

DEPLOYMENT GUIDE

VMware Cloud Foundation on FlashStack®

VMware-validated, Design-certified Partner Architecture

Contents

Introduction3

Document Reference3

VMware Cloud Foundation.....3

 FlashStack.....4

 SmartConfig.....4

FlashStack Automation with SmartConfig and VMware Cloud Foundation4

 Part 1: Deploy FlashStack with SmartConfig5

 Part 2: Deploy VMware Cloud Foundation23

 Part 3: Using Pure Storage FlashArray with VMware Cloud Foundation.....26

Support.....51

Related Information51



Introduction

Digital transformation continues to change how IT organizations provide services to support business objectives. This transformation brings new applications and integration requirements that will challenge IT infrastructure and delivery processes. More than ever, IT organizations need to be agile and have accelerated provisioning of IT services. IT teams must deploy modern tools, platforms, and infrastructure to support rapid demands.

[Pure Storage®](#) is an industry pioneer and leader that offers enterprise solutions that are simple to manage and provide a modern data experience. The Pure Storage Purity operating environment is the software-defined engine of Pure Storage FlashArray™. Purity is the driver that enables Pure FlashArray products to deliver comprehensive data services for all your traditional and modern data-center applications. Pure1® delivers cloud-based management and support of Pure Storage systems with global predictive intelligence, workload and capacity planning, analytics, and more.

This Deployment Guide provides instructions to deploy VMware Cloud Foundation (VCF) platform on FlashStack® with [SmartConfig](#) and VCF Software-Defined Data Center (SDDC). With VCF, organizations can easily deploy, manage, and support virtualized applications across private and public clouds. FlashStack solutions provide architectural blueprints to guide design considerations and planning for Day 0. VMware SDDC provides intelligent provisioning of infrastructure resources that can accelerate Day 1 and Day 2 activities. Management and Workload Domains will be up and running quickly, readily available to take on production workloads.

Document Reference

This deployment guide is based on a series of Pure Storage Knowledge Base (KB) articles. For the latest updates, visit [Pure Storage VCF on FlashStack Deployment Site](#) or [Pure Storage VCF How-To's Support Site](#).

VMware Cloud Foundation

VMware Cloud Foundation (VCF) offers an immense advantage to data center professionals in terms of simplifying Day 0 through Day 2 activities, streamlining management operations, and providing agility in deploying and decommissioning new environments within the vSphere ecosystem. Administrators simply provide imaged ESXi hosts and a few DNS/IP address entries as the input and receive a fully functional vCenter domain backed by NSX-T as the output. This domain is immediately



ready for whatever use-case or use-cases needed by the tenant organization. Additional hosts can be dynamically added or removed as requirements change and other VMware solutions such as Horizon or Tanzu can be simply layered on top of the initial deployment through the common SDDC Manager control plan. Native deployment, integration, and connectivity are also provided with members of the vRealize suite, also via SDDC Manager.

FlashStack

FlashStack solutions are engineered jointly with Cisco, to architect once and simplify deployment. Cisco Validated Designs provide architectural emphasis to cable once, upgrade in place, and scaling with no disruption. For VCF, SmartConfig maintains best practices by initializing and automating the deployment of FlashStack compute, network, and storage to the VMware vSphere ESXi environment.

SmartConfig

Cisco Validated Designs are used to record all the best practices and followed by SmartConfig to automate and simplify FlashStack deployment. SmartConfig simplifies a FlashStack deployment in four steps, as seen in Figure 1.




Figure 1. FlashStack deployment steps.

FlashStack Automation with SmartConfig and VMware Cloud Foundation

The goal of this exercise is to begin with a blank canvas (i.e. FlashStack hardware that has been racked, powered, cabled, and factory reset) and end with a fully-functional VMware environment that is ready for production workloads. This document builds the entire process with repeatable, customizable blueprints which in almost all cases completely eliminate the need for the administrator to touch any single piece of the underlying setup. Perhaps even more important is that the bulk of the steps outlined in each section of this guide can be exported to a JSON and then repeated or updated based upon unique client environment requirements over and over again.

FlashStack components to be used for your SmartConfig and VMware Cloud Foundation need to be racked, powered, and cabled up. It is required to cable all FlashStack components together according to the topology diagrams provided in the [SmartConfig Deployment Tool Getting Started Guide](#).

Below you'll find an introduction to the key technologies in play, followed by an in-depth deployment, divided into three core parts:

1. **Deploy FlashStack with ESXi via SmartConfig.** The input of this section will be factory reset Cisco and FlashArray hardware and the output will be a fully functional imaged/zoned/deployed UCS chassis with ESXi7 installed and ready for use with VMware Cloud Foundation.



- 2. **Build VMware Cloud Foundation SDDC Manager on FlashStack.** The primary input for CloudBuilder is, not ironically, the output of our work in part 1. Specifically, ESXi hosts and their underlying infrastructure, from which we will automatically deploy a Management Domain with CloudBuilder.
- 3. **Deploy a VMware Cloud Foundation Workload Domain with Pure Storage FlashArray as Principal Storage (VMFS or vVols on FC).** Options such as vVols with iSCSI are covered in additional KB articles in the VMware Cloud Foundation section of the Pure Storage support site.

Goals of FlashStack Deployment Automation

Overarching Goal: Provide a continuous process for end-to-end deployment of FlashStack with VMware Cloud Foundation. Minimize or eliminate touching a single component via an automated, repeatable and customizable framework.

FlashStack	VMware Cloud Foundation	Workload Domain
Deploy FlashStack with ESXi via SmartConfig. The input for SmartConfig is unconfigured Cisco and FlashArray hardware and the output will be a fully functional deployed FlashStack with ESXi7 installed and ready for use with VMware Cloud Foundation.	Build VMware Cloud Foundation SDDC Manager on FlashStack. Deploy four ESXi hosts as VMware Cloud Foundation Management Domain with CloudBuilder with VCF 4.0*/VVD6.0 *VCF 4.1 Required for vVols as Principal Storage	Use additional ESXi hosts for a new level of flexible Workload Domains with Pure Storage FlashArray as Principal Storage (VMFS or vVols on FC).

