPv4 address for the first interface
 - d. **Management Network Netmask**: The netmask for the first interface
12. In the **DNS** section:
 - a. In the **DNS Server list** field, enter the IP address of the **DNS Server for NSX-T Manager**.
 - b. In the **Domain Search List** field, enter the **Domain name for NSX-T Manager**.
13. In the **Services Configuration** section:
 - a. In the **NTP Server List** field, enter the IP address for **NTP server**.
 - b. Check the **Enable SSH** and **Allow root SSH logins** check boxes to enable SSH service.
 - c. Click **Next**.
14. On the **Ready to complete** window, verify the details before deployment and click **Finish** to start deployment.

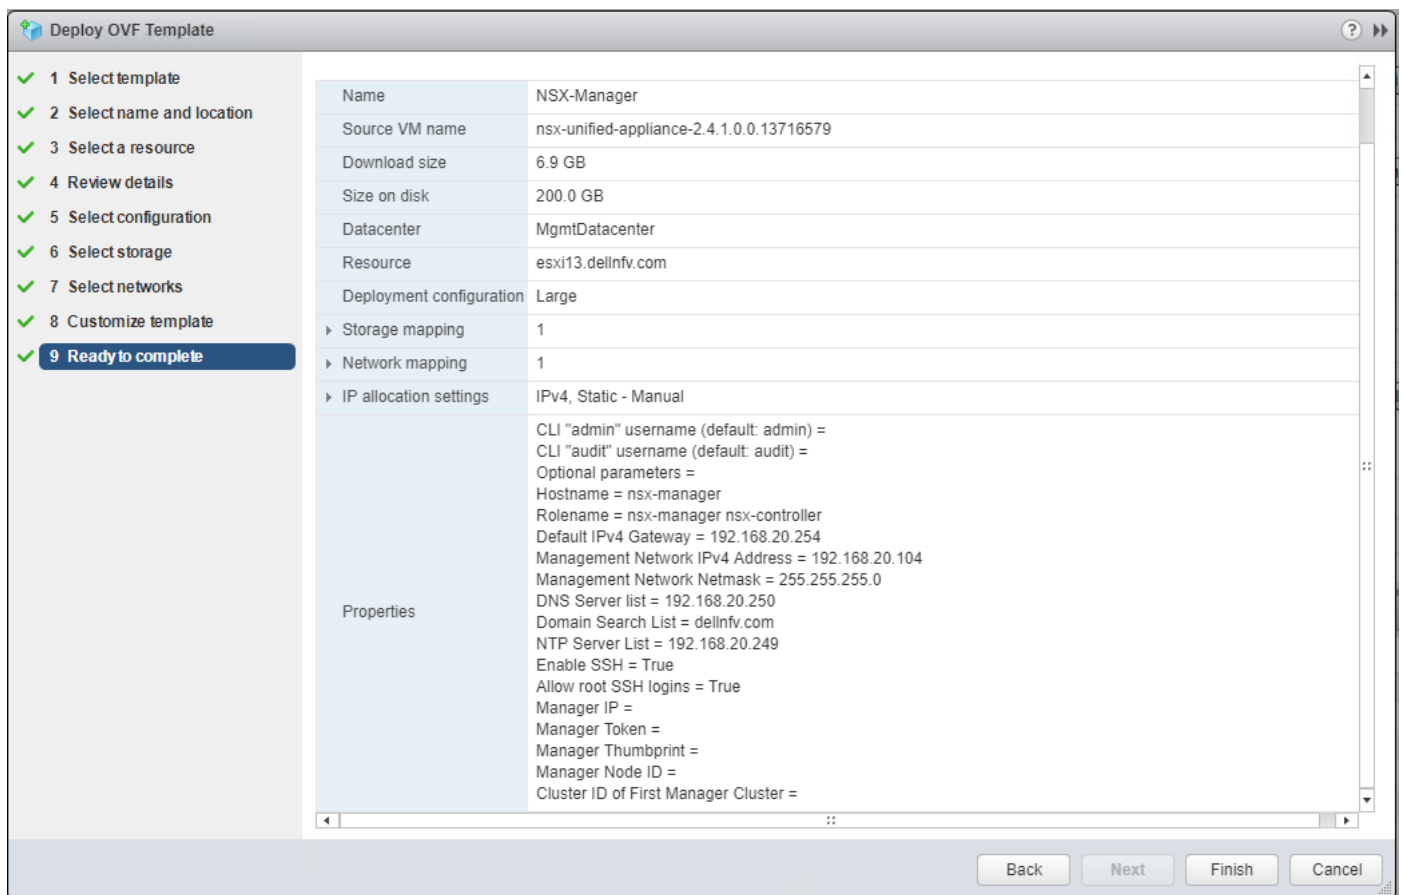


Figure 66. Ready to complete screen

15. Once the deployment is complete, perform the following steps:
 - a. Power on NSX-T Manager VM from vSphere Web Client.

- b. After the NSX-T Manager VM boots completely, connect to the NSX-T Manager GUI using the following URL: `https://<IP/FQDN of NSX-T Manager>`
- c. Review the EULA, and if you agree to the terms, check **I understand and accept the terms of the license agreement** box and click **CONTINUE**.
- d. Click **SAVE** to finish.

Add license key

About this task

This section provides the steps to assign license key to NSX-T Manager.

Steps

1. From your browser, use administrator credentials to log in to the **NSX Manager** at `https://nsx-manager-fqdn`.
2. Go to **System** and then **Licenses**.
3. Click **(+) Add**.
4. From the **Add license** screen, enter the license key then click **Add**.

Add Compute Manager for management and resource VCSA

About this task

A compute manager is an application that manages resources such as hosts and VMs. NSX-T polls compute managers to find out about changes such as the addition or removal of hosts or VMs and update its inventory accordingly.

Steps

1. From your browser, log in with admin credentials to an NSX Manager at `https://nsx-manager-ip-address`.
2. Go to **System, Fabric**, and then **Compute Managers**.
3. From the **Compute Manager** screen, click the **Add (+)** icon to add a new compute manager.

New Compute Manager ? ×

Name*

Description

Domain Name/IP Address*

Type* ▼

Username*

Password*

SHA-256 Thumbprint

Figure 67. New Compute Manager screen

4. From the **New Compute Manager** screen, enter the required details for compute manager.
 - i **NOTE:** If you do not have a compute manager thumbprint, click **ADD** to automatically detect the thumbprint of the compute manager. After clicking **ADD**, the **Invalid Thumbprint** dialog box requests that you use the new server thumbprint.
5. Click **ADD** and use the server thumbprint to create the compute manager.
6. Refresh the **Compute manager** tab to confirm that the **Status** displays as **Registered** and **Up**.
 - i **NOTE:** Repeat the steps provided in this section to create a compute manager for the management vCSA.
7. If the progress icon changes from In progress to **Not registered**, perform the following steps to resolve the error:
 - a. Select the error message and click **Resolve** One possible error message is the following: **Extension already registered at CM with ID.**
 - b. Enter the vCenter Server credentials, and then click **Resolve**. The existing registration is replaced.

Deployment of NSX-T node and cluster from NSX-T Manager

Prerequisites

- vCenter Server and vSphere ESXi hosts are successfully deployed
- NSX-T Manager is successfully deployed
- Register vSphere ESXi host to the vCenter Server
- vSphere ESXi host has the necessary CPU, memory, and hard disk resources to support 12vCPUs, 48 GB RAM, and 360 GB storage

About this task

You can deploy NSX-T Nodes using the NSX-T Manager on vSphere ESXi hosts that are managed by a vCenter Server. Two NSX-T nodes are deployed in this deployment.

Steps

1. From your browser, log in with admin credentials to an NSX Manager at <https://nsx-manager-ip-address>.
2. Go to **System**, and then **Overview**.
3. From the **Overview** page, click **ADD NODES**.
4. On the **Common Attributes** screen, select the **Compute Manager**.
Optionally, slide the **SSH and Root access** toggle switch to **Enable**.

The screenshot shows the 'Common Attributes' configuration screen in the NSX-T Manager. On the left, a sidebar shows '1 Common Attributes' selected and '2 Nodes' below it. The main area is titled 'Common Attributes' and contains the following fields:

- Compute Manager**: Mgmt-VCSA (dropdown menu)
- Enable SSH**: Yes (toggle switch)
- Enable Root Access**: Yes (toggle switch)
- Node Credentials**:
 - CLI Username**: admin
 - CLI Password**: [masked]
 - Confirm CLI Password**: [masked]
 - Root Password**: [masked]
 - Confirm Root Password**: [masked]
- DNS Servers**: 192.168.20.250 (input field)

At the bottom right, there are 'CANCEL' and 'NEXT' buttons.

Figure 68. Common Attributes screen

5. Set the **CLI password** and **Root password**.
6. Enter the **DNS Server** and **NTP Server IP** address in their respective field.
7. From the **Form Factor**, select **Medium**.
8. Click **Next**.
9. On the **Nodes** screen, enter the following details to add a **New** node:

- a. **Name**: Enter a name for the new node
- b. **Cluster** drop-down list: Select the cluster to deploy controllers
- c. **Resource** pool drop-down list: Select a resource pool if any
- d. **Host** drop-down list: Select the ESXi host on which controller is deployed
- e. **Datastore** drop-down list: Select a datastore
- f. **Network** drop-down list: Select a network for controller
- g. **Management IP/Netmask** field: Enter the network IP address or netmask IP
- h. **Management Gateway field**: Enter the gateway IP for controllers

Figure 69. New controller details screen

10. Click **FINISH**.

NOTE: Repeat the steps provided in this section to deploy another NSX-T node.

Validate NSX-T node and cluster deployment

About this task

NSX-T Controller deployment can be verified from the NSX-T Manager user interface (UI).

Steps

To validate the deployment, click **System**, and then click **Overview**.

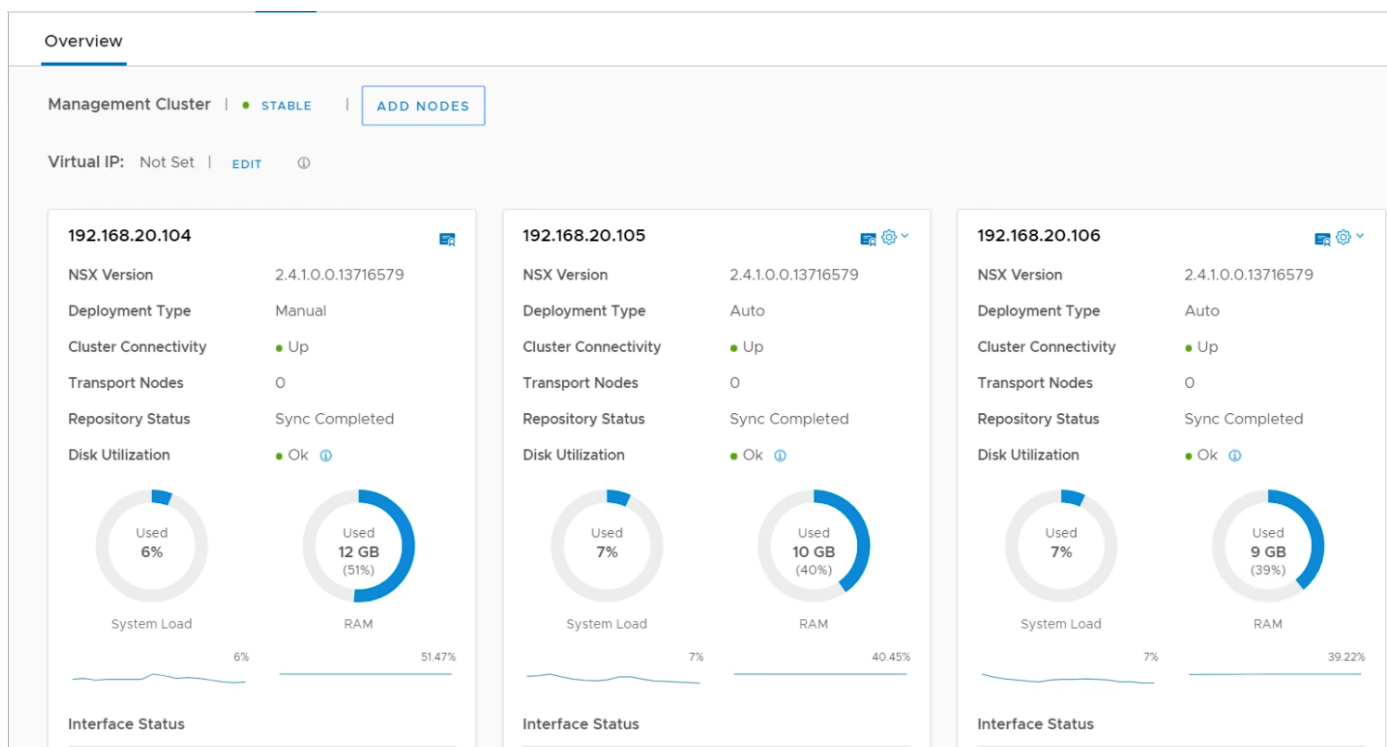


Figure 70. NSX-T Controller details page

Add Virtual IP

About this task

This section provides the steps to assign a virtual IP in NSX-T Manager.

Steps

1. From a web browser, access the **NSX Manager** at `https://nsx-manager-ip-address` and use the administrator credentials to log in.
2. Click **System**, and then **Overview**.
The **Overview** screen displays.
3. Locate the **Virtual IP** section, click **EDIT**.
The **Change Virtual IP** screen displays.
4. Enter the **Virtual IP Address** in the field that is provided and then click **Save**.

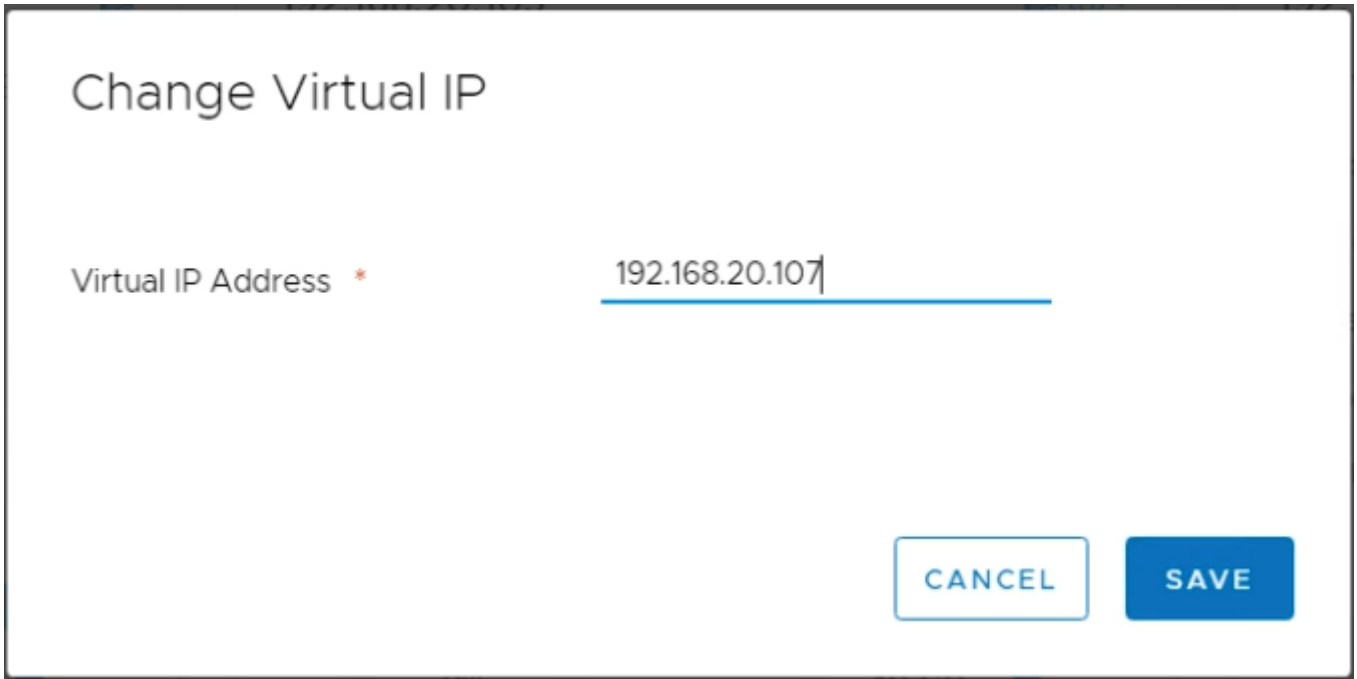


Figure 71. Change Virtual IP screen

Configure NSX-T Manager

This section covers the following steps to configure the NSX-T Manager:

- [Create transport zones](#)
- [Create uplink profiles](#)
- [Create IP pool for tunnel endpoints](#)
- [Create host transport nodes](#)

Create transport zones

About this task

This section provides the steps to add transport zones in NSX-T. The following table provides the required details used in this deployment to create transport zone:

Table 25. Transport zone details

Name	N-VDS name	Host membership criteria	Traffic type
Overlay-TZ	nvds-overlay	Standard	Overlay
Dpdk-TZ	nvds-dpdk	Enhanced Datapath	VLAN
Vlan-TZ	nvds-vlan	Standard	VLAN

NOTE: For this deployment, Intel NICs are used as it supports the N-VDS Enhanced mode feature. For the deployment with QLogic NICs, use N-VDS Standard mode as QLogic does not support N-VDS Enhanced mode. To install DPDK drivers on Intel NICs, see [Install DPDK drivers](#).

Steps

1. From a web browser, go to `https://nsx-manager-fqdn` and use the administrator credentials to log in to the **NSX Manager**.
2. Click **System, Fabric**, and then **Transport Zones**.
3. From the **Transport Zones** tab, click **(+) Add**.
The **New Transport Zone** window opens.

- Using the information provided in the [Transport zone details](#) table, enter the **Name**, **N-VDS name**, **Host membership criteria**, and **Traffic type** details.
- Click **ADD** to create transport zone.

New Transport Zone

Name *

Description

N-VDS Name *

Host Membership Criteria
 Standard (For all hosts)
 Enhanced Datapath (For ESXi hosts with version 6.7 or above)

Traffic Type
 Overlay
 VLAN

Uplink Teaming Policy Names

Figure 72. New Transport Zone screen

- Using the information provided in the [Transport zone details](#) table, repeat the steps in this section to create more transport zones.

Transport Zones						
	ID	Traffic Type	N-VDS Name	Status	Host Membership Criteria	Where Used
<input type="checkbox"/> Transport Zone ↑						
<input type="checkbox"/> Dpdk-TZ	00db...0151	VLAN	nvds-dpdk	Up	ENS	Where Used
<input checked="" type="checkbox"/> Overlay-TZ	cb86...d039	Overlay	nvds-overlay	Up	Standard	Where Used
<input type="checkbox"/> Vlan-TZ	8286...6e33	VLAN	nvds-vlan	Up	Standard	Where Used

Figure 73. Transport Zones screen

Create uplink profiles

Prerequisites

- Verify that each uplink within the uplink profile corresponds to an **Up** and **Available** physical link on your ESXi host or on the NSX-T Edge node.

About this task

An uplink profile defines policies for the links from ESXi hosts to NSX-T logical switches or from NSX-T Edge nodes to top-of-rack switches. The settings that are defined by uplink profiles may include teaming policies, active/standby links, the transport VLAN ID, and the MTU setting.

The Uplink profiles enable you to consistently configure identical capabilities for network adapters across multiple hosts or nodes. Uplink profiles are containers for the properties or capabilities that you want your network adapters to have. Instead of configuring individual properties or capabilities for each network adapter, you can specify the capabilities in uplink profiles that can be applied when you create NSX-T transport nodes.

Table 26. Uplink profile details

Name	Teaming policy	Active uplinks	Standby uplinks	Transport VLAN ID	MTU
edge-overlay-uplink-profile	Failover Order	Uplink1	-	70	1600
edge-vm-uplink-profile	Failover Order	Uplink1	-	40	1600
edge-vlan-uplink-profile	Failover Order	Uplink1	-	20	1600

Name	Teaming policy	Active uplinks	Standby uplinks	Transport VLAN ID	MTU
host-overlay-uplink-profile	LOADBALANCE_SR C_MAC	LAG1	-	70	1600
host-dpdk-uplink-profile	LOADBALANCE_SR C_MAC	LAG1	-	40	1600

NOTE: It is recommended that the Maximum Transmission Unit (MTU) settings for the Transport VLAN must be configured to support 1600 bytes.

Steps

1. From a web browser, log in to the **NSX Manager** at `https://nsx-manager-ip-address` using administrator credentials.
2. Select **System, Fabric, Profiles**.
3. From the **Uplink Profiles** tab, click **+ ADD**.
4. Using the information from the [Uplink profile details](#) table, enter the uplink profile name in the **Name** field.
5. Optionally, in the **Description** field, enter the description for Uplink profile.
6. Use the information in the [LAG details](#) table to create a LAG host overlay and host DPDK uplink profile:

Table 27. LAG details

Name	LACP mode	LACP load balancing	Uplinks	LACP time out
LAG1	Active	Source MAC address	2	Fast

- a. In the **LAGs** section, click **+ ADD**.
 - b. Refer to the [LAG details](#) table and fill the **Name, LACP Mode, LACP Load Balancing, Uplinks,** and **LACP Time Out** field.
7. In the **Teamings** section, click **+ ADD**.
 8. In **Teamings** section, use the information from the [Uplink profile details](#) table to enter the **Teaming Policy, Active Uplinks,** and **Standby Uplinks** information.
 9. Using the information from the [Uplink profile details](#) table, enter the **Transport VLAN** and **MTU** details, then click the **ADD** button.

New Uplink Profile ? ×

Name*

Description

LAGs

[+ ADD](#) [DELETE](#)

<input type="checkbox"/> Name*	LACP Mode	LACP Load Balancing*	Uplinks*	LACP Time Out
No LAGs found				

Teamings

[+ ADD](#) [CLONE](#) [DELETE](#)

<input type="checkbox"/> Name*	Teaming Policy*	Active Uplinks*	Standby Uplinks
<input type="checkbox"/> [Default Teaming]	Failover Order		

Active uplinks and Standby uplinks are user defined labels. These labels will be used to associate with the Physical NICs while adding Transport Nodes.

Transport VLAN

MTU

[CANCEL](#) [ADD](#)

Figure 74. Teamings listing

The **Uplink Profile** is successfully created.

- Using the information that is provided in the [Uplink profile details](#) table , repeat the steps in this section to create the more uplink profiles.

Uplink Profiles							
NIOC Profiles							
Edge Cluster Profiles							
Edge Bridge Profiles							
Configuration							
Transport Node Profiles							
+ ADD	EDIT	DELETE	ACTIONS				
<input type="checkbox"/>	Uplink Profile	ID	Teaming Policy	Active Uplinks	Standby Uplinks	Transport VLAN	MTU
<input type="checkbox"/>	edge-overlay-uplink-profile	1350...1b23	Failover Order	Uplink1		70	1600
<input type="checkbox"/>	edge-vlan-uplink-profile	0663...0351	Failover Order	Uplink1		20	1600
<input type="checkbox"/>	edge-vm-uplink-profile	55c2...4cfa	Failover Order	Uplink1		40	1600
<input type="checkbox"/>	host-dpdk-uplink-profile	9914...e59f	LOADBALANCE_SRC_...	LAG1		40	1600
<input type="checkbox"/>	host-overlay-uplink-profile	b34f...f081	LOADBALANCE_SRC_...	LAG1		70	1600

Figure 75. Uplink profiles

Create IP pool for tunnel endpoints

About this task

IP pool can be used for tunnel endpoints. Tunnel endpoints are the source and destination IP addresses used in external IP header to identify the ESXi hosts originating and terminating the NSX-T Data Center encapsulation of overlay frames. You can also use either DHCP or manually configured IP pools for tunnel endpoint IP addresses.

Steps

1. From a web browser, log in to the NSX Manager at `https://nsx-manager-ip-address` using administrator credentials.
2. Click **Advanced Networking & Security, Inventory**, and then **Groups**.
3. From the **IP Pools** tab, click **+ ADD**.
The **Add New IP Pool** window opens.
4. Enter the name and description for the new IP Pool in the respective fields.
5. Locate the **Subnets** section, click **+ADD**.
6. Enter the required details to add the subnets then click **ADD**.

Add New IP Pool ?

Name*

Description

Subnets

[+ ADD](#) [DELETE](#)

<input checked="" type="checkbox"/> IP Ranges*	Gateway	CIDR*	DNS Servers	DNS Suffix
<input checked="" type="checkbox"/> 192.168.7.2 - 192.168.7.99	192.168.7.254	192.168.7.0/24	192.168.20.250	dellnfv.com

Figure 76. Add New IP pool screen

The IP Pool is created successfully.

Create host transport nodes

Prerequisites

- Transport zone must be configured
- Uplink profile must be configured, or you can use the default uplink profile
- IP pool must be configured, or DHCP must be available in the network deployment
- Minimum of one unused physical NIC must be available on the host node

About this task

A transport node is a node that participates in an NSX-T Data Center overlay or NSX-T Data Center VLAN networking.

Steps

1. From a web browser, use the administrator credentials to log in to the **NSX Manager** at `https://nsx-manager-ip-address`.
2. Click **System, Fabric**, and then **Nodes**.
3. From the Host Transport Nodes tab, locate the **Managed by** drop-down list select **Res-VCSA**.
4. Expand the **Resource Cluster** hosts listing, select a host, and click **Configure NSX**.

5. On the **Host Details** screen, enter the hostname then click Next.
6. In the **Configure NSX** screen, locate the **Host Details** tab, and in the **Name** field, enter the **Transport node name** then click **Next**.
7. From the **Configure NSX** tab, locate the **Transport Zones** drop-down list, and then select the **Transport zones**.
8. From the **N-VDS Name** drop-down list, select the **N-VDS Name** created for external network, such as **nvds-dpdk**.
9. From the **Uplink Profile** drop-down list, select **host-dpdk-uplink-profile**.
10. From the **LLDP Profile** drop-down list, select **LLDP [Send Packet Enabled]**.
11. In the **Physical NICs** field, select **vmnic8** and **vmnic9** from the drop-down list for **LAG1-0** and **LAG1-1**.

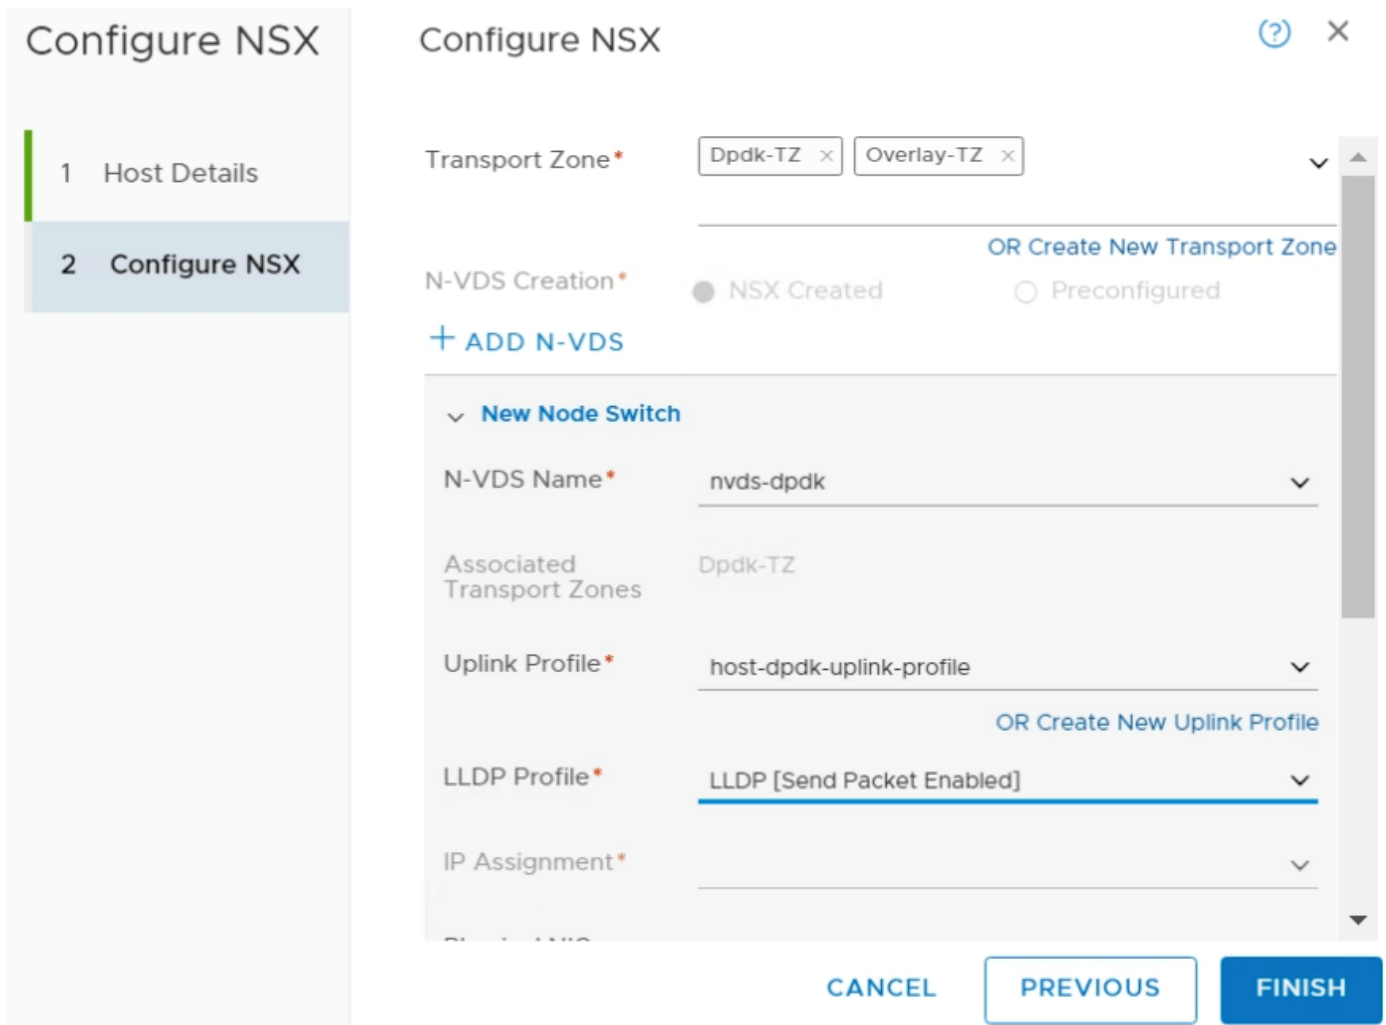


Figure 77. N-VDS screen

12. On the **N-VDS** tab, click **+ADD N-VDS** to add second N-VDS for the overlay network.
13. In the **N-VDS Name** drop-down list, select the overlay network **N-VDS**, for example, *nvds-overlay* used in this deployment.
14. From the **Uplink Profile** drop-down list, select **host-overlay-uplink-profile**.
15. From the **LLDP Profile** drop-down list, select **LLDP [Send Packet Enabled]**.
16. In the **IP Assignment** field, select **Use IP Pool** option from the drop-down- list.
17. In the **IP Pool** field select **TEP-IP-Pool** from drop-down list.
18. In the **Physical NICs** field, select **vmnic5** and **vmnic7** from the drop-down list respectively for **LAG1-0** and **LAG1-1**.
19. Repeat the steps in this section to add the more transport nodes.

Host Transport Nodes										
Managed by Res-VCSA										
CONFIGURE NSX REMOVE NSX ACTIONS View All										
Node	ID	IP Addresses	OS Type	NSX Configuration	Configuration Stat	Node Status	Transport Zones	NSX Version	N-VDS	
ResCluster (4) MoRef ID...										
esxi6.dellnfv.com	5869...e6...	192.168.100.6	ESXi 6.7.0	Configured	Success	Up	Dpdk-TZ Overlay-TZ	2.4.1.0.0.13716...		2
esxi7.dellnfv.com	fa71...995a	192.168.100.7	ESXi 6.7.0	Configured	Success	Up	Dpdk-TZ Overlay-TZ	2.4.1.0.0.13716...		2
esxi9.dellnfv.com	4b22...29...	192.168.100.9	ESXi 6.7.0	Configured	Success	Up	Dpdk-TZ Overlay-TZ	2.4.1.0.0.13716...		2
esxi8.dellnfv.com	2f3a...3e...	192.168.100.8	ESXi 6.7.0	Configured	Success	Up	Dpdk-TZ Overlay-TZ	2.4.1.0.0.13716...		2
EdgeCluster (... MoRef ID...										

