



Hewlett Packard
Enterprise

VMware Cloud Foundation (VCF) Certification on HPE Synergy

Best practices and general build information for VCF
2.3.2 use on HPE Synergy platform

Contents

Introduction and Executive Summary.....	3
Physical and Logical Layout.....	4
Solutions Components.....	4
Resources for Sizing and Use of VCF on Synergy.....	6
Key Considerations and Prerequisites.....	6
Pre-Prerequisites.....	6
HPE OneView Configuration through HPE Synergy Composer.....	7
Configuring Networks (VLANs).....	7
Configuring HPE OneView Network Sets (VLAN sets).....	9
Configuring HPE OneView Logical Interconnect Groups.....	10
Configuring HPE OneView Enclosure Group.....	16
Configuring HPE OneView Logical Enclosure.....	18
Validating HPE OneView Logical Interconnects.....	20
Configuring HPE OneView Server Profile Templates.....	21
Configuring HPE OneView Server Profiles.....	26
VMware Cloud Foundation on HPE Synergy.....	27
Deploying the VCF Imaging Server and Jump VM on Management ESXi Server.....	27
Management vSphere Standard Switch Configuration.....	27
Deploying the VMware Cloud Foundation Imaging Appliance.....	28
Configuring VCF Imaging Server.....	30
Initializing VMware SDDC Manager.....	33
VCF Bring-up Steps.....	33
Troubleshooting.....	53
Appendix – Important Links.....	56

Introduction and Executive Summary

This document is an installation guide for VMware VCF installations on the HPE Synergy hardware platform.

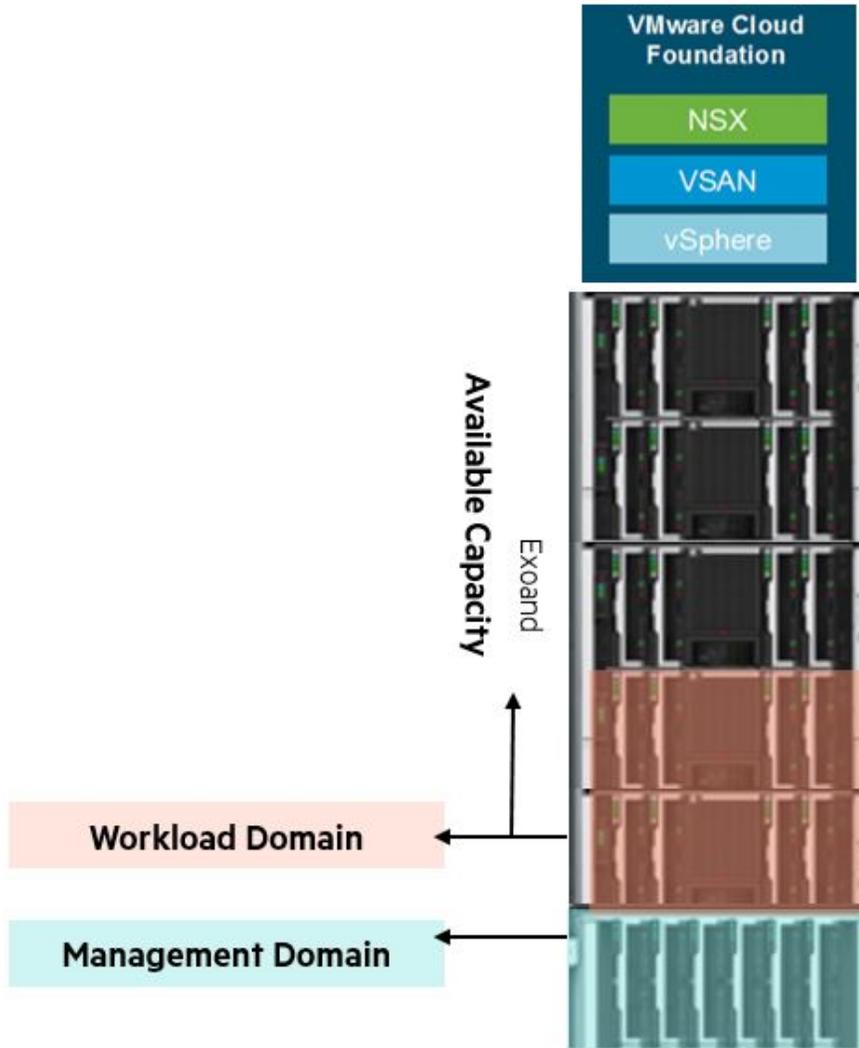
The intended users for this content are HPE, VMware, and their channel partners who are interested in setting up a lab or doing services implementations. This document is not intended to be a training guide for OneView, Synergy, or VCF. Prior knowledge of these applications and hardware is expected to understand the content herein.

VMware Cloud Foundation is an integrated hybrid cloud platform that delivers a complete set of software defined services for compute, storage, networking, security, and cloud management for the private and public cloud. Cloud Foundation drastically simplifies data center operations by deploying a standardized and validated architecture with built in lifecycle automation of the cloud stack. Cloud Foundation can also be flexibly consumed as-a-service in the public cloud (VMware Cloud on AWS, IBM, Rackspace, etc), enabling a true hybrid cloud that is based on a consistent infrastructure and operational model using common tools and processes.

HPE Synergy is a single infrastructure of pools of compute, storage, and fabric resources, along with a single management interface that allows IT to rapidly assemble, disassemble and re-assemble resources in any configuration. HPE Synergy architecture eliminates hardware and operational complexity so IT can deliver infrastructure to applications faster and with greater precision and flexibility. HPE OneView is your infrastructure automation engine built with software intelligence. It streamlines provisioning and lifecycle management across compute, storage and fabric resources in the Synergy System.

Physical and Logical Layout

Example with 3 Frames, Management Domain, and separate Workload Domain



This example can be scaled from one Synergy Frame to multiple Synergy Frames spread across multiple racks.

Solutions Components

Software components: (Starting reference during publication, Refer to HCL for current list)

Software	Version
HPE OneView for Synergy	4.00.07.02
HPE Synergy Firmware Bundle (SPP)	2018.03.00
HPE P416mi-e Firmware	1.34
HPE P416mi-e Driver	smartpqi-1.0.1.254-1OEM.650.0.0.4598673.x86_64

VMware Cloud Foundation	2.3.2
VMware ESXi Server	February 2018 HPE Custom ISO

Hardware components initial configuration:

Hardware	Detail
HPE Synergy Frame with Fabric	1-3 Frames with 2 VC SE 40Gb F8 Modules and 20 GB satellite
HPE Synergy Storage Module	D3940 (multiple depending on sizing guide)
HPE Synergy Compute Model	4x Management nodes 3 or more Production nodes HPE SY480/SY660 Gen10
HPE Synergy Compute	2x CPU - Intel Xeon, 256GB Memory each min
HPE Synergy Storage	HPE Smart Array P416i-e
HPE Synergy Network Options	Synergy 3820C 10/20Gb CNA

Resources for Sizing and Use of VCF on Synergy

- Compute, memory and storage must be vSAN certified and part of the Synergy vSAN ReadyNodes. Please work with your Hewlett Packard Enterprise sales representative to build the order by following the prescribed configurations on the Ready Nodes document in the VMware hardware compatibility list. You can access the Simple Configurator at: vSANreadynode.vmware.com/RN/RN and follow the steps outlined. For OEM: choose "Hewlett Packard Enterprise," select the model, then follow "Next steps."
- See resources below:[HPE Best Practices Guide for vSAN on Synergy](#)
- [VMware hardware compatibility list \(HCL\)](#)
- [VMware® vSAN™ Design and Sizing Guide 6.5](#)
- [VM Sizing Tool from VMWare](#)

Key Considerations and Prerequisites

Please observe the following considerations to make sure a customer is within the correct parameters.

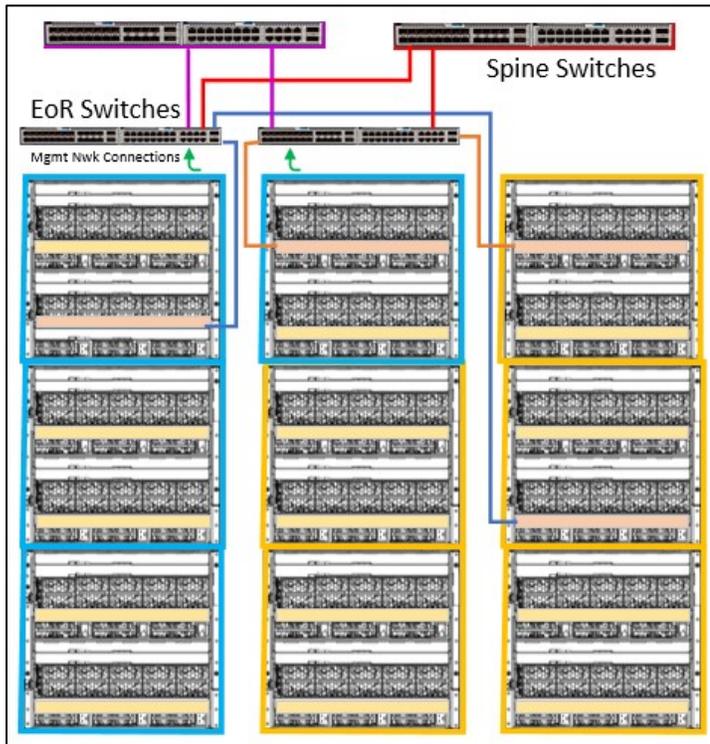
1. VMware vSAN HCL must be used as a reference for required driver and firmware levels
 - a. Check the driver/firmware combinations against what is currently being used in the latest Synergy SPP
 - b. If firmware required for HCL is not available in latest SPP then download driver from HPE support and upload to the Composer via method described in [Managing Composable Infrastructure guide](#)
2. Cache and Data drive sizes dictated by VM sizing done prior to purchasing (no set "only use these disks" in VCF)
3. Physical layout of frames and racks depending on HA and VM sizing (local to D3940) with specific drives
4. All nodes must have equivalent configurations of memory and vSAN sizing in the same cluster

Pre-Requisites

This section describes required configurations, and access credentials that must be known, prior to initiating the VCF build.

To facilitate the deployment of VMware Cloud Foundation software to Synergy Compute modules, a VMware Imaging Appliance (VIA) VM is required.

One or more HPE Synergy frames – populated with compute, memory, storage, and fabric per VMware certified specifications – must be installed and networked.



The Synergy Composer must be configured per Synergy specifications. In addition, the following OneView attributes should be preconfigured prior to proceeding with the VCF build configuration:

- Settings => Networking, Time and Locale, Licenses, and SNMP
- Settings => Address and Identifiers, including the management network address range
 - Note: The management network may be identified by any name. E.g.: **iLO_Range** in the screenshot on the next page.
- Settings => Repository (with space available)
- Enclosures
- Server Hardware
- Server Hardware Types
- Drive Enclosures

Configuration of all other required OneView parameters are covered in this document.

HPE OneView Configuration through HPE Synergy Composer

Configuring Networks (VLANs)

The VLANs configured on OneView are consumed by the Server Profiles and Interconnect. Hence they must be configured in a Network Set and Logical Interconnect Group. VMware VCF requires all VLANs used by VCF not have a VLAN ID lower than 24 or greater than 3299.

In addition to any VLANs that may already exist, create the following VLANs as **Tagged** under OneView => Networking => Networks. For a listing of the components included in the solution, see Table x in [Appendix A – Solution hardware configuration \(this section is required\)](#).

At a minimum, the following VLANs are required to bring up VMWare Cloud Foundation on Synergy.:

1. A VCF “build network” used for initial imaging and deployment
2. Vmotion Network
3. VCF-Internal Management Network
4. Datacenter Network (External Management and guests)
5. VXLAN transport VLAN
6. 1 vSAN Network per workload domain

Examples of the networks are shown in the following images. NOTE: These are **examples** and can vary depending on the customer installation.

The names do not need to be identical, but being specific will help the customer know which networks apply to what function.

The screenshot shows the VMware OneView interface. On the left, a table lists several networks. The 'Net-234' network is highlighted. On the right, the 'General' details for 'Net-234' are displayed.

Name	VLAN	Type
fc-san1		FC
fc-san2		FC
Net-230	300	Ethernet
Net-231	301	Ethernet
Net-232	302	Ethernet
Net-233	303	Ethernet
Net-234	304	Ethernet
vsan-2000	2000	Ethernet

General	
Type	Ethernet
VLAN	304
Associated with subnet ID	none
Purpose	General
Preferred bandwidth	2.5 Gb/s
Maximum bandwidth	20 Gb/s
Smart link	Yes
Private network	No
Uplink set	UplinkSet
Used by	none
Member of	1 network set

The screenshot shows the 'Review' page for a network configuration. It includes a 'General' section with various metadata and a 'Management Network' section with IP configuration details.

General	
Physical Rack Name	skrack
Company Name	vmware
Company Department	isbu
Root DNS Domain	vsphere.local
VMware Cloud Foundation Sub-Domain	vrack.vsphere.local
SSO Domain	vsphere.local

Management Network	
VLAN ID	1108
Subnet	172.21.0.0
Subnet Mask	255.255.255.0
Gateway	172.21.0.1
Primary DNS	192.168.1.104

The screenshot shows the configuration page for the 'vMotion Network'. It lists the VLAN ID, Subnet, Subnet Mask, Gateway, and Excluded IP Address Ranges.

vMotion Network	
VLAN ID	30
Subnet	172.22.0.0
Subnet Mask	255.255.252.0
Gateway	172.22.0.1
Excluded IP Address Ranges	

vSAN Network	
VLAN ID	40
Subnet	172.24.0.0
Subnet Mask	255.255.252.0
Gateway	172.24.0.1
Excluded IP Address Ranges	

VXLAN Network	
VLAN ID	50
Subnet	172.31.0.0
Subnet Mask	255.255.252.0

The screenshot shows the configuration page for the 'VXLAN Network'. It lists the VLAN ID, Subnet, Subnet Mask, Gateway, and Excluded IP Address Ranges.

VXLAN Network	
VLAN ID	50
Subnet	172.31.0.0
Subnet Mask	255.255.252.0
Gateway	172.31.0.1
Excluded IP Address Ranges	

Data Center Network	
VLAN ID	60
Subnet	172.41.0.0
Subnet Mask	255.255.252.0
Gateway	172.41.0.1
Excluded IP Address Ranges	

Data Center Uplink	
Uplink Type	L2
Uplink LAG	Disabled

Configuring HPE OneView Network Sets (VLAN sets)

Network sets are consumed only by server profiles. Network sets allow multiple VLANs to be carried on the same physical server interface. Create a network set for use from **OneView => Networking => Network Sets** and include all networks to be used by VCF.

Note: Set whichever network will be used as the “build network” for initial imaging and deployment as **Untagged**.

Networks

Name	▲	VLAN ID	Untagged	
DataCenter		100	<input type="checkbox"/>	×
Imaging		99	<input checked="" type="checkbox"/>	×
Management		101	<input type="checkbox"/>	×
vMotion		102	<input type="checkbox"/>	×
vSAN		103	<input type="checkbox"/>	×
VXLAN		104	<input type="checkbox"/>	×

There are no available networks to add.

Remove networks

Remove all

OneView
Search

Network Sets 1
All resources ▾
All labels ▾

+ Create network set

Name	▲
Net230-234	▲

Net230-234
Overview ▾
⌵

General

Preferred bandwidth 2.5 Gb/s

Maximum bandwidth 20 Gb/s

Untagged network [Net-231](#)

Used by [9 server profiles](#)
[2 server profile templates](#)

Networks

[Net-230](#) 300
 [Net-231](#) 301
 [Net-232](#) 302
 [Net-233](#) 303
 [Net-234](#) 304
 [vsan-2000](#) 2000

Configuring HPE OneView Logical Interconnect Groups

The following 2 Logical Interconnect groups must be configured from **OneView => Networking => Logical Interconnect Groups**. These will be consumed by the Logical Enclosure to configure the physical enclosures.

1. Logical Interconnect Group for 40G Ethernet Connectivity
2. Logical Interconnect Group for 12G SAS Connectivity

Adding Logical Interconnect Group for Virtual Connect SE 40Gb Ethernet Connectivity

Select **[Create logical interconnect group]**

1. Any name will suffice for the logical interconnect.
2. The Enclosure count is the number of Synergy Frames in the logical interconnect group (e.g.: networked through the physical Master Interconnect Module Pair).
3. Interconnect bay set: select **3**
 - o Interconnect bay sets are typically as follows:
 - Bay Set 1: Storage primary bays 1 & 4 and the only bays that support the SAS Connection Module
 - Bay Set 2: Storage and networking secondary bays 2 & 5
 - Bay Set 3: Networking primary bays 3 & 6
4. Redundancy: For a single frame, this is typically **Redundant**. For multiple frames, this is typically Highly Available.
 - o Redundancy selection options are as follows:
 - Highly Available: When the Master Interconnect Modules are on opposite sides of two enclosures.
 - Redundant: When a Master Interconnect Module pair are on opposite sides of a single enclosure (e.g.: bays 3 & 6)
 - Non-redundant (A side only): When a single Master Interconnect Module is used on A side (e.g.: bay 3)
5. Non-redundant (B side only): When a single Master Interconnect Module is used on B side (e.g.: bay 6)

Create Logical Interconnect Group
General ▾
?

General

Name

Logical Interconnect Group

Using the selectors below, describe the logical interconnect group to be created and then click "Select interconnects" to see the bay and interconnect choices.

Interconnect type

Enclosure count

Interconnect bay set

Redundancy

}

The interconnect type will determine configuration of interconnect choices and bay locations.

Select interconnects

Changed: Interconnect type to "Virtual Connect SE 40Gb F8 Module for Sy...

Create
Create +
Cancel

Select [Select interconnects]

Create Logical Interconnect Group

General ▾ ?

General

Name

Logical Interconnect Group

Interconnect type Virtual Connect SE 40Gb F8 Module for Synergy

Enclosure count 1

Interconnect bay set 3

Redundancy Redundant

Scope *none*

Internal 

no networks

Add uplink set

Select [Add uplink set]

Create Uplink Set

General

Name

Type

Connection mode Automatic

LACP timer

Select [Add networks]

Note: Do not mark any VLAN as native during Uplink Set creation unless that VLAN is indeed set as the PVID or Native VLAN on the upstream switch. .