Figure 2. Goals for FlashStack deployment automation.

Post-deployment, clients will enjoy the benefits of single-click upgrades for the bulk of their UCS and VMware components and the ability to dynamically scale up or down their Workload Domain deployment resources independently or collectively, based upon specific needs (e.g. compute/memory, network and/or storage).

Part 1: Deploy FlashStack with SmartConfig

For the case of our deployment covered in this guide, the component diagram and cabling in use is shown below. FlashStack connectivity is using Fibre Channel.






Figure 3. FlashStack components and cabling.

These additional items are needed before you begin Part 1:

1. UCS FIs, MDS, and Nexus components that are factory reset and in DHCP discovery mode.
2. Cisco UCS, MDS, and Nexus firmware and kickstart files.
3. Pure Storage FlashArray, initialized in one of the following methods:
 - a. Auto-Discovery FlashArray, which must be factory reset Pure Storage FlashArray, or newly shipped FlashArray with SmartConfig 1.4
 - b. Manual FlashArray with management IP on the SmartConfig
4. [ESXi 7 ISO](#).



5. Untagged native VLAN for SmartConfig (/24 minimum). Typically this is accomplished via plugging the management interfaces of the various components into a dumb switch.
6. Separate, previously deployed ESXi, Hyper-V, or other bare-metal host or other laptop/desktop where the SmartConfig and VMware Cloud Foundation OVAs can be deployed with connectivity to the above untagged native network.
7. SmartConfig OVA deployed on the above system connected to the untagged native network. SmartConfig also needs a static IP address assigned to it within the console.
8. NTP Server on the private network.
9. Minimum of three routable production-defined VLANs for use with VMware Cloud Foundation (for ESXi Management, vMotion, and VMware vSAN).
10. Edit rights over a Windows DHCP scope and DNS zone.

NOTE: Static IP addressing and forward/reverse DNS zoning will certainly work for ESXi hosts with other solutions outside of Windows. However, these two items are required if following the steps outlined in the 'Prepare ESXi Hosts for VMware Cloud Foundation' chapter at the end of this section with PowerShell.

Once the above items are verified as being available and online the first step is to navigate to the IP address that has been assigned to the SmartConfig virtual appliance.

SmartConfig Discovery

This first phase of our deployment will be done exclusively through the static IP address assigned to the SmartConfig OVA.

Enable the DHCP scope within SmartConfig so that the UCS, MDS, and Nexus components in our FlashStack can be assigned an address from the DHCP IP pool for configuration and deployment. This is accomplished by clicking on the Enable DHCP Server for Auto-Discovery radio button located in the top-right area of the GUI.





Figure 4. Enabling DHCP server for auto-discovery.

Most DHCP server fields should be automatically populated from when the networking for SmartConfig was originally assigned. Select the DHCP range you want the FlashStack components to reside within.



DHCP Settings [X]

Make sure static range do not conflict with used IP Addresses.

Network: 192.168.108.0

Subnet Mask: 255.255.255.0

Gateway: 192.168.108.1

IP Address: 192.168.108.11

DHCP & Static Range: *

Range: 12 to 64 (DHCP RANGE) | 64 to 254 (STATIC RANGE)

Cancel Submit

Figure 5. Selecting the DHCP range.

The next step is to manually add **or** use DHCP auto-discovery¹ of the FlashArray to the SmartConfig inventory so that it can be used in subsequent deployment steps.

To manually add the FlashArray, click on the hamburger button just to the right of the DHCP scope radio button and select the Add FlashArray option.





Figure 6. Adding the FlashArray.

Here we add the VIP for the FlashArray management connection and the pureuser credentials.



Add FlashArray [X]

IP Address: * 192.168.108.200

Username: * pureuser

Password: *

Figure 7. Adding the VIP for the FlashArray.

¹ Only available with SmartConfig version 1.4 and factory reset FlashArray.

While we wait for the various FlashStack components to be picked up on the DHCP scope, adding the various ISOs and kickstart scripts used to initialize, deploy and configure FlashStack can be loaded into the SmartConfig ISO library. This repository not only contains all of the various UCS and Nexus operating environments but also is where we will stage our installation of ESXi at the end of the SmartConfig phase of this deployment. The ISO library is in the same menu where we added the FlashArray in the last step:




Figure 8. ISO Library.

Multiple firmware versions for all pieces can be housed within this library and used to deploy whatever FlashStack configuration is required. The ISO library is divided into two sections; the top section shows firmware for Cisco components and provides the ability to upload additional operating systems for the current selection.




Figure 9. ISO Library firmware section.

There is also a section for the available operating systems that can be installed and their respective kickstart scripts:




Figure 10. ISO Operating system section.



The bottom section shows all previously uploaded files to the library and gives the ability to delete them if they are no longer needed or require an update:

File Name	Type	Action
nxos.7.0.3.i4.2.bin	Nexus 9k	
VMware-VMvisor-Installer-201912001-15160138.x86_64.iso	ESXi	
m9100-s5ek9-mz.7.3.0.DY1.bin	MDS	
m9100-s5ek9-kickstart-mz.7.3.0.D11.bin	MDS kickstart	
m9100-s5ek9-kickstart-mz.8.2.1.bin	MDS kickstart	
m9100-s5ek9-mz.8.2.1.bin	MDS	
ucs-6300-k9-bundle-infra.4.0.4b.A.bin	Infra Image	
ucs-k9-bundle-b-series.4.0.4b.B.bin	Blade Image	

Figure 11. Previously uploaded files.