Figure 78. Host Transport Nodes

Installation of NSX-T edge

Prerequisites

- Perform review the NSX-T Edge network requirements in the *NSX-T Edge Networking Setup Guide*
- If a vCenter Server is registered as a compute Manager in NSX-T, use the NSX-T manager UI to configure a host as an NSX-T Edge node and automatically deploy it on the vCenter Server
- Verify that the vCenter Server vSAN datastore on which the NSX-T Edge is installed, has a minimum of 120 GB storage or disk space available
- Verify that the vCenter Server cluster or host has access to the specified networks and vSAN datastore in the configuration
- Transport zones are configured
- An uplink profile is configured, or you can use the default uplink profile for bare-metal NSX-T edge nodes
- Ensure that an IP pool is configured or that it is available in the network deployment
- Verify that at least one unused physical NIC is available on the host or NSX-T edge node

About this task

NSX-T edge provides connectivity to the external networks. In this deployment four edge VMs, Edge01, Edge02, Edge03, and Edge04 are created.

Steps

- From a web browser, access the **NSX Manager** at <https://nsx-manager-ip-address> and use the administrator credentials to log in.
- Click **System, Fabric, Nodes**.
- From the **Edge Transport Nodes** tab, click **+ ADD EDGE VM**.
- In the **Name and Description** screen, enter the **Name, Host name/FQDN, Description**, in the fields provided.
- In the **Form Factor** section, select the **Medium** form factor size, then click **NEXT**.
- On the **Credentials** screen:
 - Set the **CLI password** and click to turn-on the **Allow SSH login** toggle switch.
 - Set the **Root password** and click to turn on the **Allow SSH login** toggle switch.
 - Click **NEXT**.
- On **Configure Deployment** window:
 - For the Compute Manager, select the resource compute manager.
 - Select the **Cluster, Host, and Datastore**, then click **NEXT**.

Figure 79. Configure Deployment screen

8. On the **Configure ports** screen, select **Static for the IP Assignment** and enter both the **Management IP** and **Default Gateway** information in the fields provided.
9. In the **Management Interface** section, select the **VM management network** from the drop-down list.
10. In the **Search Domain Names** field, enter the **domain name**.
11. In the **DNS Servers** field, enter the DNS server IP address.
12. In the **NTP Servers** field, enter the NTP server IP address, then click **Next**.

Figure 80. Configure ports screen

13. On the **Configure NSX** window, configure the edge transport nodes using the information in the [Edge transport nodes details](#) table:

Table 28. Edge transport nodes details

Edge VM	Transport zone	Edge switch name	Uplink profile	IP assignment	IP pool	DPDK Fastpath interfaces	
						Uplink	Connected to
Edge01	overlay-TZ	nvds-overlay	edge-overlay-uplink-profile	Use IP Pool	TEP-IP-Pool	Uplink1	Overlay-Network
	dpdk-TZ	nvds-dpdk	edge-vm-uplink-profile	-	-	Uplink1	External-Network
Edge02	overlay-TZ	nvds-overlay	edge-overlay-uplink-profile	Use IP Pool	TEP-IP-Pool	Uplink1	Overlay-Network
	dpdk-TZ	nvds-dpdk	edge-vm-uplink-profile	-	-	Uplink1	External-Network
Edge03	overlay-TZ	nvds-overlay	edge-overlay-uplink-profile	Use IP Pool	TEP-IP-Pool	Uplink1	Overlay-Network
	vlan-TZ	nvds-vlan	edge-vlan-uplink-profile	-	-	Uplink1	External-Network
Edge04	overlay-TZ	nvds-overlay	edge-overlay-uplink-profile	Use IP Pool	TEP-IP-Pool	Uplink1	Overlay-Network
	vlan-TZ	nvds-vlan	edge-vlan-uplink-profile	-	-	Uplink1	External-Network

- On the **Configure NSX** tab, use the information in the [Edge transport nodes details](#) table to select the transport zone from the **Transport Zone** drop-down list.
- Using the information in the [Edge transport nodes details](#) table, enter the information in the **Edge Switch Name, Uplink Profile, IP assignment, IP Pool,** and the **DPDK Fastpath Interfaces** details.
- Click **+ADD N-VDS** to add second N-VDS.
- Using the information from the [Edge transport nodes details](#) table, enter the **Edge Switch Name, Uplink Profile, IP assignment, IP Pool,** and **DPDK Fastpath Interfaces** details.

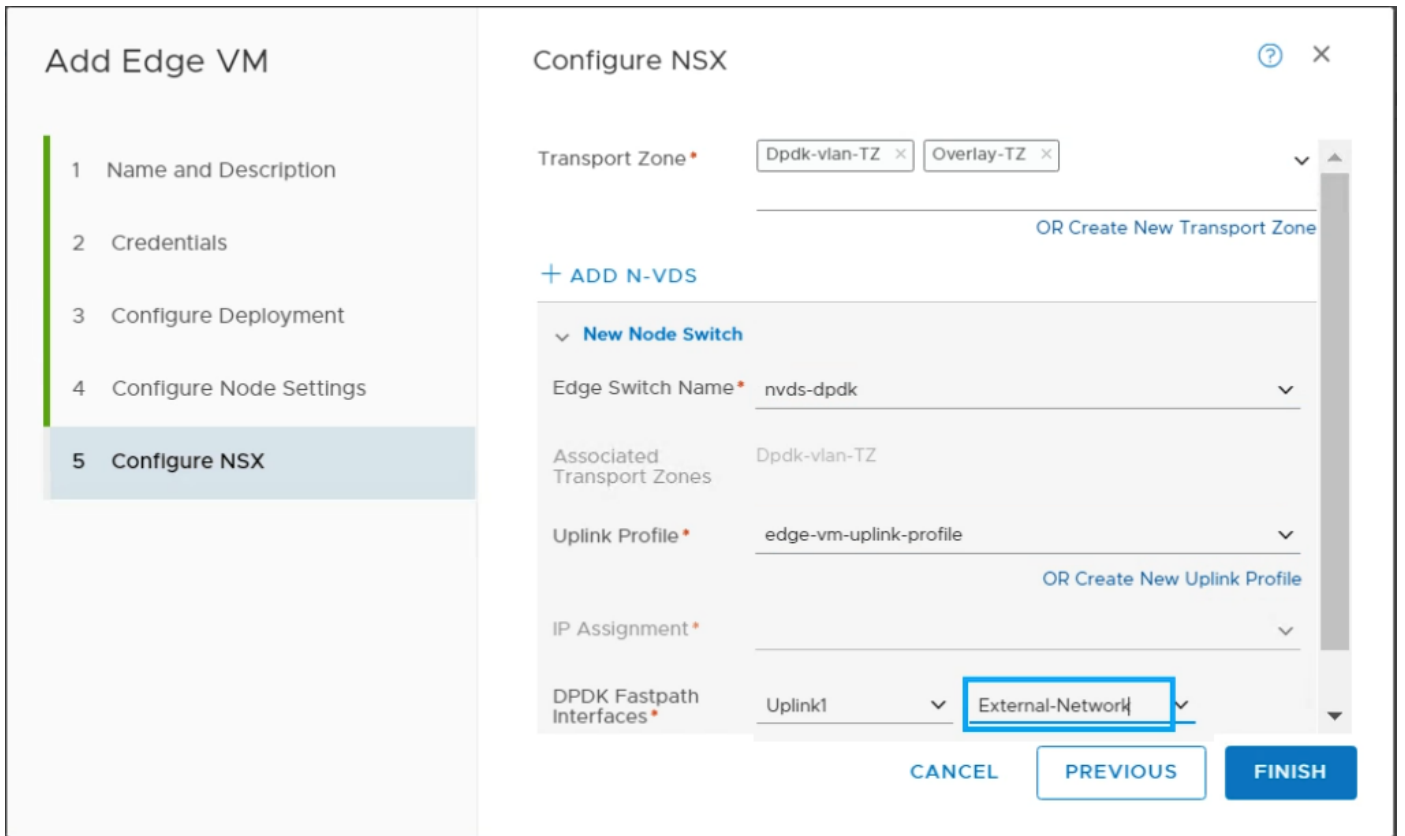


Figure 81. Configure NSX screen

14. Click **Finish**.
15. Repeat the steps in this section to deploy three additional edge VMs. The status of each edge node displays.

Host Transport Nodes **Edge Transport Nodes** Edge Clusters ESXi Bridge Clusters

+ ADD EDGE VM EDIT DELETE ACTIONS											View	All
Edge	ID	Deployment Type	Management IP	Host	Configuration Sta	Node Status	Transport Zones	NSX Version	N-VDS	Edge Cluster	Logical Routers	
edge01	5213...c126	Virtual Machi...	192.168.20.108		Success	Up	Dpdk-TZ Overlay-TZ	2.4.1.0.0.1371...	2		0	
edge02	f707...e5d1	Virtual Machi...	192.168.20.109		Success	Up	Dpdk-TZ Overlay-TZ	2.4.1.0.0.1371...	2		0	
edge03	647f...626b	Virtual Machi...	192.168.20.130		Success	Up	Vlan-TZ Overlay-TZ	2.4.1.0.0.1371...	2		0	
edge04	a199...26e2	Virtual Machi...	192.168.20.131		Success	Up	Overlay-TZ Vlan-TZ	VERSION_U...	2		0	

Figure 82. Edge VM listing

Create edge cluster

About this task

Create two edge cluster using two edge VMs in each cluster.

Table 29. Edge clusters and their participating VMs

Cluster name	Participating VM
NSX-edge-Cluster	edge01 and edge02
VCD-edge-Cluster	edge03 and edge04

Steps

1. From a web browser, use the administrator credentials to log in to the **NSX Manager** at <https://nsx-manager-fqdn>.
2. Click **System, Fabric**, and then **Nodes**.
3. From the **Edge Clusters** tab, click **Add (+)**.
4. On the **Add Edge Cluster** screen, enter the **Name**, **Description**, and select the **Edge Cluster Profile**.
5. From the **Member Type** drop-down list, select **Edge Node**.
6. Use the information from the [Edge clusters and their participating VMs table](#) to select the required edge nodes from the **Available** column and move them to the **Selected** column.
7. Click OK.

Add Edge Cluster ? ×

Name *

Description

Edge Cluster Profile × ▼

Transport Nodes

Member Type ▼

Available (2)

edge03

edge04

< BACK NEXT > 1 - 4 of 4 records

Selected (2)

edge02

edge01

Figure 83. Add Edge Cluster screen

8. Click **ADD** to create the edge cluster.
9. Repeat the steps in this section to create a second edge cluster.

Host Transport Nodes Edge Transport Nodes **Edge Clusters** ESXi Bridge Clusters

+ ADD EDIT DELETE ACTIONS ▼ Q Search

Edge Cluster	ID	Member Type	Cluster Profile	Edge Transport Nodes
<input type="checkbox"/> NSX-edge-Cluster	a41b...a4ef	Edge Node	nsx-default-edge-high-availability-pro...	2
<input type="checkbox"/> VCD-edge-Cluster	ad1e...cc7b	Edge Node	nsx-default-edge-high-availability-pro...	2

Figure 84. Cluster status screen

Create logical switches

Logical switches attach to single or multiple VMs in the network. The VMs connected to a logical switch can communicate with each other using the tunnels between hypervisors.

Prerequisites

- NSX-T Manager must be installed and configured
- A Transport zone must be configured
- Verify that fabric nodes are successfully connected to NSX-T Management Plane Agent (MPA) and NSX-T Local Control Plane (LCP)
- Verify that transport nodes are added to the transport zone
- Verify that the ESXi hosts are added to the NSX-T fabric and VMs are hosted on these ESXi
- Verify that your NSX-T Controller cluster is stable
- Verify that compatible Intel NIC drivers are available for DPDK

About this task

In this deployment three logical switches are created:

- One ENS Vlan-backed logical switch for External network connectivity using the VLAN ID of the External Network
- Two Standard overlay-backed logical switches using standard overlay transport zone
- One standard Vlan-backed logical switch for VCD by selecting a standard VLAN backed Transport Zone and using the VLAN ID of the management network


 **NOTE: QLogic drivers do not support the N-VDS Enhanced data path feature.**

Table 30. Uplink profile details

Logical switch name	Transport zone	VLAN
External_LS	dpdk-TZ	40
LS_1	overlay-TZ	-
LS_2	overlay-TZ	-
VCD_LS	vlan-TZ	20

Steps

1. From a web browser, use the administrator credentials to log in to the **NSX Manager** at <https://nsx-manager-fqdn>.
2. Click **Advanced Networking & Security, Networking**, and then **Switching**.
3. From the **Switches** tab, click **+ ADD**.
4. On the **General** tab, enter the name in the **Name**, **Transport zone**, and **VLAN ID** in their respective field using the information from the [Uplink profile details](#) table.
5. Click **ADD** to create a logical switch.

Add New Logical Switch ? X

General Switching Profiles

Name*

Description

Transport Zone* ▼

Uplink Teaming Policy Name* ▼

Admin Status Up

Replication Mode Hierarchical Two-Tier replication
 Head replication

VLAN*

VLAN Id or VLAN Trunk Spec is allowed.

Figure 85. General tab screen

6. Repeat the steps in this section to create more logical switches, as described in the [Uplink profile details](#) table.

Switches Ports Switching Profiles

+ ADD EDIT DELETE ACTIONS Search

Logical Switch ↑	ID	Admin Status	Logical Ports	Traffic Type	Config State	Transport Zone
External_LS	1cb7...1e2a	Up	0	VLAN : 40	Success	Dpdk-TZ
LS_1	b393...2f87	Up	0	Overlay : 67587	Success	Overlay-TZ
LS_2	c81c...21aa	Up	0	Overlay : 67588	Success	Overlay-TZ
VCD_LS	003b...848d	Up	0	VLAN : 20	Pending	Vlan-TZ

Figure 86. Logical switches

Create and configure tier-1 router

The tier-1 logical router must be connected to the Tier-0 logical router to get the northbound physical router access.

Prerequisites

- Verify that the logical switches are configured. See the [Create logical switches](#) section
- Verify that an NSX-T Edge cluster is deployed to perform network address translation (NAT) configuration See the [NSX-T Installation Guide](#)

About this task

This section provides the steps to create and configure tier-1 router.

Steps

- From a web browser, use the administrator credentials to log in to the **NSX Manager** at `https://nsx-manager-ip-address`.
- Click **Advanced Networking & Security, Networking**, and then **Routers**.
- From the **Routers** tab, click **+ ADD**, and then select **Tier-1 Router** from the drop-down list.
- On the **New Tier-1 Router** window, enter the **Name** as **NSX-Tier-1** and **Description**.
- From the **Edge Cluster** drop-down list select the **NSX-edge-Cluster**.
- For **Failover mode**, select **Non-Preemptive** and click **ADD**.

The screenshot shows the 'New Tier-1 Router' configuration interface. At the top, there's a title 'New Tier-1 Router' with a help icon and a close icon. Below the title, there are two tabs: 'Tier-1 Router' (selected) and 'Advanced'. The form contains the following fields and options:

- Name***: NSX-Tier-1
- Description**: An empty text area.
- Tier-0 Router**: NSX-Tier-0 (with a close and dropdown icon)
- Edge Cluster**: NSX-edge-Cluster (with a close and dropdown icon)
- StandBy Relocation**: A toggle switch set to 'Disable'.
- Failover Mode**: Two radio buttons, 'Preemptive' and 'Non-Preemptive' (selected).
- Edge Cluster Members**: An empty list area with a close and dropdown icon.

At the bottom right, there are two buttons: 'CANCEL' and 'ADD'. A mouse cursor is pointing at the 'ADD' button.

Figure 87. New Tier-1 router screen

The new tier-1 router is created.

Create router port on tier-1 router

About this task

Once the tier-1 logical router is created, you need to create a router port to connect the internal logical switch with Tier-1 Router.

Steps

1. Click the Tier-1 router, and then select **Router Ports** from the **Configuration** drop-down.
2. In the **Logical Router Ports** section, click **+ADD** to add **New Logical Router Ports**.
3. On the **New Router Port** window, enter the **Name** as **RP_1** and **Description** in the fields provided.
4. From the **Type** drop-down list select **Downlink**, and from the **Logical Switch** drop-down list select the **logical switch** as **LS_1**.

5. Select the **Attach new switch port** radio button.
6. In the **Subnets** section, enter the IP Address for the logical router port the set the **Prefix length** to **24**.
7. Click **ADD**.

New Router Port

Name*

Description

Type

URPF Mode Strict None

Logical Switch OR Create a New Switch

Logical Switch Port Attach to new switch port Attach to existing switch port

Switch Port Name

Subnets

+ ADD

<input checked="" type="checkbox"/> IP Address*	<input checked="" type="checkbox"/> Prefix Length*
<input checked="" type="checkbox"/> 172.16.50.254	<input type="text" value="24"/>

Figure 88. New Router Port for router tier-1 screen

The **New Logical router port RP_1** is added to the **Tier-1** router.

8. Repeat the above steps described in this section to create second router port , for example, RP_2, for the second LS_2 logical switch.

NSX-Tier-1 ×

Overview Configuration ▾ Routing ▾ Services ▾

Logical Router Ports

+ ADD EDIT DELETE ACTIONS ▾

Logical Router	ID	Type	IP Address/mask	Connected To	Transport Node	Relay Service	Statistics
<input type="checkbox"/>	RP_1	c33d...d...	Downlink	172.16.50.254/24	↔ LS_1 (25d7b769-5a3f-4...		
<input type="checkbox"/>	RP_2	8a66...e...	Downlink	172.16.70.254/24	↔ LS_2 (b665a4a4-2475-...		

Figure 89. Logical router ports

Configure route advertisement on tier-1 router

About this task

Configure the Route Advertisement on Tier-1 router.

Steps

1. From the Tier-1 router, click the **Routing** tab, then select **Route Advertisement** on the displayed list.
2. Click **EDIT**.
The **Edit Route Advertisement Configuration** window display.
3. Slide the **Status** slider to **Enabled**, click the **Advertise all the NSX-T Connected Routes, Advertise All Static Routes** to **Yes**, then click **SAVE**.

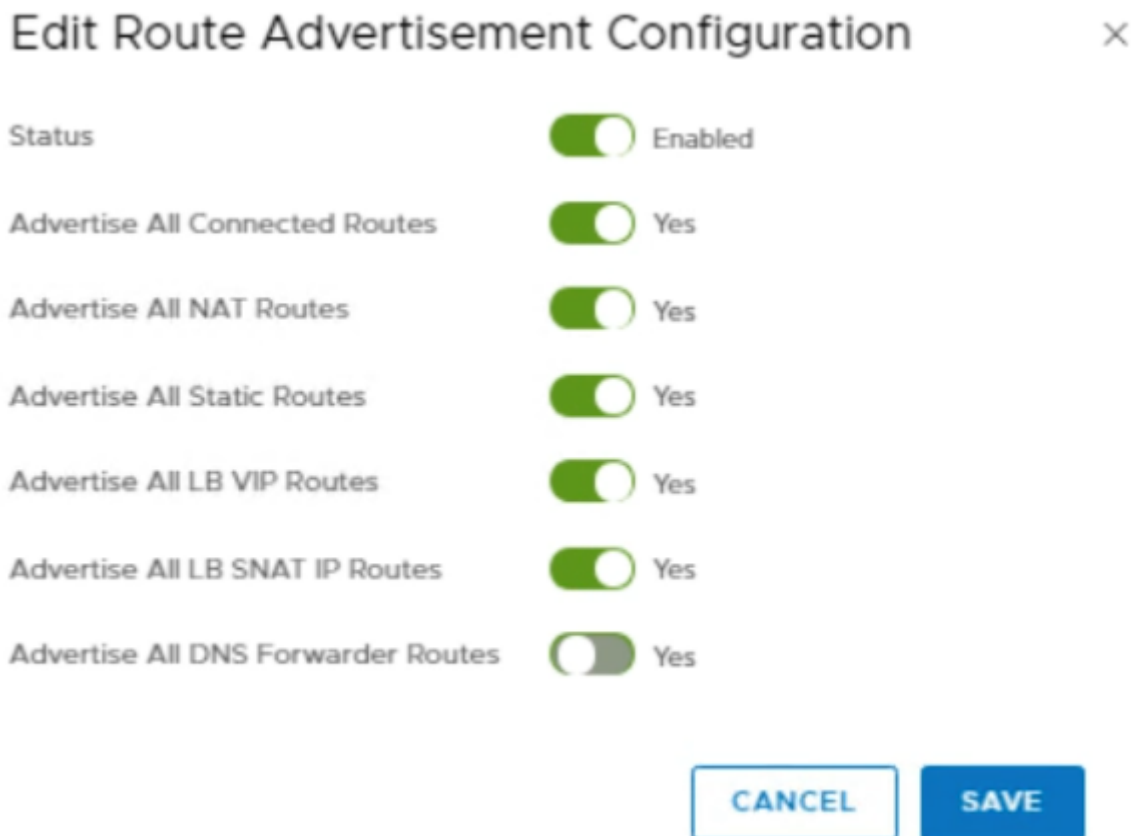


Figure 90. Edit Route Advertisement Configuration screen

Create and configure NSX-T tier 0 router

Tier-0 logical routers have downlink ports to connect to NSX-T tier-1 logical routers and uplink ports to connect to external networks.

Prerequisites

- Minimum of one NSX-T Edge is installed
- NSX-T Controller cluster is stable
- Edge cluster is configured

About this task

Create and configure NSX-T Tier 0 router.

Steps

1. From a web browser, use the administrator credentials to log in to the **NSX Manager** at <https://nsx-manager-ip-address>.
2. Click **Advanced Networking & Security, Networking**, and then **Routers**.
3. On the **Routers** tab, click **+ ADD** then select **Tier-0 Router** from the drop-down.
4. On the **New Tier-0 Router** window, enter the **Name** as **NSX-Tier-0**.
5. From the **Edge Cluster** drop-down list, select the **NSX_Edge_Cluster**.
6. For the **High Availability Mode**, select the **Active-Active** radio button, and click **ADD**.

The screenshot shows the 'New Tier-0 Router' configuration interface. At the top, there is a title 'New Tier-0 Router' and a close button. Below the title, there are two tabs: 'Tier-0 Router' (selected) and 'Advanced'. The form contains the following fields:

- Name:** NSX-Tier-0
- Description:** (Empty text box)
- Edge Cluster:** NSX-edge-Cluster
- High Availability Mode:** Active-Active (selected), Active-Standby (unselected)

At the bottom right, there are two buttons: 'CANCEL' and 'ADD'. There is also a link 'OR Create a New Edge Cluster' next to the Active-Standby radio button.

Figure 91. New Tier-0 Router screen

Connect Tier-1 router to NSX-T tier 0 router

About this task

Attach the NSX-T Tier-1 router to NSX-T Tier-0 router.

Steps

1. Go to Tier-1 Router, click the **Overview** tab, and from the **Tier-0-Connection** section click **CONNECT**.
2. On the **Connect to Tier-0 Router** screen, locate the **Tier-0-Router** drop-down list, select **Tier-0 router**, and click **CONNECT**.

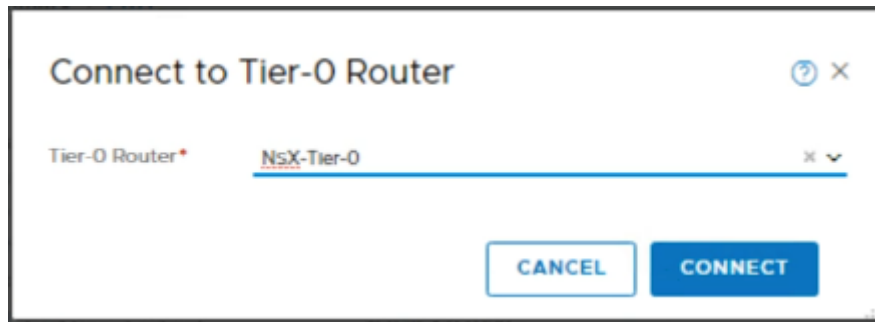


Figure 92. Connect to Tier-0 Router screen

NSX-Tier-1 and NSX-Tier-0 is connected.

Create logical router port on Tier-0 router

About this task

Once the Tier-0 logical router is created, you need to create a router port to connect the external logical switch with Tier-0 Router.

Steps

1. Click the **Tier-0** router, and then from the **Configuration** drop-down select **Router Ports**.
2. In the **Logical Router Ports** section, click **+ADD** to add **New Logical Router Ports**.
3. On the **New Router Port** window, enter the **Name** as **External-RP** and **Description** in the fields provided.
4. From the **Type** drop-down list, select **Uplink**.
5. From the **MTU** drop-down list, select **1600**.
6. From the **Transport Node** drop-down list, select the transport node.
7. From the **Logical Switch** drop-down list, select the logical switch **External_LS**.
8. Select the **Attach new switch port** radio button.
9. In the **Subnets** section, enter the **IP Address** for the logical router port the set the **Prefix length** to **24**.
10. Click **ADD**.

New Router Port

Name*

Description

Type MTU

Transport Node*

URPF Mode Strict None

Logical Switch OR Create a New Switch

Logical Switch Port Attach to new switch port Attach to existing switch port
Switch Port Name

Subnets

+ ADD

<input checked="" type="checkbox"/> IP Address*	Prefix Length*
<input checked="" type="checkbox"/> 172.16.60.10	24

Figure 93. Router Port screen

The **Logical Router Port** is created.

Redistribution on Tier-0 router

About this task

Configure the **Route Advertisement** on Tier-0 Router.

Steps

1. From the Tier-0 router, click the **Routing** tab, then select **Route Redistribution** on the **Displayed list**.
2. Click **EDIT**.
The **Edit Route Advertisement Configuration** window display.
3. Slide the **Status** slider to **Enable route redistribution configuration** and click **SAVE**.




Figure 94. Route Redistribution Configuration screen

Configure BGP on NSX-Tier-0 router

Prerequisites

- BGP should be configured on the Leaf switches.

About this task

BGP is used to exchange the network routing and reachability information between the multiple Autonomous Systems (AS) on the Internet.

Steps

1. From the Tier-0 router, click the **Routing** tab, then select **BGP** on the displayed list.
2. Click **EDIT**.
The **Edit BGP Configuration** window opens.
3. Slide the **Status** slider to **Enabled**.
4. Slide the ECMP slider to **Enabled**.
5. From the **Local AS** field, enter **65002**, and then click **SAVE**.

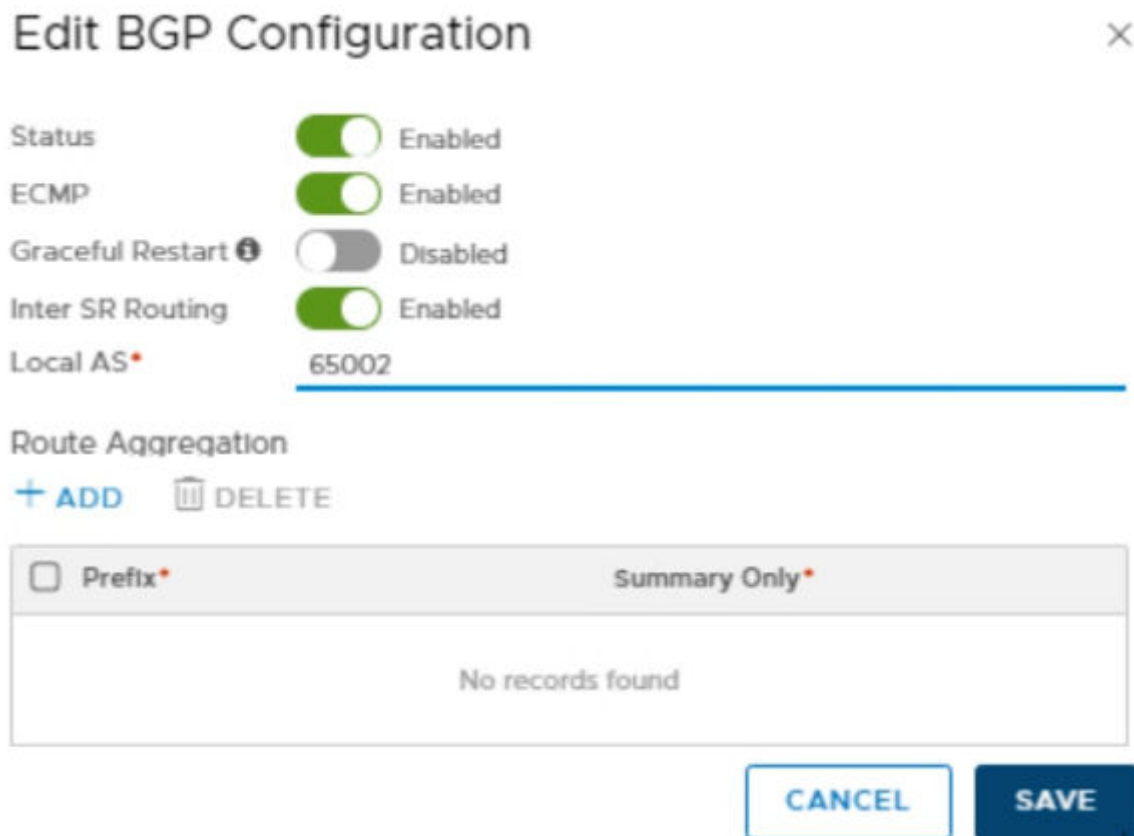


Figure 95. Edit BGP Configuration screen

Add neighbor to router NSX-Tier-0

About this task

To exchange the network routing and reachability information two BGP neighbors are created in this deployment. The [BGP neighbor details](#) table displays the required information to create neighbors for this deployment.

Table 31. BGP neighbor details

Field	Neighbor-1	Neighbor-2
Admin status	Enable	Enable
Remote AS	64502	64503
Maximum hop limit	2	2
Keep alive time (in seconds)	60	60
Hold down time (in seconds)	180	180

Steps

1. On the **BGP Configuration**, in the **Neighbors** section, click **+ ADD** to create **Neighbors**.
2. On the **New Neighbor** window, in the **Neighbor** tab:
 - a. In the **Neighbor Address** field, enter the **Leaf Router IP**.
 - b. Using the information in the [BGP neighbor details](#) table, enter the **Admin status**, **Remote AS**, **Maximum Hop Limit**, **Keep Alive Time (Seconds)**, and **Hold Down Time (Seconds)** information.

The screenshot shows the 'New Neighbor' configuration window with the following details:

- Neighbor Address:** 172.16.60.100
- Description:** (empty text box)
- Admin status:** Enabled (toggle switch)
- Maximum Hop Limit:** 2
- Remote AS:** 64502
- Keep Alive Time (Seconds):** 60
- Hold Down Time (Seconds):** 180
- Password:** (empty text box)

Buttons: CANCEL, ADD

Figure 96. New Neighbor screen

3. In the **Local Address** tab, from the **Type** drop-down list, select **Uplink**, then from the **Available** column move the uplink to **Selected** column.
4. In the **Address Families** tab, click **Add (+)** then in the **State** column click on **Edit** icon to change the state to **Enabled**.
5. Click **ADD** to create neighbor.
BGP is now configured on the NSX-Tier-0 router.
6. Repeat the steps in this section and use the information that is provided in the [BGP neighbor details](#) table to create second neighbor.