Name	▲ Type	▲ VLAN ID ▲	Native
Net-230	Ethernet	300	<input type="checkbox"/> x
Net-231	Ethernet	301	<input type="checkbox"/> x
Net-232	Ethernet	302	<input type="checkbox"/> x
Net-233	Ethernet	303	<input type="checkbox"/> x
Net-234	Ethernet	304	<input type="checkbox"/> x
vsan-2000	Ethernet	2000	<input type="checkbox"/> x

Select [Add uplink ports]

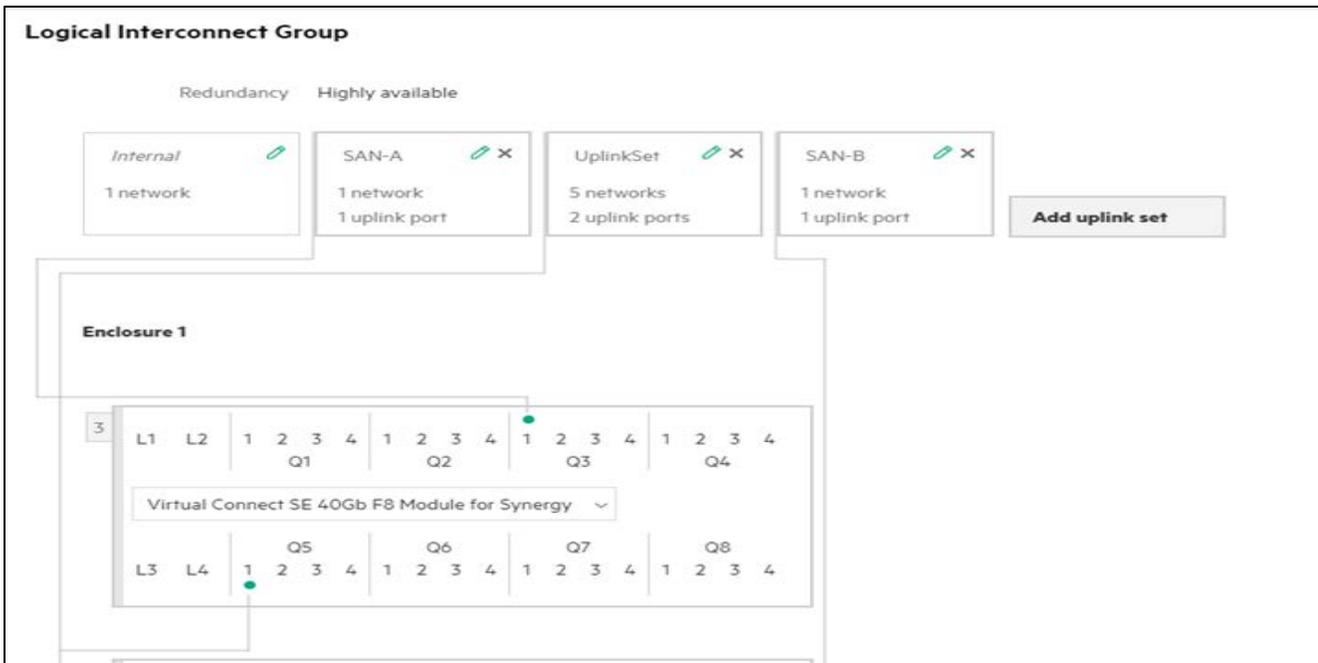
All the VLANs must be configured on the Logical Interconnect Uplink set as shown below. The Uplink set establishes the north-south connectivity from Synergy to the Customer Core switch.

Note: You must select the ports that correspond to your physical wiring.

Uplink Ports								
Interconnect Module	▲ Enclosure	▲ Bay	▲ Port	Capability	Speed	Auto-negotiation		
Virtual Connect SE 40Gb F8 Module for Synergy	1	3	Q5	Ethernet + FCoE	Auto ▾	Enabled		X
Virtual Connect SE 40Gb F8 Module for Synergy	2	6	Q5	Ethernet + FCoE	Auto ▾	Enabled		X

Notes:

1. In a 2 or more frame configuration, the second interconnect module will likely be in enclosure 2.
2. When connecting to external SANs, additional uplink sets will be required and the Logical Interconnect Map will look like:



Select [Create]

Validate the Logical Interconnect Group settings and Select [Create]

Add a Logical Interconnect Group for 12Gb SAS Connectivity

This logical interconnect group enables Server Profiles to consume storage from D3940 Storage Modules. The SAS Logical Interconnect Group applies to all enclosures in the OneView management network.

Select [Create logical interconnect group]

Give any name for the LIG

Create Logical Interconnect Group General ▾

General

Name

Logical Interconnect Group

Using the selectors below, describe the logical interconnect group to be created and then click "Select interconnects" to see the bay and interconnect choices.

Interconnect type

Enclosure count

Interconnect bay set

}

The interconnect type will determine configuration of interconnect choices and bay locations.

Select [Select interconnects]

Select [Synergy 12Gb SAS Connection Module] for Bays 1 and 4

Create Logical Interconnect Group

General ▾

General

Name

Logical Interconnect Group

Interconnect type Synergy 12Gb SAS Connection Module

Enclosure count 1

Interconnect bay set 1

Scope *none*

1

Synergy 12Gb SAS Connection Module ▾

1 2 3 4 5 6 7 8 9 10 11 12

4

Synergy 12Gb SAS Connection Module ▾

1 2 3 4 5 6 7 8 9 10 11 12

Select **[Create]**

NOTE: The steps used above all show best practices and examples. The names and exact uplinks for the Uplink Sets can be different in the customer environment and do not need to match these pictures identically.

Configuring HPE OneView Enclosure Group

Create an Enclosure Group from **OneView => Servers => Enclosure Groups**, similar to as shown in the screenshot, populating the interconnect bays with the LIG definitions created in the previous steps.

In this example the enclosure count is [3] because the environment is the maximum 20Gb interconnect satellite configuration. If starting with one enclosure as an installation then the count would be [1].

Depending on customer's specifications from earlier in the document the use of address pool or DHCP or managed externally will be selected.

As the build does not use the HPE Image Streamer there will not be a deployment settings or deployment network for that hardware in the build examples. If it exists in the customer environment then it would be added here.

Create Enclosure Group
General ▾
?

General

Name

Enclosure count

IPv4 management address configuration Use address pool Use DHCP Manage externally

IPv4 address pool	Range Name ▲	Domain	IPv4 Addresses	
	Mgmt		10.16.44.200 - 10.16.44.240	✕

Add address ranges
Remove all

OS Deployment Settings

Deployment network type

Interconnect Bay Configuration

Enclosure 1

1	Interconnect	Synergy 12Gb SAS Connection Module	
	Logical interconnect group	SAS-LIG	✕ 🔍
2	Logical interconnect group	None	✕ 🔍
3	Interconnect	Virtual Connect SE 40Gb F8 Module for Synergy	
	Logical interconnect group	ETH-LIG	✕ 🔍
4	Interconnect	Synergy 12Gb SAS Connection Module	
	Logical interconnect group	SAS-LIG	✕ 🔍
5	Logical interconnect group	None	🔍
6	Interconnect	Synergy 20Gb Interconnect Link Module	
	Logical interconnect group	ETH-LIG	✕ 🔍

Enclosure 2

Select the bay number [1] for the SAS-LIG and it will automatically populate bay number [4] as well. This must be done for every frame that has SAS connectivity to the D3940 modules inside those frames.

Select the bay number [3] for the ETH-LIG and it will automatically populate the bay number [6] as well as other [3][6] bays in different frames if more than one enclosure count selected.

Configuring HPE OneView Logical Enclosure

Create one Logical Enclosure from OneView => Servers => Logical Enclosures.

Note: The Logical Enclosure creation process will take some time since it applies the Logical Interconnect and Logical Enclosure configuration including the HPE Virtual Connect fabric and SAS Fabric. For a minimal configuration, this process could take up to 40 minutes and significantly longer with more complex configurations. Also, user interaction may be required during the creation process for firmware updates and other requirements.

Note: Depending on current SPP release it would be attached to the Logical Enclosure here in “Firmware Baseline”. The upload to the Composer should be added through the OneView => General => Firmware Bundles then “+ Add Firmware Bundle” prior to this.

Note: The names below are examples. When populating the “Enclosures” check the serial numbers of the enclosures the customer needs and use the correct Enclosure Group created in the previous step.

During the Logical Enclosure creation process, you can select the button next to create => Details to display creation process activity.

Create Logical Enclosure

Name	R10-LEC
Enclosures	CN7515010L ✕ 🔍
Enclosure group	R10-EG ✕ 🔍

Firmware

Firmware baseline	Manage manually ▼
-------------------	------------------------------------------------------------------------

Once the creation is finished, the Logical Enclosure parameters should appear similar to (NOT for every environment, just as an example):

✔ Main LE

General ▾

⌘

General

Consistency state	Consistent
Enclosure group	● Main EG
Enclosures	● CN75150484 -- Middle ● CN7515048Q -- Bottom ● CN7515010J -- Top
Logical Interconnects	● Main LE-ETH-LIG ● Main LE-SAS-LIG-2 ● Main LE-SAS-LIG-1

OS Deployment Settings

This logical enclosure is not configured for OS deployments.

Interconnects

Enclosure ● [CN75150484 -- Middle](#)

Bay ▲	Interconnect	Logical Interconnect	Installed Module	Expected Module
1	● CN75150484 -- Middle, interconnect 1	Main LE-SAS-LIG-2	Synergy 12Gb SAS Connection Module	Synergy 12Gb SAS Connection Module
2	none	not set	empty	none
3	● CN75150484 -- Middle, interconnect 3	Main LE-ETH-LIG	Synergy 20Gb Interconnect Link Module	Synergy 20Gb Interconnect Link Module
4	● CN75150484 -- Middle, interconnect 4	Main LE-SAS-LIG-2	Synergy 12Gb SAS Connection Module	Synergy 12Gb SAS Connection Module
5	none	not set	empty	none
6	● CN75150484 -- Middle, interconnect 6	Main LE-ETH-LIG	Virtual Connect SE 40Gb F8 Module for Synergy	Virtual Connect SE 40Gb F8 Module for Synergy

Enclosure ● [CN7515048Q -- Bottom](#)

Bay ▲	Interconnect	Logical Interconnect	Installed Module	Expected Module
1	none	not set	empty	none
2	none	not set	empty	none
3	● CN7515048Q -- Bottom, interconnect 3	Main LE-ETH-LIG	Synergy 20Gb Interconnect Link Module	Synergy 20Gb Interconnect Link Module
4	none	not set	empty	none
5	none	not set	empty	none
6	● CN7515048Q -- Bottom, interconnect 6	Main LE-ETH-LIG	Synergy 20Gb Interconnect Link Module	Synergy 20Gb Interconnect Link Module

Enclosure ● [CN7515010J -- Top](#)

Validating HPE OneView Logical Interconnects

Logical Interconnects will be automatically created during the Logical Enclosure creation process. Validate the Logical Interconnect configurations from **OneView => Networking => Logical Interconnects**.

Note: If there are any problems with the Logical Interconnects they will be displayed here through the menu and selecting “Activity”.

The screenshot shows the HPE OneView interface for configuring Logical Interconnects. The main title is "Logical Interconnects" with a sub-header "Main LE-ETH-LIG | Logical Interconnect". The left sidebar lists three interconnects: "Main LE-ETH-LIG", "Main LE-SAS-LIG-1", and "Main LE-SAS-LIG-2". The main content area displays the configuration for "Main LE-ETH-LIG", which is a "20Gb Interconnect Link Module".

The configuration is divided into three sections:

- CN75150484 -- Middle:** This section shows a configuration for a Synergy 20Gb Interconnect Link Module. It includes a table with columns L1 and L2, and a status indicator (green dot) for L1. Below the table, it lists "CN75150484 -- Middle, interconnect 3" with a state of "Configured" and expected/actual hardware details.
- CN75150484 -- Middle, interconnect 6:** This section shows a configuration for a Virtual Connect SE 40Gb F8 Module for Synergy. It includes a table with columns L1, L2, Q1, Q2, Q3, and Q4, and status indicators (green dots) for L1 and L2. Below the table, it lists "CN75150484 -- Middle, interconnect 6" with a state of "Configured" and expected/actual hardware details.
- CN7515048Q -- Bottom:** This section shows a configuration for a Synergy 20Gb Interconnect Link Module. It includes a table with columns L1 and L2, and status indicators (green dots) for L1 and L2. Below the table, it lists "CN7515048Q -- Bottom, interconnect 3" with a state of "Configured" and expected/actual hardware details.

Configuring HPE OneView Server Profile Templates

Typically for VMware VCF, all servers in a domain (management domain or workload domains) must have identical Server Profiles, although they may be sourced from different service profile templates.

Create a OneView Server Profile Template as follows from OneView => Servers => Server Profile Templates => [Create server profile template].

Select the appropriate hardware type and Enclosure Group created in previous steps.

Do not select any OS Deployment options if they are available (here one does not exist in the Logical Enclosure).

Select the appropriate firmware baseline and select "Firmware Only". *Using Smart Update Tools is not currently supported with VCF.*

Edit VCF-Node-Gen10 General ▾ ?

General

Name

Description

Server Profile

Server profile description

Server hardware type SY 480 Gen10 1 [Change](#)

Enclosure group Main EG [Change](#)

Affinity

OS Deployment

To define OS deployment settings, select an enclosure group configured for OS deployment.

Firmware

Firmware baseline

Force installation

Installation Method

- Firmware and OS Drivers using Smart Update Tools
- Firmware only using Smart Update Tools

Changed: Firmware baseline to "HPE Synergy Custom SPP 2017 12 20 version 2017.12.20.00"

Select [Add connection] to add two network connections, both configured to use the Network Set created earlier in this document. In this example, the network set is named "Net230-234".

Mark the first connection as PXE Primary and the second as PXE Secondary.

Add Connection	
General	
Name	vmnic0
Function type	Ethernet
Network	Net230-234
Port	Auto
Link aggregation group	None
Requested bandwidth (Gb/s)	2.5
Requested virtual functions	<input checked="" type="radio"/> None <input type="radio"/> Custom <input type="radio"/> Auto
Boot	PXE primary

Add Connection	
General	
Name	vmnic1
Function type	Ethernet
Network	Net230-234
Port	Auto
Link aggregation group	None
Requested bandwidth (Gb/s)	2.5
Requested virtual functions	<input checked="" type="radio"/> None <input type="radio"/> Custom <input type="radio"/> Auto
Boot	PXE secondary

Connections

Manage connections

ID	Name	Network	Port	Boot		
○ 1	vmnic0	<u>Net230-234</u> (network set)	Mezzanine 3:1-a	PXE primary		
	Type	Ethernet				
	MAC address	Auto				
	Requested virtual functions	None				
	Requested bandwidth	2.5 Gb/s				
	Link aggregation group	None				
○ 2	vmnic1	<u>Net230-234</u> (network set)	Mezzanine 3:2-a	PXE secondary		
	Type	Ethernet				
	MAC address	Auto				
	Requested virtual functions	None				
	Requested bandwidth	2.5 Gb/s				
	Link aggregation group	None				

Add connection

NOTE: Do not assign a LAG to networks used to provision a network via PXE boot. Without OS level drivers forming the LAG at time of the PXE request, the server will not connect properly to the PXE server. As VCF uses two connections with all uplink sets included with PXE primary/secondary, there will not be LAGs used here.

NOTE: In a SY480 Gen10 server for this build the P416i-e handles both internal and external drive configurations. For this build a separate storage controller card does not exist for the two front-facing (internal) drives and therefore the boot drive must be a single SAS drive in HBA mode (**Mixed-Mode RAID/HBA is not supported on VMware on a single controller**). The customer could also use M.2 SATA or SDcard for installation media if available.

SAS Mezzanine storage controller (P416i-e): Drives in the D3940 storage module should be configured as HBA (external logical JBOD) as follows:

1. Select the edit icon for “SAS Mezz 1 storage controller” to create the local SAS drives
2. Select Mode HBA, check “Re-initialize controller on next profile application”
3. Select [Create logical JBOD]
4. Create the Cache Tier with defined number of cache drives from the VM sizing
5. Select [Create logical JBOD] to create the capacity tier logical drive.
6. Enter the name for the capacity tier, the number of drives as determined above, and select the appropriate SAS drive type.
7. Select [Create] and [OK].
8. Your logical storage configuration should look similar to this (**NOTE:** Boot drive in HBA mode will not appear in this list).

Create Logical JBOD ?

Name

Storage Location **External**

Number of physical drives

Select drives by Drive type Size and technology

Drive type

? Changed: Drive technology to "SAS SSD"
 Create
Create +
Cancel

SAS Mezz 1 storage controller ✎

Managed by OneView

Initialization will occur on next assignment to server hardware

Name	Type	RAID Level	Number of Drives	Size GB	Drive Technology	Boot	Erase on Delete
Cache Tier	External logical JBOD	n/a	1	200	SAS SSD	n/a	Yes ✕
Capacity Tier	External logical JBOD	n/a	4	300	SAS HDD	n/a	Yes ✕

Note: VMware VIA does not support UEFI Boot Mode. Configure the Server Profile Boot Settings as shown below:

Create Server Profile Template
Local Storage ▾
?

Local Storage

Device	Capacity	RAID	Format	Actions
U1SK4	300 GB	RAID 0	NTFS	✖
External logical drive				

SAN Storage

Manage SAN Storage

Boot Settings

Manage boot mode

Boot mode: Legacy BIOS ▾

Changing the Boot Mode can impact the ability of the server to boot the installed operating system. An operating system is installed in the same mode as the platform during the installation. If the Boot Mode does not match the operating system installation, it might not boot.

Manage boot order

- 1 CD
- 2 USB
- 3 Hard disk
- 4 PXE

Drag and drop or edit rows to re-order

BIOS Settings

Manage BIOS

Advanced

Changed: Boot mode to "Legacy BIOS"

Create
Create +
Cancel

Select **[Create]**.

NOTE: Create additional server profile templates as needed for the different server profile definitions. This will include different:

1. Processor generations and versions
2. Mezzanine connectors
3. Local storage drives and allocated drives in the D3940 Storage Module

Configuring HPE OneView Server Profiles

A Server Profile will need to be added for each server used for VCF.

Create a OneView Server Profile as follows from OneView => Servers => Server Profile => [Create server profile].