After just a few minutes, we can see that all our Cisco components have been picked up by DHCP and given an address. The various FlashStack component options available for deployment are a selectable option across the top. This enables potentially many different Cisco systems to be staged and gives the deployment administrator the ability to pick only those components to be used for the build at hand. Clicking on a component will toggle it for selection, or clicking on the configuration option you want automatically will select the required underlying pieces. Our example deployment shows all components highlighted and ready for the next step: configuration. Click on the Next button at the bottom of the GUI to proceed to the next phase of SmartConfig.

Step 1: FlashStack™ Discovery

Enable DHCP Server for auto discovery: ☒

Configuration Options:

FA/PI | FA/PI | **FA/MDS/Nexus 9K** | FA/MDS/PI | FA/PI/Nexus 9K/PI

Devices

● MDS x 2 ● Nexus 9k x 2 ● FlashArray x 1 ● Fabric Interconnect x 2

Device Type	Make/Model	IP Address	Serial Tag
Fabric Interconnect	Cisco UCS-FI-6332-16UP	192.168.108.22	SAL204N780
Fabric Interconnect	Cisco UCS-FI-6332-16UP	192.168.108.21	SAL2023RLBY
Nexus 9k	Cisco N9K-C9372PX	192.168.108.23	SAL938P4L8
Nexus 9k	Cisco N9K-C9372PX	192.168.108.20	SAL9465TWH
MDS	cisco MDS - 1cl	192.168.108.20	JPG951044
MDS	cisco MDS - 1cl	192.168.108.19	JPG 010005U
FlashArray	FAm50i2	192.168.108.200	M_SERIES_PCH1M250059

Figure 12. Components ready for configuration.

SmartConfig Configuration

The configuration phase of FlashStack deployment is where you input specific environmental constructs including management IP addresses, operating system versions, kickstart scripts, and VLANs to be used on the production network post-deployment. The beauty of this phase is that you potentially only need to do it once as you can export all the values you used as a JSON at the end of the deployment and then import it for use on subsequent deployments, letting you automate them. A sample of the JSON used for this guide is available at the following GitHub page [here](#).



DEPLOYMENT GUIDE

The top area of the Basic Manual Configuration section of the Configuration window includes some basic networking information; the administrative password for all Cisco gear; an IP range for KVM IP addresses to the B and/or C-series blades, which FlashArray is to be used with the deployment; and lastly what Operating System (in our case ESXi7) to be installed on top of the UCS servers along with the ESXi kickstart script. The kickstart script for ESXi is important as it sets some key variables so these hosts can easily be used for eventual use with VMware Cloud Foundation.

For more detail on all SmartConfig Configuration fields, please follow [this link](#).

Step 2: FlashStack™ Configuration

Type: ☒ Manual Configuration ☐ Import Configuration

General Information

Mgmt Subnet Mask: 255.255.255.0 Default Gateway: 192.168.108.1

NTP Server: 192.168.108.15 DNS Server: 192.168.108.10

Admin Password: Confirm Password: FlashArray: snl-m502-b08-21

Domain Name: Domain Name

Fabric Interconnect

Compute Type: ☒ Rack Server ☒ Blade Server Operating System: VMware-VMvisor-Installer-7.0.0-15843807-x86_64.iso

Virtual IP Address: 192.168.108.70 Kick Start: KS-cfg

System Name: S UCS Firmware: KVM Console IP Address Range: 192.168.108.69

Mgmt IP Address: 192.168.108.69

Fabric Setup: ☒ A ☐ B

Figure 13. FlashStack Configuration.

The lower section of the basic configuration screen includes firmware version and kickstart selection for the Nexus 9K switches and MDS, as well as management IP address assignment including FlashArray:

MDS

Switch Selection: ☒ UPRISOR Switch Selection: ☒ UPRISOR

Switch Name: mds-a Switch Name: mds-b

IP Address: 192.168.108.67 IP Address: 192.168.108.68

System Image: m9100-v5ek9-mz.8.2.1.bin Kickstart Image: m9100-v5ek9-kickstart-mz.8.2.1.bin

Nexus

Switch Selection: ☒ SALISOR Switch Selection: ☒ SALISOR

Switch Name: nexus-a Switch Name: nexus-b

IP Address: 192.168.108.65 IP Address: 192.168.108.66

Switch Image: nexus.7.0.3.14.2.bin

Figure 14. FlashStack configuration screen continued.

FlashArray (PCHFL16250059)

FlashArray Name: FlashArray01 Virtual IP Address: 192.168.108.200

Controller 0 IP Address: 192.168.108.201 Controller 1 IP Address: 192.168.108.202

Organization Name: Pure Storage Your Name: Kyle Grossmiller

Your Title: Solution Architect Sender Domain: purestorage.com

Alert Email Address(s)**: admin@purestorage.com Timezone: America/Los_Angeles

**Comma separated

Figure 15. FlashStack configuration screen continued.



The advanced section of the Manual Configuration phase gives us the ability to specify what production VLANs, port channels, and SAN connections to be assigned throughout the hardware stack. By default these values align with a Cisco Validated Design, but in our example deployment we are using VLAN2140 for ESXi management, 2137 for vMotion, and 2138 for vSAN (denoted below as ‘Application VLAN’) so they are changed as shown below. More VLANs can be added here for specific application usage.

Step 2: FlashStack™ Configuration

Configuration Type
Type: ☒ Manual Configuration ☐ Import Configuration

Advanced Configuration

Mgmt VLAN id: *	2140	Native VLAN id: *	2
vMotion VLAN id: *	2137	Application VLAN id: *	2138
Application VLAN id: *	2136	Application VLAN id: *	2143
In-Band Interface VLAN id: *	108	Uplink PortChannel F/A: *	151
Uplink PortChannel F/B: *	152	SAN PortChannel F/A: *	1
SAN PortChannel F/B: *	2	VSAN F/A: *	101
VSAN F/B: *	102		

Figure 16. FlashStack advanced configuration options.

With these values filled out we next move on to the Initialization Phase of the deployment. In just a few screens we are already more than halfway done with our SmartConfig deployment!