NSX-Tier-0 X

Overview Configuration **Routing** Services

BGP Configuration | EDIT

Status ● Enabled
 ECMP ● Enabled
 Graceful Restart ● Disabled
 Inter SR Routing ● Enabled
 Local AS 65002
 Route Aggregation 0

Neighbors

User System

+ ADD EDIT DELETE ACTIONS

IP Address	Local Address	ID	Admin status	Maximum Hop	Remote AS	Address Family	BFD	Keep Alive	Hold Down
<input type="checkbox"/> 172.16.60.3	172.16.60.10	22b2...5d20	● Enabled	2	64503	1	Disabled	60	180
<input type="checkbox"/> 172.16.60.2	172.16.60.10	5999...d662	● Enabled	2	65402	1	Disabled	60	180

Figure 97. Neighbors section

Create and configure VCD-Tier1 router

Prerequisites

- Minimum of one NSX-T Edge is installed
- NSX-T Controller cluster is stable
- Edge cluster is configured

About this task

The VCD Tier-1 logical router is a stand-alone router, and it does not have any downlink or connection with Tier-0 router. It has a service router but no distributed router. The VCD Tier-1 logical router has a centralized service port (CSP) to connect with a Load Balancer.

Steps

1. From a web browser, use the administrator credentials to log in to the **NSX Manager** at <https://nsx-manager-fqdn>.
2. Click **Advanced Networking & Security, Networking**, and then **Routers**.
3. On the **Routers** tab, click **+ ADD** then select **Tier-1 Router** from the drop-down.
4. From the **New Tier-1 Router** screen, enter the **Name** and **Description**.
5. From the **Edge Cluster** drop-down list, select **VCD_Edge_Cluster**.
6. For **Failover Mode**, select **Non-Preemptive** and click **ADD**.

The screenshot shows the 'New Tier-1 Router' configuration interface. At the top, there's a title 'New Tier-1 Router' with a help icon and a close icon. Below the title, there are two tabs: 'Tier-1 Router' (selected) and 'Advanced'. The form contains several fields: 'Name' with the value 'VCD-Tier-1', an empty 'Description' text area, 'Tier-0 Router' (empty), 'Edge Cluster' with the value 'VCD-edge-Cluster', 'StandBy Relocation' set to 'Disable' (toggle off), 'Failover Mode' with 'Non-Preemptive' selected (radio button), and 'Edge Cluster Members' (empty). At the bottom right, there are two buttons: 'CANCEL' and 'ADD'.

Figure 98. Tier-0 Router screen

The VCD Tier 1 Router is created.

Create logical router port on VCD tier-1 router

About this task

Create a logical router port on VCD-Tier-1 router and connect with logical switch.

Steps

1. Click the **VCD Tier-1** router, and then from the **Configuration** drop-down, select **Router Ports**.
2. In the **Logical Router Ports** section, click **+ADD** to add **New Logical Router Ports**.
3. On the **New Router Port** screen, enter the **Name** as **VCD-RP** and **Description** in the fields provided.
4. From the **Type** drop-down list select **Centralized**, and then from the **Logical Switch** drop-down list, select **VCD-LS**.

5. Select the **Attach new switch port** radio button, and in the **IP Address/mask** field enter the **IP Address or mask**, and then click **ADD**.

New Router Port

Name* VCD-RP

Description

Type Centralized MTU 1600

URPF Mode Strict None

Logical Switch VCD-LS

Logical Switch Port Attach to new switch port Attach to existing switch port

Switch Port Name

Subnets

+ ADD DELETE

IP Address*	Prefix Length*
<input checked="" type="checkbox"/> 192.168.20.252	<input checked="" type="checkbox"/> 24

CANCEL ADD

Figure 99. New Router Port screen

The Logical Router Port is created.

Configure Route Advertisement on VCD Tier-1 router

About this task

Configure Route Advertisement on VCD-Tier-1 Router.

Steps

1. From the VCD-Tier-1 router, click the **Routing** tab, then select **Route Advertisement**.
2. Click **EDIT**.
The **Edit Route Advertisement Configuration** window display.
3. Slide the **Status** slider to **Enabled**, click the **Advertise all the NSX-T Connected Routes**, slide the **Advertise All Static Routes** switch to **Yes**, then click **SAVE**.

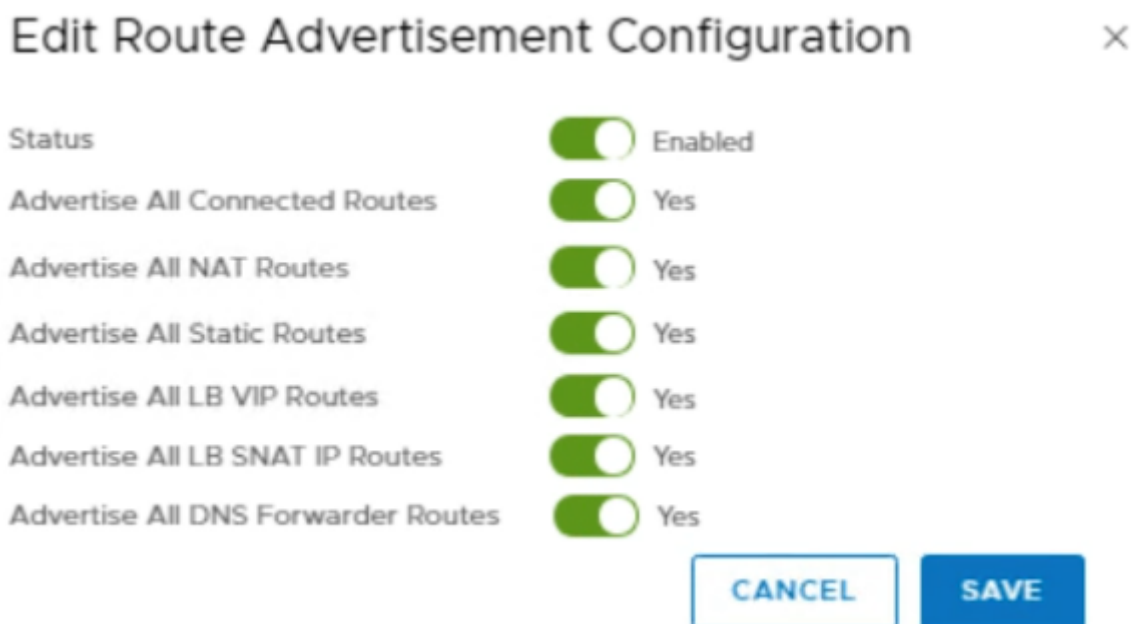


Figure 100. Route Advertisement Configuration screen

Create and configure the Load Balancer

Prerequisites

- Verify that **vCD Cell1**, **vCD Cell 2**, and **vCD Cell 3** are up and running. See *Installation and configuration of vCloud Director* to configure it.

About this task

The NSX-T logical load balancer provides the high-availability services and distributes the network traffic load between the servers. Only the Tier-1 router supports the NSX-T load balancer. One load balancer can be linked with only a Tier-1 logical router.

Steps


1. From a web browser, use the administrator credentials to log in to the **NSX Manager** at <https://nsx-manager-ip-address>.
2. Click **Advanced Networking & Security, Networking**, and then **Load Balancing**.
3. On the **Load Balancer** tab, click **+ Add**.
4. On the **Add Load Balancer** screen, enter the load balancer name as **VCD_LB** and provide a description.
5. Select the **Load balancer virtual server size** and click **OK** to create load balancer.

Add Load Balancer ? X

Name * VCD-LB

Description

Load Balancer Size *

Select from one of the three available choices of size for the Load Balancer 

SMALL	MEDIUM	LARGE
<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
Virtual Servers: 20 Pool Member: 300	Virtual Servers: 100 Pool Member: 2000	Virtual Servers: 1000 Pool Member: 7500
CPU: 2 Memory: 4GB	CPU: 4 Memory: 8GB	CPU: 12 Memory: 16GB

Error Log Level * INFO

Figure 101. Add Load Balancer screen

Attach load balancer with VCD-Tier1 router

About this task

Once the NSX-T load balancer is created, it is required to link it with the VCD-Tier1 router to use the high-availability services.

Steps

1. Click **Advanced Networking & Security, Networking**, and then **Load Balancing**.
2. From the **Load-Balancing** tab, select the load balancer that you created in **Create and configure the Load Balancer** section.
3. From the **Actions** drop-down list select **Attach to a Logical Router** option.
4. Select **VCD-Tier1** that you created in the *Create and configure VCD-Tier1 router* section and click **OK**.

Attach to a Logical Router



Select the Router to which the Load Balancer VCD_LB is to be attached. Only Tier-1 Routers in 'Active Standby' are currently supported. Note: The Load Balancer can only be Enabled if it had a Virtual Server associated with it.

Tier-1 Logical Router *

VCD-Tier1

CANCEL

OK

Figure 102. Attach to a Virtual Server screen

Create a health monitor for load balancer

About this task

Once the load balancer is created, create a health monitor to test whether a server is available. This health monitor does the different tests to monitor servers health. The health monitor starts health checks once:

- The server pool is added to the load balancer
- Load balancer is linked to VCD tier-1 router

Steps

1. Click **Advanced Networking & Security, Networking**, and then **Load Balancing**.
2. On the **Monitors** tab, click **+ Add** to create a new active health monitor.
3. On the **Monitor Properties** screen, enter the name and provide as **TCP_VCD** a description for health monitor.
4. From the **Health Check Protocol** drop-down list, select **LbTcpMonitor** option.
5. In the **Monitoring Port** field, enter **443** as **Port number**.
6. Keep the default options for **Monitoring interval (sec)**, **Fall Count**, **Rise Count**, and **Timeout Period (sec)**, then click **Next**.

Field	Value
Name *	TCP_VCD
Description	
Health Check Protocol *	LbTcpMonitor
Monitoring Port	443
Monitoring Interval (sec) *	5
Fall Count *	3
Rise Count *	3
Timeout Period (sec) *	15

Figure 103. Monitor properties

7. Review the health check configuration settings and click **Finish**.

The Active health monitor is created successfully.

Add a server pool for load balancing

About this task

Server pool is made of multiple servers that are configured and running on the same environment.

Steps

1. Click **Advanced Networking & Security, Networking**, and then **Load Balancing**.
2. From the **Server Pools** tab, click **+ Add** to create a server pool.
The **General properties** screen displays.
3. Enter a name as **VCD_IP** and a description for the load balancer pool.
4. From the **Load Balancing Algorithm** drop-down, select **ROUND_ROBIN** for the **Server pool**.
5. Keep the **default** option for **TCP Multiplexing** and **Maximum Multiplexing Connections** then click **Next**.

The screenshot shows the 'Add New Server Pool' configuration window. On the left is a sidebar with four steps: 1 General Properties (selected), 2 SNAT Translation, 3 Pool Members, and 4 Health Monitors. The main area is titled 'General Properties' and contains the following fields:

- Name**: VCD_IP
- Description**: (empty text box)
- Load Balancing Algorithm**: ROUND_ROBIN (dropdown menu)
- Advanced Properties** (expanded):
 - TCP Multiplexing**: Disabled (toggle switch)
 - Maximum Multiplexing Connections**: 6

At the bottom right, there are two buttons: 'CANCEL' and 'NEXT'.

Figure 104. General Properties screen

6. From the **SNAT Translation** window, set the **Translation Mode** to **Auto Map**, and click **Next**.
7. On the **Pool Members** window, select the **Membership Type** to **Static**.
8. In the **Static Membership** section, click **+ ADD** and add three pool members:
 - a. In the **Name** column, enter the pool member name.
 - b. In the **IP** column, enter the IP address of VCD Cell 1.
 - c. In the **State** column, select **Enabled** from the drop-down list.
 - d. Click **+ Add** in the **Static Membership** section and create the remaining two vCD Cells as pool members.
 - e. Click **Next**.

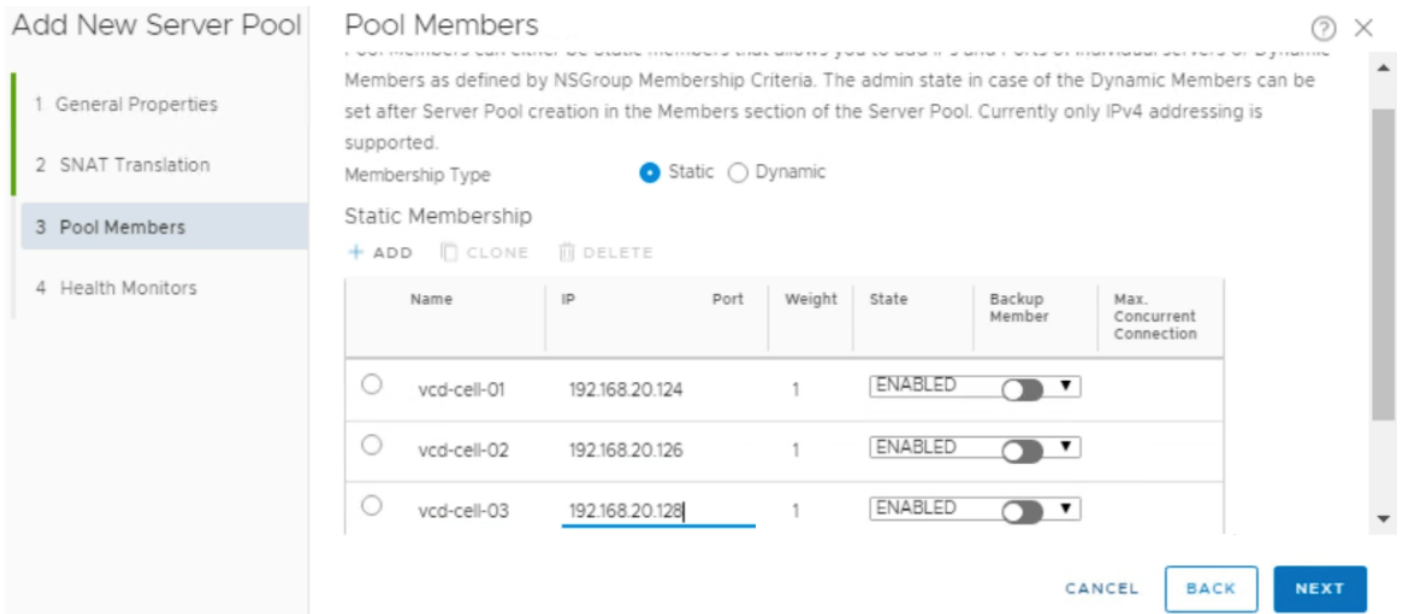


Figure 105. Pool Members screen

9. On the **Health Monitors** screen, in the **Enter the Minimum Active Members** field, enter the number of active health monitors. In this deployment.

NOTE: In this deployment, the 1 active health monitor is used.

10. Select the **Active Health Monitor** that you have created in [Create a health monitor for load balancer](#) section.

11. Click **Finish** to add server pool.

Create a virtual server

About this task

Virtual servers receive all the client connections and distribute them among the servers. A virtual server has an IP address, a port, and a protocol.

Steps

1. Click **Advanced Networking & Security, Networking**, and then **Load Balancing**.
2. From the **Virtual Servers** tab, click **+ ADD** to add a new virtual server. The **General Properties** screen displays.
3. Perform the following selections:
 - a. In the **Name** field, enter **VCD_IP**.
 - b. Enter a brief description in the **Description** box.
 - c. From the **Load Balancer Application Profile** section, select the **Layer 4** radio button.
 - d. From the **Application Profile** drop-down list, select the **nsx-default-lb-fast-tcp-profile** option.
 - e. Use the slider to set the **Access Log** option to **Enabled**, and then click **Next**.

Figure 106. General Properties screen

4. On the **Virtual Server Identifiers** screen:
 - a. In the **IP Address** field, enter the VCD-Tier-1 router centralized port IP address.
 - b. In the **Port** field, enter **443, 80**.
 - c. Keep the default values for **Protocol**, and click **Next**.

Figure 107. Virtual Server Identifiers screen

5. From the **Server Pool** screen, select the server pool that you have created in the [Add a server pool for load balancing](#) section, then click **Next**.
6. On the **Load Balancing Profiles** screen, select the **Source IP** to **nsx-default-source-ip-persistence-profile**, and then click **Finish**.

Attach the virtual server to the load balancer

About this task

The virtual server receives the client traffic and then distributes it between the servers. Attach the virtual server to load balancer to enable the high-availability services to distribute the network traffic load between the servers.

Steps

1. Click **Advanced Networking & Security, Networking**, and **Load Balancing**.
2. From the **Load Balancers** tab, select the load balancer that you created in [Create and configure the Load Balancer](#) section.
3. From the **Actions** drop-down list, select the **Attach to a Virtual Server** option.

4. Select virtual server that you have created in [Create a virtual server](#) section, and click **OK**.

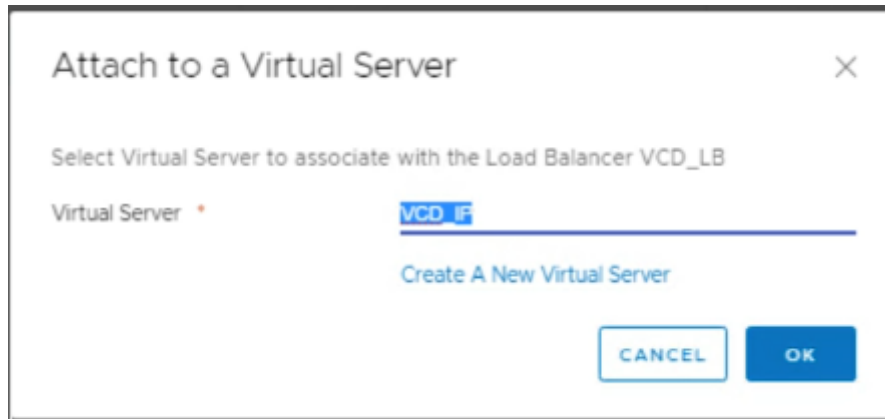


Figure 108. Attach to a Virtual Server

Configure vCloud Director

VMware vCloud Director (vCD) is a VIM component and works on top of other VIM components, vCenter Server, and NSX-T Manager. The vCloud Director is deployed on the Management pod. vCloud Director connects with:

- vCenter Server to manage the workloads
- NSX Manager associated with tenant networking

The vCloud Director server is grouped by deploying three vCD instances to create vCD cells: one vCD cell is used as Primary cell, and remaining two are used as Standby cells. These cells are attached with the NSX-T load balancer for high availability. An NFS server instance is created to provide the temporary storage for upload or download the catalog items that are published externally.

Installation of NFS server

Prerequisites

- A virtual machine running CentOS7 with following configuration:
 - 8GB Memory
 - Disk space: 1 TB
 - vCPU:1
 - vNIC:1
- Management pod should be configured and it should have internet connectivity
- DNS entries must be added in the DNS server for all the vCD cells

About this task

You must deploy and configure an NFS server accessible to all the servers of vCD server group.

Steps

1. Log in as a root user into Linux CentOS virtual machine and open the terminal.
2. Run the following command to install package nfs-utils:
`yum install nfs-utils`
3. Run the following command to start NFS-related services:
`systemctl start nfs-server`
4. Run the following command and make an export directory:
`mkdir /opt/vcd-share`
5. Run the following command to restart the NFS server:
`systemctl restart nfs-server`
6. Append the following line within the `/etc/exports` file:
`/opt/vcd-share *(rw, sync, no_root_squash)`
7. Stop the firewall then turn it off using the following commands:

```
systemctl stop firewalld
chkconfig firewalld off
```



NOTE: Verify that the NFS server is correctly configured by running the following command:

```
# showmount -e <NFS_IP>
Output: Export list for <NFS_IP>:
/opt/vcd-share*
```

Installation and configuration of vCloud Director

The vCloud Director servers consist of one or more vCD cells. These vCD cells are created by deploying vCloud Director Appliances. The process to install and configure the vCloud Director creates vCD cells. Each server in the group runs number of services that the vCloud Director Cell calls. These cells have a common database and connect with vCenter Server, ESXi hosts, and NSX-T Manager.

In this deployment, three vCD Cell is deployed, such as **VCD-Cell01**, **VCD-Cell02**, and **VCD-Cell03**. **VCD-Cell01** are used as the primary cell and **VCD-Cell02** and **VCD-Cell03** are used as the stand by cells. Deployment size for both the primary and standby cells must be the same. For example, you can use one primary-small and two standby-small cells, or one primary-large and two standby-large cells for HA cluster. For this deployment, one primary-large and two standby-large cells are used.

Prerequisites:

- vCenter Server must be up and running
- AD-DNS, and NTP server should be up and running
- DNS entries must be added in the DNS server for all the vCD cells
- DRS Automation option on the vCenter Server cluster that is used for vCD deployment must be set to Fully Automated

 **NOTE:** See the [Enable vSphere DRS](#) section for information about setting the DRS automation.

Deployment and configuration of vCD Cell 01

About this task


Deploy the vCloud Director Cell 01.

Steps

1. Using the VMware vSphere Web Client, log in to the Management vCenter.
2. Right-click the **Management Datacenter**, and then click **Deploy OVF Template**. The **Select template** window opens.
3. Enter the download URL or click **Browse** to locate the .OVA file on your computer, then click **Next**. The **Select name and location** screen displays.
4. In the field provided, enter the **Name**, select the **Location**, and then click **Next**. The **Select a resource** screen displays.
5. Select the **ESXi** to deploy **vCD cell 01** and click **Next**.
6. From the **Review details** screen, review the settings that are selected then click **Next**.
7. Use the scroll bar to review the information on the **Accept license agreement** screen and if you agree, click **Accept** and click **Next**.
8. On the **Select configuration** screen, select the type of deployment configuration from the **Configuration** drop-down list and click **Next**.

 **NOTE:** For this deployment, **Primary large configuration is used**.

9. On the **Select storage** screen:
 - a. Locate the **Select virtual disk format** drop-down list, and select **Thin provision**.
 - b. From the **VM storage policy** drop-down list, select **vSAN Default Storage Policy**.
 - c. Select the **vSAN datastore** and click **Next**.
10. On the **Select networks** screen, select the appropriate networks, then click **Next**.
11. On the **Customize template** screen, locate the **VCD Appliance Settings** section and complete the following fields:
 - a. **NTP Server:** Enter the NTP server IP address.
 - b. **Initial root password:** Set the root password.
 - c. **Expire Root Password upon First Login:** Clear the checkbox to disable the password expiration on first root login.
 - d. **Enable SSH service in the appliance:** Check the box to enable SSH services in the appliances.
 - e. **NFS mount for transfer file location:** Enter the NFS Server Share folder path.

 **NOTE:** This is the export directory path that you have created in the [Installation of NFS server](#). This shared folder path must be in the following format: `<NFS-Server-IP>:/<Share Folder Path>`

For example: `192.168.20.122:/opt/vcd-share`

12. On the **Customize template** screen, locate the **VCD Configure - Required only for primary appliances** section, and fill the following fields:
 - a. **vCloud DB password for the vCloud user:** Set the password for vCloud Database user.
 - b. **Admin User Name:** Enter the username for the system administrator or use the default name.
 - c. **Admin Full Name:** Enter the full name of vCD system administrator.
 - d. **Admin user password:** Set the system administrator user password.
 - e. **Admin email:** Enter the email ID of administrator user.
 - f. **System name:** Enter the system name or use keep the default name.
 - g. **Installation ID:** Enter the installation ID for vCD cell 01, or use keep the default ID.
13. On the **Customize template** screen, locate the **Networking Properties** section, and enter the following fields:
 - a. **Default Gateway:** Enter the IP address of default gateway for vCD Cell 01.
 - b. **Domain Name:** Enter the domain name for vCD Cell 01.
 - c. **Domain Search Path:** Enter the domain search path for vCD Cell 01.
 - d. **Domain Name Servers:** Enter the DNS IP address.
 - e. **eth0 Network IP Address:** Enter the IP address for eth0 network interface.
 - f. **eth0 Network Netmask:** Enter the netmask IP for eth0 network interface.
 - g. **eth1 Network IP Address:** Enter the IP address for eth1 network interface.
 - h. **eth1 Network Netmask:** Enter the netmask IP for eth1 network interface.

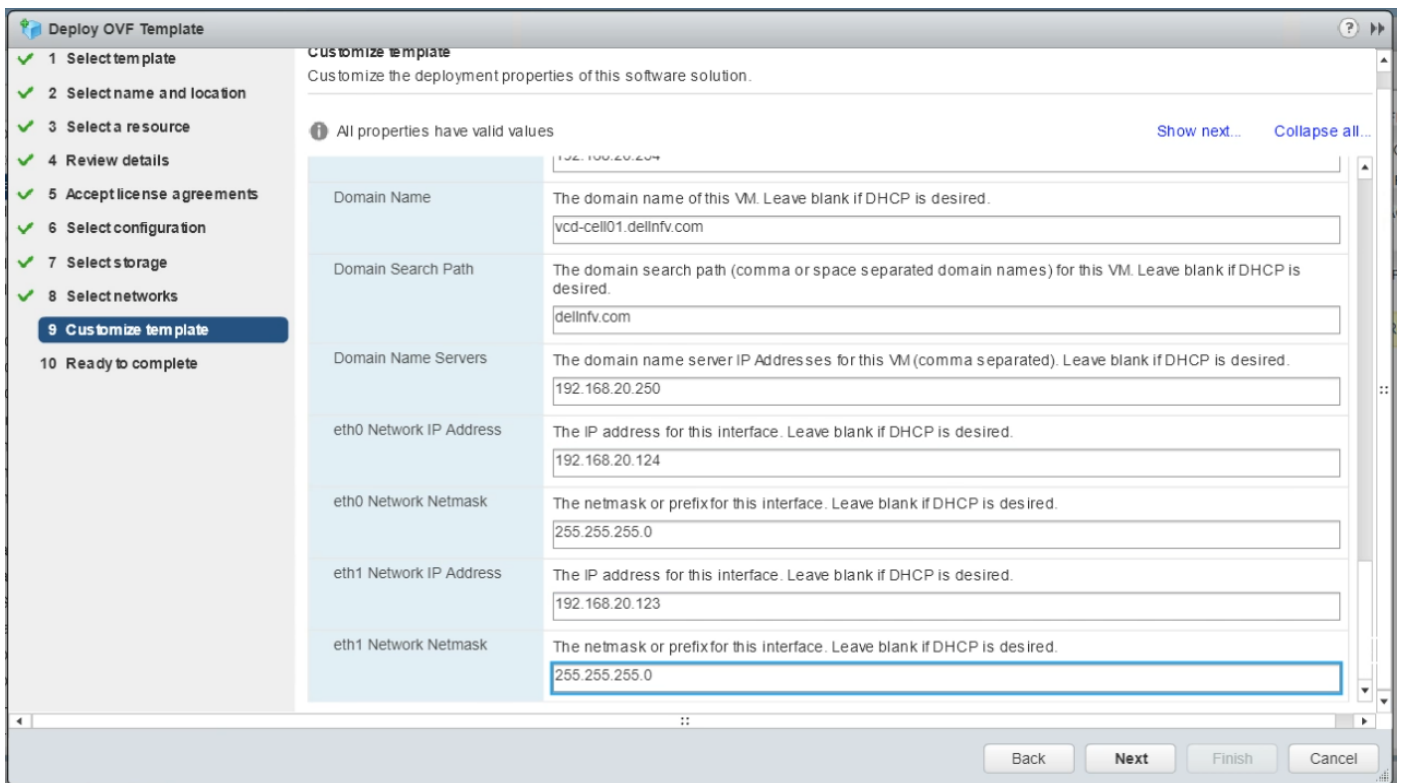


Figure 109. Network Properties screen

14. Click **Next**.
15. From the **Ready to complete** screen, review the provided configuration details, and then click **Finish** to deploy **vCD Cell 01**.

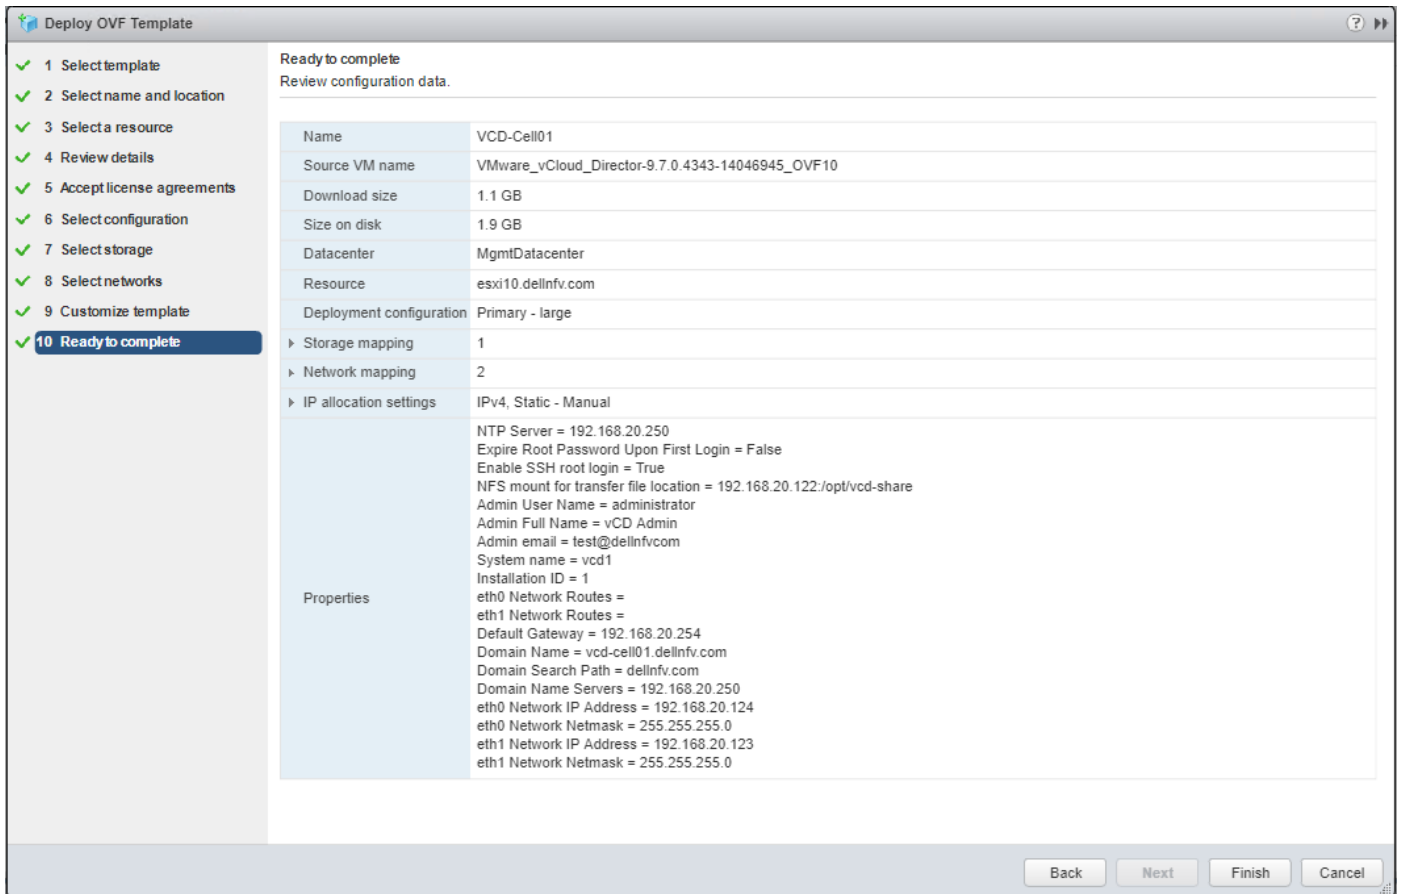


Figure 110. Ready to Complete screen

Assign license to vCD Cell 01

About this task

Assign license to vCD Cell 01.

Steps

1. From a web browser, log in to **vCD cell 01** at <https://<<vCD-Cell-01-fqdn>>/cloud>
2. Click **Administrator, System Settings, License**.
3. On the **License** page, enter the vCD license key in the **Serial Number** field.
4. Click **Apply** to save the license key.



Deployment of vCD Cell 02

About this task

Deploy vCD Cell 02 using the steps provided in this section.

Steps

1. Use the VMware vSphere Web Client to log in to the Management vCenter.
2. Right-click the **Management Cluster**, and then click **Deploy OVF Template**.
3. From the **Select template** screen, enter the download URL or click **Browse** to locate the .OVA file on your computer, and then click **Next**.
The **Select name and location** screen displays.
4. Enter the **Name**, **Location**, and then click **Next**.