Enter the General information for the specific server as follows:

The screenshot shows the 'Create Server Profile' interface in HPE OneView. At the top, there is a header 'Create Server Profile' with a 'General' dropdown menu and a help icon. Below the header is a yellow warning box with the following text: 'On next assignment to server hardware, the local storage controllers marked for re-initialization will have their logical drives deleted making existing data inaccessible. It is strongly suggested that you back up any data on existing logical drives on these controllers before applying a profile with this option selected. If you wish to preserve any existing logical drives on these controllers, deselect the re-initialize storage checkbox.' Below the warning box is the 'General' tab with the following fields:

- Name: VCF-Node-1
- Description: (empty)
- Server profile template: VCF-Node-Gen10
- Server hardware: unassigned (with a search icon and a close button)
- Show empty bays:
- Server hardware type: SY 480 Gen10 1
- Enclosure group: Main EG
- Affinity: Device bay (dropdown menu)

The default entries should be fine for all remaining fields. Select the appropriate server to be used in VCF and validate the entries and select **[Create]**. During the create process, select the circular button to the left of the Create task bar and select Details to see the creation details.

Repeat these steps for each server to be used by VCF. The end result should depend on how many profiles are initially created. (VCF needs 4 hosts for "management" and a minimum of 3 for a workload domain, resulting in 7 servers required at a minimum).

VMware Cloud Foundation on HPE Synergy

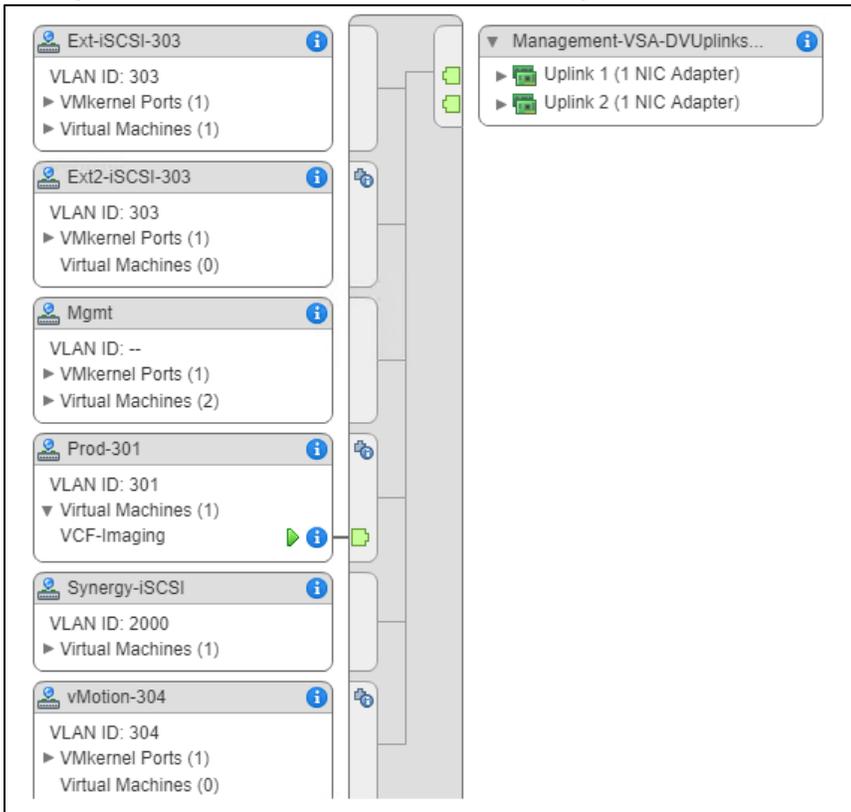
Cloud Foundation is an integrated hybrid cloud platform that delivers a complete set of software defined services for compute, storage, networking, security, and cloud management for the private and public cloud. Cloud Foundation drastically simplifies data center operations by deploying a standardized and validated architecture with built in lifecycle automation of the cloud stack. Cloud Foundation can also be flexibly consumed as-a-service in the public cloud (VMware Cloud on AWS, IBM, Rackspace, etc), enabling a true hybrid cloud that is based on a consistent infrastructure and operational model using common tools and processes.

Deploying the VCF Imaging Server and Jump VM on Management ESXi Server

While VCF will deploy its own internal vCenter, in order to stage all the components for install, an “imaging station” needs to be created. To image servers to be used for VMware Cloud Foundation, VMware provides the “VIA OVF” imaging appliance as well as the Cloud Foundation software bundle. There are many supported methods to image a collection of hosts, such a laptop running VMware Workstation plugged into the infrastructure.

In this example, there is a vSphere cluster residing on a set of Apollo servers plugged into the same upstream switch as the Synergy Frame. The same VLAN that will be used for imaging (and is marked as Native in the Synergy Network Set) is configured as a port group on this cluster.

Management vSphere Standard Switch Configuration



Deploying the VMware Cloud Foundation Imaging Appliance

Deploying the VMware Imaging Appliance (VIA) VM. Deploy the VMware OVF in the port group backed by the VLAN to be used for PXE-based Imaging which was configured on the HPE Synergy OneView and the network switch in the [Configuring VLANs on network switch](#) section. Follow the steps in the “VMware Cloud Foundation VIA User Guide”, 2.3.1 version pages 34-36.

Enabling EMS Mode on the VIA Appliance

By Default, the VIA Imaging appliance is configured to deploy an entire rack worth of infrastructure, including Top-of-Rack switches. In a HPE Synergy based deployment, the Top-of-Rack-Switch deployment options are not required. VIA has a hidden option to bypass all switch configuration and use externally managed switches called “EMS Deployment.” For Synergy deployments, once the Imaging Appliance VM has been deployed, it is necessary to enable the “EMS Deployment” option.

To enable the EMS Deployment Option, perform the following steps.

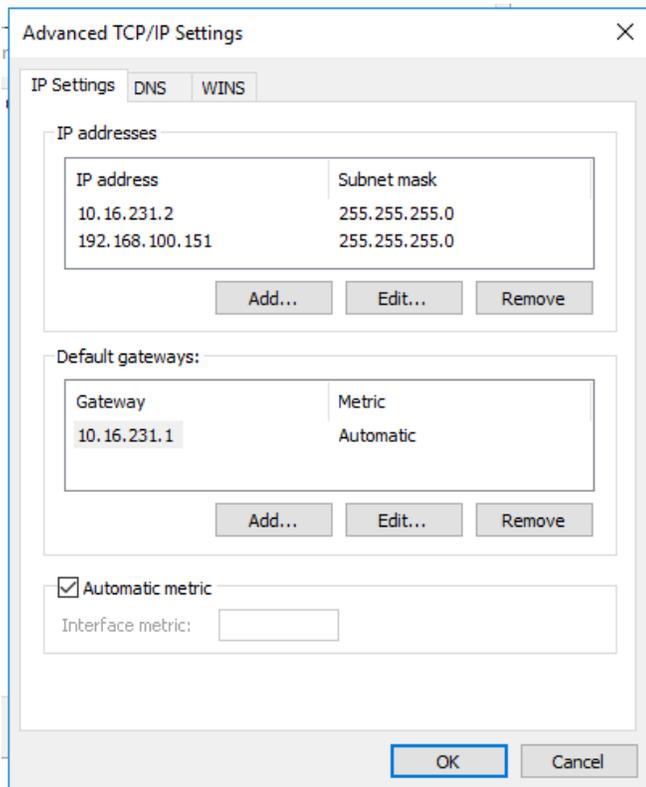
1. Open a VMware console to the VIA Imaging Appliance.
2. Alt-F2 to get to a CLI Prompt.
3. Login as **root** Password **root123**
4. Edit the file `/opt/vmware/evorack-imaging/config/via.properties`
5. Append `via.enable.ems=true` to the end of this file and save.
6. Stop the VIA services by issuing `service via-service stop`.
7. Wait a few moments, then restart the services by issuing `service via-service start`

Note: Ignore the “No Networking Detected” warnings on the VIA Imaging Host’s console. This is normal.

Deploying the Windows Jump VM

Depending on which method used to deploy and stand up Cloud Foundation, it may be convenient to deploy a utility VM to assist with running the VIA Imaging Appliance and the first-time bring up wizard of VMWare Cloud Foundation. In the lab in which this guide was built, the build network (VLAN 301) also has a routed IP network associated with it (10.16.231/24). The VIA appliance, however, deploys with a pre-configured address of 192.168.100.2, and will leverage the 192.168.100.0/24 address scheme throughout the deployment process.

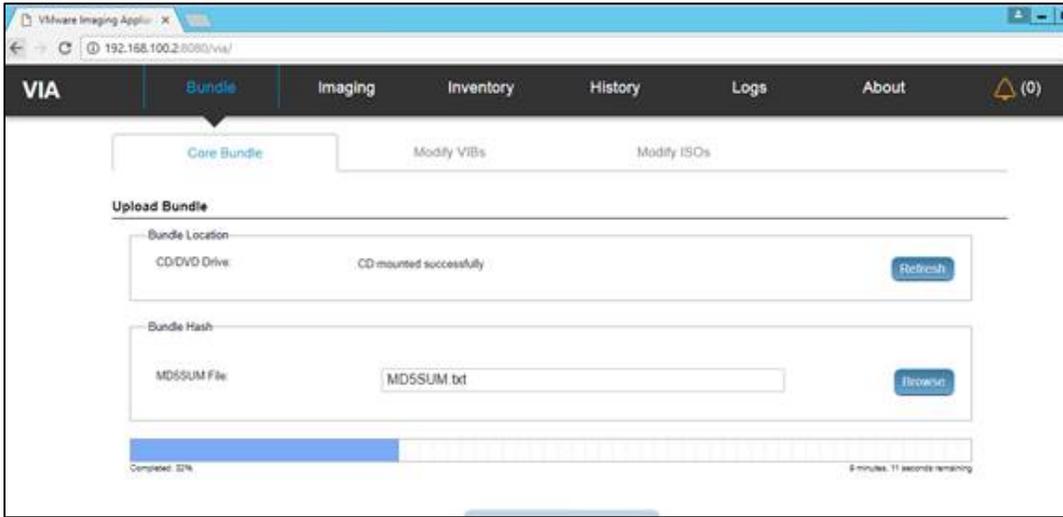
By going into the “Advanced TCP/IP Settings” of the Jump VM, we can set a secondary IP address in the 192.168.100.0/24 network on the same interface. A secondary IP address on the “192.168.100” network that the VCF Imaging Host uses is easily configured in the Advanced TCP/IP menu of the VM’s NIC properties.



Note: In order to avoid IP address conflicts during the imaging and initial bring-up phase, please use an IP address in the 192.168.100.151-199 range. Once the IP address is set, and launch a browser session and navigate to <http://192.168.100.2:8080/via/>

1. From the vCenter console, mount the VCF Bundle directly to the VIA Imaging Appliance VM.
2. Go back to the web session established and "refresh" under core bundle and validate that it sees the CD mount.
3. Browse / select the appropriate MD5SUM for the bundle.

4.



5. Press "upload bundle" - watch the bundle upload. It will take 15-45 minutes for it to finish.
6. Wait for it to finish before going to the next steps.
7. Once the upload finishes, click "Activate Bundle" Wait for it to show "activated"

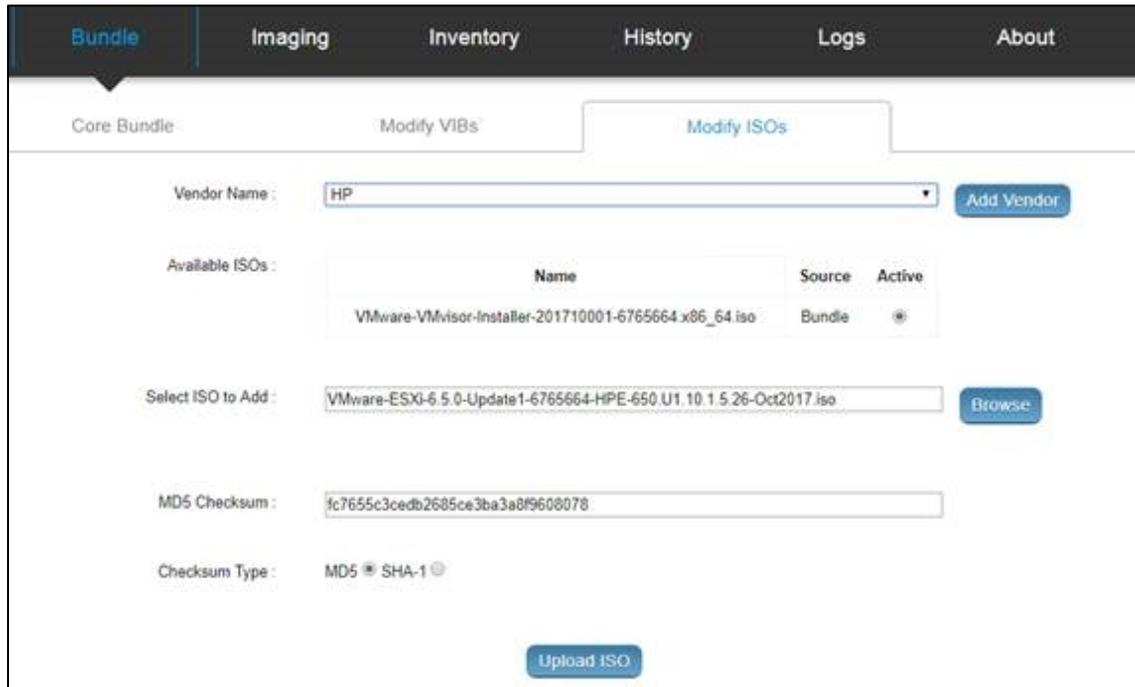
Configuring VCF Imaging Server

Uploading the HPE Custom ISO

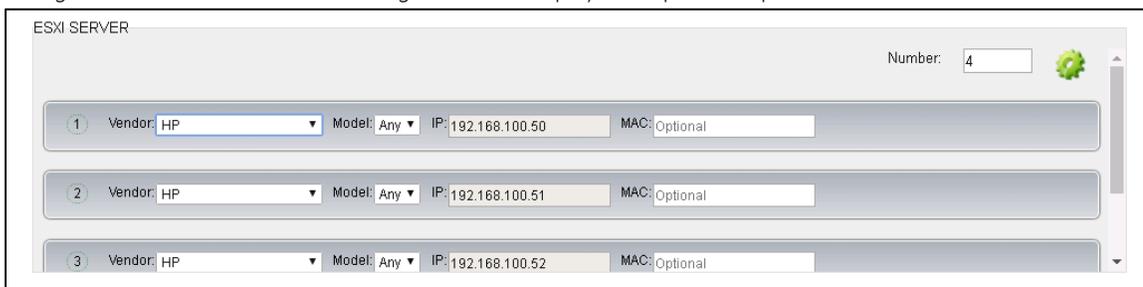
The steps in this section follow along with the "VMware Cloud Foundation VIA User Guide", 2.2 version pages 40-48.

1. Upload the latest HP custom ESXi image that is current with VMware requirements.
 - a. Obtain the compatible HPE ESXi Custom Image to be used that is supported by VCF and HPE Synergy
 - b. If a current driver exists that is required but not in current HPE ISO then upload driver to VIA after current ISO.
2. Switch to the "Modify ISO" tab in the VIA browser session.
3. Select Vendor "HP" from vendor pull down.
4. Browse to current HPE custom ISO.
5. Add in the MD5 sum.
6. Upload ISO.

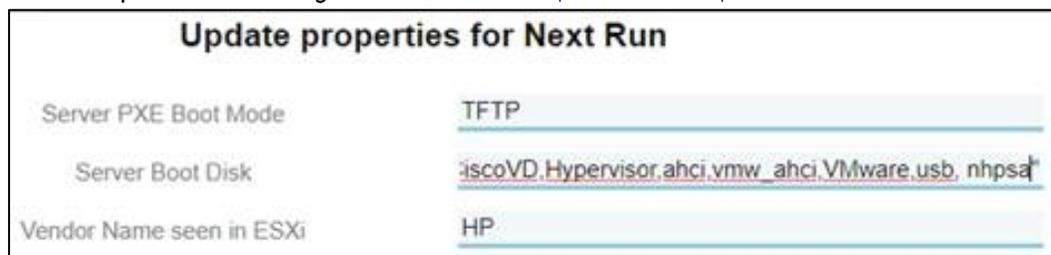
7.



- 8. For any additional VIBs such as updated storage drivers upload them in the same menu above.
- 9. **Important:** Mark the new Uploaded HPE ISO as "active" in this screen, or servers will not find a NIC when they PXE boot.
 - a. Go back to main bundle and then back to "ISOs" and make sure it is still marked as ACTIVE. (Imaging App is sensitive)
- 10. After uploading the HPE media and marking it active, switch to the Imaging tab. Under the imaging tab change the imaging type to be "VMWARE CLOUD FOUNDATION EMS DEPLOYMENT"
 - a. Change the vendor to HP and click the "gear" icon to display the "Update Properties for Next Run" screen



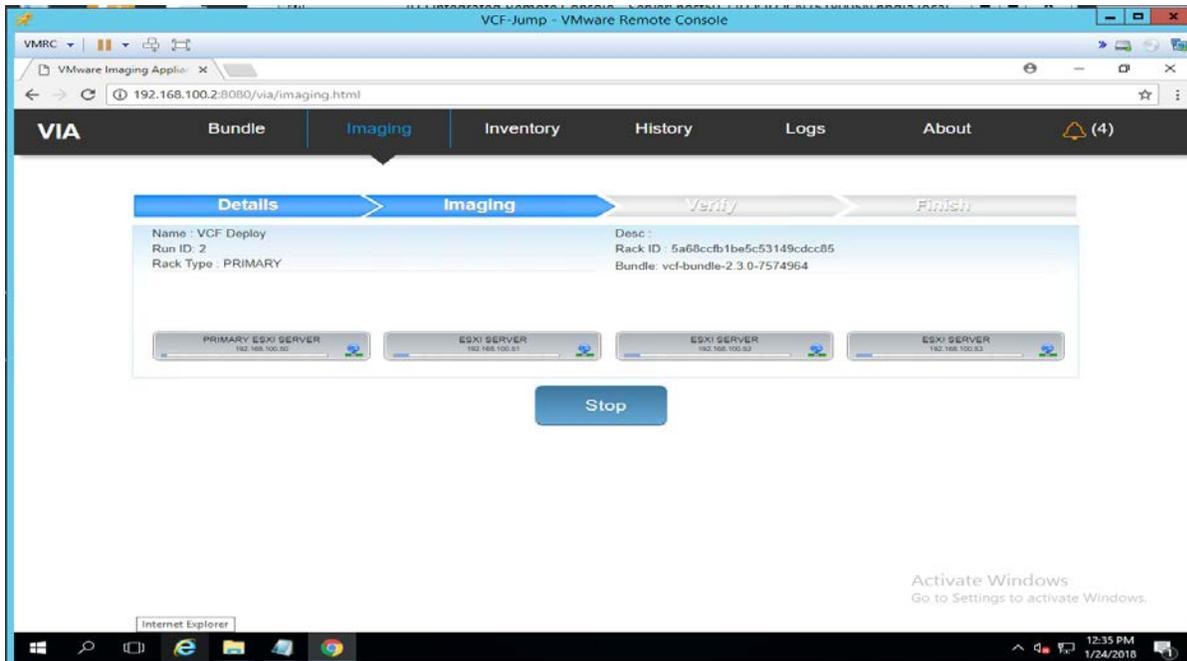
- b. If using local "boot media" you have to add or replace for the proper media driver in this screen.
 - i. **NOTE:** For Gen10 and the P416i, the driver is "smartpqi"
 - ii. **Gen 9 compute modules using the P542 will use "nhpsa"**. The example below is for a Gen9.