SmartConfig Initialization

This phase of the process is exceedingly simple - we click on the ‘Initialize’ button at the bottom of the screen and all of the items we put into the Configuration phase previously are loaded and configured onto the various Cisco devices (management IP, upload and/or upgrade selected firmware and kickstart scripts, etc..).

FlashStack™ SmartConfig

Getting Started | Help | About

1 Discovery 2 Configuration 3 Device Initialization 4 Deployment

Step 3: FlashStack™ Device Initialization

#	Device Type	Make/Model	Serial Tag	Status
1	Fabric Interconnect	Cisco UCS-FI-6332-16UP	SAL204N7R0	Configuration in progress
2	Fabric Interconnect	Cisco UCS-FI-6332-16UP	SAL2023RLBY	Configuration in progress
3	Nexus 9K	Cisco N9K-C9372PX	SAL193BP4L8	Upgrading to Version 7.3.3A.2
4	Nexus 9K	Cisco N9K-C9372PX	SAL19465TWH	Upgrading to Version 7.3.3A.2
5	MDS	cisco MDS-101	JPG1951044	Upgrading to Version 8.2.1
6	MDS	cisco MDS-101	JPG2013005J	Upgrading to Version 8.2.1

Initialize

Activate Windows
Go to Settings to activate Windows.

© 2018 Pure Storage Inc. Version 1.3.20190220003

Figure 17. Loading configuration items into Cisco devices.



SmartConfig Deployment

This final step within SmartConfig will:

- Complete our Nexus switch configuration.
- Create and associate service profiles with our UCS hosts.
- Configure and zone the MDS switches.
- Create a boot from SAN policy.
- Create a boot LUN for ESXi on the Pure Storage FlashArray.
- Create a host group and a shared Fibre Channel volume on the FlashArray for immediate use post-deployment.

For use cases that match the CVD, simply click on the Deploy button within the basic workflow and all steps will complete automatically. Imported JSON configurations will also run without the need for user interaction.

However, just as in the Configuration phase of the SmartConfig process, there are advanced options available to help customize items that may not match a Cisco Validated Design. As with the Configuration phase, these customizations can be captured and exported as a JSON post-deployment and used to automate subsequent FlashStack deployments. For this example, we will show the individual steps needed to customize our deployment for VMware Cloud Foundation within our lab.

VMware Cloud Foundation requires that the Management, vMotion, and vSAN VLANs are all aggregated on ESXi vmnic0 and vmnic1 physical ports as part of a distributed switch. To allow for this, we first switch into 'Advanced Mode' within the Deployment phase of SmartConfig to make a couple of changes to the Nexus switch port channels and later to the UCS service profile management vNICs.

To get into the Advanced mode of the Deployment phase, click on the 'Advanced' button towards the top-right of the SmartConfig GUI. Doing so will unlock the ability to edit the individual Deployment workflows as shown below.




Figure 18. Advanced mode in the deployment phase.



Clicking on the edit button for the Configure Nexus Switches workflow gives us a graphical representation of the steps for each deployment action.

The change that we need to make specific to VMware Cloud Foundation is to include the additional VLANs that it requires as part of the upstream port channels on both of our Nexus switches. So that will require editing these following four workflows in the same fashion.

First select the Task Input button for the first Nexus workflow:




Figure 19. Selecting the Task Input for the first Nexus workflow.

Then, switch to the Advanced tab (#1), click the + sign to add an additional allowed VLAN (#2); and add vMotion, vSAN, and the native VLAN to the allowed list (#3). Note that these VLANs were set and are shown as an alias from the Configuration phase.




Figure 20. Adding an additional allowed VLAN to the allowed list.

Once all of the additional VLANs have been added like in the below example, click on Save (#4).




Figure 21. Saving the allowed additional VLANs.

Repeat this process for the three remaining upstream Port Channel workflows.

Now that the proper VLANs have been added, we can proceed with running the Nexus workflow by pushing on the play button.

For our use case, the MDS Switch workflow requires no editing, so that can be run with the default values once the Nexus workflow completes.



In a similar fashion to the Nexus setup, we need to add our additional VLANs to the management vNIC template (A&B) of the UCS service profile so that the VLANs can communicate across that interface when VMware Cloud Foundation is deployed.

To update this, click on the edit button in the Configure UCS Fabric Interconnect workflow.

Step 4: FlashStack™ Deployment

Workflows			
#	Name	Description	
1	Configure Nexus Switches	Initial Nexus 9K Configuration and Port Channel Setup	✓
2	Configure MDS Switches	Initial MDS Configuration and Port Channel, VSAN Setup	Rollback ✓
3	Configure UCS Fabric Interconnects	Deploy UCS Configuration for Servers, LAN, SAN, Boot Policy	▶ 

Figure 22. Edit button in the Configure UCS Fabric Interconnect workflow.

In the spawned graphical workflow hierarchy, expand Configure UCS LAN Connectivity.




Figure 23. Expanding the Configure UCS LAN Connectivity workflow.

We will be editing the workflows of Create Management vNIC Template (A) and (B).




Figure 24. Editing the workflows of Create Management vNIC Template (A) and (B).

As with the Nexus setup, select the Advanced tab (#1), but this time we will be using the dropdown menu (#2) to add the additional VLANs that we require (#3). Save the updated configuration (#4) and repeat the process on the other vNIC Template (B) workflow.




Figure 25. Adding additional VLANs.



This represents the last required customized step needed for the Deployment phase to complete. All subsequent workflows can be run with default values.

Once the final workflow has been completed, all of the customizations that we made previously can be downloaded as a JSON file via the Export Configuration button shown below. A sample JSON from this above example can be downloaded from our project GitHub page [here](#).

More savvy users can edit this JSON file directly for their needs and import for automated SmartConfig deployment.




Figure 26. Exporting the configuration as a JSON file.