5. From the **Select a resource** screen, select the ESXi to deploy **vCD cell 02** and click **Next**. The **Review details** screen displays.
6. Review the settings that are selected then click **Next**.
7. From the **License Agreement** screen, review the license agreement terms and if you accept then check the **I accept all license agreements** box and click **Next**.
8. On the **Select configuration** window, select the type of deployment configuration from the **Configuration** drop-down list and click **Next**.
 -  **NOTE:** For this deployment, **Standby large configuration for vCD-cell02** is used.
9. On the **Select storage** screen:
 - a. From the **Select virtual disk format** drop-down list, select **Thin provision**.
 - b. From the **VM storage policy** drop-down list, select **vSAN Default Storage Policy**.
 - c. Select the **vSAN datastore** and click **Next**.
10. From the **Select networks** screen, select the appropriate networks, then click **Next**.
 -  **NOTE:** For this deployment, **VM-Mgmt-Network** is used.
11. On the **Customize template** screen, locate the **VCD Appliance Settings** section and complete the following fields:
 - a. **NTP Server:** Enter the NTP server IP address.
 - b. **Initial root password:** Set the root password.
 - c. **Expire Root Password upon First Login:** Clear the check box to disable the password expiration on first root login.
 - d. **Enable SSH root login:** Check the box to enable SSH services in the appliances.
 - e. **NFS mount for transfer file location:** Enter the NFS Server Share folder path. This is the export directory path that you have created in Step 4 of [Installation of NFS server](#). This share folder path must be in the following format: <NFS-Server-IP>:/<Share Folder Path>
 For example: 192.168.20.122:/opt/vcd-share
12. On the **Customize template** screen, locate the **Networking Properties** section, and complete the following fields:
 - a. **Default Gateway:** Enter the IP address of default gateway for vCD Cell 02.
 - b. **Domain Name:** Enter the domain name for vCD Cell 02.
 - c. **Domain Search Path:** Enter the domain search path for vCD Cell 02.
 - d. **Domain Name Servers:** Enter the DNS server IP address.
 - e. **eth0 Network IP Address:** Enter the IP address for eth0 network interface.
 - f. **eth0 Network Netmask:** Enter the netmask IP for eth0 network interface.
 - g. **eth1 Network IP Address:** Enter the IP address for eth1 network interface.
 - h. **eth1 Network Netmask:** Enter the netmask IP for eth1 network interface.
13. Click **Next**.
14. On the **Ready to complete** screen, review the provided configuration details then click **Finish**.
15. Repeat the above steps to deploy **vcd-cell-03**.

vCD integration with vCenter

About this task

Integrate the VMware vCenter Server with vCD to use vCenter resources with vCD.

To integrate vCD with vCenter and NSX-T, perform the following steps:

Steps

1. From a web browser, log in to **vCD cell 1** at `https://<<vCD-Cell-01-fqdn>>/provider`.
2. Click the **Main menu** icon then select **vSphere Resources** from the list.
3. From the left navigation panel, click **vCenters**, and then click the **Add-on vCenters** page.
4. On the **vCenter** screen, enter the following information:
 - a. **Name:** Enter the resource vCenter name.
 - b. **Description:** Enter a brief description.
 - c. **URL:** Enter the resource vCenter URL/FQDN.
 - d. **User name:** Enter the username of resource vCenter.

- e. **Password:** Enter the password of entered user.
- f. Select the **vSphere Web Client URL** radio button, and then enter the resource vCenter server URL/FQDN.

Figure 111. Name this vCenter screen

- 5. On the **Connect to NSX Manager** screen, move the **Configure Settings** switch to disable it, then click **Next**.
- 6. On the **Ready to Complete** screen, review the provided information and click **Finish**.
The Resource vCenter is connected with the vCD.

vCD integration with NSX-T

About this task

Integrate the VMware NSX-T with vCD to use its resources with vCD.

Follow the below steps to integrate vCD with vCenter and NSX-T.

Steps

1. From a web browser, log in to **vCD cell 1** at `https://<<vCD-Cell-01-fqdn>>/provider`.
2. Click the **Main menu** icon then select **vSphere Resources** listing.
3. From the left navigation panel, click **NSX-T Managers**, then click the **Add-on NSX-T Managers** page.
4. On the **Register NSX-T Manager** screen, enter the following information in the fields provided:
 - a. **Name:** NSX-T Manager name
 - b. **Description:** Brief description
 - c. **URL:** URL/FQDN of NSX-T Manager
 - d. **User name:** User name of NSX-T Manager
 - e. **Password:** Assigned password of the username entered
5. Click **Save** to register NSX-T Manager.

Figure 112. Register NSX-T Manager screen

Creating a session token for vCD

Generate a vCD session token to integrate vCD with vCenter Server and NSX-T manager.

Prerequisites

NOTE: For information about generating a vCD session token to integrate vCD with vCenter Server and NSX-T manager, see the [VMware API Reference Guide](#).

- Download and install Postman on the deployment VM. For more information, see [Postman Documentation](#).
- Open the Postman application, go to **Settings** and turn-off the **SSL Certificate Verification**.

About this task

Generate the session token to run APIs.

Steps

1. On the deployment VM, open the Postman application.
2. POST a request to the vCD login URL and enter the vCD administrator credentials into the Authorization header of the request.

```
url = https://<FQDN>/api/sessions
Method = POST
Authorization
Type- Basic Auth
HEADERS
Key          VALUE
Accept      application/*+xml;version=31.0;
```

NOTE: The values provided in the above example only for reference purposes, update the values as per your requirement.

3. Update the value for the parameter above using the following table:

Table 32. Parameter description

Parameter	Description
FQDN	Enter the FQDN for vCloud Director

The **Response status 200 OK** message means that the session code is generated successfully.

4. In the headers section of the response, note the value of `x-vcloud-authorization` field. This value is used as a session token in all other API calls.

Retrieve VIM server details

About this task

Post a GET request on vCD to retrieve the VIM server details.

Steps

1. On the deployment VM, open the **Postman** application.
2. Paste the following parameters in the **Postman Headers**:

```
url: https://<FQDN>/api/admin/extension/vimServerReferences
Method: GET
Header:
x-vcloud-authorization: Use the value fetched from the session API.
Accept application/*+xml;version=31.0;
```

3. Update the values for the parameters above using the following table:

Table 33. Parameter description

Parameter	Description
FQDN	Enter the FQDN for vCloud Director
x-vcloud-authorization	Enter the session ID received from the Creating a session token for vCD section

The **Response status 200 OK** message displays.

4. From the received response, make note of the **href**, **name**, and **ID** of the **vCenter Server** as shown in the following example:

 **NOTE: This information is required when creating the provider VDC.**

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<vmext:VMWVimServerReferences xmlns="http://www.vmware.com/vcloud/v1.5"
xmlns:vmext="http://www.vmware.com/vcloud/extension/v1.5" xmlns:ovf="http://
schemas.dmtf.org/ovf/envelope/1" xmlns:vssd="http://schemas.dmtf.org/wbem/wscim/1/cim-
schema/2/CIM_VirtualSystemSettingData" xmlns:common="http://schemas.dmtf.org/wbem/wscim/1/
common" xmlns:rasd="http://schemas.dmtf.org/wbem/wscim/1/cim-schema/2/
CIM_ResourceAllocationSettingData" xmlns:vmw="http://www.vmware.com/schema/ovf"
xmlns:ovfenv="http://schemas.dmtf.org/ovf/environment/1" xmlns:ns9="http://www.vmware.com/
vcloud/versions" type="application/vnd.vmware.admin.vmwVimServerReferences+xml">
  <Link rel="up" href="https://192.168.20.124/api/admin/extension" type="application/
vnd.vmware.admin.vmwExtension+xml"/>
  <vmext:VimServerReference href="https://192.168.20.124/api/admin/extension/vimServer/
3cd1ac67-88de-4c4b-8e1a-d171e322d8d1" id="urn:vcloud:vimserver:3cd1ac67-88de-4c4b-8e1a-d171e322d8d1"
name="ResVCSA" type="application/vnd.vmware.admin.vmwvirtualcenter+xml"/>
</vmext:VMWVimServerReferences>
```

Update VIM server

About this task

Create a PUT request on the postman to update VIM server. For more information, see the [VMware API Reference Guide](#).

Steps

1. From the deployment server, open the **Postman** application.

- Paste the following parameters in the **Postman Headers**:

```
url = https://<FQDN>/api/admin/extension/vimServer/{ID}
Method = PUT
Authorization
Type- Basic Auth
HEADERS
Key          VALUES
Accept      application/*+xml;version=31.0;
Content-Type application/vnd.vmware.admin.vmwvirtualcenter+xml;version=31.0;
Content-Length 596
x-vcloud-authorization Use the value fetched from the sessions api
```

Update the values for the table above using the parameters in the following table:

Table 34. Uplink profile details

Parameter	Description
FQDN	Enter the FQDN for vCloud Director
ID	Enter the VIM Server ID received from the response of Retrieve VIM Server Details
Content-Length	Enter the total number of characters available in Body section
x-vcloud-authorization	Enter the session ID received from the Creating a session token for vCD section

- From the **Body** tab, select the **Raw** radio button, and paste the following parameters to create a **PUT** request to register VIM Server.

```
BODY:
<?xml version="1.0" encoding="UTF-8"?>
<vmext:VimServer
  xmlns:vcloud="http://www.vmware.com/vcloud/v1.5"
  xmlns:vmext="http://www.vmware.com/vcloud/extension/v1.5"
  name="ResVCSA-Name">
  <vmext:Username>Administrator@resvsphere.local</vmext:Username>
  <vmext>Password>*****</vmext>Password>
  <vmext:Url>https://FQDN:443</vmext:Url>
  <vmext:IsEnabled>true</vmext:IsEnabled>
  <vmext:IsConnected>true</vmext:IsConnected>
  <vmext:UseVsphereService>true</vmext:UseVsphereService>
</vmext:VimServer>
```

Update the values for above parameter meters using the following table:

Table 35. Uplink profile details

Parameter	Description
Name	Resource vCenter Server name
Username	Resource vCenter Server administrator username
Password	Password for use with the assigned username
URL	FQDN for resource vCenter Server

NOTE: Any change in the body section requires an update to the Content-Length in the header section. Depending on the number of characters you add or delete in the Body section, update the Content-Length in the Header section by the same amount.

- Keep the remaining parameters set at **Default** and **POST** the request.

NOTE: The Response status 202 Accepted display status means that the vCenter Server is updated successfully.

Retrieve the list of available resource pool

About this task

You can retrieve the list of available resource pools available on the vCenter server to create a provider VDC. To retrieve the list, create a **GET** request. For more information, see the [VMware API Reference Guide](#).

Steps

1. On the deployment server, open the postman application.
2. Paste the following parameters to create a GET request to retrieve the list of available resource pool:

```
url = https://<FQDN>/api/admin/extension/vimServer/<ID>/resourcePoolList
Method = GET
HEADERS
Key          VALUES
Accept       application/*+xml;version=31.0;
Content-Type application/vnd.vmware.admin.resourcePoolList+xml;
x-vcloud-authorization Value as obtained from sessions api
```

3. Update the values for the parameters above using the following table:

Table 36. Uplink profile details

Parameters	Description
FQDN	Enter the FQDN for vCloud Director
ID	Enter the VIM Server ID received from the response of Retrieve VIM Server Details
x-vcloud-authorization	Enter the session ID received from the Creating a session token for vCD section

NOTE: The values that are provided in the example above are for reference only. Update the values as required for your configuration.

The **Response status 200 OK** message displays and a list of available resource pools displays.

Retrieve NSX-T Manager instance details

About this task

Post a GET request on vCD to retrieve the NSX-T Manager details. To retrieve the VIM server details, perform the following steps:

Steps

1. On the deployment VM, open the **Postman** application.
2. Paste the following parameters in the **Postman Headers**:

```
url: https://<FQDN>/api/admin/extension/nsxtManagers
Method = GET
Header:
x-vcloud-authorization: Use the value fetched from the session API.
Accept application/*+xml;version=31.0;
```

3. Using the parameters in the following table to update the values for the parameters above:

Table 37. Parameter description

Parameter	Description
FQDN	Enter the FQDN for vCloud Director
x-vcloud-authorization	Enter the session ID received from the Creating a session token for vCD section

The **Response status 200 OK** message displays.

- From the received response, make note of the **Name**, **href**, and **ID** of the **NSX-T Manager**, as shown in the example below. This information is required when creating the provider VDC.

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<vmext:NsxTManagers xmlns="http://www.vmware.com/vcloud/v1.5" xmlns:vmext="http://
www.vmware.com/vcloud/extension/v1.5" xmlns:ovf="http://schemas.dmtf.org/ovf/envelope/1"
xmlns:vssd="http://schemas.dmtf.org/wbem/wscim/1/cim-schema/2/
CIM_VirtualSystemSettingData" xmlns:common="http://schemas.dmtf.org/wbem/wscim/1/common"
xmlns:rasd="http://schemas.dmtf.org/wbem/wscim/1/cim-schema/2/
CIM_ResourceAllocationSettingData" xmlns:vmw="http://www.vmware.com/schema/ovf"
xmlns:ovfenv="http://schemas.dmtf.org/ovf/environment/1" xmlns:ns9="http://www.vmware.com/
vcloud/versions">
  <Link rel="add" href="https://192.168.20.124/api/admin/extension/nsxtManagers"
type="application/vnd.vmware.admin.nsxTmanager+xml"/>
  <Link rel="up" href="https://192.168.20.124/api/admin/extension" type="application/
vnd.vmware.admin.vmwExtension+xml"/>
  <vmext:NsxTManager name="nsxManager1" id="urn:vcloud:nsxtmanager:c9ae8923-1b4e-49f6-beff-
aff6518c8d" href="https://192.168.20.124/api/admin/extension/nsxtManagers/c9ae8923-1b4e-49f6-beff-
aff6518c8d" type="application/vnd.vmware.admin.nsxTmanager+xml">
    <Description>NSX-T Manager</Description>
    <vmext:Username>admin</vmext:Username>
    <vmext:Url>https://192.168.20.104</vmext:Url>
  </vmext:NsxTManager>
</vmext:NsxTManagers>
```

Create a provider VDC

Prerequisites

- A vCenter Server instance must be available to provide a resource pool and storage information to provider VDC

About this task

A provider VDC is a collection of compute, memory, and storage resources from a vCenter Server instance. For network resources, a provider VDC uses NSX-T Data Center. A provider VDC provides resources to organization VDCs. For more information, see the [VMware API Reference Guide](#).

Steps

- On the deployment VM, open the **Postman** application.
- Paste the following parameters in the **Postman Headers**.

```
url = https://<FQDN>/api/admin/extension/providerVdcParams
Method = POST
Authorization
Type- Basic Auth
HEADERS
Key          VALUES
Accept      application/*+xml;version=31.0;
Content-Type application/vnd.vmware.admin.createProviderVdcParams+xml;
x-vcloud-authorization Use the value fetched from the sessions api
```

Update the values for the parameters above using the following table:

Table 38. Uplink profile details

Parameter	Details
FQDN	Enter the FQDN for vCloud Director
x-vcloud-authorization	Enter the session ID received from the Creating a session token for vCD section

- In the **Body** tab, select the **Raw** radio button, and paste the following parameters to create a POST request to create VDC provider.

```
BODY:
<?xml version="1.0" encoding="UTF-8"?>
```

```

<vmext:VMWProviderVdcParams
xmlns="http://www.vmware.com/vcloud/v1.5"
xmlns:vmext="http://www.vmware.com/vcloud/extension/v1.5"
name="nsxTPvdc1">
<vmext:ResourcePoolRefs>
  <vmext:VimObjectRef>
    <vmext:VimServerRef
      href="https://192.168.20.124/api/admin/extension/vimServer/3cd1ac67-88de-4c4b-8e1a-
d171e322d8d1"/>
    <vmext:MoRef>resgroup-10</vmext:MoRef>
    <vmext:VimObjectType>RESOURCE_POOL</vmext:VimObjectType>
  </vmext:VimObjectRef>
</vmext:ResourcePoolRefs>
<vmext:VimServer
href="https://192.168.20.124/api/admin/extension/vimServer/3cd1ac67-88de-4c4b-8e1a-
d171e322d8d1"
id="urn:vcloud:vimserver:3cd1ac67-88de-4c4b-8e1a-d171e322d8d1"
name="ResVCSA-name"
type="application/vnd.vmware.admin.vmwvirtualcenter+xml"/>
<vmext:NsxTManagerReference
href="https://192.168.20.124/api/admin/extension/nsxtManagers/c9ae8923-1b4e-49f6-beff-
afff65518c8d"
id="urn:vcloud:nsxtmanager:c9ae8923-1b4e-49f6-beff-afff65518c8d"
name="nsxManager-name"
type="application/vnd.vmware.admin.nsxTmanager+xml"/>
<vmext:HighestSupportedHardwareVersion>vmx-7</vmext:HighestSupportedHardwareVersion>
<vmext:IsEnabled>true</vmext:IsEnabled>
<vmext:StorageProfile>*</vmext:StorageProfile>
</vmext:VMWProviderVdcParams>

```

Update the values for above parameter using the following table:

Table 39. Uplink profile details

Parameter	Description
Name	Enter the name of provider VDC
For ResourcePoolRefs	
VimServerRef href	Provide the VimServerRef hyperlink received from the Retrieve VIM server details response
MoRef	Provide the MoRef value received from the Retrieve the list of available resource pool response
VimObjectType	Provide the VimObjectType value received from the Retrieve the list of available resource pool response
ForVimServer	
VimServerRef href	Provide the VimServerRef hyperlink received from the Retrieve VIM server details response
ID	Provide the vCenter Server ID received from the Retrieve VIM server details response
Name	Enter the resource vCenter server name received from the Retrieve VIM server details response
For NSX_ManagerReference	
Name	Enter the NSX-T Manager name received from the Retrieve NSX-T Manager instance details response
href	Provide the NSX-T Manager link received from the Retrieve NSX-T Manager instance details response
ID	Provide the NSX-T Manager ID received from the Retrieve NSX-T Manager instance details response
HighestSupportedHardwareVersion	Enter the supported VMX hardware version

- Keep the remaining parameters default and POST the request.

Response status 201 Accepted display means that the Provider VDC is created successfully.

Create an organization

About this task

This section provides steps to create organization in vCloud Director environment.

Steps

1. From a web browser, use the administrator credentials to log in to vCD cell 1 at `https://<<vCD-Cell-01-fqdn>>/provider`.
2. On the **Organizations** page, click **Add**.
3. From the **New Organization** screen, locate the **Organization name** field and enter the organization name.
NOTE: The name that is provided in this field is a unique identifier that displays as a part of URL that organization users use to log in to the organization.
4. In the **Organization full name** field, enter the organization name.
5. In the **Description** field, provide a description for the organization.
6. Click **Create** to create organization.
NOTE: If required, repeat the above steps that are provided in this section to create more organizations.

Create a new Organization VDC

About this task

You allocate resources to an organization by creating an organization virtual data center that is partitioned from a provider Virtual Data Center (VDC). A single organization can have multiple organization virtual data centers.

Steps

1. From a web browser, use the administrator credentials to log in to vCD cell 1 at `https://<<vCD-Cell-01-fqdn>>/provider`.
2. From the left navigation panel, select **Organization VDCs** and then click **New**.
3. On the **General** screen:
 - a. Enter the name and description of the resource.
 - b. Select the **Enable the Organization VDC** box then click **Next**.
4. On the **Organization** screen, select the organization to assign the resource and click **Next**. The **Provider VDC** screen displays.
5. Select the provider VDC to assign the resource and click **Next**.
6. From the **Allocation Model** screen, select the **Allocation pool** radio button and click **Next**.
7. On the **Configure Allocation Pool Model** screen, keep the default settings and click **Next**.
8. On the **Storage Policies** screen, select **Thin provisioning**, then select the storage policies and click **Next**.
9. On the **Network Pool** screen, move the **Use Network Pool** slider to **Disabled**, and then click **Next**.
10. On the **Ready to Complete** screen, review the settings and click **Finish**.
NOTE: If required, repeat the steps in this section to create more organization VDCs.

Name	Status	State	Allocation Model	Organization	Provider VDC	vCenter
org-vdc-1	Enabled	Allocation Pool	Test-Dell	nsxTPvdc1	ResVCSA	
org-vdc-2	Enabled	Allocation Pool	Test-Dell	nsxTPvdc1	ResVCSA	

Figure 113. Ready to Complete screen

Create new catalog

About this task

A newly created organization does not have a catalog in it. A catalog is required to store vApp templates, media files, and catalog items that are used as building blocks to create their own vApps.

Steps

1. On the **Organization VDC** page, click on the name of organization VDC to view the details then click **Open in Tenant Portal**.
2. On the **Tenant Portal**, click the **Main menu** icon, and then select **Libraries** from the displayed list.
3. From the left navigation panel, click **Catalogs**, and then click **New**.
4. On the **Create Catalog** screen, locate the **Name** field enter the catalog name.
5. Click Create.

NOTE: If required, repeat the above steps provided in this section to create more catalogs.

Create Catalog ×

Name this Catalog

A catalog allows you to share vApp Templates and media with other users in your organization. You can also have a private catalog for vApp Templates and media that you frequently use.

Name *

Description

Pre-provision on specific storage policy

CANCEL OK

Figure 114. Create Catalog screen

Create vApp Templates

Prerequisites

- Make sure that you have all the OVF files to vApp Template.

NOTE: Verify that these OVF or OVA files do not have any network adapter attached to it while creating the vApp Template. Once vApp Templates are created you can add a network adapter to the template.

About this task

The vApp templates are VM images that are preloaded with the OS, application, or data. These templates ensure that VMs are consistently configured across an entire organization. vApp templates are added to catalogs. In this deployment we will be creating two vApp templates: one vApp template for windows vApps and Second vApp Templates for CentOS vApps.

Steps

1. From the **Organization VDC** page, click on the name of organization VDC to view the details then click **Open in Tenant Portal**.
2. On the **Tenant Portal**, click the **Main menu** icon, and then select **Libraries** from the displayed list.
3. From the left navigation panel, click **vApp Templates** then click **Add**.
4. On the **Select Source** screen, select **Browse** radio button, click the **Upload** icon and select the all of the vApp files from your local. Then click **Next**.
5. On the **Review details** screen, review the settings selected then click **Next**.

6. On the **Select vApp Template Name** screen:
 - a. In the **Name** field, enter the vApp template name.
 - b. In the **Description** field, enter a brief description of vApp template.
 - c. In the **Catalog** field, select the vApp template catalog from the drop-down list.
 - d. Click **Next**.
7. On the **Ready to Complete** screen, review the provided settings, and then click **Finish** to create vApp template.

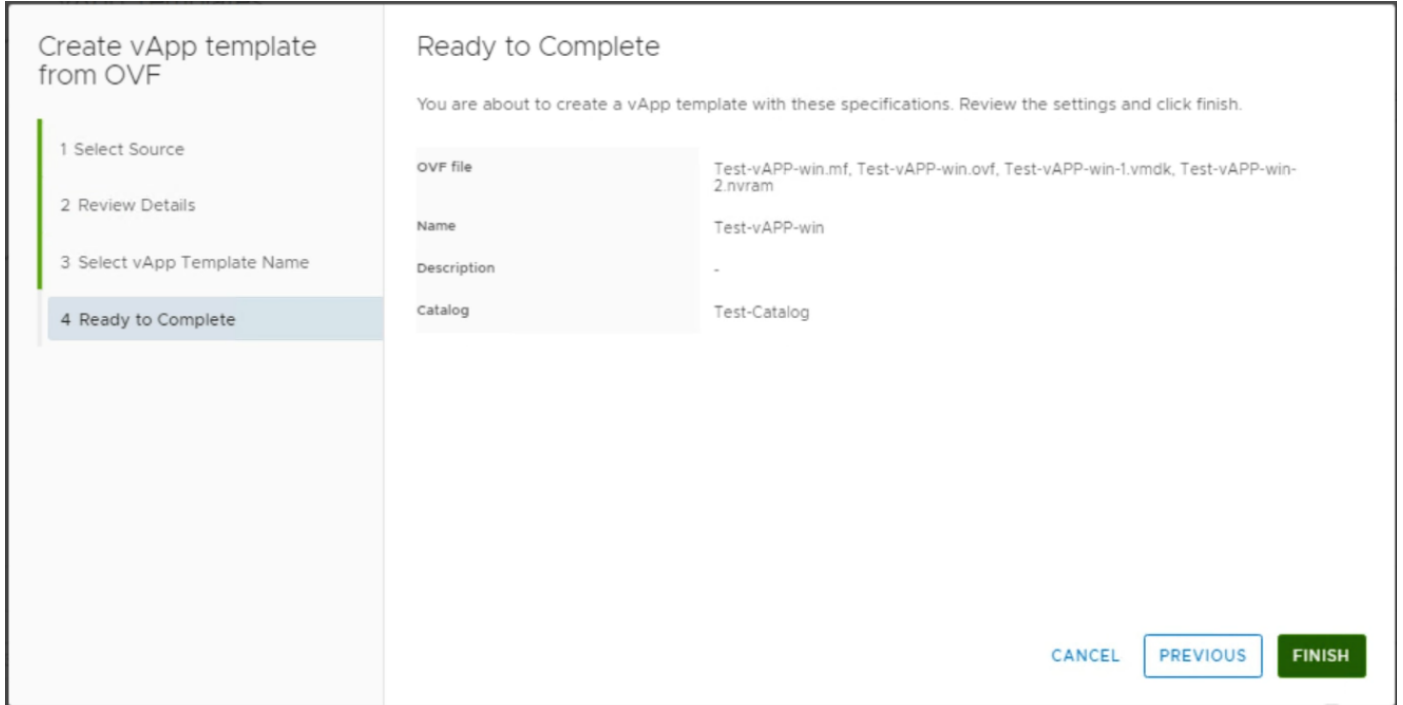


Figure 115. Ready to Complete screen

NOTE: If required, repeat the steps in this section to create more vApp Templates.

Create vApp

About this task

A vApp are VMs to communicate over a network and to use resources and services in a deployed environment.

Steps

1. On the **Tenant Portal**, click the **Main menu** icon then select **Datacenters** from the list.
2. From the left navigation panel, click **vApps**, and then click **NEW VAPP**.
3. On the **New vApp** window, in the **Name** field enter the vApp name.
4. In the **Description** field, enter a brief description about vApp.
5. Click **Create**.

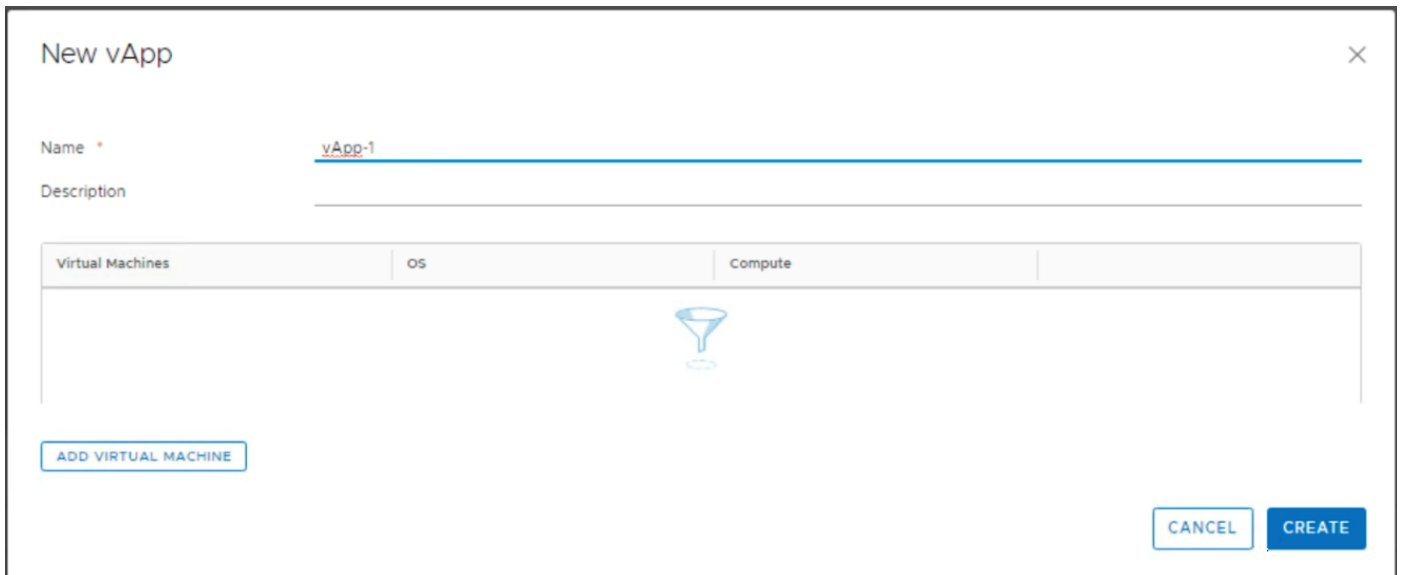


Figure 116. New vApp screen

NOTE: If required, repeat the steps provided in this section to create more vApps.

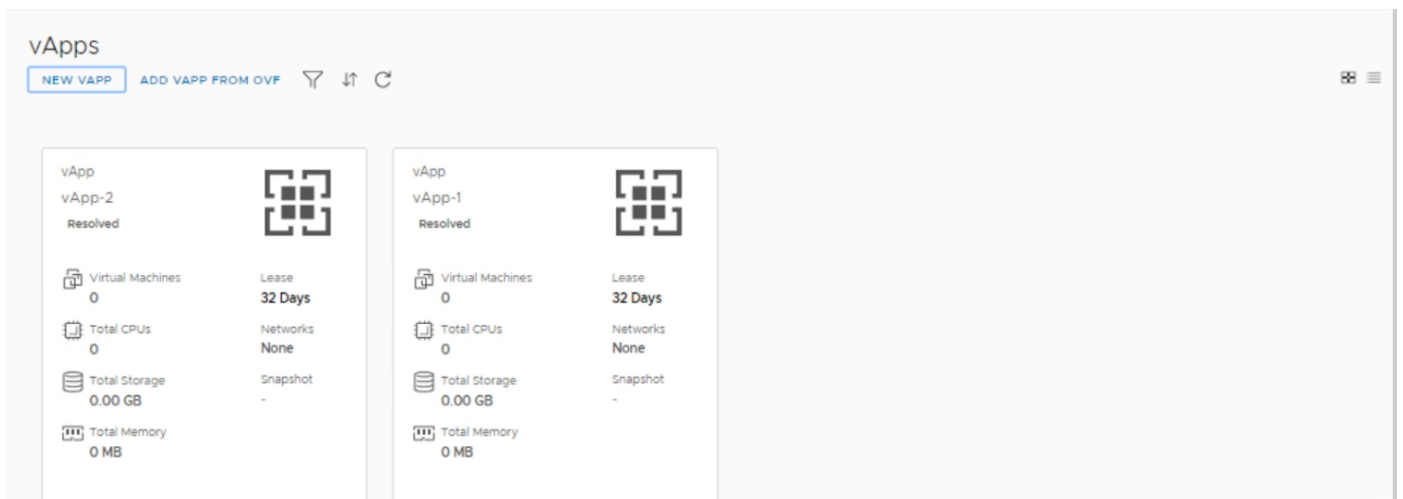


Figure 117. vApps screen

Create virtual machine for vApp template

About this task

This section provides steps to create virtual machine for vApp template in vCloud Director environment.

Steps

1. On the **Tenant Portal**, click the **Main menu** icon, and then select **Datacenters** from the displayed list.
2. From the left navigation panel, click **Virtual Machines** then click **New VM**.
3. On the **New VM** window, locate the **Name** field enter the new VM name.
4. In the **Computer Name** field, enter the computer name for the VM.
5. In the **Description** field, enter a brief description for new VM.
6. In the **Type** field, select the **From Template** radio button.
7. In the **Templates** section select the vApp template for VM.
8. Click **OK** to create VM.