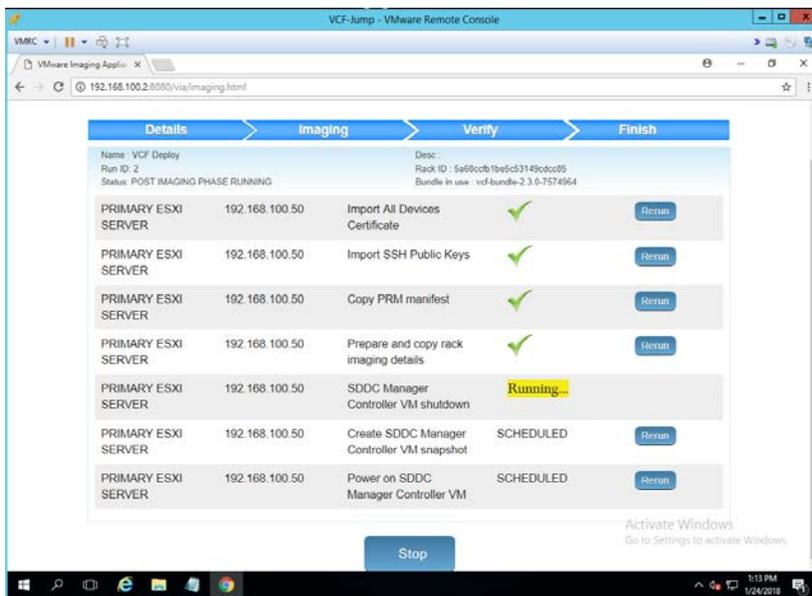
- c. Select how many servers that will be imaged in this run.

d. Select “start imaging” – wait about 30-40 seconds for all imaging services to start in the background

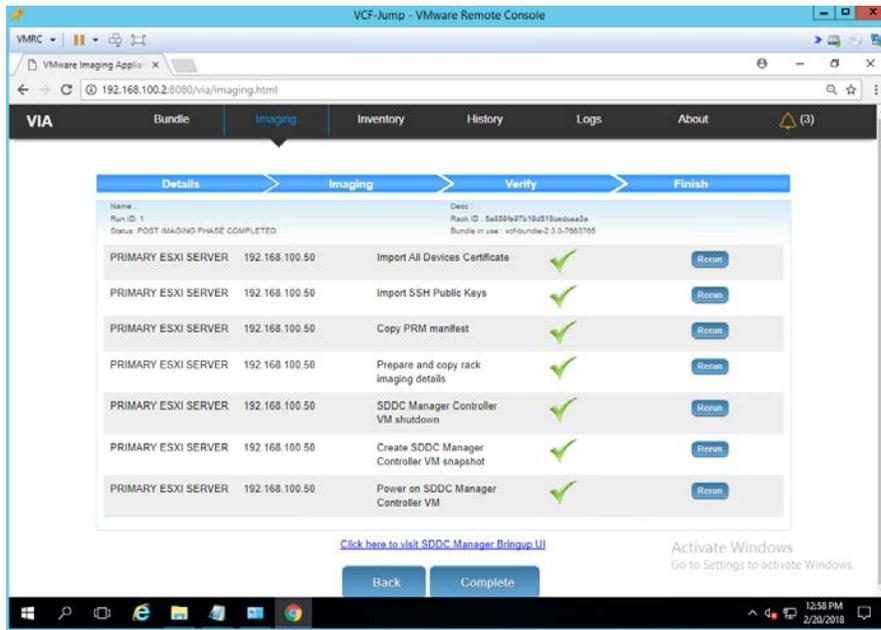
11. Power on ESXi servers and watch boot via ILO remote console. They should go into PXE booting. The first server to power on will be the "core" server and the customer should wait 30 seconds before powering on additional servers.
12. Once all servers are booting / deploying, switch back to the VIA Browser session -- the customer should see all servers in progress. For imaging detail on a specific server, click one of the boxes corresponding to the server.



13. After all servers complete the initial imaging process, the screen will automatically refresh and run through various post-imaging verification tests.



14. Once all servers are fully imaged, and all validation has been performed then check with the following screen.



15. Choose complete, or alternatively, click the link to continue to the SDDC Manager Bring Up UI.

Initializing VMware SDDC Manager

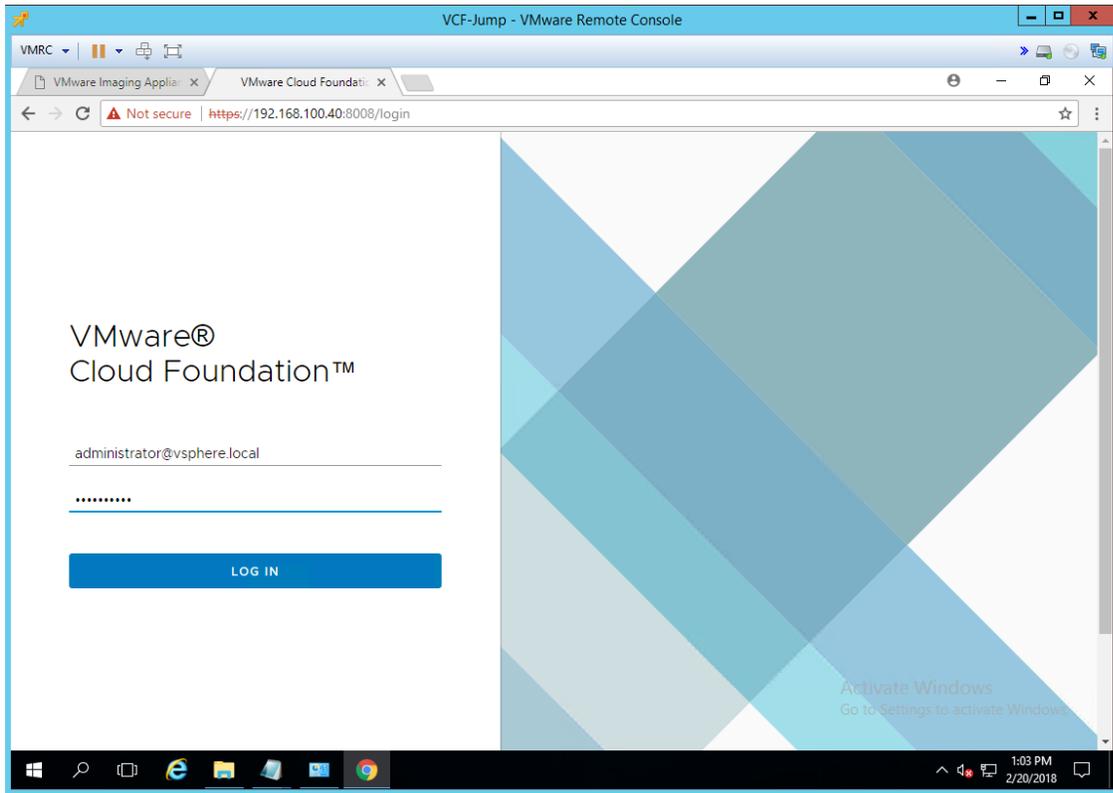
Before You Begin

1. VCF relies heavily on DNS and runs its own DNS server. Please have administrative access to DNS and be prepared to create a DNS Delegation to the VCF sub-domain.
2. For VLAN and IP Scheme used in the PoC, refer to the “Network Switch Configuration” section of this document.
3. Before starting VI WLD deployment the administrator will need an extra VLAN/subnet for the vSAN which needs to be configured on the vSAN networking screen during VI WLD deployment configuration. If the default vSAN VLAN is used then deployment will fail.

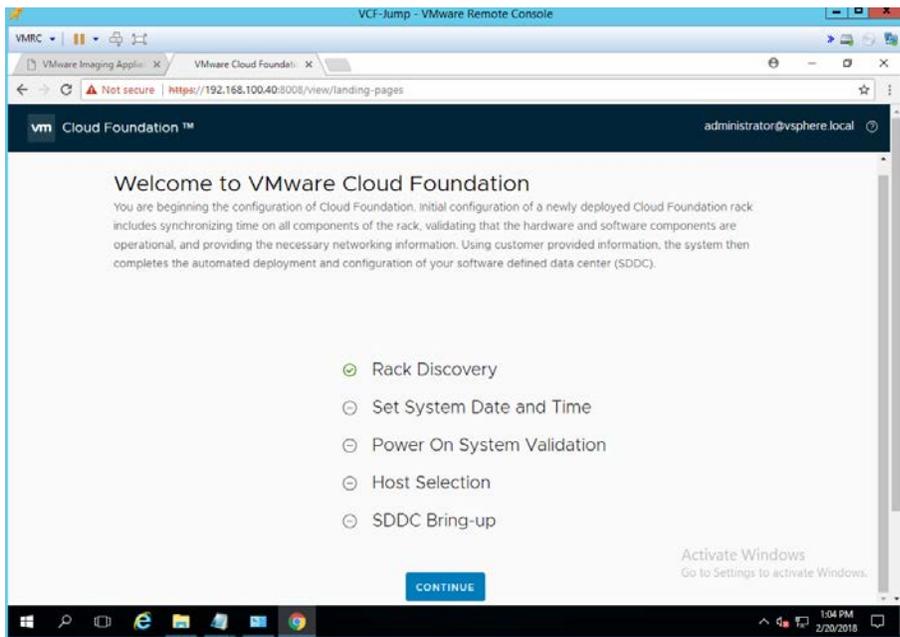
VCF Bring-up Steps

Logging into the SDDC Manager BringUp UI

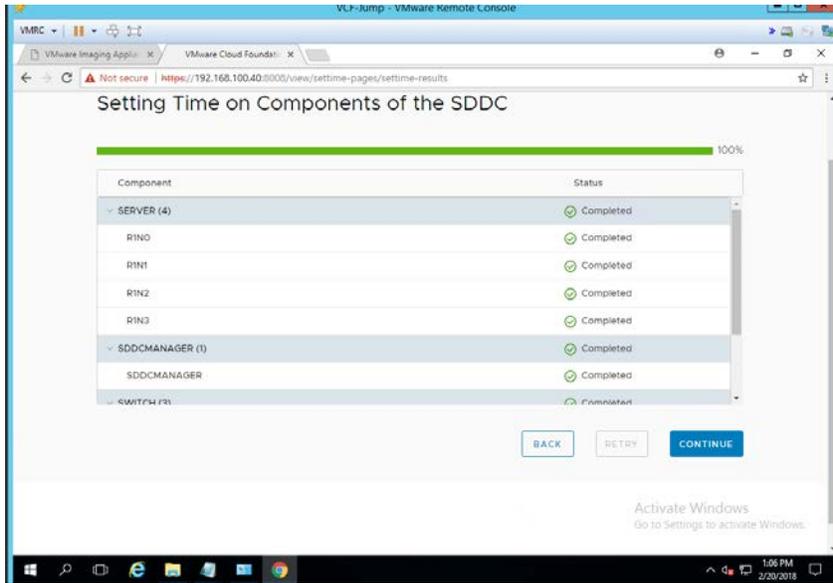
1. 1. Login to the BringUP UI using the credentials administrator@vsphere.local and password of “vmware1234”



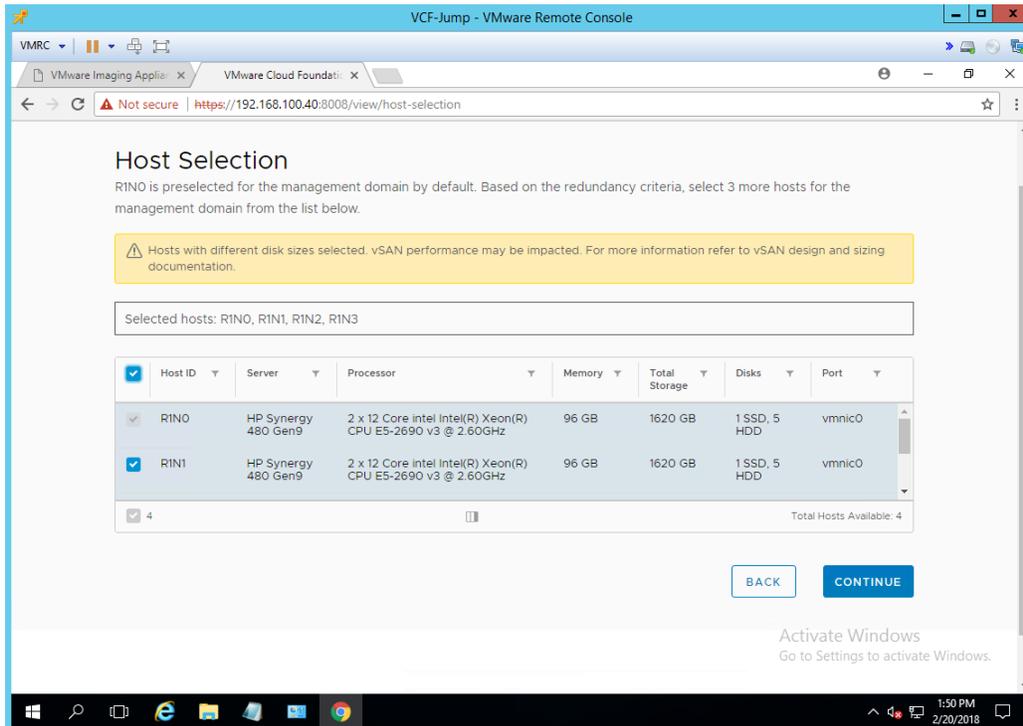
2. The BringUP UI will first discover the installed components and then prompt to move to the next steps.



3. Set the system date and time
4. Time will be synchronized across all components



5. After time synchronization, VCF SDDC Manager will validate all the servers are ready to continue with deployment.
 - a. If the “power On System Validation” hangs on ToR switch validation. Clear the alerts and retry.
6. Select 4 hosts to create the Management domain and choose “Continue”. Ignore the warning on “different disk sizes” this error is false.

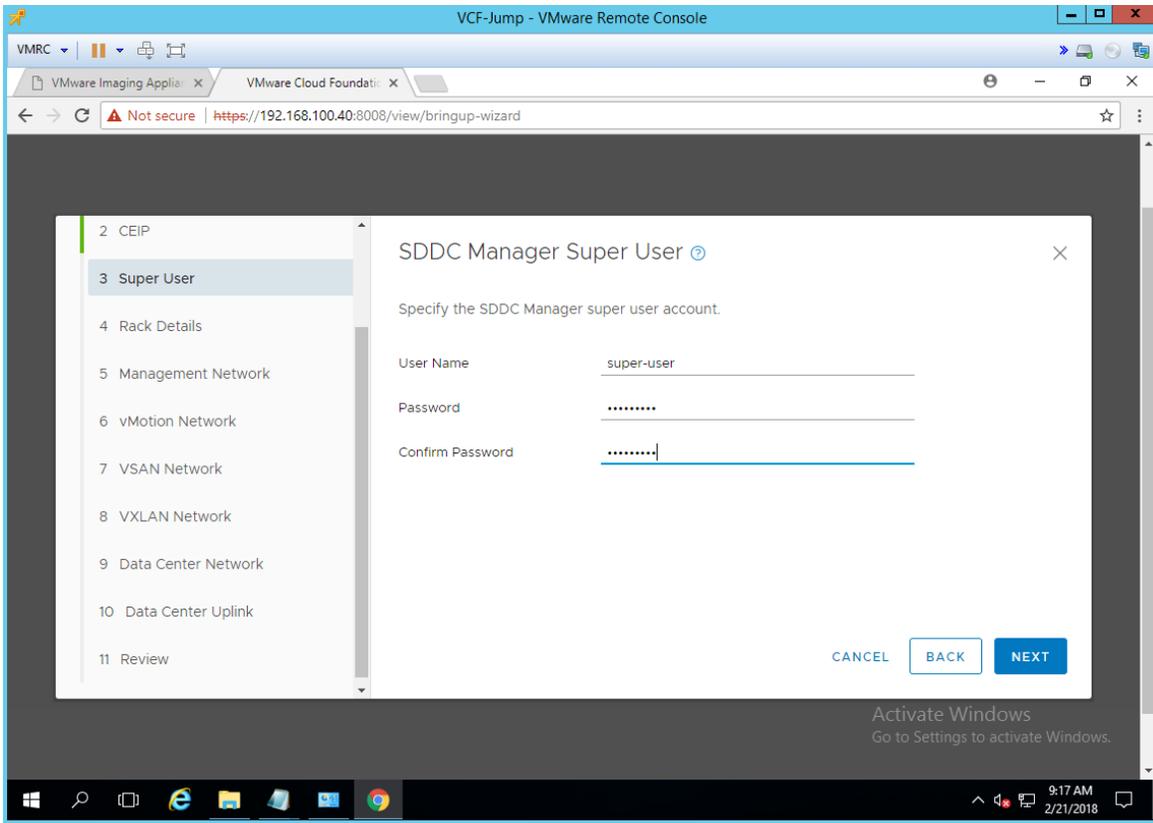


7. Begin deployment and Agree to the EULA

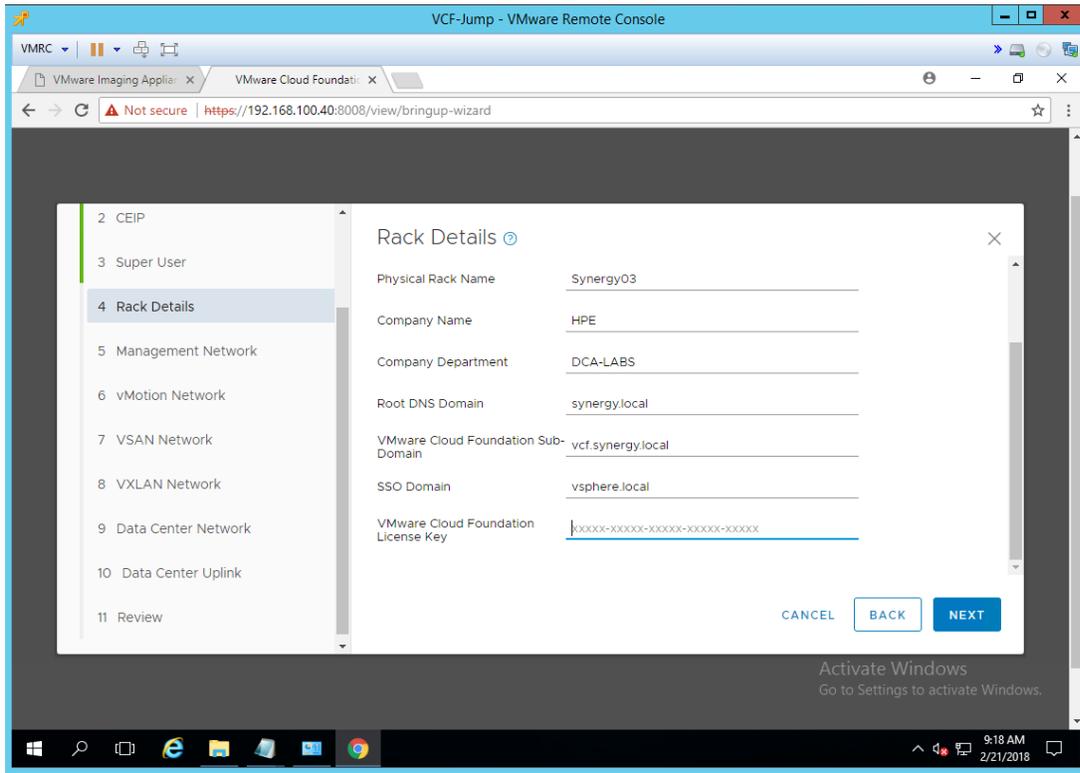
Begin deployment of the SDDC

START DEPLOYMENT

8. Choose to Opt-in to VMware's Customer Experience Program (CEIP).
9. Create a new super-user account (It can't be "administrator"). For this example, the username is "super-user".