Finalizing ESXi for VMware Cloud Foundation

The ESXi kickstart script deployed within SmartConfig is available at [this GitHub repository](#). This script handles a majority of the steps needed to prepare the ESXi servers for use with VMware Cloud Foundation.

Once SmartConfig deployment completes, the kickstart script:

- Sets a root password.
- Installs ESXi onto the selected Pure Storage FlashArray via boot from SAN.
- Sets a VLAN for both the Management and VM networks.
- Enables ssh and the ESXi shell.
- Sets and enables NTP.

One item that the kickstart script cannot handle (since it is not able to be run on a per-host basis) is assigning each ESXi host a hostname, DNS entry, and static IP address.

There are numerous ways to assign a management IP and give a DNS name to each host, but for this guide, we decided to create a DHCP reserved range that is based upon each host's MAC address. This aligns with our theme of minimizing or eliminating the need to individually touch servers.

To use DHCP in this fashion, we first need to obtain the MAC address for each UCS/ESXi host in the chassis. This can be pulled from the Nexus switches (via the `show mac address-table` command) or, via UCSM which we will show here.



After logging in to the UCS Manager GUI, we navigate to Servers > LAN > MAC Identity Assignment tab. We can then filter for the management MAC addresses of interest for our eight ESXi hosts.




Figure 27. Navigating the MAC Identity Assignment tab.

From the Advanced Filter panel, we remove all non-management MAC addresses as follows:




Figure 28. Removing all non-management MAC addresses.

Then we export the filtered data to a CSV file:




Figure 29. Exporting MAC addresses to a CSV file.



The exported CSV includes the MAC addresses under the 'ID' column.




Figure 30. MAC addresses in the ID column.

For the creation of our reserved DHCP range and forward/reverse DNS entries we have created a PowerShell script that is available on GitHub [here](#). This script will take three inputs, two of which must be created prior to usage on a Windows server: a DHCP range and a DNS zone.

The third item for the script is a CSV file (example [here](#)) that contains the MAC address, IP address (within the DHCP scope) and desired ESXi hostnames:

```
IP,MAC,FQDN
10.21.140.226,00-25-b5-91-1a-15,sc-esxi1
10.21.140.227,00-25-b5-91-1a-12,sc-esxi2
10.21.140.228,00-25-b5-91-1a-0f,sc-esxi3
10.21.140.229,00-25-b5-91-1a-0c,sc-esxi4
10.21.140.230,00-25-b5-91-1a-09,sc-esxi5
10.21.140.231,00-25-b5-91-1a-06,sc-esxi6
10.21.140.232,00-25-b5-91-1a-03,sc-esxi7
10.21.140.233,00-25-b5-91-1a-00,sc-esxi8
```

Figure 31. IP address, MAC address and ESXi hostnames in the CSV file.

When invoked, users input those three pieces of information for the script to run:






Figure 32. Inputting the information to run the script.

After script completion and refreshing the DHCP and DNS management windows, we can see that our DHCP reserved scope has been created along with forward/reverse DNS entries for each host within the CSV file.

10.21.140.226	Pointer (PTR)	SC-ESXi1.puretec.purestorage.com.
10.21.140.227	Pointer (PTR)	SC-ESXi2.puretec.purestorage.com.
10.21.140.228	Pointer (PTR)	SC-ESXi3.puretec.purestorage.com.
10.21.140.229	Pointer (PTR)	SC-ESXi4.puretec.purestorage.com.
10.21.140.230	Pointer (PTR)	SC-ESXi5.puretec.purestorage.com.
10.21.140.231	Pointer (PTR)	SC-ESXi6.puretec.purestorage.com.
10.21.140.232	Pointer (PTR)	SC-ESXi7.puretec.purestorage.com.
10.21.140.233	Pointer (PTR)	SC-ESXi8.puretec.purestorage.com.

Figure 33. The created DHCP reserved scope.






Figure 34. DHCP scopes.

The final step of this phase in our deployment is to reset the management interfaces on the ESXi hosts so that they can pick up their new hostname and management IP. One way to accomplish this is to run the below script via connecting to a fabric interconnect with ssh.

NOTE: This script will **immediately reboot** a UCS host, so make certain that you are connected to the correct Fabric Interconnect before running this script.

```

Reboot_UCS_Servers
1  scope server 1/8
2  cycle cycle-immediate
3  commit-buffer
4  exit
5  scope server 1/7
6  cycle cycle-immediate
7  commit-buffer
8  exit
9  scope server 1/6
10 cycle cycle-immediate
11 commit-buffer
12 exit
13 scope server 1/5
14 cycle cycle-immediate
15 commit-buffer
16 exit
17 scope server 1/4
18 cycle cycle-immediate
19 commit-buffer
20 exit
21 scope server 1/3
22 cycle cycle-immediate
23 commit-buffer
24 exit
25 scope server 1/2
26 cycle cycle-immediate
27 commit-buffer
28 exit
29 scope server 1/1
30 cycle cycle-immediate
31 commit-buffer
32 exit

```

Figure 35. Resetting the management interfaces on the ESXi hosts.



Once the ESXi hosts come back online after reboot, we can see that this example server has been assigned the desired management IP address and hostname. We are now able to proceed and deploy VMware Cloud Foundation with CloudBuilder.




Figure 36. The example server with the desired management IP address and hostname.

Part 2: Deploy VMware Cloud Foundation

You'll need the below additional items before you begin part 2:

- Separate, previously deployed ESXi, Hyper-V, or other bare metal host or other laptop/desktop where the CloudBuilder OVA can be deployed with management network connectivity to the FlashStack ESXi hosts
- Four vSAN-ready ESXi hosts built in the previous phase of this document for Management Domain
- NTP Server that can communicate with CloudBuilder and underlying ESXi hosts
- CloudBuilder OVA deployed on the above host or laptop with a static IP address
- Minimum of three routable production VLANS for use with VMware Cloud Foundation (Management, vMotion, and VMware vSAN)
- Required VMware Cloud Foundation license keys.