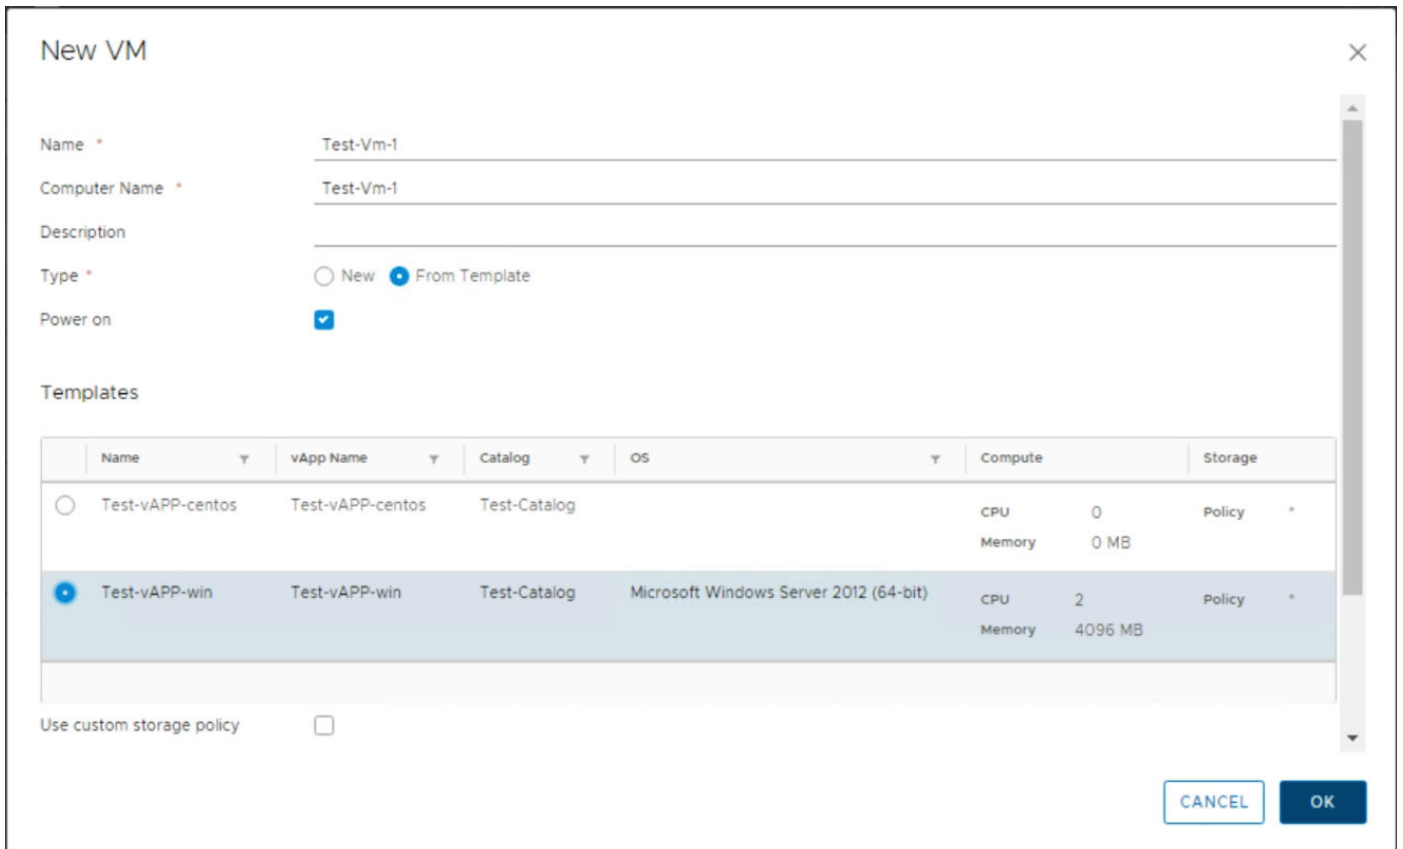


Figure 118. New VM screen

NOTE: If required, repeat the steps provided in this section to create more VM.

Add a network to organization VDC

About this task

Add a network to organization VDC in vCloud Director environment.

Steps

1. On the **Tenant Portal**, click the **Main menu** icon, select **Datacenters** from the displayed list.
2. From the left navigation panel, click **Networks**, and then click **Add**.
3. On the **Network Type** screen, select **Imported** radio button to use existing NSX-T logical switches and click **Next**.
4. On the **NSX-T Logical Switch** screen, select the **NSX-T logical switch** and click **Next**.
5. From the **General** screen:
 - a. In the **Name** field, enter the name of the Organization VDC network.
 - b. In the **Gateway CIDR** field, enter the CIDR of logical switch.
 - c. In the **Description** field, enter a brief description of the network.
 - d. Click **Next**.
6. On the **Static IP Pools** screen, in the **Static IP Pools** field, enter the Static IP range, click **Add**, and then click **Next**.
7. On the **DNS** screen:
 - a. In the **Primary DNS** field, enter the **DNS IP** address.
 - b. In the **DNS suffix name** field, enter the domain name.
 - c. Click **Next**.
8. On the **Ready to Complete** screen, review the provided information and click **Finish**.

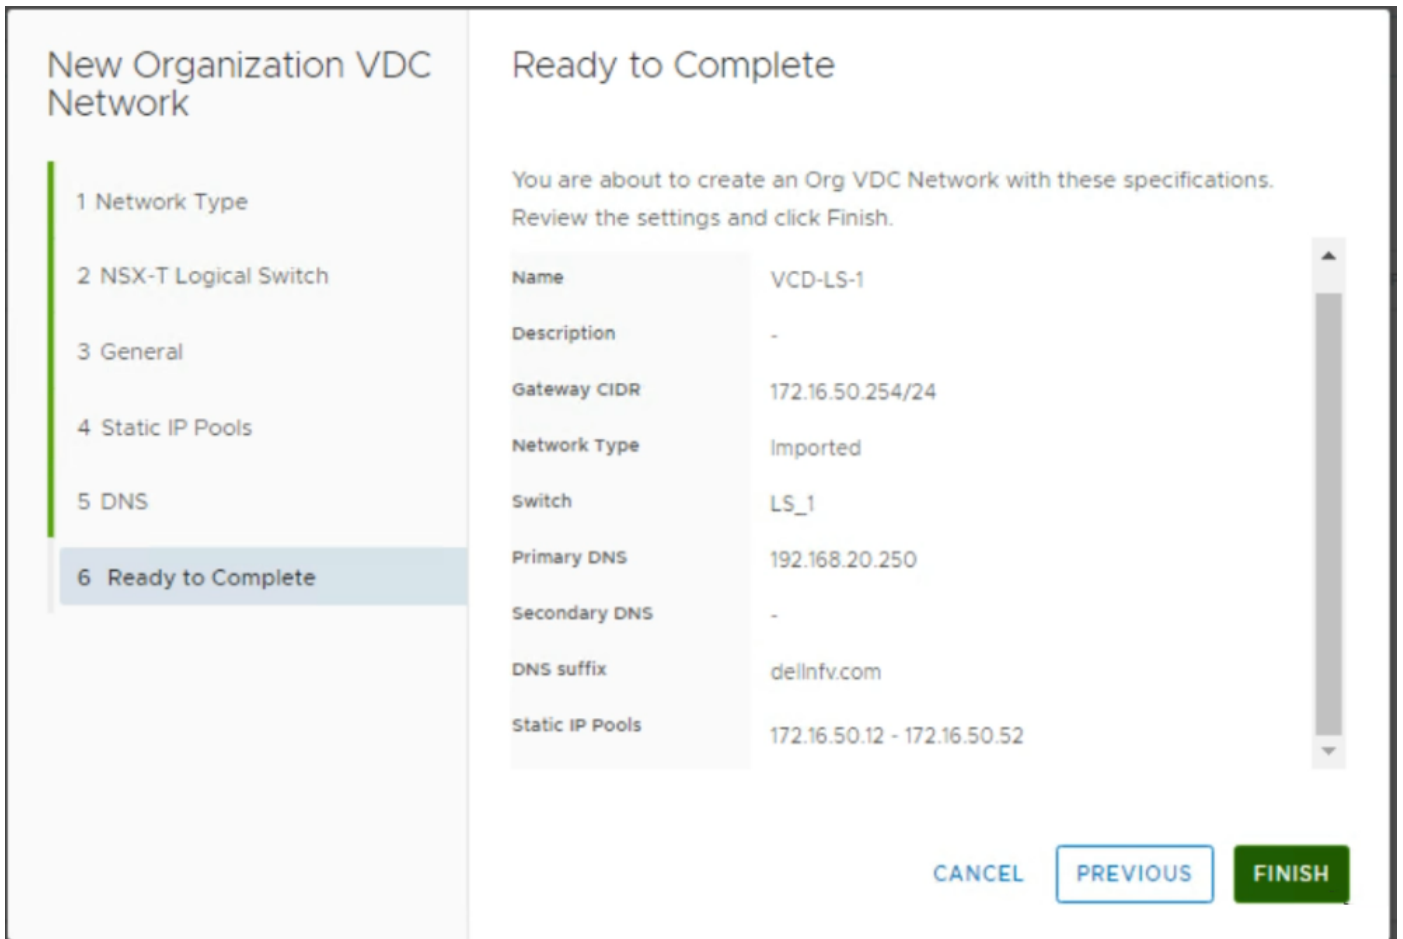


Figure 119. Ready to complete screen

NOTE: If required, repeat the steps in this section to add more networks.

Networks

ADD

Name	Status	Gateway CIDR	Network Type	Connected To	IP Pool Consumed	Shared
VCD-LS-1	✓	172.16.50.254/24	Imported	-	0%	-
VCD-LS-2	✓	172.16.70.254/24	Imported	-	0%	-

Figure 120. Networks listing

Add Network to vApp

About this task

Add networks to a vApp in vCloud Director environment.

Steps

1. On the **Tenant Portal**, click the **Main menu** icon then select **Datacenters** from the list.
2. From the left navigation panel, click **vApps**.
3. In the **vApps** page, locate the desired **vApp**, click the **Actions** drop-down list and select **Add Network**.
4. On the **Add Network** window, select the **OrgVDC Network** radio button in the **Type** field.

5. From the listing of networks, select the network to add with vApp, and then click **Add**.

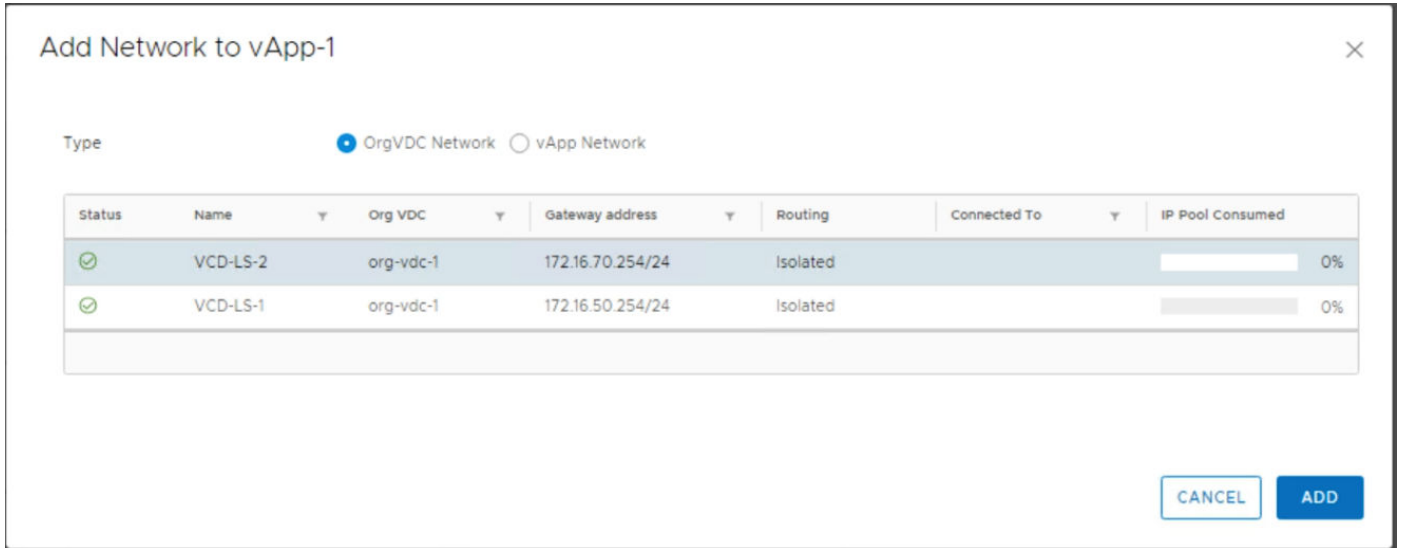


Figure 121. Add Network to vApp screen

NOTE: If required, repeat the steps in this section to add more networks to vApp.

Add Network to Virtual Machine

About this task

This section provides steps to add networks to virtual machine in vCloud Director.

Steps

1. On the **Tenant Portal**, click the **Main menu** icon then select **Datacenters** from the displayed list.
2. From the left navigation panel, click **Virtual Machines**. The **Virtual Machines** screen displays.
3. Select the desired VM and then click **Details**.
4. In the **Hardware** section, click **Add in the NICs** sub-section.
 - a. From the **Network** drop-down list select the organization VDC network.
 - b. Click to select the **Connected** check box.
 - c. From the **IP Mode** drop-down list, select the **Static - IP Pool** option.
 - d. Click **Save**.

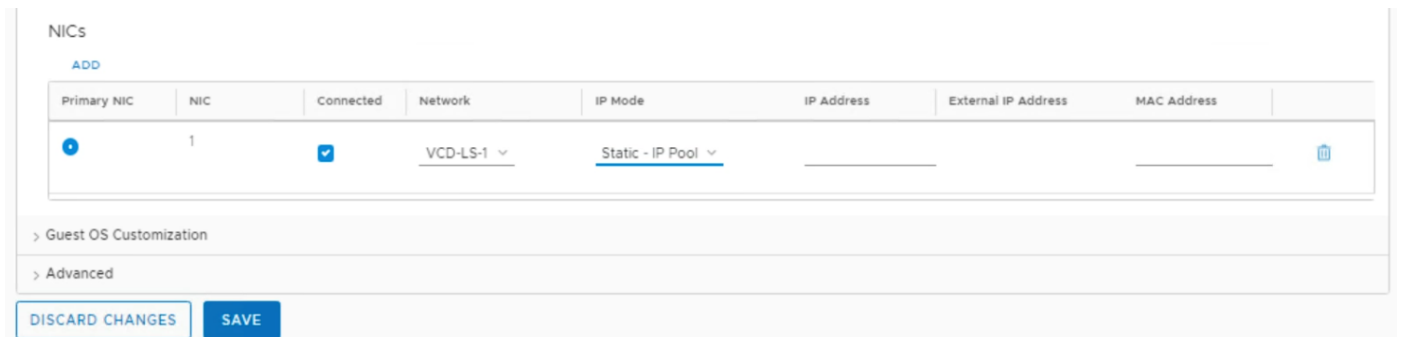


Figure 122. Adding network to VM screen

Add VM to a vApp

About this task

This section provides the to add VMs to vApp.

Steps

1. On the **Tenant Portal**, click the **Main menu** icon, and then select **Datacenters**.
2. From the left navigation panel, click **vApps**.
The **vApps** screen displays.
3. Locate the desired vApp, click the **Actions** drop-down list, and select **Add VM**.
4. On the **Add VMs** , click the **ADD VIRTUAL MACHINE** button.
5. From the **New VM** screen, locate the **Name** field and enter the VM name.
6. In the **Computer Name** field, enter the computer name.
7. In the **Description** field, enter a brief description for VM.
8. In the **Type** field, select the From Template radio button.
9. In the **Templates** section select the vApp template for VM.
10. Click **OK**.

New VM

Name *

Computer Name *

Description

Type * New From Template

Power on

Templates

	Name	vApp Name	Catalog	OS	Compute	Storage
<input type="radio"/>	Test-vAPP-centos	Test-vAPP-centos	Test-Catalog	CentOS 7 (64-bit)	CPU 2 Memory 4096 MB	Policy *
<input checked="" type="radio"/>	Test-vAPP-win	Test-vAPP-win	Test-Catalog	Microsoft Windows Server 2012 (64-bit)	CPU 2 Memory 4096 MB	Policy *

Use custom storage policy

Figure 123. New VM window screen

11. Click **ADD** to create.

NOTE: If required, repeat the steps in this section to add more VM to vApp.

Move a VM to vApp

About this task

This section provides steps to move a VM to vApp.

Steps

1. On the **Tenant Portal**, click the **Main menu** icon then select **Datacenters**.
2. From the left navigation panel, click **Virtual Machines**.
3. On the **Virtual Machines** screen, locate the desired VM, click the **Actions** drop-down list, and select the **Move to** option.
4. On the **Select Destination vApp** screen, select the vApp, and then click **Next**.
5. From the **Configure Resources** screen, select the following from the **NICs** section:
 - a. From the **Network** drop-down list select the organization VDC network.
 - b. Click to place a check in the **Connected** check box.
 - c. From the **IP Mode** drop-down list select the **Static - IP Pool** option.
 - d. Click **Next**.
6. On the **Ready to Complete** screen, click **Done** to move the VM to vApp:

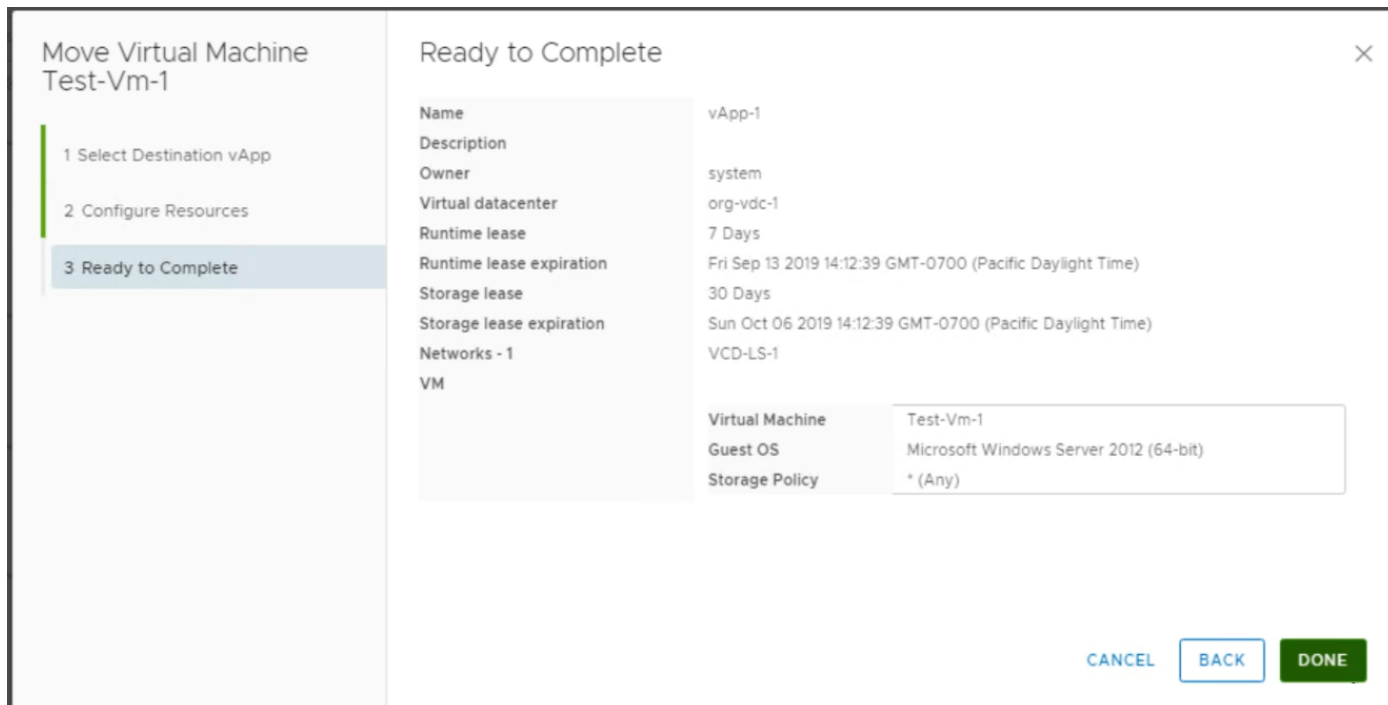


Figure 124. Ready to complete screen

NOTE: If required, repeat the steps in this section to move VM to vApp.

VMware vRealize Log Insight deployment and configuration

Dell EMC Ready Solution bundle uses the VMware vRealize Log Insight (vRLI) to collect the log data from ESXi hosts, it also connects with vCenter servers to collect the log data of server events, tasks, and alarms.

In this deployment, vRLI is deployed in a single cluster configuration that consists of three nodes:

- Master
- Worker
- Witness

Prerequisites:

- ESXi 6.7 U2 server is up and running
- AD-DNS and NTP is up and running
- Management and Resource VCSAs are up and running
- Manual creation of forward and reverse lookup entries for all vRealize Log Insight instances on DNS server are added prior to deployment

Deploy the vRealize Log Insight virtual appliance

About this task

Deploy the vRLI virtual appliances using the steps provided in this section.

Steps

1. Log in to the Management vCenter using the VMware vSphere Web Client.
2. Right-click the **Management Datacenter**, then click **Deploy OVF Template**.
3. On the **Select template** window, enter the download URL or click **Browse** to locate the .OVA file on your computer, then click **Next**.
4. On the **Select name and location** window, enter the **Name**, select the **Location** then click **Next**.
5. On the **Select a resource** window, select the ESXi to deploy vRLI, and click **Next**.
6. Review the settings that are selected then click **Next**.
7. Use the scroll bar to review the information in the **Accept license agreement** section, if you agree, click **Accept**, and click **Next**.
8. On the **Select Configuration** page, select the size of the vRealize Log Insight virtual appliance based on the size of the environment for which you intend to collect logs, then click **Next**.
9. On the **Select storage** screen:
 - a. From the **Select virtual disk format** drop-down list, select **Thin provision**.
 - b. From the **VM storage policy** drop-down list, select **vSAN Default Storage Policy**.
 - c. Select the datastore, and click **Next**.

NOTE: Deploy the vRealize Log Insight virtual appliance with thick provisioned eager zeroed disks whenever possible for better performance and operation of the virtual appliance.

10. In the **Select networks** screen, select the appropriate networks.
11. On the **Customize template** screen, configure the Networking Properties for the vRLI virtual appliance.

NOTE: If you do not provide network settings such as an IP address, DNS server, and gateway information, vRLI uses DHCP to set those settings.

NOTE: Do not specify more than two domain name servers. If you specify more than two domain name servers, the configured domain name servers are ignored within the vRealize log.

12. On the **Customize template** page, select **Other Properties**, set the root password for the vRealize Log Insight virtual appliance, then click **Next**.
13. Review the settings in the **Ready to complete** screen, and click **Finish** to deploy the vRLI VM.

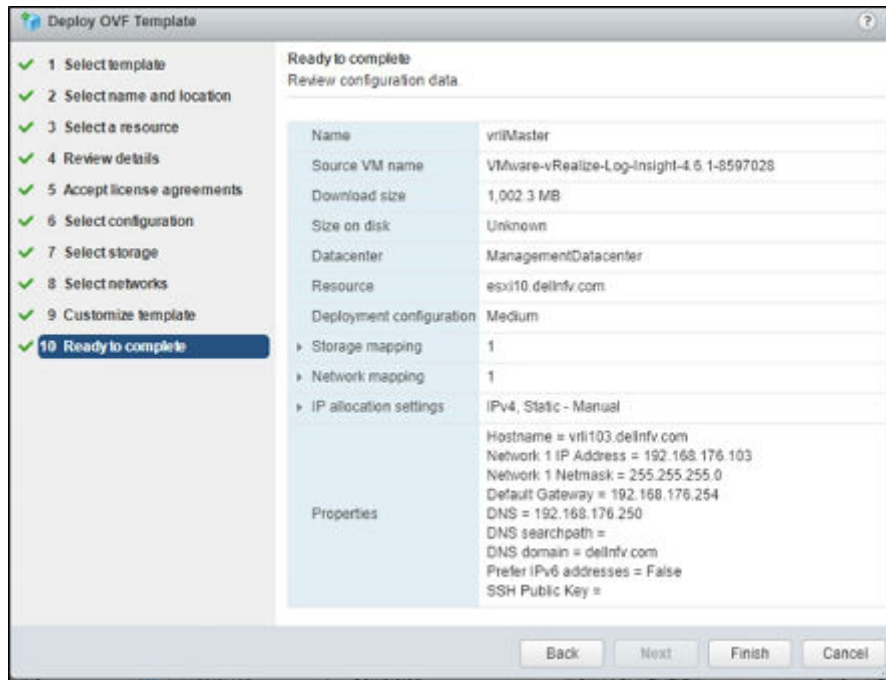


Figure 125. Review Configuration data screen

NOTE: Repeat the steps in this section two more times to deploy worker nodes, and to create a cluster of three appliances, with only the Master node turned on.

Configure the root SSH password for vRLI virtual appliance

About this task

NOTE: The steps in this section are optional and only required if the root SSH password is not set at the time of vRealize Log Insight .OVA file deployment.

By default, the SSH connection to the virtual appliance is disabled. You can configure the root SSH password from the VMware Remote Console or when you deploy the vRLI virtual appliance. You can also enable SSH and set the root password from the VMware Remote Console.

Before configuring the root SSH password for vRLI virtual appliance, verify that the vRealize Log Insight virtual appliance is deployed and running.

Steps

1. In the **vSphere Client inventory**, click the **vRealize Log Insight virtual appliance**, and open the **Console** tab.
2. Go to a command line by following the key combination specified on the splash window.
3. In the console, type `root`, and press **Enter**.
4. Leave the password field empty and press **Enter**. The **Password change requested. Choose a new password.** message displays in the console.
5. Leave the old password empty and press **Enter**.
6. Enter a new password for the root user, then press **Enter**.
7. Enter the new password again for the root user, and press **Enter**.

NOTE: The password must consist of at least eight characters, and must include at least one upper case letter, one lower case letter, one digit, and one special character. You cannot repeat the same character more than four times.

The **Password changed** message displays.

Master node configuration

About this task

Configure the master node using the steps provided in this section.

Steps

1. Go to the vRLI Web user interface at <https://<vRLI_Host_IP_/FQDN>.
2. From the **Setup** window, click **Next**.
3. On the **Choose Deployment Type** window, click **Start New Deployment**.
4. From the **Admin Credentials** screen:
 - a. In the **Email** field, enter the admin email ID.
 - b. In the **New password** field, enter the admin password.
 - c. In the **Confirm new password** field, reenter the password to confirm.
 - d. Click **Save and Continue**.
5. Enter the license key, click **Add License**, and click **Save and Continue**.
6. On the **General Configuration** page, enter an email address in the field that is provided to receive system notifications from vRealize Log Insight.
7. Optionally, you can participate in the **Customer Experience Improvement Program** by selecting the **Join the VMware Customer Experience Program** check box. Otherwise, leave the check box blank.
8. Click **Save and Continue**.
9. On the **Time Configuration** page, set how time is synchronized on the VRLI appliance by selecting NTP server then enter the IP address for NTP server and click **Test**.
10. Once the test is successful, click **Save and Continue**.
11. On the **SMTP Configuration** window, keep the default settings, and click **Skip**.
12. On the **Setup complete** screen, click **Finish** to complete the setup.

Worker node configuration

About this task

NOTE: Configure a minimum of three nodes in a vRLI cluster to provide ingestion, configuration, and user space high availability.

Configure the worker node using steps provided in this section.

Steps

1. Power on the second Log Insight Appliance and wait for the configuration process to complete.
2. Go to the second Log Insight URL, for example, `https://,vRLI 2nd Host IP/FQDN`.
3. From the **Setup** window, click **Next**.
4. On the **Choose Deployment Type** window, click **Join Existing Deployment**.
5. Enter the IP address or hostname of the vRLI master node and click **Go**.
6. Select **Click here** to access the **Cluster Management** page to be redirected to the **Master Node vRLI login** page.
7. Click **Allow** button on the **Cluster Management** page for the new worker node appliance to join.
8. Repeat the steps in this section to configure a third Log Insight Appliance.

Enable Integrated Load Balancer

Prerequisites


- Verify that all vRLI nodes and the specified Integrated Load Balancer IP address are on the same network
- DNS records have been configured for the IP addresses

About this task

The Integrated Load Balancer (ILB) ensures that incoming Ingestion traffic is accepted by vRLI even if some vRLI nodes become unavailable. The ILB also balances incoming traffic fairly among available vRLI nodes. vRLI clients, using both the Web user interface and ingestion (through Syslog or the Ingestion API), should connect to the vRLI using the ILB address.

Steps

1. Login to the Master node Log Insight web UI with Admin login, for example, <https://<Log Insight FQDN>>
2. From the upper-right menu, click the **Administration** then select **Cluster**.
3. On the **Cluster** page, click **+NEW VIRTUAL IP ADDRESS**.
4. On the **New Virtual IP** window, enter the **IP** and **FQDN for ILB**, then click **Save**.
5. Refresh the **Cluster Management** page to confirm that the LB Status displays as **Available**.

 **NOTE:** You can configure multiple virtual IP addresses. Click **+NEW VIRTUAL IP ADDRESS** and enter the IP Address in the field provided. This option also allows you to enter the FQDN and tags.

Integrate vRLI with AD

About this task

Use the steps provided in this section to integrate vRLI with AD.

Steps

1. Navigate and login to vRLI, for example, [<https://<Log Insight FQDN>](https://<Log Insight FQDN>)
2. From the upper-right menu, click the **Administration** then select **Authentication**.
3. On the **Active Directory** tab configure as follows:
 - **Enable Active Directory support:** slide the toggle switch to ON
 - **Default Domain:** Enter the relevant domain name
 - **Username:** user must admin rights
 - **Password:** Password for above user
 - **Connection Type:** Can be Standard or can be set to Custom for testing specific ports.
 - **Require SSL:** Check if SSL required.
4. Click **Test Connection** to validate the settings.
5. Once the connection is validated successfully, click **Save**.

Integrate vRLI with VMware vCenter

About this task

Integrate vRLI with VMware vCenter to pull the tasks, events, and alerts.

Steps

1. Go to the URL for the Master Node Log Insight sever, for example, <https://Log Insight FQDN/IP>
2. From the upper-right corner of the window, click **Administration**, select **vSphere**, then click **+ Add vCenter Server**.
3. Enter the **Hostname (IP/FQDN)** for the **Management vCenter** and the user credentials to connect to the vCenter Server system, and click **Test Connection** to verify the connection.
4. Once tested successfully, click **Save**.

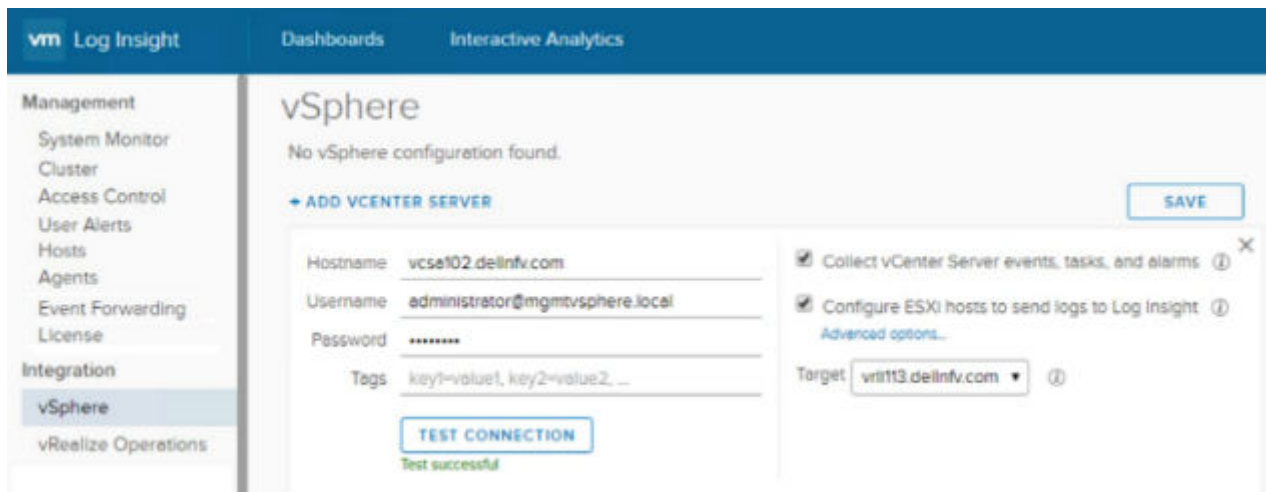


Figure 126. VMware vSphere integration log in screen

- Repeat the above steps to configure the integration of the resource vCenter.