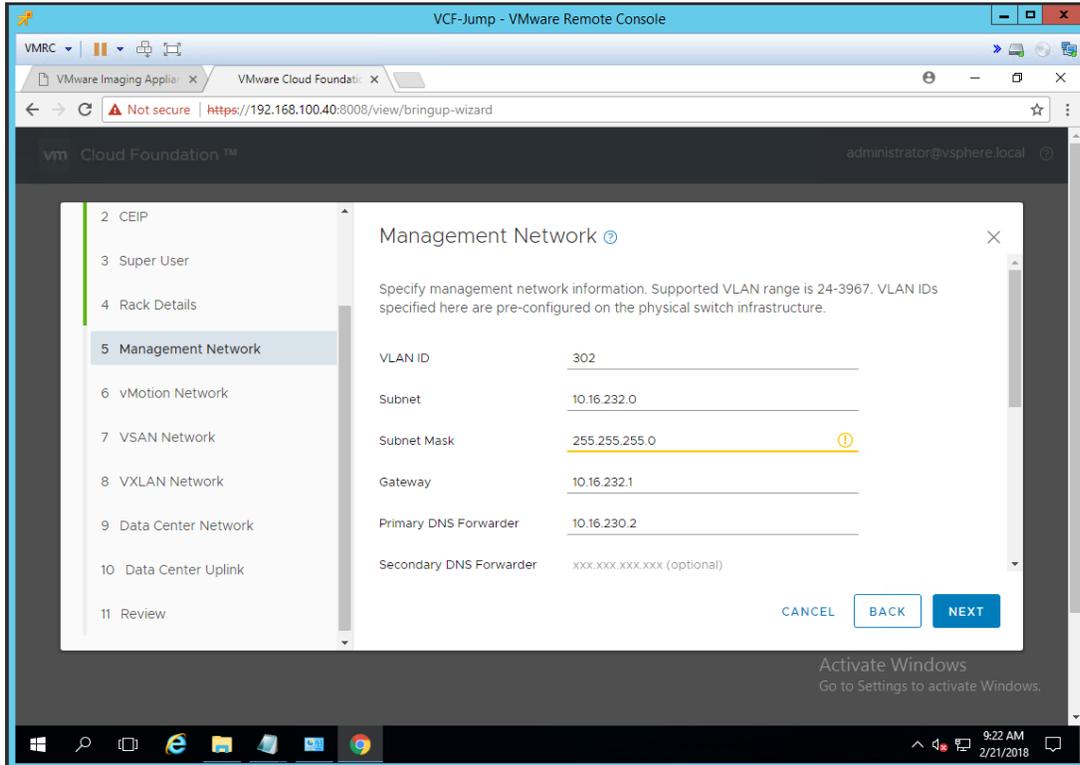


10. Enter the Rack Details specific to the customer environment.

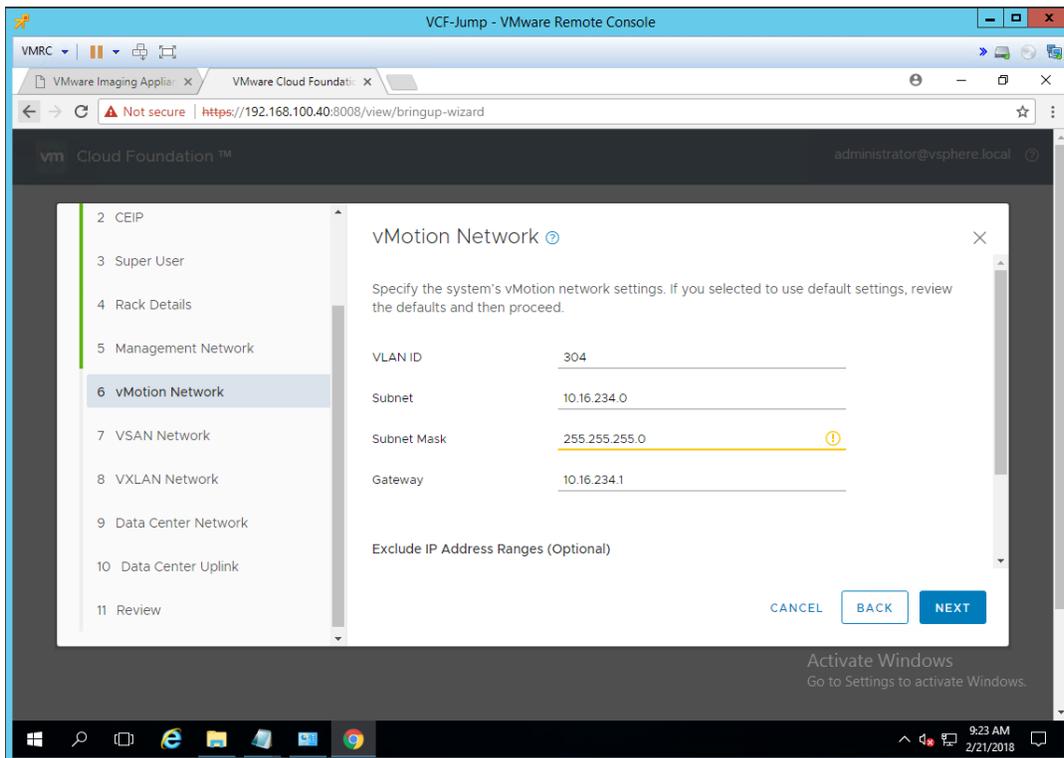


11. Enter the management network VLAN and IP network information.

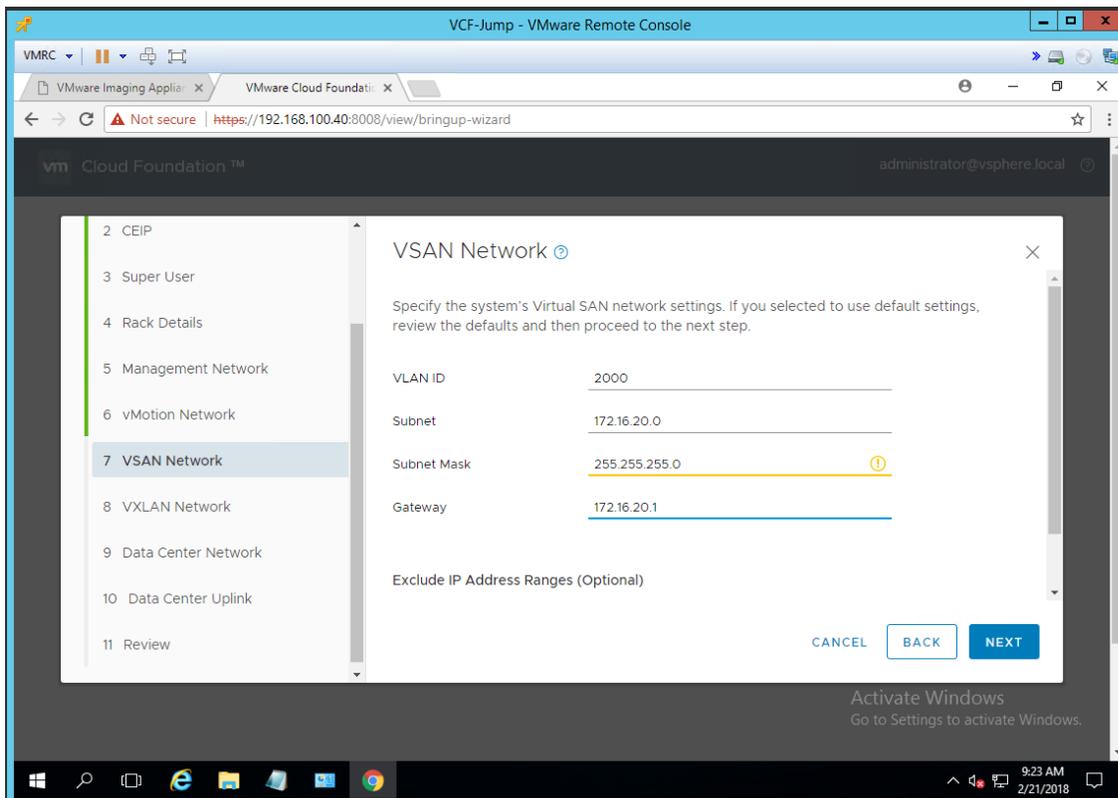
12.



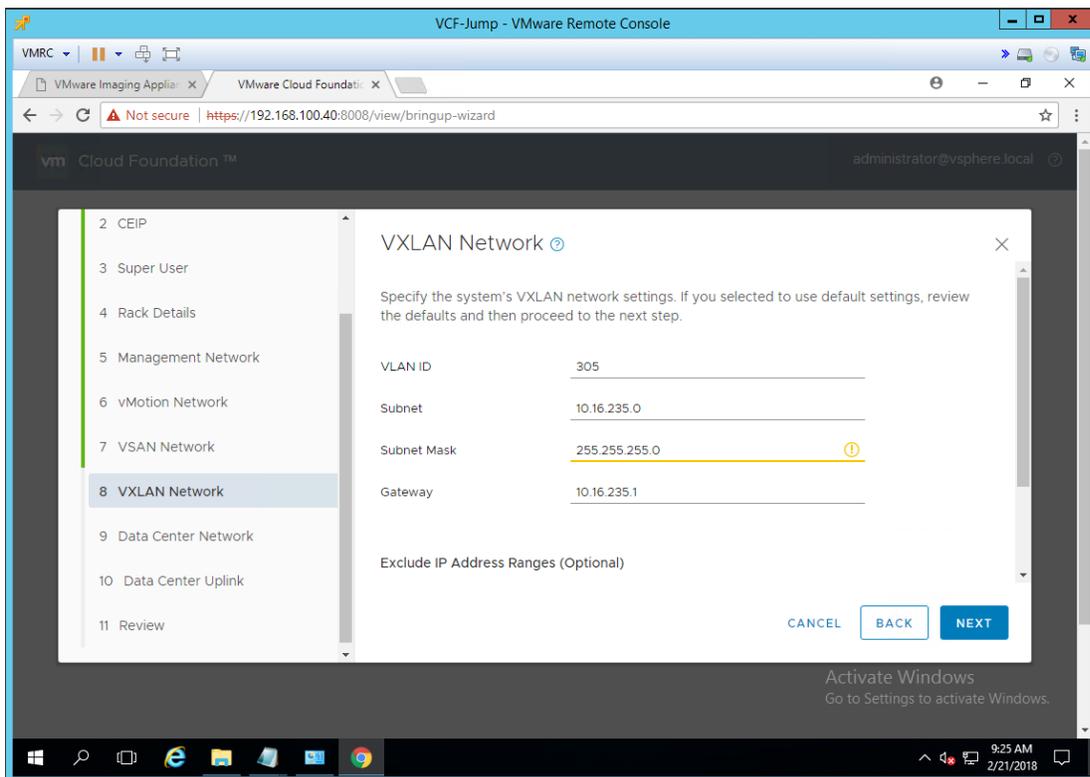
13. Enter the information for the internal vMotion network



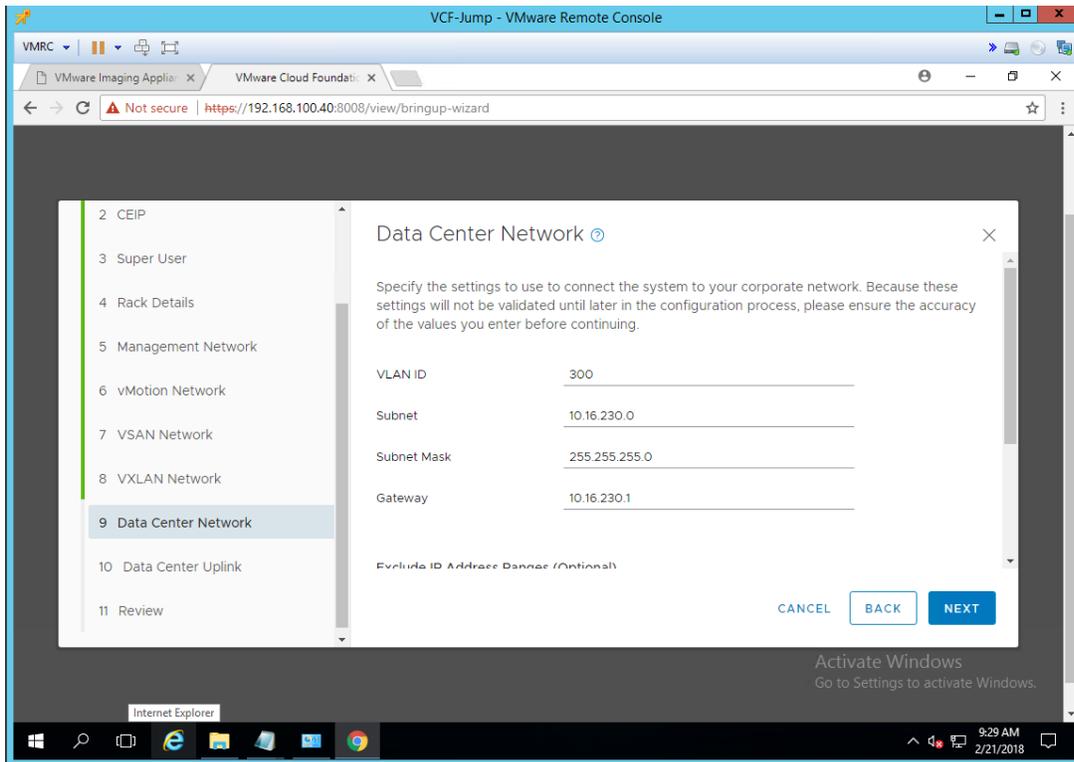
14. Enter the vSAN network information



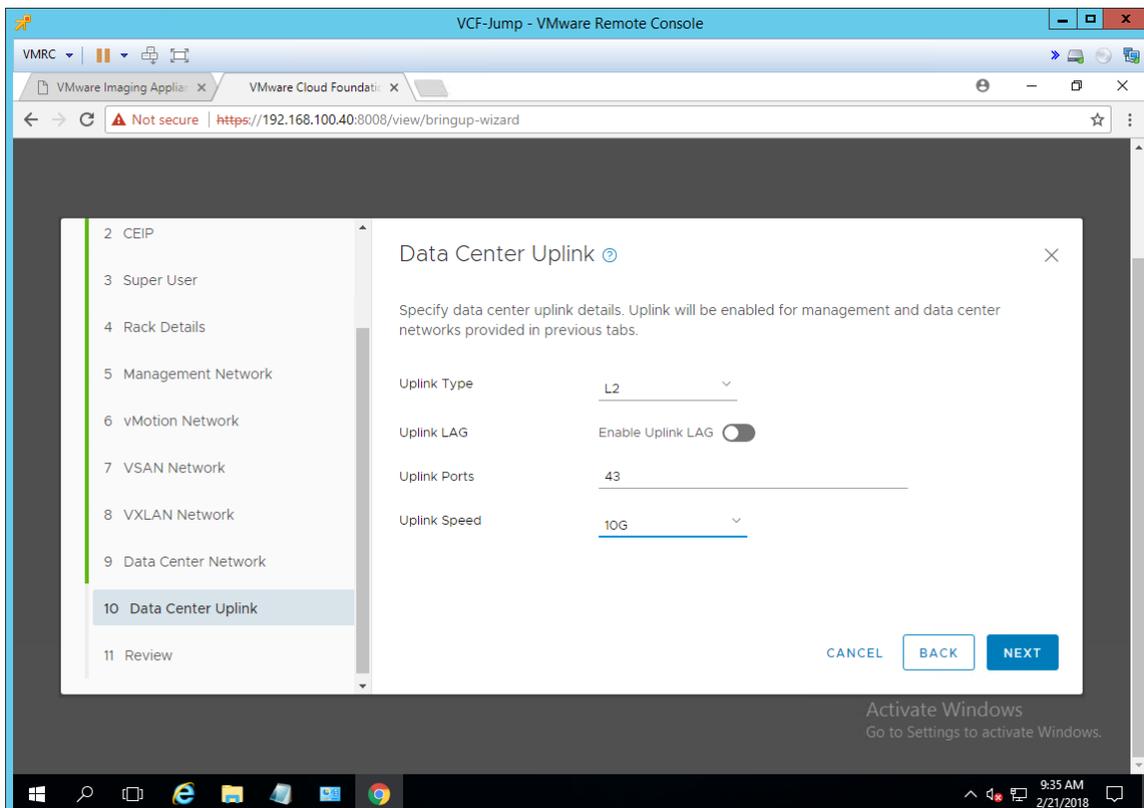
15. Enter the VXLAN transport network information



16. Enter the network information for the Datacenter Network. The Datacenter network is used for access to the Workload Domain and it needs to be routed also.