Build the Cloud Foundation Management Domain

VMware Cloud Foundation deployments all begin in the same way: building a Management Domain with vSAN via the CloudBuilder OVA. This takes a VMware Validated Design (VVD) as input. This methodology ensures consistency in the critical first phase of a greenfield deployment by requiring users to input a predefined set of values into an excel spreadsheet or JSON file. This, in turn, minimizes the chance of a mistake and directs users to correct errors before deployment kicks off. Since this deployment guide is based upon VMware Cloud Foundation 4.0, we will be using VVD6 for our sample architecture.

The first step is to log in to the CloudBuilder VM with the administrator credentials provided when the VM was built. Once logged in, there is a more exhaustive list of prerequisites specific to VMware Cloud Foundation that need to be reviewed and acknowledged before moving forward.






Figure 37. VMware Cloud Foundation prerequisites.

Once the prerequisites have been reviewed, and the platform type has been specified (VMware VVD) the end-user can download the input specification as an excel spreadsheet or as a JSON file.

This comprehensive deployment specification needs to be filled out with items such as:

- VLANs to be used
- Hostnames for both appliances and ESXi hosts (DNS entries must be made before Management Domain deployment)
- VMware product license keys
- Appliance users and passwords
- DNS/NTP servers

The screenshot shows an Excel spreadsheet titled 'Management Workloads'. The table has columns for 'Proposed VM Name', 'Application', 'Operating System', 'vCPU s', 'vRAM', and 'Storage Total GB'. It lists various VMs like 'sc-vcenter-mgmt', 'sc-nsx-cluster', and 'sc-nsx-1' through 'sc-nsxt-eg2'. At the bottom, there are summary rows for 'Virtual Infrastructure Totals', 'Management Workload Domain Calculations', and 'Total Resources' with a 30% free calculation.

Proposed VM Name	Application	Operating System	vCPU s	vRAM	Storage Total GB
sc-vcenter-mgmt	vCenter Server	Virtual Appliance	4	19	462.08
sc-nsx-cluster	NSX-T Manager Cluster	Virtual Appliance	-	-	-
sc-nsx-1	NSX-T Virtual Appliance Node #1	Virtual Appliance	6	24	200
sc-nsx-2	NSX-T Virtual Appliance Node #2	Virtual Appliance	6	24	200
sc-nsx-3	NSX-T Virtual Appliance Node #3	Virtual Appliance	6	24	200
SC-NSXT-EN1	NSX-T Edge Node #1 - North-South	Virtual Appliance	4	8	200
SC-NSXT-EG2	NSX-T Edge Node #2 - North-South	Virtual Appliance	4	8	200
sc-sddc	SDDC Manager Appliance	Virtual Appliance	4	16	816
Virtual Infrastructure Totals			34	123	2278.08 Thick Provisioned
Management Workload Domain Calculations					
			vCPU	vRAM	Storage
Total Resources			34	123	2278.08
Total with 30% free					2962

Figure 38. Management workloads sheet.



In a larger environments some of above needed items might require communication with multiple groups (e.g. network, security, and software procurement teams) to gather all of the input data needed. However, once all of this data has been amassed, the input specification will be uploaded to CloudBuilder, checked, and then the SDDC and Management Domain deployment will kick off after verification.




Figure 39. SDDC deployment verification.

As with Cisco Validated Designs, VMware Validated Designs are blueprints that automate formerly repetitive tasks and will hold up the entire Management Domain stack on the end-users behalf, bringing VMware online much faster than the manual alternative.

As the deployment process continues, users can see each step of the deployment process. If any errors are encountered, they will be displayed within the GUI, providing context on when and where an issue was encountered. Detailed logs of the bring-up process can be found by connecting via ssh to the CloudBuilder VM itself and reviewed for more details when necessary.




Figure 40. Steps in the deployment process.

SDDC Manager and Management Domain deployment times can vary greatly depending on how many ESXi hosts are in the management cluster, ESXi host resources, and overall network connectivity, but once completed CloudBuilder will have fulfilled its role and SDDC Manager can be launched.






Figure 41. Launching the SDDC manager.

With the Management Deployment phase now completed in the preceding steps, we will now walk through a few options for using Workload Domains with Pure Storage.

Part 3: Using Pure Storage FlashArray with VMware Cloud Foundation

As we covered in the [introduction to VMware Cloud Foundation](#), Workload Domains represent three or more ESXi hosts aggregated together in one or more clusters under a single vCenter instance. Once deployed, they are then managed by VMware Cloud Foundation administrators within SDDC Manager. These are units of compute, network, memory, and storage that can be rapidly expanded and contracted, upgraded, and orchestrated via integrated connectivity to the vRealize suite. Once deployed, Workload Domains can be assigned to one or more groups of tenant organizations who can then deploy, manage, and use the VMs and applications required for their respective use case(s). A key differentiator between Management Domains and Workload Domains is that Workload Domains allow for other types of Principal Storage besides vSAN.

A brief description of the differences between Principal Storage and Supplemental Storage and how it relates to VCF is needed. Fortunately, it is very easy to distinguish between the two storage types. *Principal Storage* is any storage type that you can connect directly to your Workload Domain as a part of the setup process within SDDC Manager. It is currently comprised of vVols², vSAN, NFS, and VMFS on Fibre Channel (as of October, 2020). *Supplemental Storage* simply means that you connect your storage system to a Workload (or Management) Domain after it has been deployed. Examples of the Supplemental Storage include the iSCSI VMFS and NVMe-oF³ protocols.

Workload Domain Deployment with VMFS on FC as Principal Storage

The below flowchart represents the overall steps needed to deploy a VCF Workload Domain with Fibre Channel.

² Requires VMware Cloud Foundation 4.1+ and vSphere 7.0.1+.

³ Requires VMware vSphere 7.0+.






Figure 42. Steps needed to deploy a VCF Workload Domain with Fibre Channel.