Configure vRLI to send notifications to vRealize Operations Manager

Prerequisites

- ROps Manager VM should be turned on and configured properly
- vROps FQDN should be registered with DNS

About this task

You can configure vRLI to send alert notifications to vRealize Operations (vROps) Manager. Integrating vRLI alerts with vROps Manager allows you to view all information about your environment in a single user interface. You can send notification events from multiple vRLI instances to a single vROps Manager instance.

Steps

- Go to the URL for the Log Insight server, for example, <https://<Log Insight FQDN>>.
- From the upper-right corner of the window, click **Administration**, and then select **vRealize Operations**.
- Enter the **FQDN/IP** in the **Hostname** field for the vROps, then enter user credentials to connect with vROps server system.
- Click to place a check in the **Enable alerts integration** box.
- Click **Test Connection** to verify the connection.
- Once tested successfully, click **Save**.

Add Log Insight content packs

The Content Pack Marketplace is where you can access content packs for VMware and non-VMware products. Content Packs include domain-specific queries, alerts, dashboards, field extractions, and agent group templates for their associated products. A content pack is not required to ingest logs from a specific product. Essentially, the content pack makes it easier and faster to find critical log data by selecting and alerting admin to common issues that are present in the ingested log data. As a result, troubleshooting and root cause analysis efforts take less time.

NOTE: Ensure that you download the latest and compatible Content Packs for vSAN, vROps, NSX-T, VCD, vRO, and vSphere from VMware Marketplace.

Offline update for content pack

Prerequisites

- Content pack needs to be downloaded for the import
- Components which requires the content packs: vSAN, vROps, NSX-T, VCD, vRO, and vSphere
- The browser time and server time must be in same time zone as UTC to forward the log with different products

About this task

This section provides the steps to update the content pack offline.

Steps

1. Go to the URL for the Log Insight sever, such as `https://<Log Insight FQDN>`
2. From the upper-right corner of the window, select **Content Packs**, and then click **Import Content Pack**.
3. Browse for the downloaded **Content Pack** and select **Import**.
4. Click **OK** to complete Content Pack import.
5. Repeat the steps in this section to install content packs for the remaining components.

Online update for content pack

Prerequisites

- Internet connectivity
- Components that require content packs: vSAN, vROps, NSX-T, VCD, vRO, and vSphere

About this task

This section provides steps to update content pack online.

Steps

1. Content packs are available for many of the components used in the DELL NFV and can be imported into any instance of the Log Insight.
2. Go to the URL for the Log Insight sever, for example, `https://<Log Insight FQDN>`
3. From the upper-right corner of the screen, click **Content Packs**.
4. From the **Customer Pack Marketplace**, click **Marketplace**.
5. Select a content pack.
6. Select the license agreement and click the **Install** button to install the content pack.
7. Click **OK** to complete the VSAN setup instructions.
8. Repeat the steps in this section for the remaining vROps, NSX-T, vCD, vRO, and vSphere Content Packs.

vRLI integration with vCD

Prerequisites

- vCD content pack for vRLI should be installed. See the Add Log Insight content packs section to install

About this task

Integrate vRLI with vCD to view the operational and health status of vCD environment.

Steps

1. Click **Administration**, and then click **Agents**.
2. From the drop-down list, select **Cloud Director Cell Servers**.

3. Click the **COPY TEMPLATE** button.
4. On the **Copy Agent Group** window, enter the agent group name, and click **Copy**.
5. Specify a filter then click the **Save New Group** button.
The **Agent Group** is created successfully.
6. Download the log insight agent on the each vCD cells:
 - a. Click **Administration**, and then click **Agents**
 - b. At the end of the page click **Download Log Insight Agent Version 4.8**.

The **LinuxRPM** file of the Log Insight agent starts to download.
7. Install the downloaded **LinuxRPM** file of log insight agent to each vCD cells:
 - a. Copy the LinuxRPM file of log insight agent to the `.tmp` folder on vCD cell.
 - b. SSH to the vCD cell with root user.
 - c. Run the following command on SSH to install the log insight agent LinuxRPM file:

```
rpm -i VMware-Log-Insight-Agent-4.8.0-13020979.noarch_192.168.20.113.rpm
```

The agent installation begins.

8. Once the installation is complete, configure the installed log insight agent:
 - a. Go to the `etc` director, and run the following command:

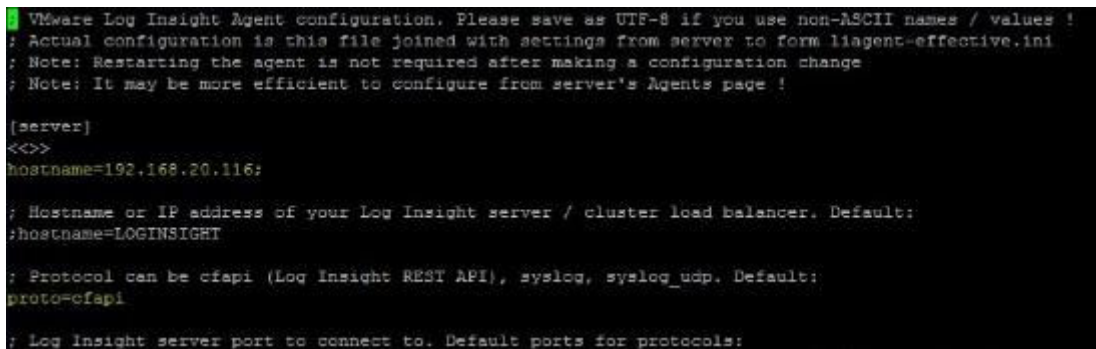
```
vi liagent.ini
```

OR

```
vi /var/lib/loginsight-agent/liagent.ini
```

- b. Using the downloaded agent from log insight, verify that the Log Insight hostname is present. If it is not, add the hostname:

```
Uncomment proto=cfapi
```



```
VMware Log Insight Agent configuration. Please save as UTF-8 if you use non-ASCII names / values !
; Actual configuration is this file joined with settings from server to form liagent-effective.ini
; Note: Restarting the agent is not required after making a configuration change
; Note: It may be more efficient to configure from server's Agents page !

[server]
<<>>
hostname=192.168.20.116:

; Hostname or IP address of your Log Insight server / cluster load balancer. Default:
;hostname=LOGINSIGHT

; Protocol can be cfapi (Log Insight REST API), syslog, syslog_udp. Default:
proto=cfapi

; Log Insight server port to connect to. Default ports for protocols:
```

Figure 127. Host name and uncommented protocol screen

- c. Append the vRLI agent configuration:
 1. From the vRLI click **Installed Content Packs, VMware vCloud Director**, and then click **Agent Groups**.
 2. Copy the **Configuration**.
 3. In the VCD Cell SSH, append the configuration to the `liagent.ini` file.

```
[update]
; Do not change this parameter
package_type=rpm

; Enable automatic update of the agent. If enabled:
; the agent will silently check for updates from the server and
; if available will automatically download and apply the update.
;auto_update=yes
[filelog|vcd-essential]
directory=/opt/vmware/vcloud-director/logs
include=vcloud-container-debug*;upgrade*;vmware-vcd-support*;watchdog*
event_marker=(\d{2})\d{4})-\d{2}-\d{2}\s\d{2}:\d{2}:\d{2},\d{3}\s
tags={"vmw_product":"vcd"}
```

Figure 128. Append the vRLI agent screen

9. Run the following command to restart the log insight agent services:
`service liagentd restart`
10. Run the following command to restart the vCD services:

```
vmware-vcd restart
```

OR

```
service vmware-vcd restart
```

vRLI integration with vRO

Prerequisites:

- vRO must be installed. Refer the Installation of vRO
- vRO must be installed. See [Installation of vRO](#).
- vRO must be configured to forward logs to vRLI. See [Configure vRealize Orchestrator to forward logs to vRLI](#).
- vRO content pack must be installed on vRLI

Integrate vRLI with vCD to view the operational and health status of vCD environment.

Integrate vRLI with vRO

Steps

1. From a web browser, open and log in to the **Log Insight Sever**.
2. Click **Administration**, and then **Agents**.
3. From the drop-down list, select **vRealize Orchestrator**.
4. Click the **COPY TEMPLATE** button.
5. On the **Copy Agent Group** window, enter the agent group name, and click **Copy**.
6. Specify a filter then click the **Save New Group** button.
 The **Agent Group** is created successfully.
7. Refresh the page and select the newly vRO agent group then click **Dashboards** to view the vRO dashboard.

vRealize Orchestrator

vRealize Orchestrator (vRO) is a development and process-automation platform and contains a workflow library and a workflow engine. This allows administrators to create and run workflows to automate the orchestration processes. Orchestrator provides a standard set of plug-ins including a plug-in for vCenter Server and vRealize automation to allow you to orchestrate the tasks in the different environments. In this deployment, one instance of vRO will be deployed.

Installation of vRO

About this task

This section provides steps to vRealize Orchestrator.

Steps

1. Log in to the Management vCenter using the VMware vSphere Web Client.
2. Right-click on the **Management Datacenter**, then click **Deploy OVF Template**.
3. From the **Select template** screen, enter the download URL or click **Browse** to locate the .OVA file on your computer, then click **Next**.
4. On the **Select name and location** screen, enter the **Name** and select the **Location** then click **Next**.
5. On the **Select a resource** screen, select the **ESXi to deploy vRealize Orchestrator** and click **Next**.
6. From the **Review details** window, review the settings selected then click **Next**.
7. Use the scroll bar to review the information in the **Accept license agreement** window and if you agree, then click **Accept** and click **Next**.
8. On the **Select storage** screen:
 - a. From the **Select virtual disk format** drop-down list, select **Thin provision**.
 - b. From the **VM storage policy** drop-down list, select **vSAN Default Storage Policy**.
 - c. Select the datastore and click Next.
9. In the **Select networks** window, select the appropriate networks, then click **Next**.
10. On the **Customize template** window, in the **Application** section fill the following fields:
 - a. **Initial root password**: Set the root password
 - b. **Enable SSH service in the appliance**: Check the box to enable SSH services in the appliances
 - c. **Hostname**: Enter the hostname or FQDN for this VM
11. On the **Customize template** window, locate the **Network properties** section, fill the following fields:
 - a. **Default Gateway**: Enter the IP address of default gateway for vRealize Orchestrator
 - b. **Domain Name**: Enter the domain name for vRealize Orchestrator
 - c. **Domain Search Path**: Enter the domain search path for vRealize Orchestrator
 - d. **Network 1 IP Address**: Enter the IP address for vRealize Orchestrator
 - e. **Network 1 Netmask**: Enter the netmask IP for vRealize Orchestrator
 - f. Click **Next**
12. On the **Ready to complete** window, review the provided configuration details then click **Finish** to deploy vRealize Orchestrator.

Configure NTP in vRO

About this task

This section provides steps to configure NTP in vRO.

Steps

1. From your browser, log in to vRO Appliance Configuration page with administrator credentials at: `https://<<RO-IP>>:5480/`
2. Click **Admin**, and then click **Time Settings**.
3. In the **Time Server** field, enter the IP address of NTP.




Figure 129. Time Settings tab

4. Click **Save Settings** to save the NTP settings.

Configure Orchestrator Server with vSphere Authentication

About this task

Configure the vSphere authentication method in orchestrator to use the vRealize Orchestrator appliances.

Steps

1. In a web browser, login to Orchestrator Access Control Center with administrator credentials at: `https://<<orchestrator_server_IP_or_DNSname>>:8283/vco-controlcenter`
2. On the Host Settings page, click **Change**.

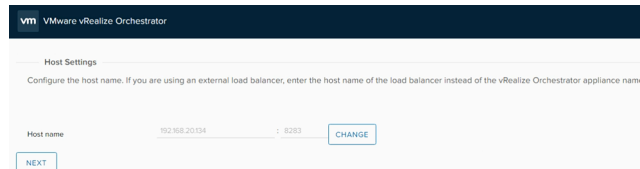


Figure 130. Host Settings

3. In the **Host Name** field, enter the host name or FQDN of vRO VM then click **Apply**.
4. Click **Next**.
5. On the **Authentication Provider** page, from the **Authentication** mode drop-down list, select **vSphere**.
6. In the **Host address** field, enter the host name of resource vCenter, then click **Connect**.
7. Review the **Certificate information** then click **Accept Certificate**.
8. In the **User name** field, enter the user name of resource vCenter admin.
9. In the **Password** field, enter the password for resource vCenter admin user.
10. In the **Default tenant** field, enter the resource vCenter tenant name, then click **Register**. For this deployment, `resvsphere.local` is used as default tenant name.
11. In the **Admin group** field, enter the name of an admin group then click **Search**. For this deployment `resvsphere.local \ComponentManager.Administrators` is selected.

Figure 131. Configure Authentication Provider

12. Click **Save Changes** to save the configuration settings.
13. On the **Test Login** tab, login with resource vCenter administrator credentials then click **Test** to validate the connection.

Updating the vRO using ISO

Prerequisites

- Create a backup of existing vRO appliance.

About this task

Update the vRO using ISO file.

Steps

1. Download the updated vRO iso file from [here](#).
2. Mount the ISO file to the **vRO VM**.
3. In the web browser, log in to the **vRO Appliance management** page at <https://<IP or FQDN>:5480>
4. Click **Update**, and then **Settings**.
5. Change the **Update Repository** to **Use CDROM Updates** then click **Save Settings**.
6. Click **Update**, and then **Status**.
7. Click **Check Updates**.
8. When the update is display, click **Install Updates**.

Configure vRO plug-in for vSphere Web Client

The vRO will be integrated with resource vCenter server instance. Two workflows are required to run in the vRO Orchestrator Client to integrate it with vCenter Server:

- [Add a vCenter Server instance to vRO](#)
- [Register vRealize Orchestrator as a vCenter](#)

Add a vCenter Server instance to vRO

About this task

This section provide steps to add vCenter Server instance to vRO.

Steps

1. From your browser, log in to vRO at <https://<<vRO-fqdn>>>
2. On the **Orchestrator Appliance Home** page, click **Start the Orchestrator Client** to create and manage workflows.
3. From the left navigation panel, click **Library, Workflows**.

4. In the **Search for** field, search **Add a vCenter Server** instance.
5. Click **Run on the Add a vCenter Server instance workflow**.
6. On the **Set the vCenter Server instance properties** tab:
 - a. In the **IP or host name of the vCenter Server instance to add** field, enter the host name or FQDN of resource vCenter.
 - b. In the **HTTPS port of the vCenter Server instance** field, enter the resource **vCenter instance port number**. For this deployment default port number, 443 is used.
 - c. In the **Location of the SDK that you use to connect to the vCenter Server instance** field, enter a path for SDK to connect with resource vCenter server instance. For this deployment default path, /sdk is used.
 - d. Check the **Will you orchestrate this instance?** check box if you want to orchestrate the vCenter Server instance.
 - e. Check the **Do you want to ignore certificate warnings** check box if you want to ignore certificates warnings for the vCenter Server instances. If you select **Yes**, the vCenter Server instance certificate is accepted silently and the certificate is added to the trusted store.

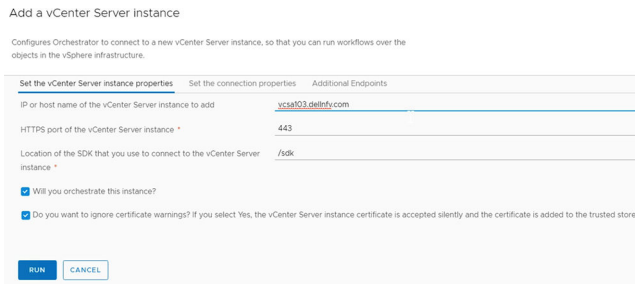


Figure 132. Set the vCenter Server instance properties tab

7. On the **Set the connection properties** tab:
 - a. Check the **Do you want to use a session per user method to manage user access to the vCenter Server system?** check box. This option creates a new session to vCenter Server.
 - b. In the **User name** field of the user that Orchestrator uses to connect to the vCenter Server instance field, enter the Administrator user name of resource vCenter Server.
 - c. In the **Password** field of the user that Orchestrator uses to connect to the vCenter Server instance field, enter the Administrator password of resource vCenter Server.
 - d. In the **Domain name** field, enter the domain name for Orchestrator.

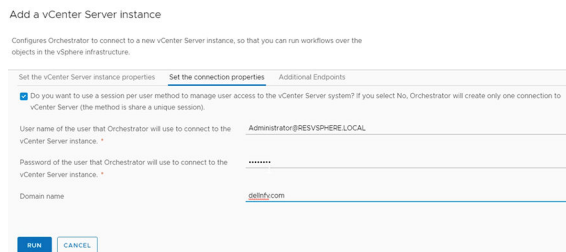


Figure 133. Set the vCenter Server connection properties tab

8. Click **Run** to establish the connection.

Register vRealize Orchestrator as a vCenter Server extension

About this task

To register vRealize Orchestrator as a vCenter server extension:

Steps

1. On the vRealize Orchestrator, click the left navigation panel, **Library**, and then **Workflows**.
2. In the **Search for** field, search **Register vCenter Orchestrator as a vCenter server extension**.
3. Click **Run on the Register vCenter Orchestrator as a vCenter workflow**.
4. On the **Register vCenter Orchestrator as a vCenter extension** page:

- a. In the **vCenter Server instance to register Orchestrator with** field, select the **resource vCenter**.
- b. Click **Run**.




Figure 134. Register vCenter Orchestrator as a vCenter extension page

5. Once the workflow is completed, reboot the resource vCenter server integrated with vRO.
6. Re-login to the resource vCenter Server and verify that vRO plug-in is present by clicking **Home** and then **Inventories**.

Configure vRealize Orchestrator to forward logs to vRLI

About this task

You can configure each vRO to forward logs to the vRealize Log Insight.

Steps

1. In a web browser, login to Orchestrator Control Center with administrator credentials at: `https://<<orchestrator_server_IP_or_DNSname>>:8283/vco-controlcenter`
2. On the **Home** page, under the **Log** section, click **Logging Integration**.
3. On the **Logging Integration** page, set the following properties:
 - a. Move the **Enable logging to a remote log server** slider to allow vRLI to collect logs from vRO.
 - b. In the **Type** field, select **Use Log Insight Agent** radio button.
 - c. In the **Host** field, provide the vRLI host name.
 - d. In the **Port** field, set the port number to 9000.
 - e. In the **Protocol** drop-down list, select the protocol to **cfapi**.
4. Click **Save**.

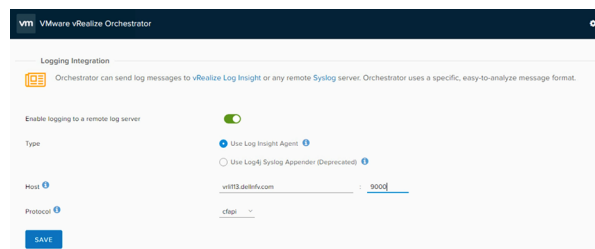


Figure 135. vRLI Logging Integration page

VMware vRealize Operations Manager deployment and configuration

The vRealize Operations (vROps) Manager delivers intelligent operations management with application-to-storage visibility across physical, virtual, and cloud infrastructures. Using policy-based automation, operations teams automate key processes and improve IT efficiency.

As part of the vROps Manager deployment, three nodes will be deployed as follows:

- Master
- Data
- Data used as a master replica node

NOTE: The vROps deployment covered in this document has been done for Medium configuration. The vROps OVA template deployment is to be done three nodes – data, master, and replica. The vRealize Operations Manager UI can be used to add the Management vCenter and the Resources vCenter.

Deployment prerequisites for vRealize Operations Manager

You can create a single node and configure it as a master node or create a master node in a cluster to handle additional data. All vRealize Operations Manager installations require a master node. With a single node cluster, administration and data functions are on the same master node. A multiple-node vRealize Operations Manager cluster contains one master node and one or more nodes for handling additional data.

Prerequisites:

- ESXi 6.7 U2 server is up and running
- AD-DNS and NTP is up and running
- vCenter, vSAN, vRLI is installed, configured, and running
- Manual creation of forward and reverse lookup entries that are completed for all vROps instances on DNS server before deploying them

Deploy vRealize Operations Manager

About this task

The first stage is the deployment of the OVA File as a vRealize Operations Manager. To deploy vROps Manager:

Steps

1. Log in to vCenter using vSphere web client, right-click on the **vCenter server** and select **Deploy OVF Template**.
2. Enter the **Name** and select the **Location** click **Next**.
3. On the **Select a resource** window, select an **ESXi for vROps**.
4. On the **Review details** window, review the entered information then click **Next**.
5. Read and if you agree then accept the license agreement and click **Next**.
6. On the **Select configuration** window select the configuration size based on the size of the environment, then click **Next**.
7. On the **Select storage** window, select **virtual disk format**, **VM storage policy**, and **vsanDatastore store** then click **Next**.
8. In the **Select networks** window, select the appropriate networks for vROps, then click **Next**.
9. On the **Customize template** window, configure the **Networking Properties** for the vROps virtual appliance, and click **Next**.
10. On the **Ready to complete** window, review the configuration data and click **Finish**.


 **NOTE:** Repeat the steps in this section two more times to deploy more nodes.

Configuration of vRealize Operations Manager

About this task


Configure the vROps manager using steps provided in this section.

Steps

1. Go to the FQDN or IP address of the node that will be the master node of vRealize Operations Manager.
2. Log in to **vRealize Operations Manager** and click **New Installation**.
3. From the **Getting started** window, review the information, then click **Next**.
4. Set the **Administrator password** and click **Next**.
5. Choose the appropriate certificate, then click **Next**.
 **NOTE:** If certificates need to be installed, select the **Install a certificate option and browse to the selected file**.
6. On the **Deployment Settings** window, enter the **Cluster Master Node Name** then the **NTP Server FQDN/IP address** for the environment, click **Add**, then click **Next**.
7. Keep the default settings on the **Add Nodes** window, and click **Next**.
8. On the **Ready to complete** window, review the entered information then click **Finish** to complete initial set up.

Add data nodes to scale out vRealize Operations Manager

About this task

 **NOTE:** Add two data nodes using the steps in this section. The second data node works as the replica of the master node.

This section provides steps to add data node in vROps.

Steps

1. Log in to the new node of vRealize Operations Manager, then select **Expand an Existing Installation**.
2. On the **Expand Existing Cluster** window, review the information then, click **Next**.
3. On the **Enter node settings and cluster information** window, enter the name of the node in the **Node name** field, and select the **Data** from the **Node Type** drop-down.
4. In the **Master node IP address or FQDN** field, enter the FQDN or IP address of the master node, then click **VALIDATE** to validate master node connection.
5. Verify the displayed master node certificate information and if correct check the **Accept this certificate** box, and click **Next**.
6. In the **Username and Password** window verify that the vROps administrator username is admin, enter the vROps Manager Admin password, then click **Next**.
7. On the **Ready to Complete** screen, verify the configuration details and click **Finish**.

Add master replica node

About this task

This section provide steps to add master replica node.

Steps

1. Log in to new node of vRealize Operations Manager, then select **Expand an Existing Installation**.
2. On the **Expand Existing Cluster** window, review the information then, click **Next**.

3. On the **Enter node** settings and cluster information window, enter the name of the node in the **Node name** field, and select the **Data** from the. **Node Type** drop-down.
4. In the **Master node IP address or FQDN** field, enter the FQDN or IP address of the master node, then click the **VALIDATE** button to validate master node connection.
5. Verify the displayed master node certificate information and if correct check the **Accept this certificate** box and click Next.
6. On the **Username and Password** window verify that the vROps administrator username is admin, enter the vROps Manager Admin password, then click **Next**.
7. On the **Ready to Complete** screen, verify the configuration details and click **Finish**.

Enable High Availability for clusters

Prerequisites

- Ensure the configuration and operation of the [Add master replica node](#) is complete

About this task

Enable the high availability mode for vROPs clusters.

Steps

1. From a web browser, use your administrator credentials to log in to the vROps Manager GUI.
2. From the **System Status** screen, select the node and click the **Enable** button in the **High Availability** field.




Figure 136. Enable HA settings

3. On the **Enable High Availability** screen, select the **Master node**, click to place a check in the **Enable High Availability** for the cluster box, then click **OK**.
 - NOTE:** While performing Cluster Configuration, if the process stops responding when at the **Waiting for Analytics** screen, perform the following steps:
 - a. Do not reboot any of the vROps nodes or stop the cluster configuration process.
 - b. Make sure NTP server is running and all the vROps nodes are configured to use same NTP server.
 - c. Synchronize time between all the vROps nodes and NTP server.
 - d. Ensure that the time difference between vROps nodes must not be greater than 30 seconds.
 - e. After a few seconds, the cluster configuration proceeds automatically.

Start cluster

About this task

See [vRealize Operations Manager Analytics Cluster](#) if the process fails to start with the status of **Waiting for Analytics**.

Steps

1. On the **System Status** window click the **Start vRealize Operation Manager** button, to start the **vROps Manager**.
2. Click **Yes**. The vROps master node deploys changes to an **Online** state.
3. Open the vROps user interface appliance by entering the UI URL: `https://fqdn or ip of vROps` and log in to the portal with default local user ADMIN.
4. Click **Finish**.

Product license

About this task

Add license to vROps using following steps.

Steps

1. Log in to the vROps Manager with admin credentials.
2. After logging in, you are directed to the **vROps Configuration** page, then click **Next**.
3. Review the information provided within the EULA and if you agree to the terms, check the **I accept the terms of this agreement** check box, then click **Next**.
4. In the **Product Key** field, enter a valid product license key, then click **Next**.
5. Review the information provided in the **Customer Experience Improvement Program** window. To participate in the program, click to place a check in the **Join the VMware Customer Experience Improvement Program** check box to participate, then click **Next**.
6. On the **Ready to complete** window review the selected settings, then click **Finish**.


vROps integration with other components

Activate vCenter, vSAN, and vRLI Management packs

About this task

You are required to activate the vSphere, vSAN, and vRLI Management packs in the vROps GUI in-order to integrate them with vROps. Sometimes, these Management packs are automatically activated and you can view them on Solutions window to configure. If they are not activated and you are not able to view these packs on Solution window then follow the below steps to activate:

Steps

1. In your web browser log in to vROps Manager GUI with admin credentials.
2. Click **Administration, Solutions, Repository**.
3. On the **Repository** screen, activate the required packages.
 **NOTE: For this deployment, VMware vSphere, VMware vSAN, and VMware vRealize Log Insights packages are required to activate.**

Once these packages are activated, you will be able to see their management packs in the Solution window to configure them.


Integrate vROps with VMware vCenter

About this task

You are required to add an adapter in the vROps manager for both management and resource vCenter instances to integrate vRPOs with vCenter server.

-  **NOTE: Perform the steps in this section for both the management and resource clusters.**

Steps

1. Log in to vROps Manager Web GUI with admin credentials.
2. Click **Administration, Solutions**, and then **Configuration**.
3. On the **Solutions** screen, select **vSphere Management solutions**, then click the **Configure** icon.
4. On the **Manage Solution** window, enter the adapter **Display** name, and **Description** in the fields provided.
 **NOTE: To add more adapter click the + (Add) icon in the left navigation panel.**

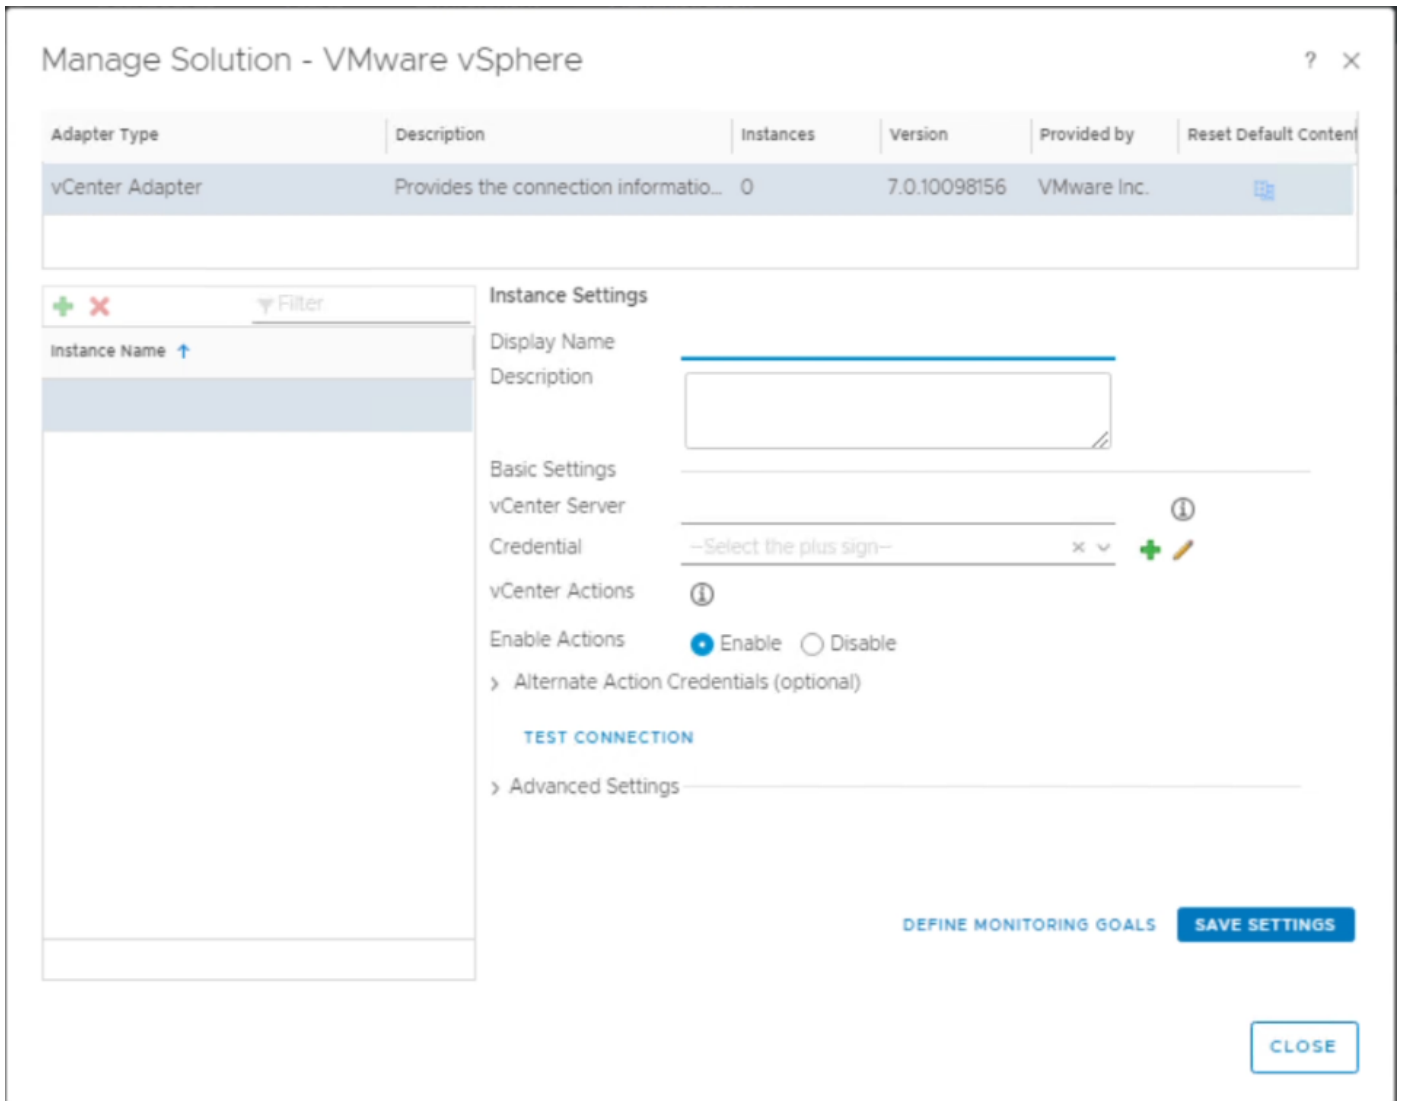


Figure 137. Manage Solution window

5. Click the **Add (+)** icon. The **Manage Credential window** displays.
 6. On the **Manage Credential** window, enter the required **Credential name**, **User Name**, and **Password** in the provided fields, then click **OK**.
 7. On the **Manage Solution** window, click **Test Connection** to initiate the communication with vCenter and match the thumb print and certificate.
Once the test is verify, the **Review and Accept Certificate** screen.
 8. Click **Accept** to acknowledge the certificate. Once you acknowledge the certificate the **Test connection successful** message displays, then click **OK**.
 9. Click **Define Monitoring Goals**, and select the appropriate options and click **Save**.
- NOTE:** Keep the default settings and then click **Save**.