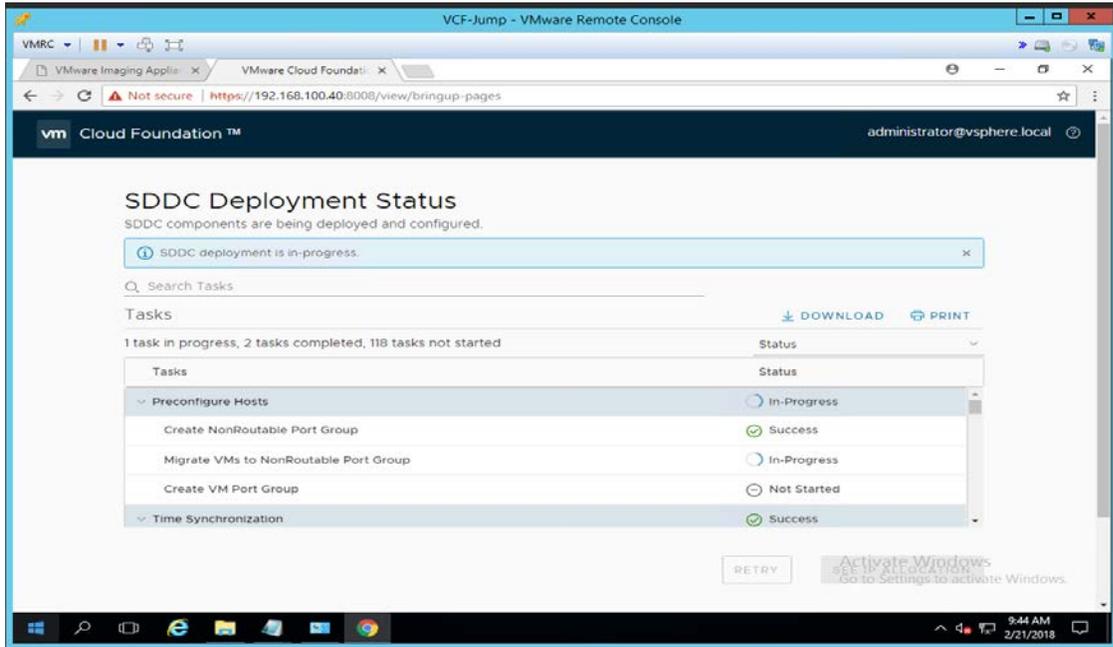


- At the Datacenter Uplink screen, select an uplink type of “L2” – and put in “uplink port” value between 43 and 46. These settings are not used in Synergy deployments, but values must be inserted to move past this screen.

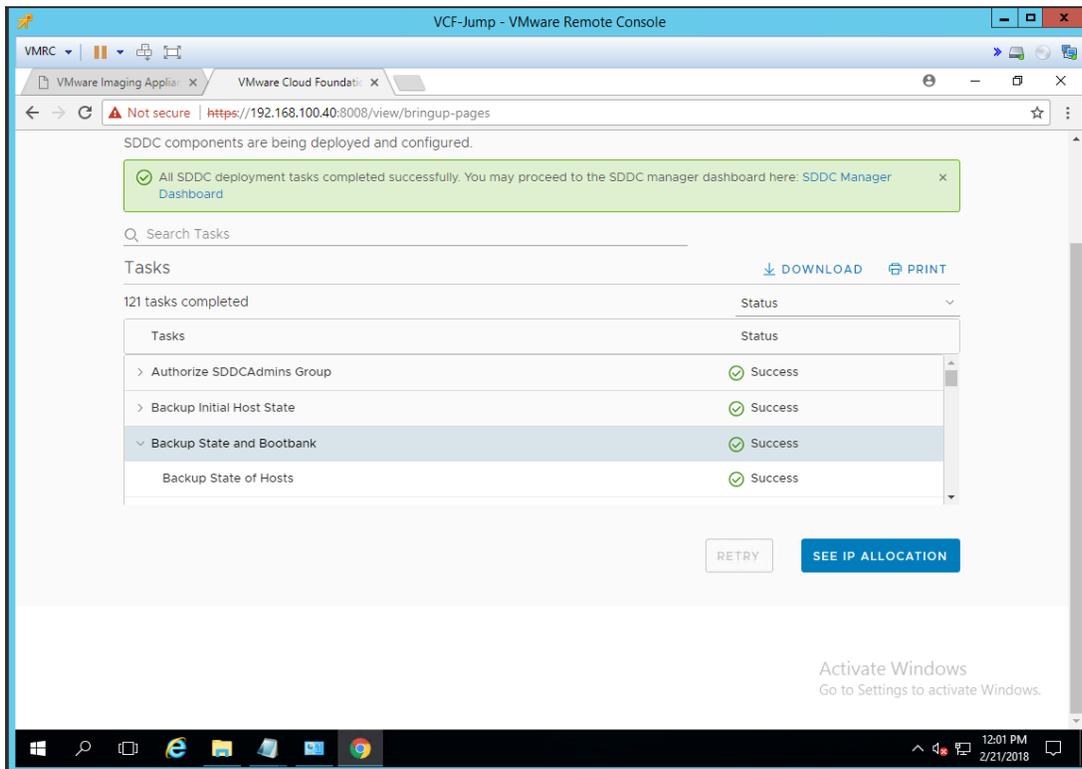


- After selecting “Finish,” the SDDC BringUp process will begin and will take some time to complete.

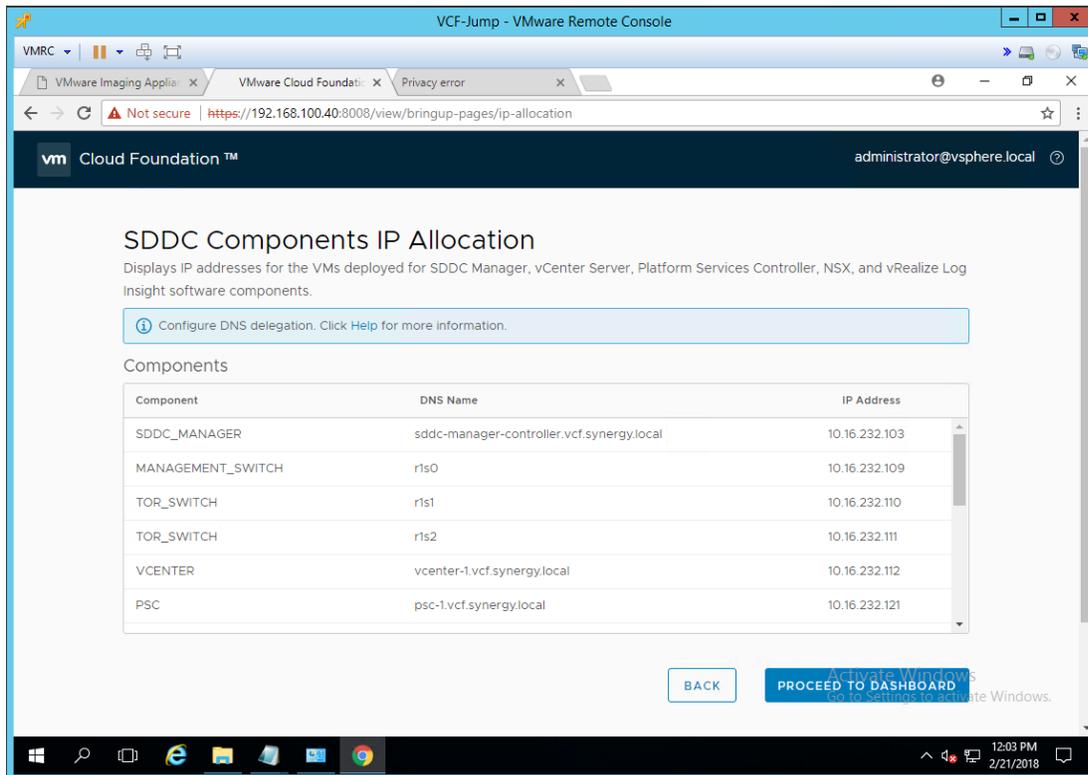
19.



20. The entire deployment will take roughly 2 hours.



21. Upon completion, check the "See IP Allocation" and record the value for the SDDC Manager VM.



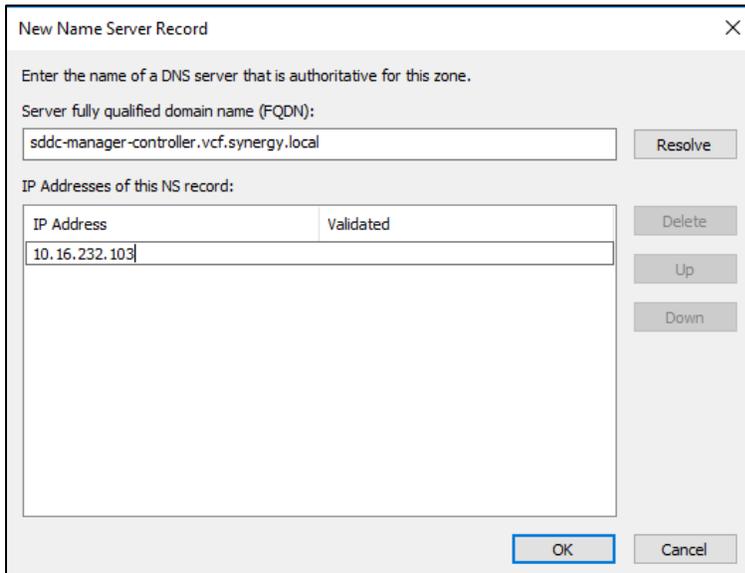
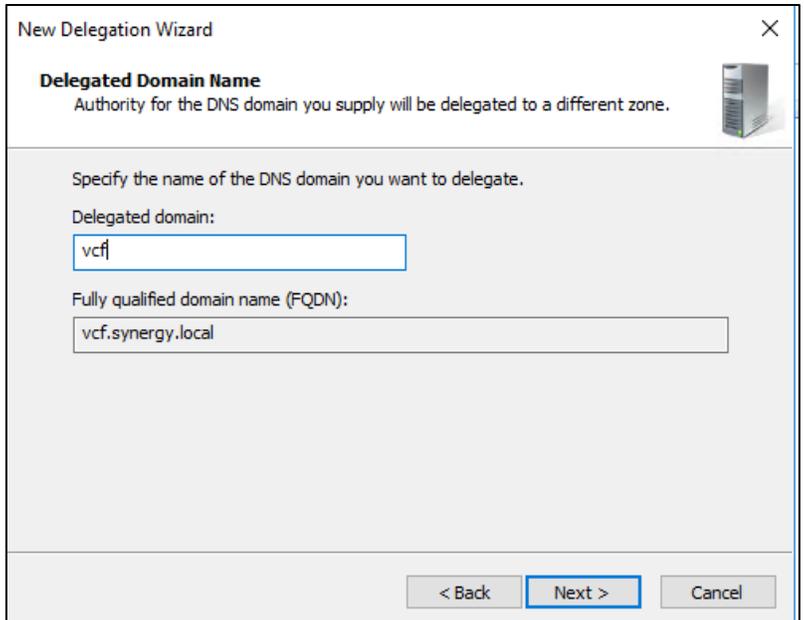
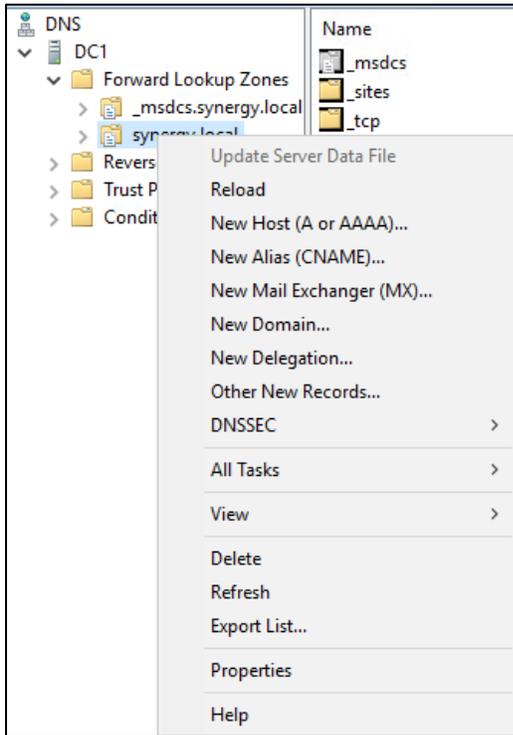
Configure DNS Delegation

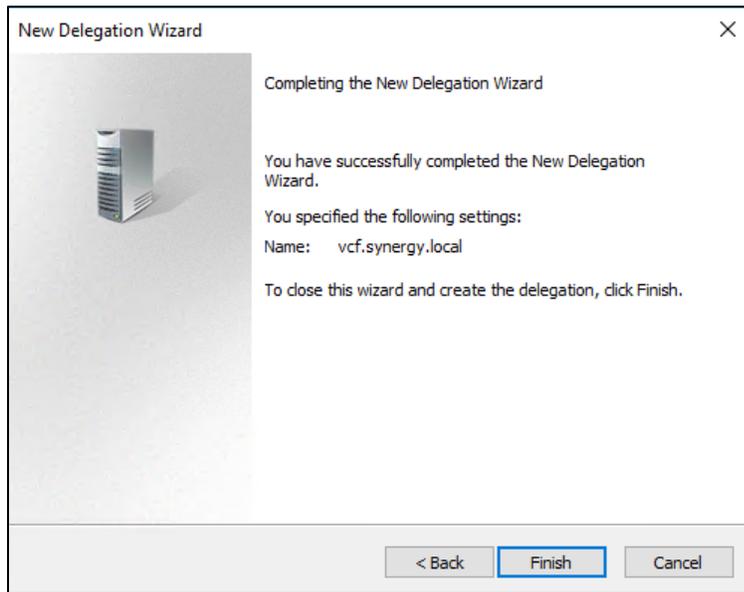
For quick validation the DNS server of the jump vm or the laptop from where the SDDC manager will be accessed from, can be updated to point to the SDDC Manager Controller vm.

SDDC Manager uses Unbound (a DNS server software) for name resolution during the Cloud Foundation bring-up. The customer must now configure the corporate DNS server to delegate zone control for the Cloud Foundation domain to SDDC Manager.

For example, if the corporate domain is synergy.local, and the Cloud Foundation Sub Domain is subdomain.synergy.local, the corporate DNS server must be configured to delegate control of subdomain.synergy.local to SDDC Manager.

1. Install DNS on your server by adding a new role through Server Manager and selecting DNS.
2. Ensure that your jump server uses the local DNS for name resolution.
3. Configure the primary zone (vmware.corp) as a zone managed by Windows DNS.
4. Right-click the zone and select New Delegation.
5. Enter the name of the sub-domain (subdomain in our example).
6. In the Server fully qualified domain name (FQDN) field, type the IP address of SDDC Manager and click Resolve.
7. Click OK.
8. The new zone appears as a delegated zone under your primary domain.





9. Test a ping to the platform services controller VM.

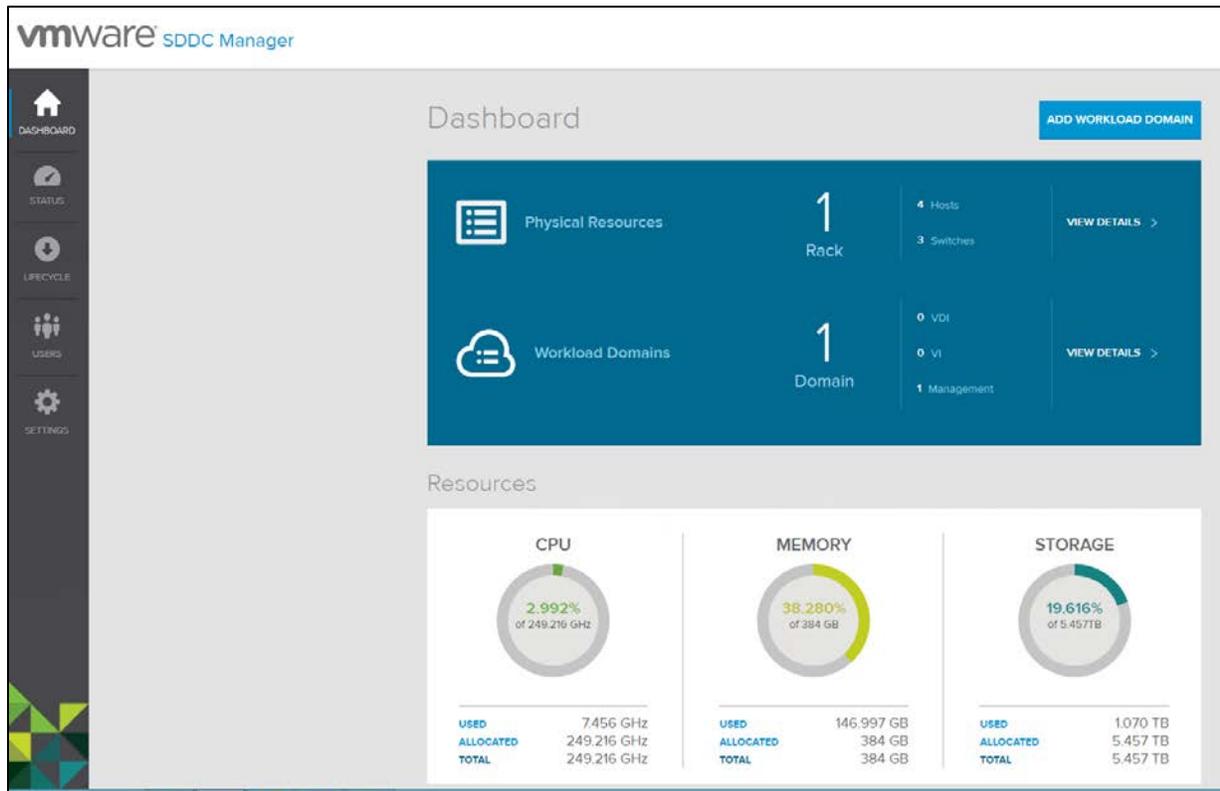
```
C:\>ping psc-1.vcf.synergy.local

Pinging psc-1.vcf.synergy.local [10.16.232.121] with 32 bytes of data:
Reply from 10.16.232.121: bytes=32 time<1ms TTL=63

Ping statistics for 10.16.232.121:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\>_
```

10. Once DNS Delegation tasks are complete go back to the BringUP UI and click "Proceed to Dashboard"



The VCF Management Domain is now deployed.

At this point, additional servers can be imaged, if they were not imaged previously.