The first step in this process is creating a host group comprised of the ESXi hosts you want to use in your Workload Domain, creating and then connecting a volume to the host group, and finally formatting a VMFS volume for use. This pre-work can be accomplished in a few ways but the easiest is with the Pure Storage [PowerShell module](#) with the `Initialize-PfavCfWorkloadDomain` cmdlet.

This cmdlet:

- Takes in a comma-separated list of ESXi FQDNs/IPs, their credentials, a datastore name, a size, and a FlashArray connection.
- Connects directly to each ESXi host, gets their FC WWNs, and creates a host object on the FlashArray for each.
- Creates a host group and adds each host.
- Creates a new volume of the specified size.
- Rescans one ESXi host and formats it with VMFS.
- Rescans the remaining hosts.
- Disconnects from the hosts.
- If any step fails it will clean up anything it did.

Usage:

```

$faCreds = get-credential
New-PfaConnection -endpoint <FlashArray FQDN/IP> -credentials $faCreds -ignoreCertificateError -
defaultArray
$creds = get-credential
Initialize-PfavCfWorkloadDomain -esxiHosts <array of IPs or FQDNs> -credentials $creds -datastoreName
<datastore name> -sizeInTB <datastore size> -fc
  
```

Log into one of the ESXi hosts either through the Direct Connect UI or CLI and navigate to Storage.






Figure 43. Navigating to Storage.

Rescan the vHBA adapters on the host to pick up the newly created Pure Storage volume.




Figure 44. Picking up the new Pure Storage volume.

Log in to the other ESXi hosts and perform a rescan. Confirm that the new VMFS datastore is available on all hosts to be part of the Workload Domain before proceeding.




Figure 45. Confirming the new VMFS datastore is available.

Log in to the VMware Cloud Foundation SDDC Manager instance. Click on Adding a Workload Domain, and then select the VMFS on FC storage type.

Storage Selection


Select the type of storage you would like to use for this Workload Domain.

- ☐ vSAN
Configure vSAN based workload domain.
- ☐ NFS
Configure NFS based workload domain.
- ☒ VMFS on FC
Configure Fibre Channel based workload domain.

CANCEL BEGIN

Figure 46. Selecting the VMFS on FC.

Provide basic cluster identifying information.




The screenshot shows the 'VI Configuration' window with the 'Name' tab selected. The left sidebar lists steps 1 through 8, with '1 Name' highlighted. The main area contains three input fields: 'Virtual Infrastructure Name' with the value 'FC-vCF', 'Cluster Name' with the value 'FC-vCF', and 'Organization Name' with the value 'FC-vCF'.

Field	Value
Virtual Infrastructure Name	FC-vCF
Cluster Name	FC-vCF
Organization Name	FC-vCF

Figure 47. Providing basic cluster identifying information.

Provide Workload Domain vCenter specifications. Note that forward and reverse DNS lookup entries for vCenter must be made before Workload Domain construction.



The screenshot shows the 'VI Configuration' window with the 'Compute' tab selected. The left sidebar lists steps 1 through 8, with '2 Compute' highlighted. The main area is titled 'vCenter' and contains six input fields: 'vCenter IP Address' (10.21.143.140), 'vCenter FQDN' (vcsa-workload1.puretec.purestorage), 'vCenter Subnet Mask' (255.255.255.0), 'vCenter Default Gateway' (10.21.143.1), 'vCenter Root Password' (masked with dots), and 'Confirm vCenter Root Password' (masked with dots).

Field	Value
vCenter IP Address	10.21.143.140
vCenter FQDN	vcsa-workload1.puretec.purestorage
vCenter Subnet Mask	255.255.255.0
vCenter Default Gateway	10.21.143.1
vCenter Root Password	
Confirm vCenter Root Password	

Figure 48. Provide Workload Domain vCenter specifications.

There are two options for networking with VMware Cloud Foundation: NSX-V and NSX-T. Please note that the NSX-T installation package must be downloaded to the SDDC Manager repository before that option will become functional. Similar to vCenter, all DNS forward and reverse lookup entries must also be made before deployment. In this example deployment, we use NSX-T.



VI Configuration

1 Name

2 Compute

3 Networking

4 Storage

5 Host Selection

6 License

7 Object Names

8 Review

Networking ?

NSX Platform

☐ NSX-V
☒ NSX-T

Overlay Networking

VLAN ID ⓘ

2143

NSX-T Manager

NSX-T Manager Cluster IP ⓘ

10.21.143.141

NSX-T Manager Cluster FQDN ⓘ

nsx-t-cluster.puretec.purestorage.c

NSX-T Manager 1 IP Address ⓘ

10.21.143.142

NSX-T Manager 1 FQDN ⓘ

nsx-t-mgmt1.puretec.purestorage.c

NSX-T Manager 2 IP Address ⓘ

10.21.143.143

NSX-T Manager 2 FQDN ⓘ

nsx-t-mgmt2.puretec.purestorage.c

NSX-T Manager 3 IP Address ⓘ

10.21.143.144

NSX-T Manager 3 FQDN ⓘ

nsx-t-mgmt3.puretec.purestorage.c

NSX-T Manager Subnet Mask ⓘ

XXX.XXX.XXX.XXX

NSX-T Manager Default Gateway ⓘ

XXX.XXX.XXX.XXX

NSX-T Manager Admin Password ⓘ

.....

Confirm NSX-T Manager Admin Password

.....

Figure 49. NSX-T networking.

Select at a minimum three ESXi hosts to be members of the Workload Domain cluster. Additional hosts can be added and removed later (provided you keep a minimum of three hosts). This greatly enhances the ability to scale workloads up or down based on real-time business needs.

VI Configuration

1 Name

2 Compute

3 Networking

4 Storage

5 Host Selection

6 License

7 Object Names

8 Review

Host Selection ?