Define Monitoring Goals ✕

Please answer the following list of questions to create a new default policy or Save to modify the existing default policy. To adjust advanced settings of the default policy or create a new policy, proceed to Administration > Policies Page.

Which objects do you want to be alerted on in your environment?

[Learn More](#)

- Infrastructure objects except for Virtual Machines
- Virtual Machines only
- All vSphere objects

Which type of alerts do you want to enable? (Select all that apply)

[Learn More](#)

- Health alerts that usually require immediate attention.
- Risk alerts indicating that you should look into any problems in the near future
- Efficiency alerts indicating that you can reclaim resources.

Enable vSphere Hardening Guide Alerts?

[Learn More](#)

- Yes
- No

CANCEL
SAVE

Figure 138. Define Monitoring Goals

10. On the **Success** window, the default policy has been successfully configured message displays, then click **OK**.
11. Click **Save settings** for the adapter instance to be successfully saved, then click **Close**. The **vRealize Operations Manager** starts to collect information about vCenter and its linked hosts and services.
12. Click the **Home** icon to display information.

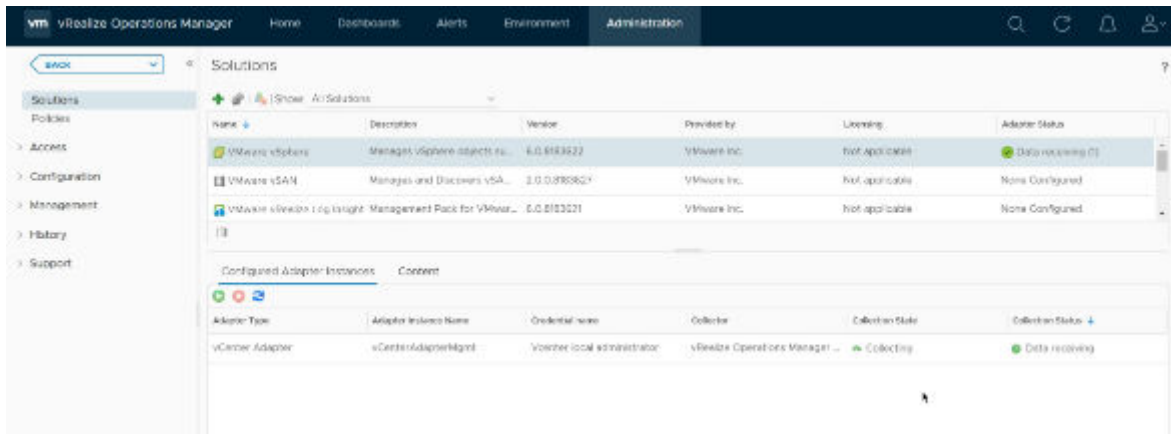


Figure 139. Solutions screen

vROps is deployed, configured and integrated with vCenter Server managing Management Cluster.

NOTE: Repeat the steps in this section to integrate resource vCenter Server with vROps.

vROps integration with AD

About this task

Follow the steps provides in this section to integrate vROps with AD.

Steps

1. Log in to the vRealize Operations Manager user interface with administrator privileges.
2. On the **Administration** screen, then click **Authentication Sources** in the left navigation panel.
3. Click the **Add (+)** to add an Authentication source.
4. On the **Add Source for User Group Import** screen, enter the AD domain details in the required fields, then click **Test** to validate settings.
Once the settings are validated, the **Test connection was successful** message displays.
5. Click **OK** to close the window.
The **Active Directory** is added.

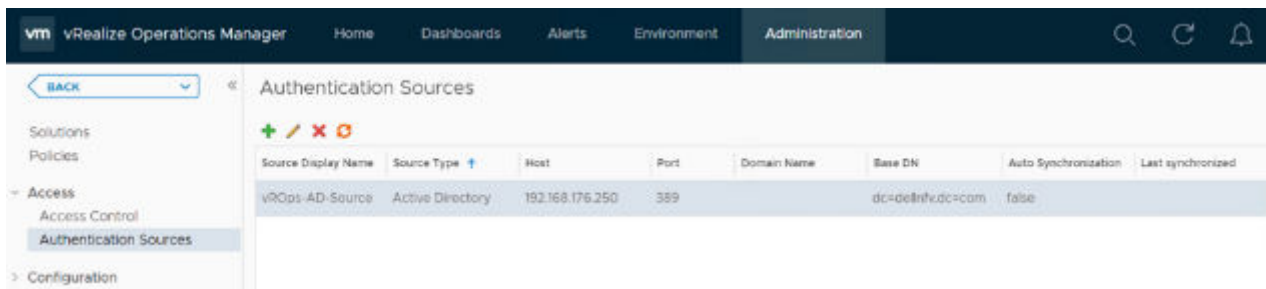


Figure 140. Authentication Sources screen

Importing users to AD

About this task

Once the AD is integrated with vROps import the AD users.

Steps

1. From the **Administration** screen, click **Access Control** in the left navigation panel.
2. Click the **User Accounts** tab, then click the **Import User** icon.

3. In the **Import From** field, select the **AD Domain** previously added.
4. In the **Search string** field, enter a full or partial group name to search for, then click **Search**.
5. On the **Display search results** screen, select a group and click **Next**.
6. Click the **Objects** tab and from the **Select Role** drop-down list, assign a relevant role to the user group.
7. Check the **Allow access to all objects in the system** check box, and click **Finish**.
Users are added in the AD with the necessary permissions.
8. Log in to vROps from one of the accounts imported and verify that the selected permissions are accessible.

NOTE: Users can also be added using the **Importing Users from User Accounts** function.

vROps integration with vRLI

About this task

This section provides steps to integrate vRLI with vROps manager.

Steps

1. Log in to vROps Manager Web GUI with admin credentials.
2. Click **Administration**, **Solutions**, and then click **Configuration**.
3. In the **Solution** window, select **Management pack for VRLI**, and then click the **Configure** icon.
4. In the **Manage Solution** window, enter the **Display name**, and the Log insight server IP then click **Test Connection** to validate.

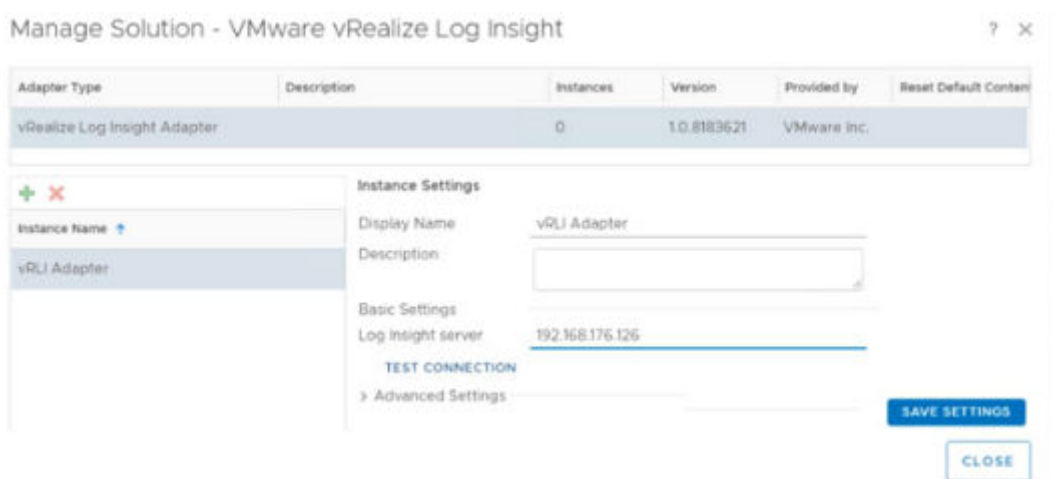


Figure 141. Manage Solution screen

Once the settings are validated, the **Test connection was successful** message displays.

5. Click **OK** to close the window.
6. Click **Save Settings**, then click **Close**. vROps is configured and is able to open Log Insight from vROps.
7. Log in to vRLI to configure the connection to vROps.
8. Click **Administration**, then click **vRealize Operations** and enter the information for vROps.
9. Click **Test Connection** and then **Save**. This allows Log Insight to communicate with vROps.
10. From vROps, click the **Log Insight** tab in the left navigation panel to open **Log Insight from vROps**.
11. To view Log Insight logs inside vROps for a particular Virtual Machine, search for the VM and click to select it.

vROps integration with NSX-T

About this task

vROps Management pack for NSX-T can be download from the VMware Marketplace. This management pack for NSX-T provides inventory and health monitoring for other NSX-T components, such as NSX-T Controller clusters, Edge clusters. Logical routers, transport zones, transport nodes, and load balancers. It also provides a dashboard of NSX-T topology graph, environment overview, and top alerts.

NOTE: Download the Management Pack for NSX-T from the VMware Marketplace. Perform the steps in this section for the NSX-T Manager.

Steps

1. Log in to vROps Manager Web GUI with admin credentials.
2. Click **Administration, Solutions,** and then **Repository**
3. In the **Repository** window, click the **Add (+)** icon in the other **Management Packs** section to upload the PAK file for NSX-T.
4. Click **Browse** and go to the location where the PAK file is located, then click **Upload**.
5. Once the upload is complete, click **Next**.
6. Review the EULA information and if you agree to the terms, click to select the **I accept the terms of this agreement** box then click **Next**.
7. After the solution is installed, click **Finish**.
8. Once the installation is complete, select the **NSX-T management pack** in the **Solution** window and click the **Configure** icon to configure the Solution Adapter instance.
9. On the **Manage Solutions** window, locate the **Instance settings** section, enter the required instance details for the NSX-T Manager.
10. In the **Credential** field, click the **(+) Add** icon. The **Manage Credential** window displays.
11. On the **Manage Credential** window, enter the NSX-T manager credentials then click **OK**.
12. Click the **Test Connection** button.
13. Once the test validation is complete, click **OK**.
14. Click **SAVE SETTINGS** to save the details of the NSX-T Adapter and click **CLOSE**.

vROps integration with vSAN

About this task

By integrating the vROps with vSAN, you can use the provided vROps dashboard to evaluate, manage, and optimize the performance of vSAN objects and vSAN-enabled objects in your vCenter Server system.

Steps

1. Log in to vROps Manager Web GUI with admin credentials.
2. Click **Administration, Solutions, Configuration.**
3. From the **Solutions** window, select **Solution pack for vSAN**, then click the **Configure** icon.
4. In the **Manage Solution** screen, enter the **Display name, Description,** and **vCenter server IP** in the fields provided.

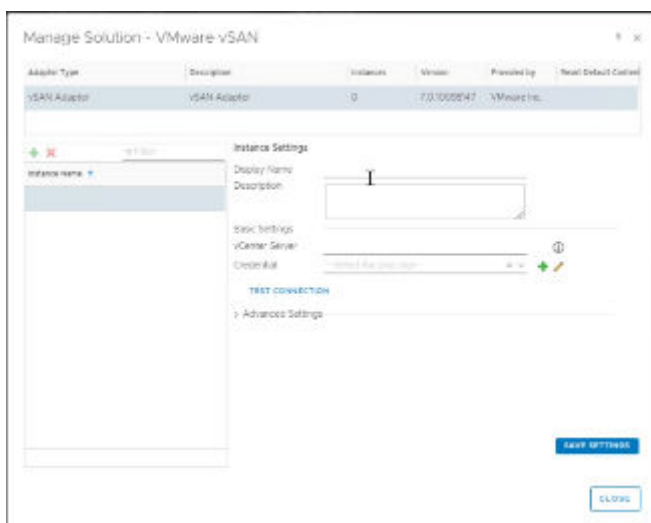



Figure 142. Instance settings

5. In the **Credential** field, click the **(+) Add** icon. The **Manage Credential** window displays.

6. On the **Manage Credential** screen, enter the vCenter credentials then click **OK**.
7. From the **Manage Solution** screen, click the **Test Connection** button to validate.
8. Once the connection is validated, the **Review and Accept Certificate** window displays.
9. Review the certificate details and click **Accept** to accept the certificate.
10. Click **OK** to acknowledge the test results.
11. Click **Save Settings** then click **Close**.
 -  **NOTE: Repeat the steps in this section to add the resource and edge clusters.**
12. Return to the **Home** screen and verify that the VSAN has dedicated dashboard items available on vROps.

vROps integration with vCD

About this task

vROps Management pack for vCloud Director can be download from the VMware Marketplace. vROps uses this pack to monitors the vCloud Director health and sends the early warnings, alerts.

-  **NOTE: Download the Management Pack for vCD from the VMware Marketplace. Perform the steps in this section for the vCD.**

Steps

1. Log in to vROps Manager Web GUI with admin credentials.
2. Click **Administration**, **Solutions**, and then **Repository**.
3. In the **Repository** window, click the **Add (+)** icon in the other **Management Packs** section to upload the PAK file for vCD.
4. Click **Browse** to go to the location where the PAK file is located, then click **Upload**.
5. Once the upload is complete, click **Next**.
6. Review the EULA and if you agree to the terms, click to place a check in the **I accept the terms of this agreement** check box then click **Next**.
7. After the solution is installed, click **Finish**.
8. Once the installation is complete, select the vCD management pack in the Solution window, and click the **Configure** icon to configure the **Solution Adapter** instance.
9. On the **Administration** window, click the **Solutions** tab, select **Solution pack for vCD**, then click the **Configure** icon.
10. From the **Manage Solution** screen, enter the **Display name**, **Description**, and **vCloud Director name/IP for vCD Cell 1** in the fields provided.

Adapter Type	Description	Instances	Version	Provided by	Reset Default Content
vCloud Adapter		0	5.1.0.13477673	VMware Inc.	

+ X Filter

Instance Name ↓

Instance Settings

Display Name

Description

Basic Settings

vCloud Director Host

Auto Discovery ▼

Organization

Filter By Provider VDCs List

Filter By Organizations List

Credential x ▼ + ✎

TEST CONNECTION

SAVE SETTINGS

CLOSE

Figure 143. Manage Solution screen

11. In the **Credential** field, click the **(+)** icon. The **Manage Credential** window displays.
12. On the **Manage Credential** window, locate and click on the **Credential kind** drop-down list, select **vCloud System Credential**, and enter the **vCD Cell1** credentials then click **OK**.
13. On the **Manage Solution** screen, click the **Test Connection** button to validate. Once the connection is validated, the **Review and Accept Certificate** message displays.
14. Review the certificate details and click **Accept** to accept the certificate.
15. Click **OK** to acknowledge the test results.
16. Click **Save Settings**.
17. In the left navigation pane, click the **(+)** icon to add **vCD Cell 2**.
18. Enter the **Display name**, **Description**, and **vCloud Director name/IP for vCD Cell 2** in the fields provided.
19. In the **Credential** field, click the **(+)** icon. The **Manage Credential** window displays.
20. On the **Manage Credential** screen, enter the **vCD Cell2** credentials then click **OK**.
21. On the **Manage Solution** window, click the **Test Connection** button to validate. Once the connection is validated, the **Review and Accept Certificate** message displays.
22. Review the certificate details and click **Accept** to accept the certificate.
23. Click **OK** to acknowledge the test results.
24. Click **Save Settings**.
25. In the left navigation pane, click the **(+)** icon to add **vCD Cell 3**.
26. Enter the **Display name**, **Description**, and **vCloud Director name/IP for vCD Cell 3** in the fields provided.
27. In the **Credential** field, click the **(+)** icon. The **Manage Credential** window display.
28. On the **Manage Credential** window, enter the **vCD Cell 3** credentials then click **OK**.
29. On the **Manage Solution** window, click the **Test Connection** button to validate. Once the connection is validated, the **Review and Accept Certificate** message displays.
30. Review the certificate details and click **Accept** to accept the certificate.
31. Click **OK** to acknowledge the test results.
32. Click **Save Settings**, then click **Close**.

vROps integration with vRealize Orchestrator

About this task

vROps Management pack for vRealize Orchestrator can be download from the VMware Marketplace. vROps uses this pack to monitors the health of vRO and sends the early warnings, alerts.

 **NOTE:** Download the Management Pack for vRO from the VMware Marketplace. Perform the steps in this section for the vRO.

Steps

1. Log in to vROps Manager Web GUI with admin credentials.
2. Click **Administration**, **Solutions**, and then **Repository**.
3. In the **Repository** screen, click the **Add (+)** icon in the other **Management Packs** section to upload the PAK file for vRO.
4. Click **Browse** to go to the location where the PAK file is located, then click **Upload**.
5. Once the upload is complete, click **Next**.
6. Review the EULA and if you agree to the terms, click to place a check in the **I accept the terms of this agreement** box then click **Next**.
7. After the solution is installed, click **Finish**.
8. Once the installation is complete, select the vRO management pack in the **Solution** window and click the **Configure** icon to configure the **Solution Adapter instance**.
9. On the **Administration** screen, click the **Solutions** tab, select management pack for vRO, then click the **Configure** icon.
10. From the **Manage Solution** screen, enter the **Display name**, **Description**, **vRealize Orchestrator name/IP**, and **Port number** in the fields provided.
11. Expand the **Advanced Settings**, then from the **Collectors and Group** drop-down list, select **vRealize Operations Manager Collector-vrops-master** option.

Adapter Type	Description	Instances	Version	Provided by	Reset Default Cont...
vRealize Orchestrator Adapter	vRealize Orchestrator Adapter	1	3.0.13077040	VMware Inc.	

+ X Filter

Instance Name ↑

VRO

Instance Settings

Display Name: VRO

Description:

Basic Settings

vRealize Orchestrator Host: vro.dellinfv.dellinfv.com ⓘ

Port: 8281

Auto Discovery: true ▾

Credential: x ▾ + ✎

TEST CONNECTION

Advanced Settings

Collectors/Groups: vRealize Operations Manage ▾ ⓘ

SAVE SETTINGS

Figure 144. Manage Solution screen

12. In the **Credential** field, click the **(+)** icon. The **Manage Credential** window displays.
13. On the **Manage Credential** window, locate and click on the **Credential kind** drop-down list, select **vRO System Credential**, and enter the vRO credentials then click **OK**.
14. On the **Manage Solution** window, click the **Test Connection** button to validate. Once the connection is validated, the **Review and Accept Certificate** message displays.
15. Review the certificate details and click **Accept** to accept the certificate.
16. Click **OK** to acknowledge the test results.
17. Click **Save Settings**, then click **Close**.

vSphere Replication

Prerequisites

- Download the vSphere Replication iso image file and mount it on the system
- Management pod must be up and running
- DNS entry for vSphere Replication must be added in the DNS server.

About this task

vSphere Replication is an alternative to storage-based replication. It protects virtual machines from partial or complete site failures by replicating the virtual machines between the following sites:

- From a source site to a target site
- Within a single site from one cluster to another
- From multiple source sites to a shared remote target site

Steps

1. Log in to the Management vCenter using the VMware vSphere Web Client.
2. Right-click on the **Management Datacenter**, then select **Deploy OVF Template**.
3. On the **Select template** screen, enter the download URL or click Browse to locate the .OVA following files on your computer, then click **Next**.
 - vSphere_Replication_OVF10.cert
 - vSphere_Replication_OVF10.mf
 - vSphere_Replication_OVF10.ovf
 - vSphere_Replication_support.vmdk
 - vSphere_Replication_system.vmdk
4. On the **Select name and location** window, enter the **Name** and select the **Location to deploy vSphere Replication** then click **Next**.
5. On the **Select a resource** screen, select the ESXi to deploy vSphere Replication, and click **Next**.
6. On the **Review details** screen, review the settings selected then click **Next**.
7. Use the scroll bar to review the information in the **Accept license agreement** screen and if you agree, then click **Accept**, and click **Next**.
8. On the **Select configuration** screen, select the type of deployment configuration from the **Configuration** drop-down list and click **Next**.
For this deployment, 4vCPU is used.
9. On the **Select storage** screen:
 - a. From the **Select virtual disk format** drop-down list, select **Thin provision**.
 - b. From the **VM storage policy** drop-down list, select **vSAN Default Storage Policy**.
 - c. Select the **vSAN datastore** and click **Next**.
10. In the **Select networks** screen:
 - a. Select the appropriate networks. For this deployment, VM-Mgmt-Network is used.
 - b. From the **IP allocation** drop-down list, select **Static - Manual**.
 - c. Click **Next**.
11. On the **Customize template** screen, in the **Application** section fill the following fields:
 - a. In the **Password** field, set the root account password.
 - b. In the **NTP Servers** field, enter the NTP server IP address.
 - c. In the **Hostname** field, enter the host name for vSphere Replication VM.
12. On the **Customize template** screen, in the **Networking Properties** section, fill the following fields:
 - a. **Default Gateway**: Enter the IP address of default gateway for vSphere Replication.
 - b. **Domain Name**: Enter the domain name for vSphere Replication.

- c. **Domain Search Path:** Enter the domain search path for vSphere Replication.
- d. **Domain Name Servers:** Enter the DNS IP address.
- e. **Management Network IP Address:** Enter the IP address for Management network.
- f. **Management Network Netmask:** Enter the netmask IP for Management network.

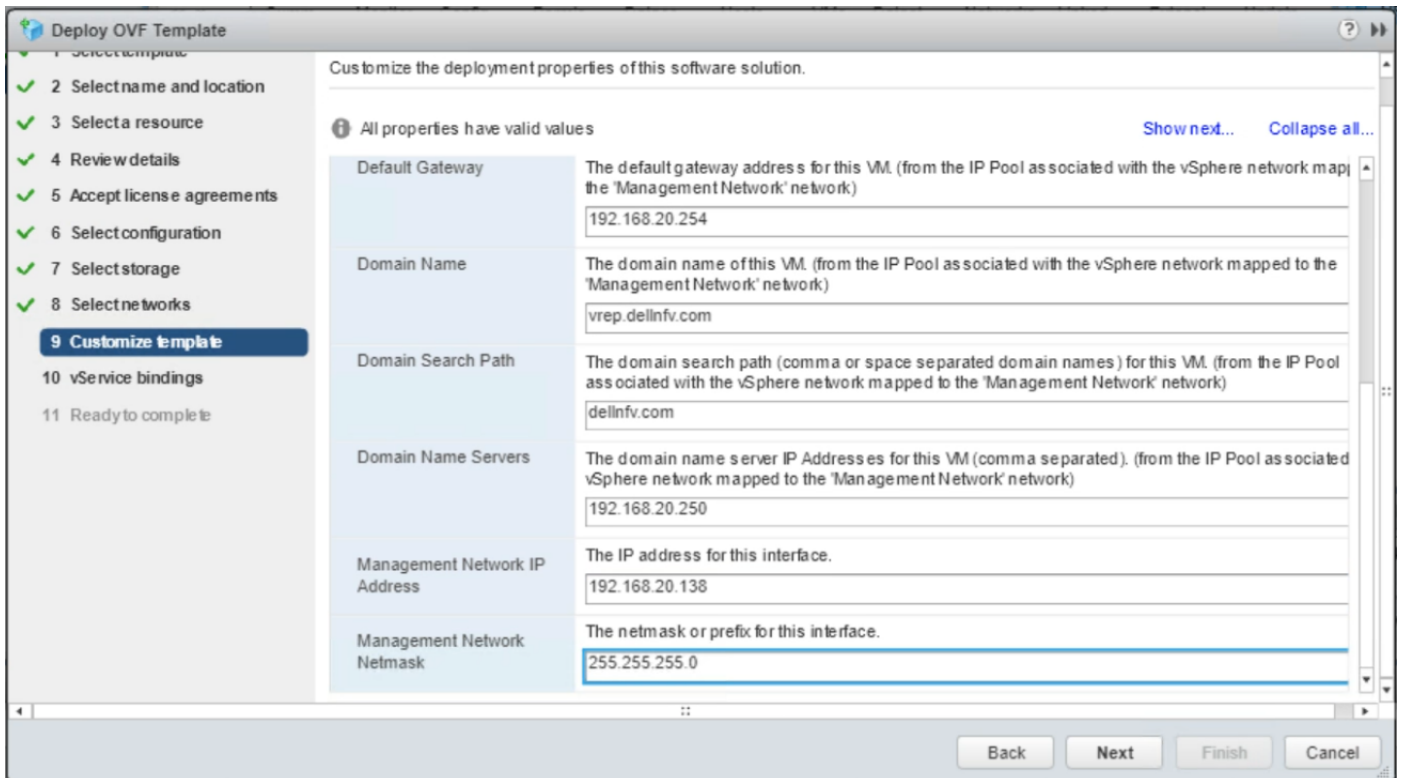


Figure 145. Customize Template screen

13. Click **Next**.
14. On the **vService bindings** screen, verify the binding status is up, and click **Next**.
15. On the **Read to complete** screen, review the provided configuration details then click **Finish** to deploy vSphere Replication VM.
16. Once the deployment is completed, on the select the **vSphere Replication VM**.
17. In the right navigation panel, locate the **Actions** drop-down list, then select **Edit Settings**.
18. On the **Edit Settings** window, add a network adapter for **Replication Network**:
 - a. From the **New device** drop-down list, select **Network** then click **Add**.
 - b. From the **New Network** drop-down list, select the **Replication Network**.
 - c. Check the **Connect** box.
19. Click **OK** and then **Power-on the VM**.

Configuring vSphere Replication

About this task

Configure the vSphere Replication using steps provided in this section.

Steps

1. Log with the root credentials in to vSphere Replication Web Interface at: <https://<<vSphere-Replication-FQDN/IP>>:5480>
2. Click **Network**, and then **Address**.
3. In the **eh1 info** section, from the **IPv4 Address Type** drop-down list select **Static**.
4. In the **IPv4 Address** field, enter the **vSphere Replication VM IP address**.
5. In the **Netmask** field enter the **netmask IP** for vSphere Replication VM.
6. Click **Save Settings**.

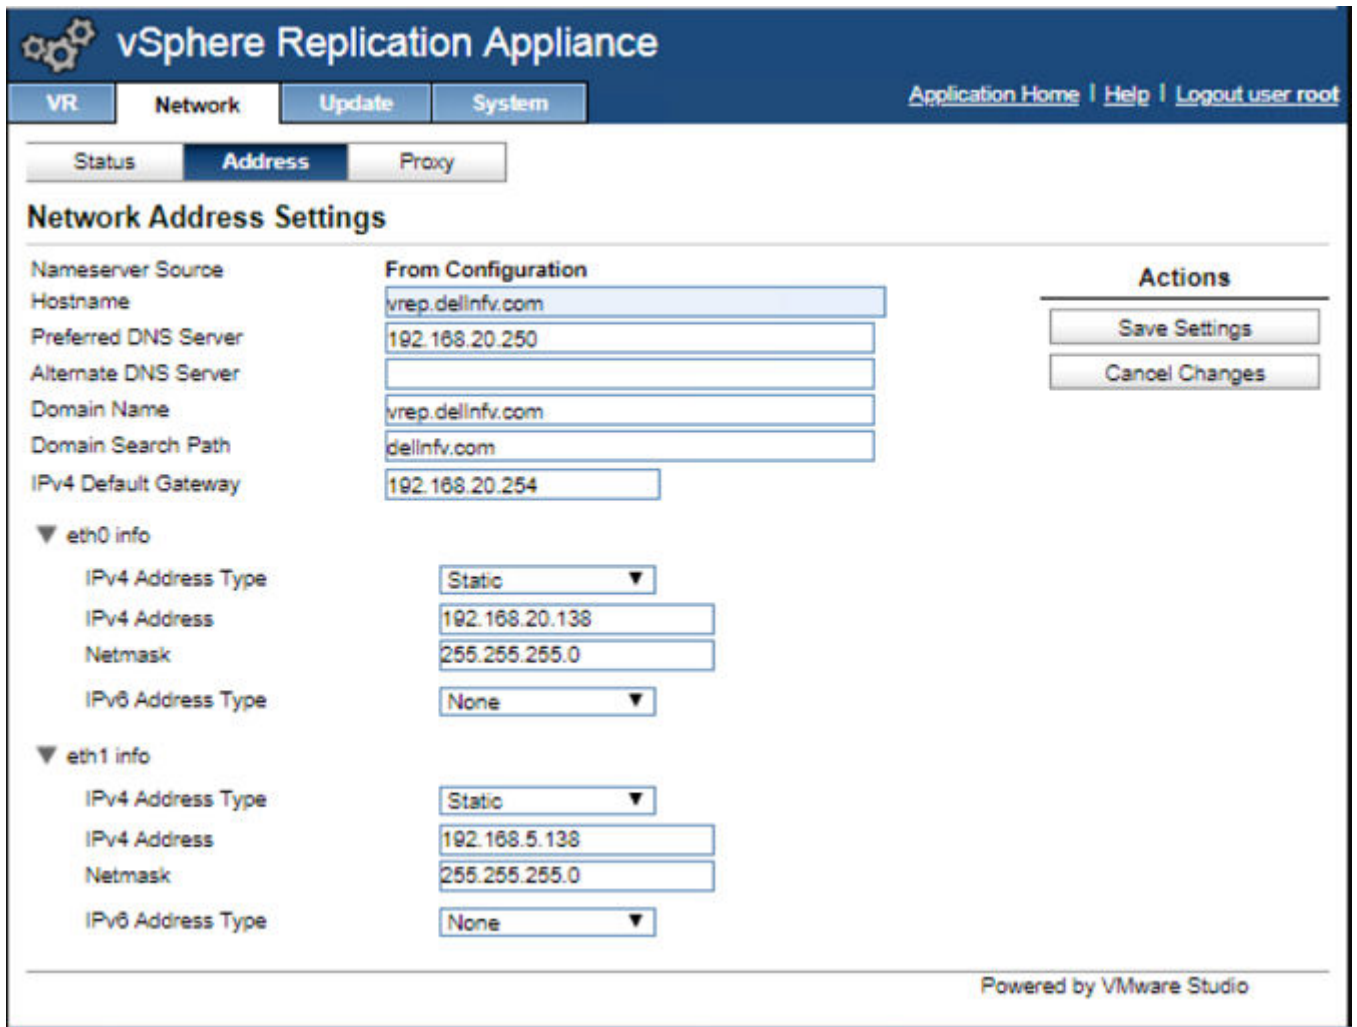


Figure 146. Network Address screen

7. Click **VR**, and then **Configuration**.
8. In the **IP Address for Incoming Storage Traffic** field, enter the IP address of network adapter used by Replication Server for incoming replication data.
9. Click **Network Setting**.
The **VR network settings changed successfully** message displays.

vSphere Replication Appliance

Application Home | Help | Logout user root

VR | Network | Update | System

Getting Started | **Configuration** | Security | Support

Startup Configuration

VR network settings changed successfully

Configuration Mode:

- Configure using the embedded database
- Manual configuration
- Configure from an existing VRM database

Actions

- Save and Restart Service
- Unregister VRMS
- Reset Embedded Database

LookupService Address:

SSO Administrator:

Password:

VRM Host:

VRM Site Name:

vCenter Server Address:

vCenter Server Port:

vCenter Server Admin Mail:

IP Address for Incoming Storage Traffic:

SSL Certificate Policy

Accept only SSL certificates signed by a trusted Certificate Authority
(You must click the 'Save and Restart Service' button after changing this setting)

Install a new SSL Certificate

Generate a self-signed certificate

Upload PKCS#12 (*.pfx) file No file chosen

Service Status

Figure 147. Target site screen

10. In the **LookupService Address** field, enter the FQDN of the vCenter Server.
11. Enter the Management vCenter SSO administrator username in the **SSO Administrator** field.
12. In the **Password** field, enter the SSO administrator password.
13. Click **Save** and **Restart Service**.
14. Review the SSL Certificate then click **Accept**.
15. Log-out and Log-in from vSphere Replication and Management vCenter to save the changes.