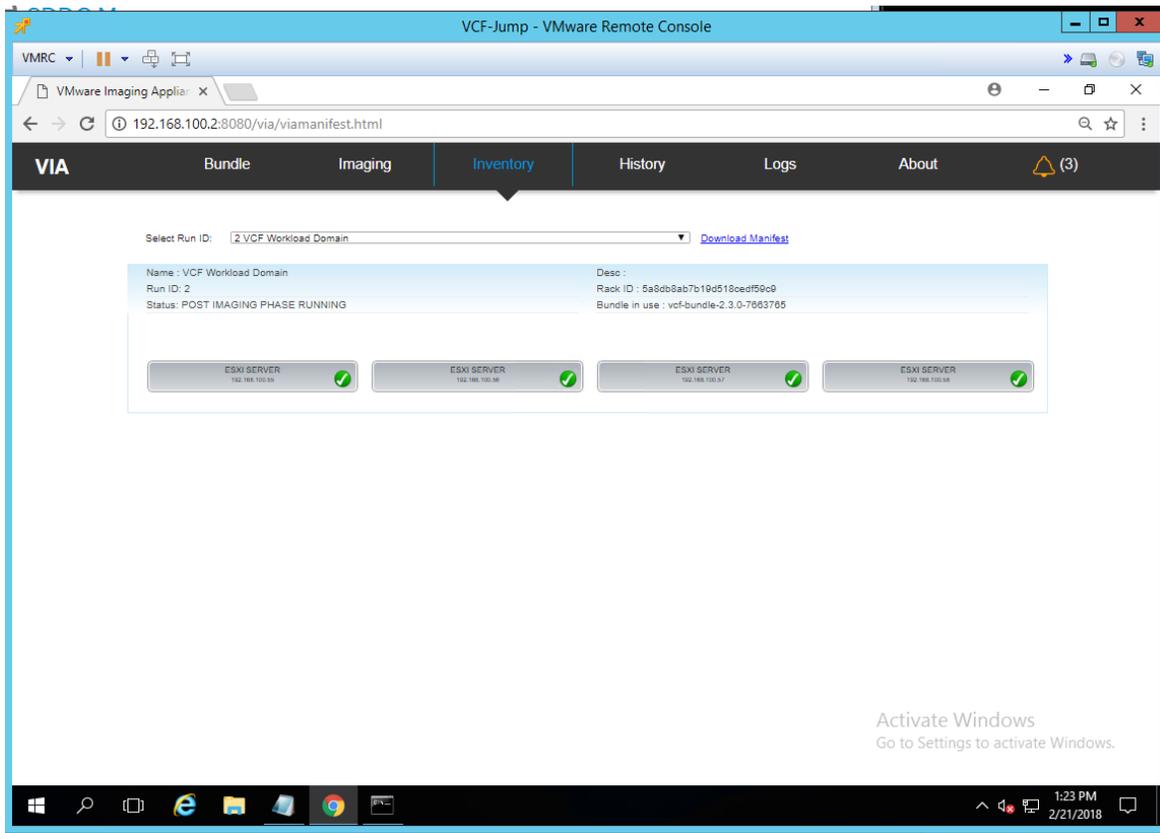
NOTE: To image additional servers, go back to the "VIA", choose "individual deployment", do not select primary esxi host, and select the number of servers then continue in a similar fashion to the first servers that were imaged.

Importing additional imaged hosts to VCF's inventory

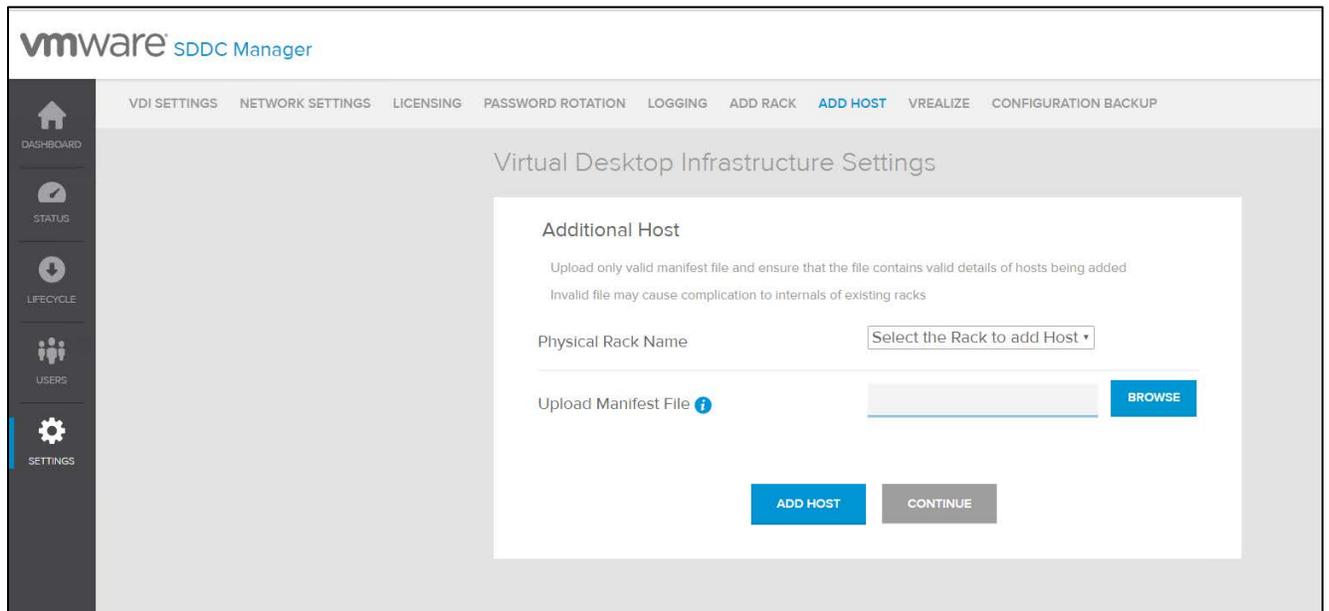
If additional hosts were imaged separately, it is important, post-imaging, to save the "manifest" and import it into VCF.

Following a successful image run of additional nodes, navigate to the "Inventory" tab in the VIA Imager.

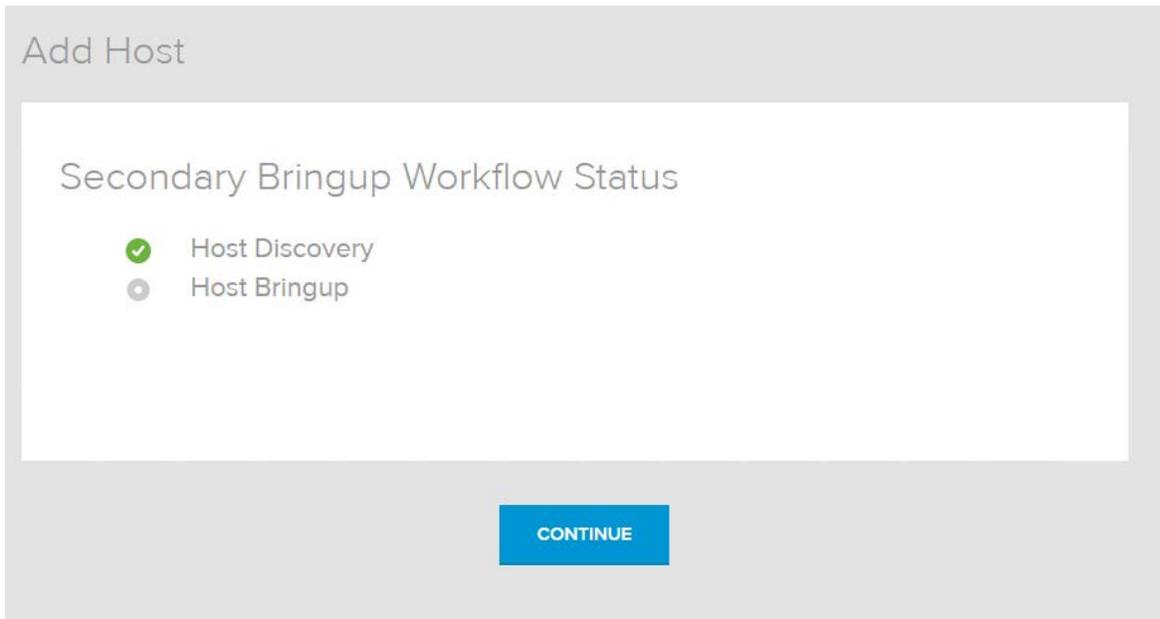
1. Choose "Download Manifest" and save the downloaded *.tgz file.



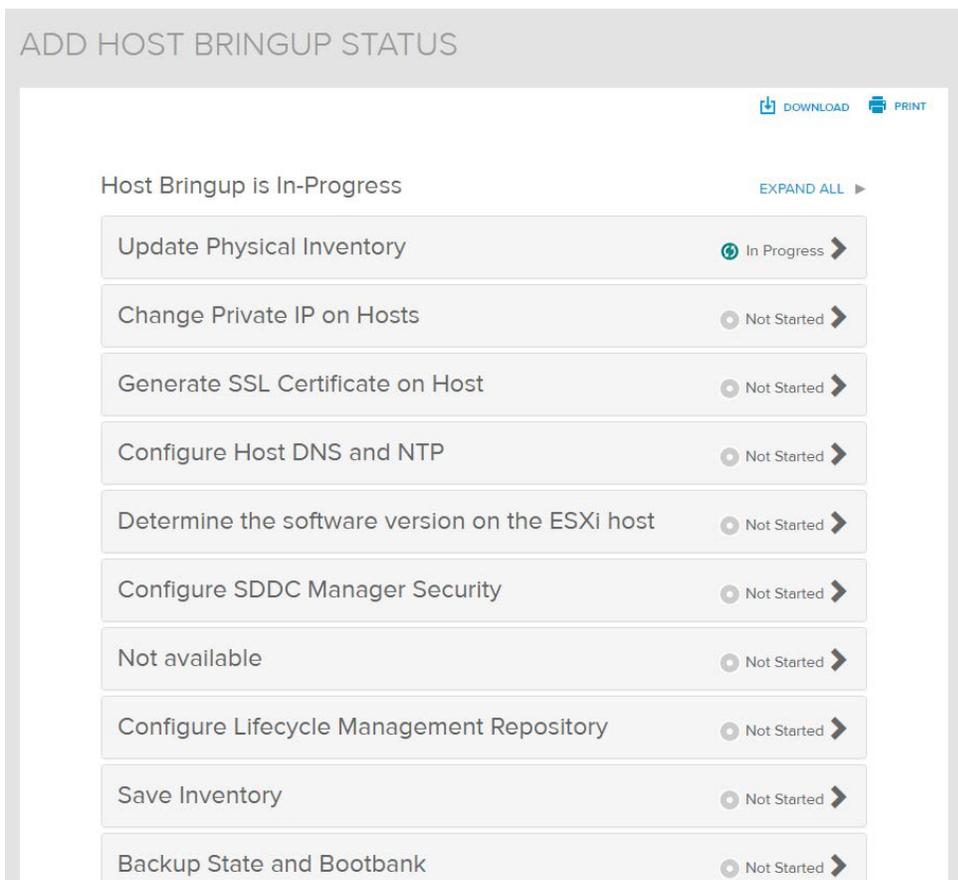
2. In the SDDC Manager Controller UI, navigate to “Settings” and then “Add Host”.



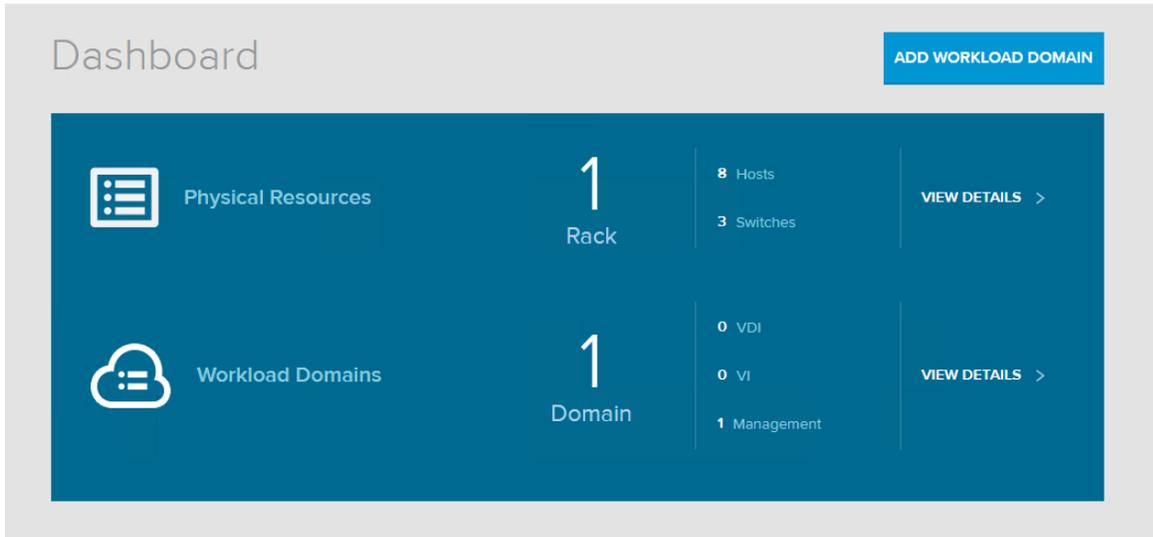
- 3.
4. Choose the appropriate Physical rack, browse to the TGZ file saved in the previous step and choose “add host”.
5. The screen will pause for a moment on “Host Discovery” and when prompted, choose “Continue”.



- The Host Bring Up will now begin and the progress can be monitored.

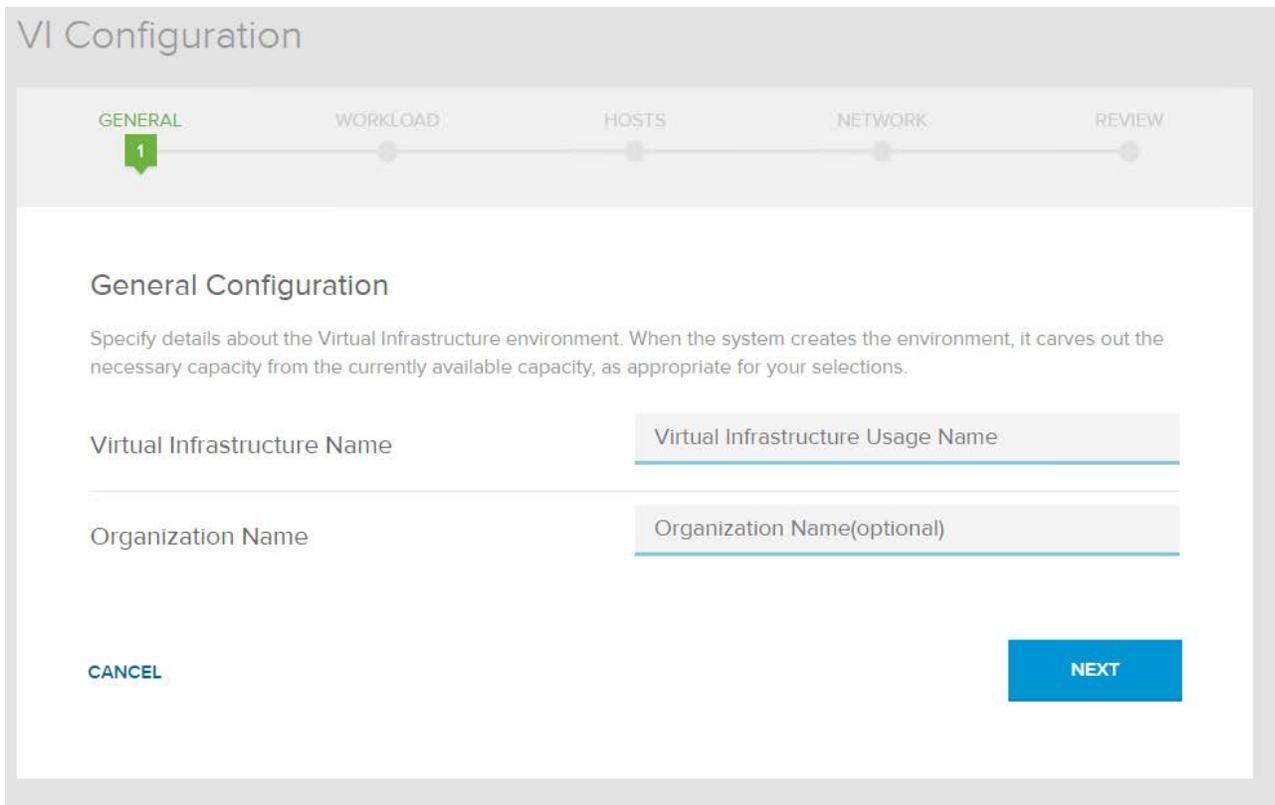


- Once the discovery finishes, return to the Dashboard and the additional hosts will be available.

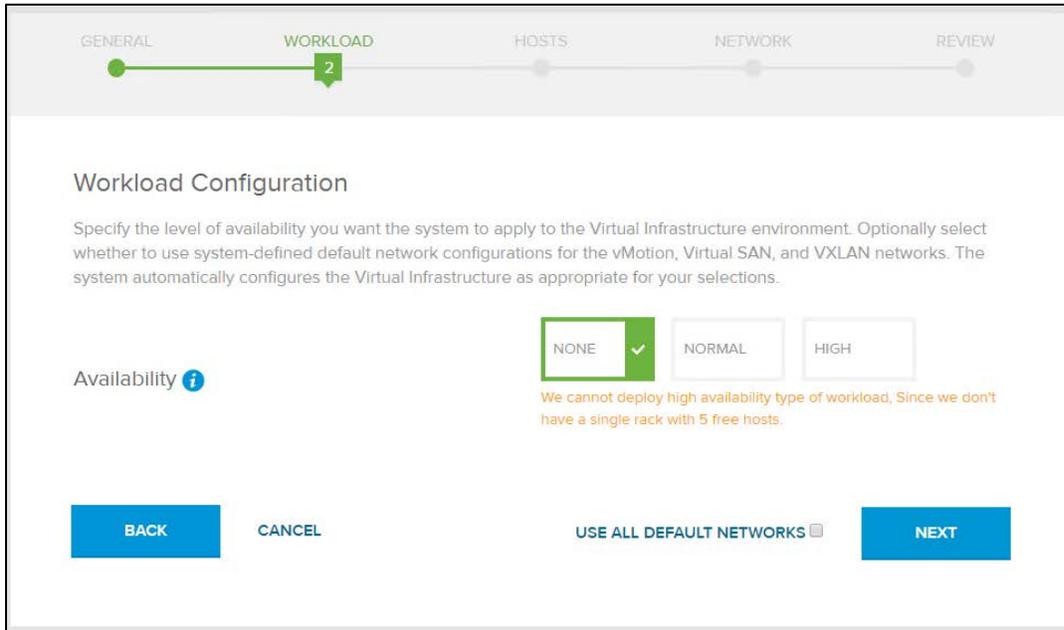


Adding a Workload Domain to VCF

1. Click the Add Workload Domain button.
2. Choose the appropriate type of Workload Domain (VDI or VI).
 - a. **NOTE:** In this example, it is a “Virtual Infrastructure” workload domain.



3. Choose the Workload Configuration appropriate to your needs. Note: Do not choose “Use all Default Networks” during this wizard. Leave this button blank.



4. Select available imaged hosts to use in this workload domain. (Minimum of 3 required)

Host Selection

Calculated vSAN FTT=0.. A host failure with current configuration will lead to data loss. Increasing availability will solve this. ✕

Hosts with different memory size selected.. HA failover between hosts may be impacted. For more information on HA constraints refer to product documentation. ✕

Hosts with different CPU types selected.. vMotion between hosts may not work as Enhanced vMotion(EVC) is not enabled. For enabling EVC manually please refer to product documentation. ✕

Hosts with different disk sizes selected.. vSAN performance may be impacted. For more information refer to vSAN design and sizing documentation. ✕

Total Selected Resources : Hosts - 4, Cores - 116 CORES, Memory - 704 GB, Disk - 6.48 TB, CPU - 274.404 GHZ

Show only selected Host(s) CLEAR CUSTOM FILTERS CLEAR SELECTIONS

HOST ID ESXi VERSION	PORT #	HOSTS	MEMORY (GB)	PROCESSOR	TOTAL STORAGE (GB)	DISKS	STORAGE TYPE
<input checked="" type="checkbox"/> R1N5 unknown	vmnic0	HP Synergy 480 ...	96	intel Intel(R) Xeon(R)...	1620	1xHP LOGICAL V...	HYBRID
<input checked="" type="checkbox"/> R1N6 unknown	vmnic0	HP Synergy 480 ...	96	intel Intel(R) Xeon(R)...	1620	1xHP MO0200JE...	HYBRID
<input checked="" type="checkbox"/> R1N7 unknown	vmnic0	HP Synergy 480 ...	256	intel Intel(R) Xeon(R)...	1620	1xHP MO0200JE...	HYBRID
<input checked="" type="checkbox"/> R1N8 unknown	vmnic0	HP Synergy 480 ...	256	intel Intel(R) Xeon(R)...	1620	1xHP MO0200JE...	HYBRID

5. For the next several screens, review the Management, vMotion, vXLAN VLAN and IP address items. For all networks except for the vSAN network, an individual “use defaults” button can be selected. For vSAN, however, it’s important to use a discrete vSAN VLAN ID

per workload domain. If this is a new VLAN ID, it will be necessary to switch to the Synergy Composer and verify that the VLAN ID is included in the Network Set on the profiles. If it is not, it can be added at this time.

Management Configuration

Review the management network settings that will be used for this Virtual Infrastructure environment. The management network settings are read-only because the deployed Virtual Infrastructure environment will use the system's already-configured management network.

VLAN ID <i>i</i>	<input type="text" value="302"/>
Subnet <i>i</i>	<input type="text" value="10"/> - <input type="text" value="16"/> - <input type="text" value="232"/> - <input type="text" value="0"/>
Subnet Mask <i>i</i>	<input type="text" value="255"/> - <input type="text" value="255"/> - <input type="text" value="255"/> - <input type="text" value="0"/> <small>Subnet mask bit length is recommended to be less than or equal to 22</small>
Gateway <i>i</i>	<input type="text" value="10"/> - <input type="text" value="16"/> - <input type="text" value="232"/> - <input type="text" value="1"/>
DNS <i>i</i>	<input type="text" value="10"/> - <input type="text" value="16"/> - <input type="text" value="232"/> - <input type="text" value="103"/>
Secondary DNS <i>i</i>	<input type="text" value="xxx"/> - <input type="text" value="xxx"/> - <input type="text" value="xxx"/> - <input type="text" value="xxx"/>
NTP <i>i</i>	<input type="text" value="time.nist.gov"/>
Exclude IP Address Ranges <i>i</i>	START IP ADDRESS <input type="text" value="xxx"/> - <input type="text" value="xxx"/> - <input type="text" value="xxx"/> - <input type="text" value="xxx"/>
	END IP ADDRESS <input type="text" value="xxx"/> - <input type="text" value="xxx"/> - <input type="text" value="xxx"/> - <input type="text" value="xxx"/> + CLICK TO ADD

VI Configuration

GENERAL WORKLOAD HOSTS **NETWORK** REVIEW

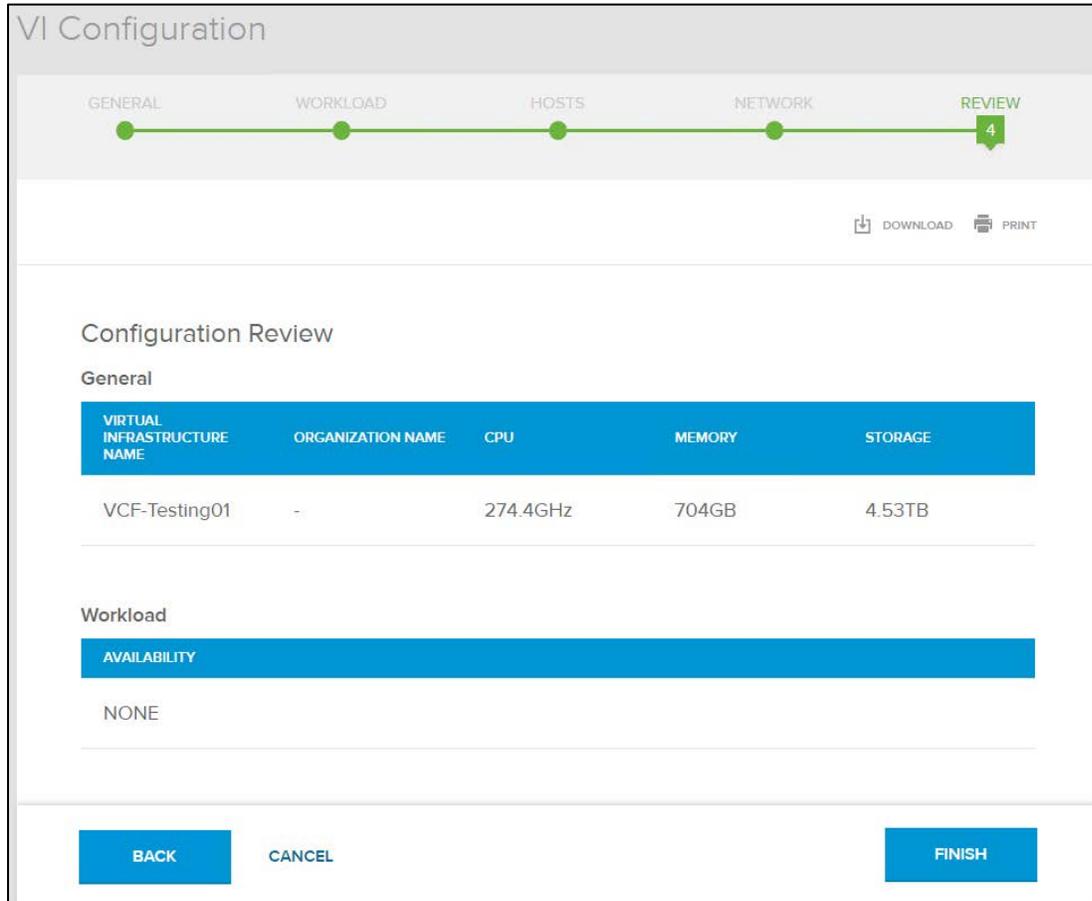
MANAGEMENT vMotion Virtual SAN VXLAN DATA CENTER

Data Center Connections

Data Center Connections won't be validated until later in the process. Please review them carefully.

Data Center Network <i>i</i>	<input type="text" value="PUBLIC"/>
Configuration Name <i>i</i>	PUBLIC
VLAN ID <i>i</i>	300
Subnet <i>i</i>	10.16.230.0
Subnet Mask <i>i</i>	255.255.255.0
Gateway <i>i</i>	10.16.230.1

6. Select the datacenter network “Public” from the drop-down menu. This was created during the initial bringup. If a new network is needed, first provision that on OneView before creating it on SDDC Manager.
7. Review the configuration and choose finish.



8. VCF management and 1 workload domain is now finished and the system can now be consumed.

NOTE: Consumption of and use cases for VCF are out of scope for this paper. The intended purpose is to get the customer started on their installation and first initialization of VCF and HPE Synergy. Consult VMware services for further help with VMware Virtual Cloud Foundation.

Troubleshooting

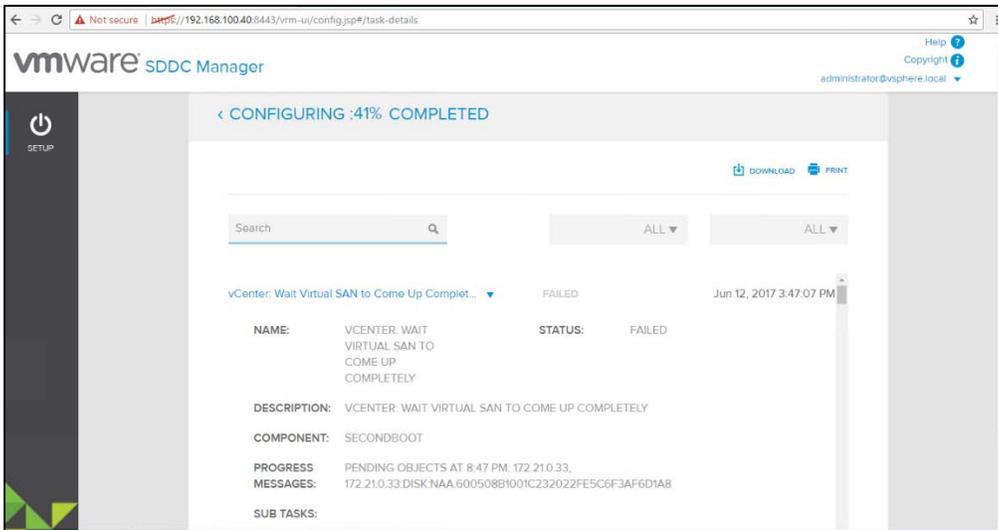
Issue 1: UEFI Boot does not start PXE Installation

Resolution: Use Legacy Boot Mode

Use Legacy Boot Mode on OneView Service Profile. UEFI is not validated with VCF at the moment. This is described in the “Configuring HPE OneView Server Profile Templates” section of this document.

Issue 2: VRM Error: SSD in use - 4 of 6, pending 172.21.0.33, 172.21.0.33:DISK:naa.600508b1001c232022fe5c6f3af6d1a8

Resolution: Possible Sever Profile storage configuration inconsistency



Ensure all Server Profile have identical Storage configuration.

Issue 3: vSAN Datastore configuration fails and vSAN Datastore capacity displayed is 0GB even though devices are present on the P416mi-e Controller.

Resolution: This issue may occur if the disks were previously used for another vSAN configuration or metadata exists such as prior RAID.

1. If the ESXi Server can list all the devices on the Storage Adapter for P416mi-e Controller, it is possible that the disks were used by some other vSAN installation. This requires cleanup of the disks.
2. To clean up the disks run login to the ESXi Server using SSH and run the following commands.

List all devices and note their device name. Note: Ensure that the local boot disk must not be wiped or ESXi installation will be lost

`esxcfg-mpath -L` (Sample output below)

```
[root@rack-1-n4:~] esxcfg-mpath -L
vmhba0:C0:T0:L1 state:active naa.600508b1001c2a4ea24dd090a5e3187a vmhba0 0 0 1 NMP active san sas.5001438037642fb0
vmhba1:C0:T0:L1 state:active naa.600508b1001c2eed7a694ed948ea8c2c vmhba1 0 0 1 NMP active local sas.500143803755bac0
vmhba1:C0:T0:L2 state:active naa.600508b1001c090e0eff971e58f1f710 vmhba1 0 0 2 NMP active san sas.500143803755bac0
vmhba1:C0:T0:L3 state:active naa.600508b1001cbea6085875a5a8770220 vmhba1 0 0 3 NMP active san sas.500143803755bac0
vmhba1:C0:T0:L4 state:active naa.600508b1001c06567b8cc02891a9a85b vmhba1 0 0 4 NMP active san sas.500143803755bac0
vmhba1:C0:T0:L5 state:active naa.600508b1001c760da13c96f3f4721ca9 vmhba1 0 0 5 NMP active san sas.500143803755bac0
[root@rack-1-n4:~] █
```

Delete all partition in the D3940 Disk Drive

`for I in 1 2 3 4 5 6 7 8 9; do partedUtil /dev/disks/ naa.600508b1001c760da13c96f3f4721ca9 $I; done`

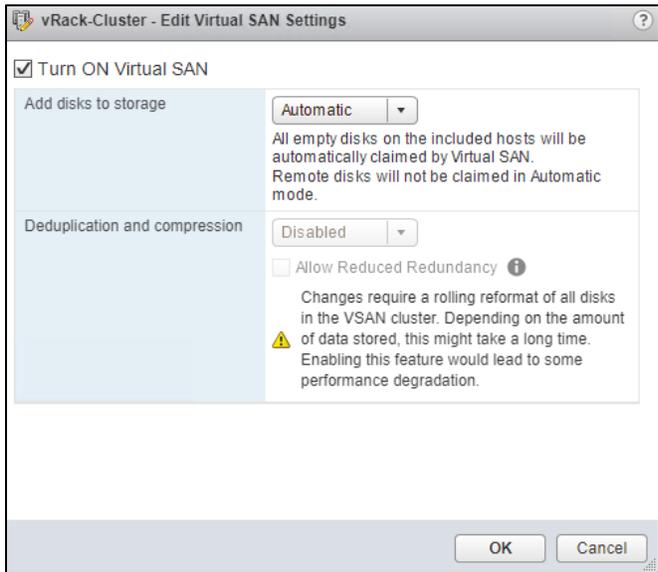
#Rescan all storage adapter

`esxcfg-rescan -A`

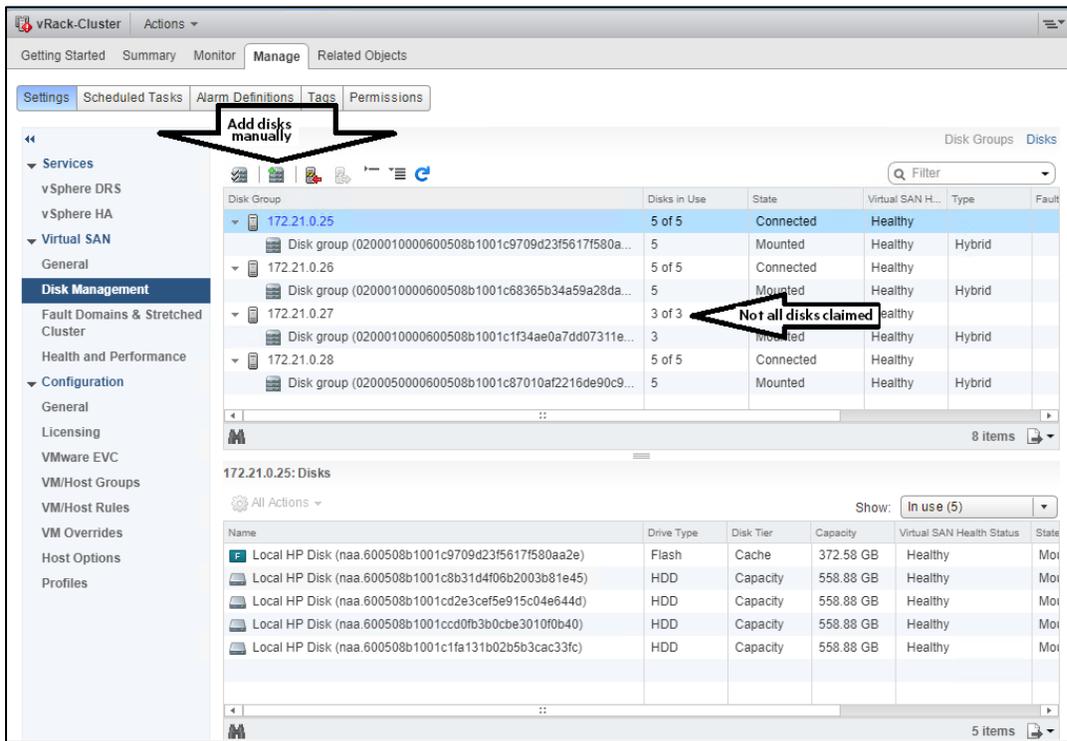
Issue 4: SDDC Manager Initial Setup fails to bring up the vSAN Cluster successfully

Resolution: Occurs if one of the host is unable to claim disks and create diskgroup in automatic mode.

1. Perform steps in "Issue 3: vSAN Datastore configuration may fail. vSAN Datastore capacity displayed is OGB even though devices are present on the P542D Controller" to clean up the vSAN Disks.
2. Disable Automatic Claim on the vSAN Cluster.



3. Select host which does not have enough disks. Click on add disks SSD and HDD. vSAN will automatically add the disks in the appropriate Tier.



4. Re-enable Automatic Mode for the vSAN cluster.

Issue 6: vCenter “Wait for vSAN” to come up completely

Resolution: This may occur if the ESXi host is unable to detect or mount the disks in a disk group.

Follow the documented steps for [Issue 5: SDDC Manager Initial Setup fails to bring up vSAN Cluster successfully](#).

Appendix – Important Links

VMware Documentation

1. [VMware VCF Install Files](#)
2. [VMware VCF Imaging Appliance deployment and Install Guide](#)
3. [VMware VCF Bring-Up Guide](#)
4. [VMware vSAN Disk Requirements KB 2106708](#)

HPE Documentation

1. [HPE Synergy System \(1 Frame\) Setup](#)
2. [HPE Synergy with VMware vSAN Best Practices Guide](#)

Learn more at <http://www.hpe.com/info/synergy-docs>



Sign up for updates



© Copyright 2018 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Add trademark acknowledgments as needed. For trademark policy, see <https://hpe.sharepoint.com/sites/F4/OGC/Pages/tradeack.aspx>

ENW, June 2018, Rev. 2