Configure vSphere Replication connection

About this task

This section provides steps to configure vSphere Replication for Site Recovery.

Steps

1. Log in to the Management vCenter using the VMware vSphere Web Client.
2. In the **Navigator tree**, select the **vSphere Replication VM** and click **Home**, and then **Site Recovery**.
3. Verify that the **vSphere Replication** displays the **OK** status.

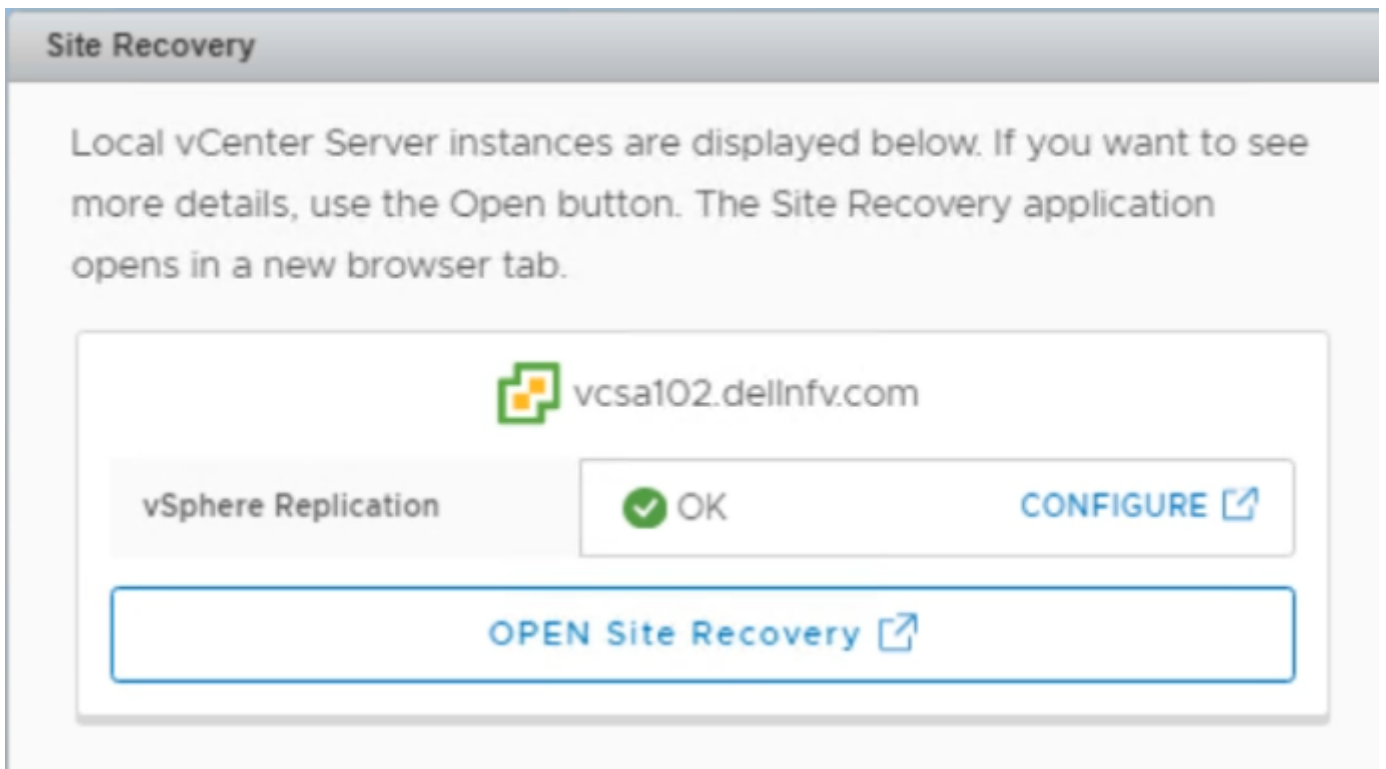


Figure 148. vSphere Replication status

4. Click **Open Site Recovery**.
The **Site Recovery** opens in new tab.
5. On the **Site Recovery** page, from the **Menu** drop-down list, select **Replications** within the same vCenter Server then click the **FQDN/Host name of the vCenter**.
6. On the **Replications** tab, click **+ New**.
7. On the **Virtual machines** screen, check the box of each virtual machines that you want to protect, and click **Next**.
8. On the **Target site** screen:
 - a. On the **Select the vSphere Replication server that will handle the replication**, select **Manually select vSphere Replication Server** radio button.
 - b. Select the **vSphere Replication VM**.
 - c. Click **Next**.

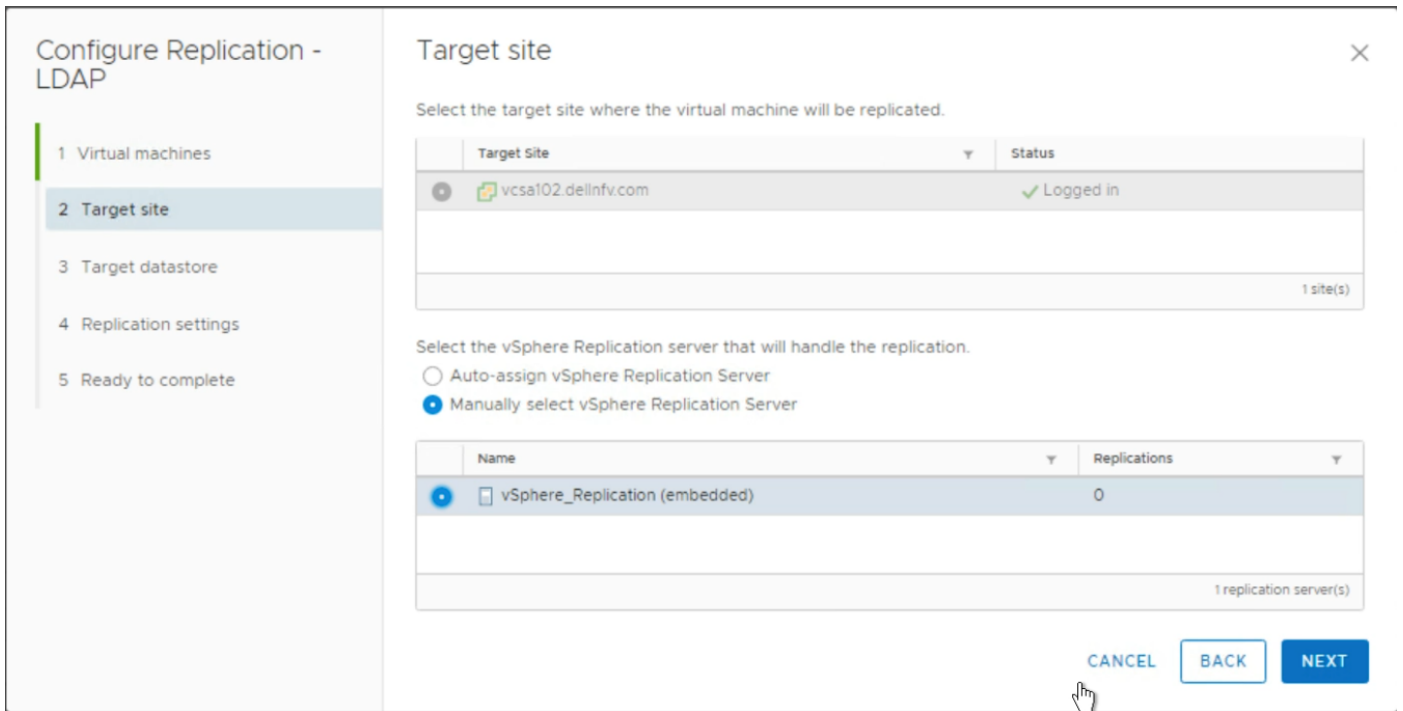


Figure 149. Target site screen

9. On the **Target datastore** screen:
 - a. From the **Disk format** drop-down list select **Same as source**.
 - b. From the **VM storage policy** list, select **vSAN Default Storage Policy**.
 - c. Select the **vSAN datastore**.
 - d. Click **Next**.
10. On the **Replication settings** screen:
 - a. In the **Recovery point Object (RPO)** field, use the slider to set the acceptable period for which data can be lost in the case of a site failure.

NOTE: A RPO determines the maximum data loss that you can tolerate. For example, an RPO of 1 hour seeks to ensure that a virtual machine loses the data for no more than 1 hour during the recovery. For more information, refer the [VMware vSphere Replication Documentation](#).
 - b. Check the **Enable point in time instances** box to save the multiple replication instances that can be converted to snapshots of source VM during recovery.
 - c. Then, enter the number of instance to keep and validity to store this instances.
 - d. Click **Next**.
11. On the **Ready to complete** screen, review the provided settings and click **Finish**. Once the configuration is complete, the status changed to **OK**.




Figure 150. Status updated to OK

Set up anti-affinity rules

The Affinity rule setting establishes a relationship between two or more VMware VMs and hosts. Anti-affinity rules tell the vSphere hypervisor platform to keep virtual entities separated.

Anti-affinity rules are created to keep two or more VMs separated on different hosts.

NOTE: The anti-affinity rules can be created for management and edge cluster VMs. The creation of a separate anti-affinity rule for the VMs listed is recommended for the respective cluster. Dell EMC recommends the execution of this section only after every component has been deployed and configured.

To create anti-affinity rules for management and edge pod components, see the *Create an anti-affinity rule* section.

The [Management VM cluster list](#) table displays a list of management cluster VMs that must be kept on different hosts using an anti-affinity rule. For example, create an anti-affinity rule for management VCSA to always keep its three VMs (Active, passive and witness node VM) on different hosts.

Table 40. Management VM cluster list

Rule name	Create rule for virtual machines		
Mgmt VCSA	VCSA-Mgmt-Active	VCSA-Mgmt-Passive	VCSA-Mgmt-witness
Resource VCSA	VCSA-Res-Active	VCSA-Res-Passive	VCSA-Res-witness
vRLI_Rule	vRLI-master	vRLI-worker1	vRLI-worker2
vROPS_Rule	vROPS-master	vROPS-data	vROPS-replica
NSX_Rule	Nsx-Manager	Nsx-1	Nsx-2
vCD_Rule	vCD_Cell1	vCD_Cell2	vCD_Cell3

The following is a list of edge cluster VMs that must be kept on different hosts using an anti-affinity rule:

Table 41. Edge cluster list

Rule name	Create rule for VMs	
Edge_Rule01	Edge01	Edge02
Edge_Rule02	Edge03	Edge04

Prerequisites:

- All components of the management cluster are deployed
- All components of the resource cluster are deployed

Create an anti-affinity rule

Prerequisites

All the management VM should be deployed

About this task

Create anti-affinity rules using steps provided in this section.

Steps

1. Go to the VMware vCenter GUI, click the desired cluster, such as ManagementCluster, then select the **Configure** tab.
2. From the **Configuration** window, locate the **VM/Host Rules** and click the **Add** button.

3. In the **Create VM/Host Rule** dialog box, enter a **Name** for the rule, for example Mgmt VCSA, and check the **Enable rule** box.
 4. From the **Type** drop-down menu, select **Separate Virtual Machines**, then click **Add**.
 5. From the **Add Rule Member** window, select the virtual machines to keep on different hosts then click **OK**.
 6. Click **OK** to create the rule.
- NOTE:** Repeat the steps provided in this section to create the anti-affinity rule for the management, resource, and edge cluster VMs.

Enable vSphere DRS

Prerequisites

Make sure that anti-affinity rules are created for VMs

About this task

This section provides steps to enable vSphere DRS on Management cluster, Resource cluster, and Edge cluster.

Steps

1. On the VMware vSphere Web Client, go to the Management cluster, click the **Configure** tab, **Services**, **vSphere DRS**, and then click **Edit**.
2. On the **Edit cluster** settings window, check the **Turn ON vSphere DRS** box.
3. Set the **DRS Automation to Fully Automated**.
4. Set the **Power Management to Off**, and then set the **Advanced Options to None**.
5. Click **OK**.
6. Verify that the **vSphere DRS** shows the **Turned ON** status.




Figure 151. vSphere DRS status window

7. Repeat the steps provided within this section on the edge and resource clusters to enable the DRS.

Enabling vSphere availability

About this task

Perform the following steps on management, edge and resource clusters to enable vSphere availability:

Steps

1. In the VMware vSphere Web Client, go to the **management cluster**.
2. Click the **Configure** tab, then go to **Services**, click **vSphere Availability**, then click the **Edit** button.
3. Click to place a check in the **Turn ON vSphere HA** and **Turn on Proactive HA** boxes.
4. In the **Failures and Responses** section and select **Enable Host Monitoring**, then click **OK**.
5. Repeat the steps in this section, on the resource and edge cluster.

Forwarding logs to vRLI

Forwarding vROps log to vRLI

Prerequisites

- Configured vRLI server to receive the logs

About this task

You can configure vROps VM to send log messages to a remote logging server.

Steps

- Log in to the vROps master VM.
- From the top navigation panel, click **Administrator** and select the **Management** tab.
- Click **Log forwarding**.
- Select the box: **output logs to external log server**.
- Select the other for Log Insight Servers:
 - In Host give FQDN of vrlI master
 - Port:9000
 - Protocol: cfapi

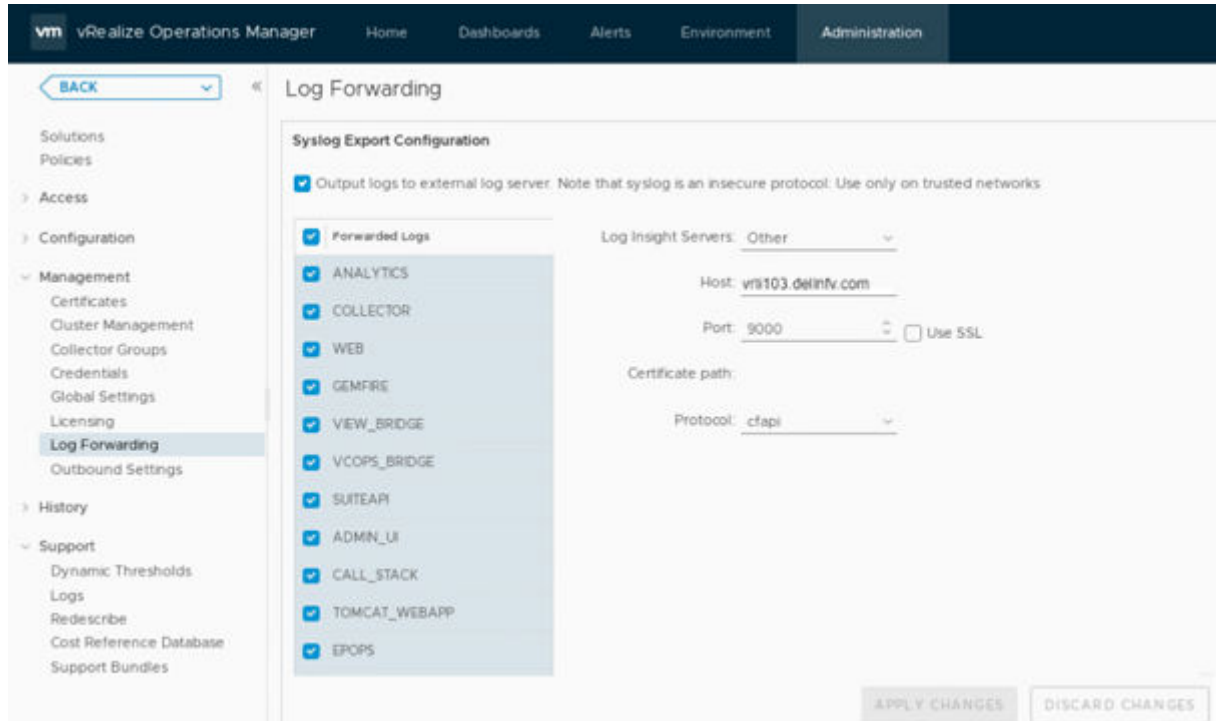


Figure 152. Log forwarding

Forwarding vSAN logs to vRLI

About this task

This section provides steps to forward logs from vSAN to vRLI.

Steps

1. Open the vRealize Log Insight UI in a browser.
2. From the vRealize Log Insight UI, click **Admin** then click **Content Packs**.
3. Click **Interactive Analytics**.
4. Click the **Bell** icon and select **Manage Alerts**.
5. Select the alerts that are vSAN related. In the search box of the **Alerts** dialog box, enter vSAN as a search phrase.

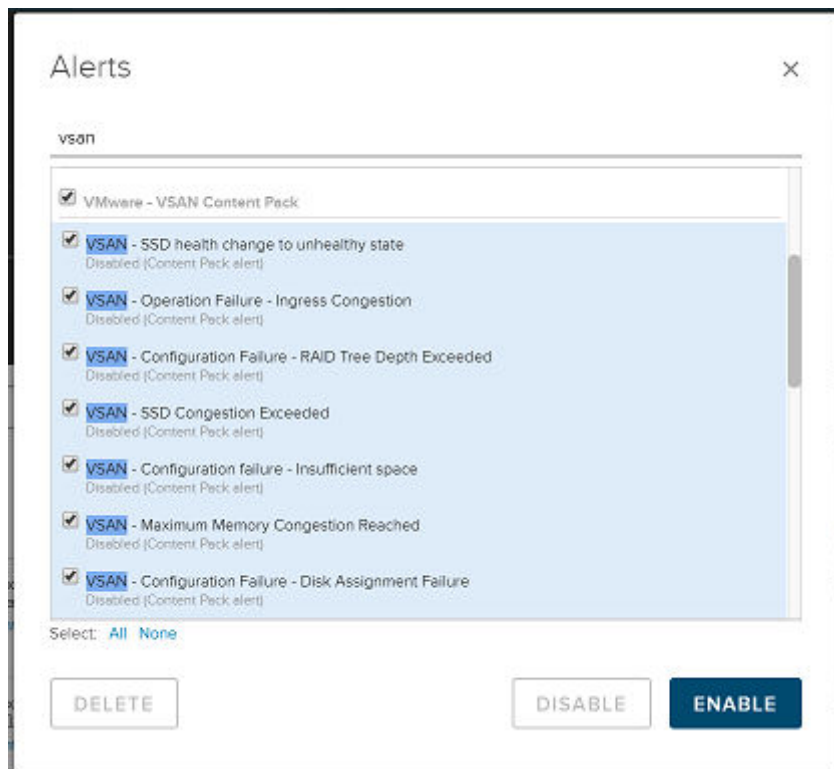


Figure 153. Alerts

6. Select the alerts VMware-vSAN content pack box and then click **Enable**.
7. In the **Enable Alerts** dialog box provide **Email ID**, **fallback object**, and **Criticality**.

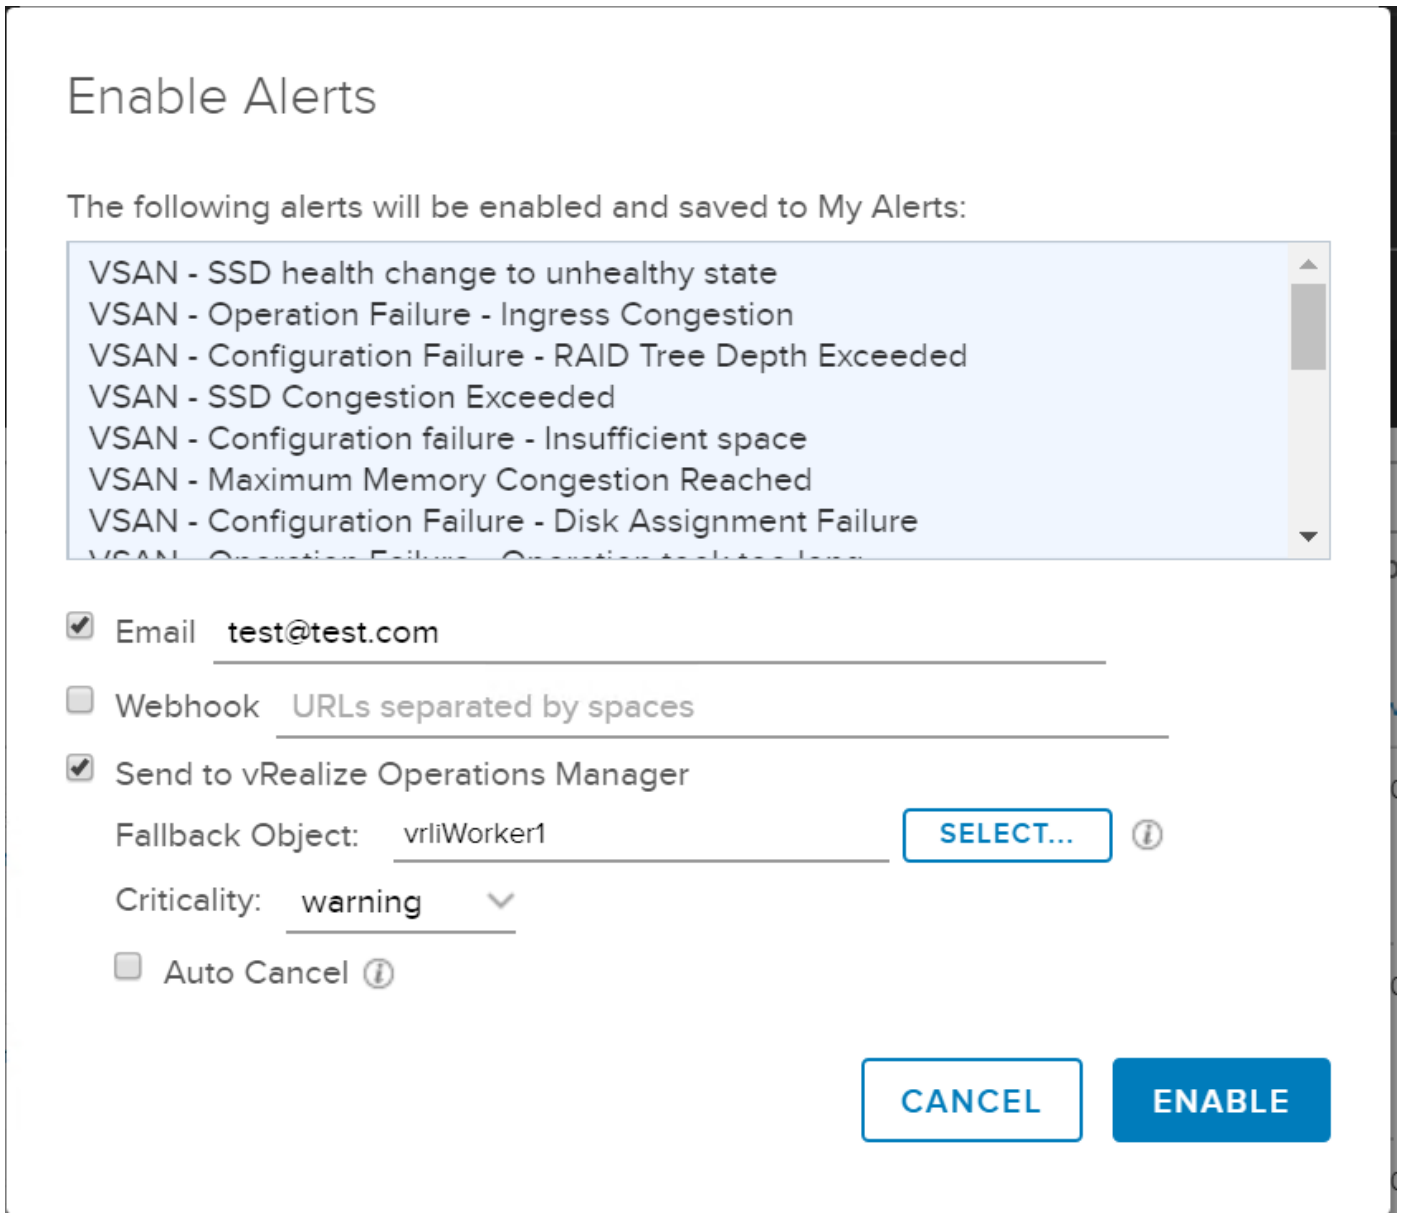


Figure 154. Enable alerts

8. Click **Enable**.

Forwarding logs from vCD to vRLI

About this task

This section provides the steps to forward logs from vCloud Director to vRLI.

Steps

1. Open your browser and log in to vCD Cell1 with administrator credentials.
2. From the **Administrator** screen, locate the **System Settings** section and click **General**.
3. Enter the **vRLI Master node IP** in **Syslog server 1**.
4. Enter the **vRLI virtual IP** in **Syslog server 2**.
5. Click **Apply** to save the change.

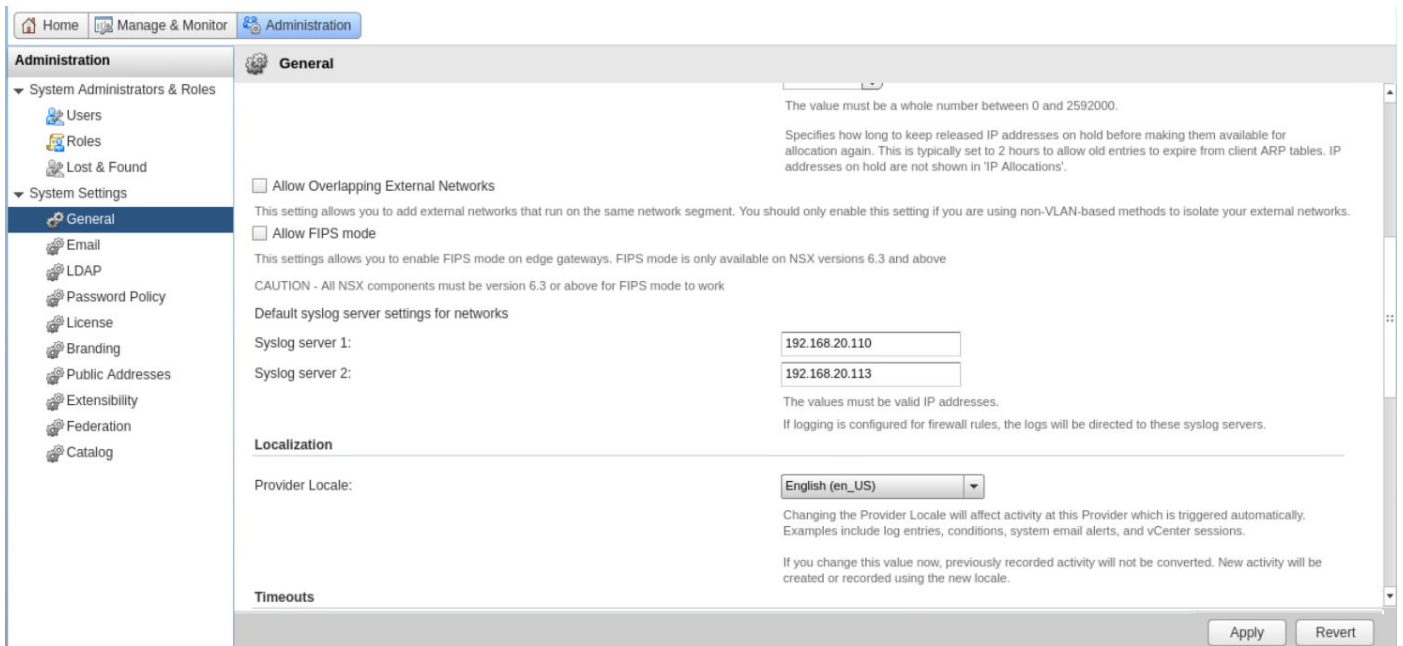


Figure 155. Syslog server for vCD

Configure syslog server for NSX-T

Prerequisites

- Configured vRLI server to receive the logs

About this task

You can configure NSX-T appliances to send log messages to a remote logging server.

Steps

- From the CLI, log in to the NSX-T – Manager using admin credentials.
- Run the following command to configure a log server and the types of messages to send to the log server.

NOTE: Multiple facilities or message IDs can be specified as a comma delimited list, without spaces. See the [Log message IDs table](#) for a list of available log message IDs.

```
set logging-server <hostname-or-ip-address[:port]> proto <proto> level <level> [facility <facility>] [messageid <messageid>] [certificate <filename>] [structured-data <structureddata>]
```

NOTE: For more information about this command, see the [NSX-T CLI Reference section](#).

The command can be run multiple times to add multiple logging server configurations. For example:

```
nsx> set logging-server 192.168.20.113 proto udp level info facility syslog messageid SYSTEM,FABRIC structured-data audit=true
nsx> set logging-server 192.168.20.113 proto udp level info facility auth,user
```

- To view the logging configuration, enter the get logging-server command. For example:

```
nsx-manager> get logging-servers
192.168.20.113 proto udp level info facility syslog messageid SYSTEM,FIREWALL structured-data audit="true"
192.168.20.113 proto udp level info facility syslog messageid SYSTEM,MONITORING structured-data audit="true"
192.168.20.113 proto udp level info facility syslog messageid SYSTEM,DHCP structured-data audit="true"
192.168.20.113 proto udp level info facility syslog messageid SYSTEM,ROUTING structured-
```

```

data audit="true"
192.168.20.113 proto udp level info facility syslog messageid SYSTEM,SWITCHING structured-
data audit="true"
192.168.20.113 proto udp level info facility syslog messageid SYSTEM,FIREWALL-PKTLOG
structured-data audit="true"
192.168.20.113 proto udp level info facility syslog messageid SYSTEM,- structured-data
audit="true"
192.168.20.113 proto udp level info facility syslog messageid SYSTEM,SYSTEM structured-
data audit="true"
192.168.20.113 proto udp level info facility syslog messageid SYSTEM,GROUPING structured-
data audit="true"

```

NOTE: Repeat the steps in this section for the NSX-T Controller and NSX-T Edge to configure remote logging on each node individually.

Log Message IDs

About this task

In a log message, the message ID field identifies the type of message. You can use the messageid parameter in the set logging-server command to filter which log messages are sent to a logging server. Table 40 displays the list of available log Message IDs.

Message ID	Examples
FABRIC	<ul style="list-style-type: none"> Host node Host preparation Edge node Transport zone Transport node Uplink profiles Cluster profiles Edge cluster Bridge clusters and endpoints
SWITCHING	<ul style="list-style-type: none"> Logical switch Logical switch ports Switching profiles switch security features
ROUTING	<ul style="list-style-type: none"> Logical router Logical router ports Static routing Dynamic routing
FIREWALL	<ul style="list-style-type: none"> Firewall rules Firewall connection logs
FIREWALL-PKTLOG	<ul style="list-style-type: none"> Firewall connection logs Firewall packet logs
GROUPING IP sets	<ul style="list-style-type: none"> Mac sets NSGroups NSServices NSService groups VNI Pool IP Pool
DHCP	<ul style="list-style-type: none"> DHCP relay
SYSTEM	<ul style="list-style-type: none"> Appliance management (remote syslog, ntp, etc)

Message ID**Examples**

MONITORING

- Cluster management
- Trust management
- Licensing
- User and roles
- Task management
- Install (NSX-T Manager, NSX-T Controller)
- Upgrade (NSX-T Manager, NSX-T Controller, NSX-T Edge and host-packages upgrades)
- Realization
- Tags

- SNMP
- Port connection
- Traceflow
- All other log messages

Reference documentation

For additional system information, see the following documents:

[Dell EMC Ready Solution vCloudNFV3.0 OpenStack - Deployment Manual Operations Guide](#)

[Dell EMC PowerEdge R640 Installation and Service Manual](#)

[Dell EMC PowerEdge R740 Installation and Service Manual](#)

[Dell EMC PowerEdge R740xd Owner's Manual](#)

[VMware vCloud NFV Reference Architecture v3.0](#)

[VMware API Reference Guide](#)

[NSX-T Installation Guide](#)

[Postman Documentation](#)

[vCloud Director 9.7 Documentation](#)