As a best practice, VMware recommends deploying ESXi hosts with similar or identical configurations across all cluster members, including similar or identical storage configurations. The minimum configuration required for FC is 3 hosts. For more detail, please check product documentation.

Selected resources: 120 Cores, 766.45 GB Memory, 3,577.03 GB Storage

Show only selected hosts ☐

RESET FILTER

CLEAR SELECTION

☒	FQDN	Network Pool	Memory	Dirty Host
☒	sn1-m5-ch1-06.puretec.purestorage.com	vMotion-Only ⓘ	255.48 GB	
☒	sn1-m5-ch1-08.puretec.purestorage.com	vMotion-Only ⓘ	255.48 GB	
☒	sn1-m5-ch1-07.puretec.purestorage.com	vMotion-Only ⓘ	255.48 GB	

☒ 3

Figure 50. Host selection.

Enter the VMFS datastore name that was created earlier in this article.





Figure 51. Entering the VMFS datastore name.

Provide license files for vCenter and NSX-T.

Confirm the auto-generated Object Names from the earlier steps in the deployment do not conflict with any existing VMware Objects.



Figure 52. Confirming object names.

Review the deployment settings and click 'Finish' to build the environment.

Once the Workload Domain has finished deployment, you now have a multitude of options for how to connect and administer Pure Storage alongside it.

Some options are:

- [Pure Storage vSphere Plug-In](#)
- [VMware vRealize Automation](#)
- [VMware vRealize Operations](#)
- [Pure Storage GUI and/or CLI](#)



Workload Domain Deployment with vVols on FC as Principal Storage

In VMware Cloud Foundation version 4.1, vVols have taken center stage as a Principal Storage type available for Workload Domain deployments. This inclusion in one of VMware's products of focus should eliminate any doubt that vVols is not an important area of investment for VMware and their ecosystem partners. This technical KB will walk through the steps required to deploy a Workload Domain using Fibre Channel and vVols with the Pure Storage FlashArray.

Prerequisites:

- VMware Cloud Foundation Management Domain deployed with VCF/Cloud Builder version 4.1
- FlashArray with Fibre Channel connectivity
- FlashArray running Purity 5.3.10+

NOTE: Do not use Purity version 5.3.9 as there is a bug which prevents vVols deployment from completing--this is resolved in 5.3.10

- Three or more ESXi hosts with the following characteristics:
 - Fibre Channel connectivity and zoning completed
 - ESXi version 7.0.1 or above
 - Setup for use in VCF 4.1 per VMware's Documentation
 - Added as a host object to the FlashArray with their respective WWNs assigned
 - Hosts should not have any shared VMFS datastores attached to them (a private volume like boot from SAN is fine, though)

With the above prerequisites confirmed, the process of building a Workload Domain using vVols and FC can be broken down into these steps that we will detail in the remainder of this document:

- Register FlashArray VASA Provider in SDDC Manager.
- Create a Host Group and attach it to Pure Storage Protocol Endpoint on FlashArray.
- Commission ESXi Hosts within SDDC Manager for Workload Domain.
- Create the Workload Domain with the vVols Storage Type.
- Complete VASA registration with Pure Storage vSphere Plugin post-Workload Domain Deployment.

Add FlashArray VASA Provider in SDDC Manager

A cornerstone for building vVols-based Workload Domains is registering a Storage Provider to the Workload Domain vCenter instance during the deployment process. Storage Providers leverage VASA (VMware API for Storage Awareness) as the key that enables vCenter to deploy vVols and the numerous benefits that they provide with a FlashArray. As of VMware Cloud Foundation 4.1, VASA Storage Providers can be added and managed in SDDC Manager by going to the Administration > Storage Settings menu.






Figure 53. Storage Settings in the Administration menu.

Within the Storage Settings menu, we next select the **+Add VASA Provider** to open the wizard for that task.




Figure 54. Opening the Add VASA Provider wizard.



VASA Providers

Add VASA Provider

Ensure that all the required details are entered for adding a VASA provider

Name ⓘ 1 →

URL ⓘ 2 →

User Credentials ⓘ

User Name 3 →

Password 4 →

[ADD ADDITIONAL USER CREDENTIALS](#)

Storage Containers ⓘ

Container Name 5 →

Container Type 7 →

[REMOVE](#)

[ADD ADDITIONAL STORAGE CONTAINERS](#)

6 →

Figure 55. Adding VASA providers.

The fields required for adding a VASA provider are broken out individually below to show what information is needed to register the FlashArray in SDDC Manager.

1. Provide a descriptive **name** for the VASA provider. It is recommended to use the FlashArray name and append it with -ct0 or -ct1 to denote which controller the entry is associated with.
2. Provide the URL for the VASA provider. This cannot be the management VIP of the array. Instead, this field needs to be the management IP address associated with one of the controllers. The URL also is required to have the VASA port and version.xml appended to it. The format for the URL is: `https://<IP of FlashArrayController>:8084/version.xml`
3. Give a FlashArray username with the arrayadmin role. The procedure for how to create such a user can be found [here](#). While the pureuser account can be used, we recommend creating and using a separate FlashArray local user for VASA operations.
4. Provide the password for the FlashArray username to be used.
5. Container Name must be vVol container. Note that this value is case-sensitive. This is the default container; customized names will come in a future release.
6. For Container Type, select FC from the drop-down menu to use Fibre Channel.



- Once all entries are completed, click Save.

This completes the SDDC Manager VASA Registration component of the process and we can now proceed to the next step.

NOTE: Each FlashArray has two VASA providers—one on each controller. SDDC manager only offers the ability to register a single VASA provider when deploying a new workload domain. So you should register one VASA provider with SDDC manager per FlashArray (which VASA provider is up to you—consistency, however, is good, so always choose CT0 for instance), and ensure that post-deployment the second VASA provider is registered. Instructions on that process are below.

Create and Attach PE to FlashArray Host Group

A protocol endpoint (PE) is a logical I/O proxy that establishes the data path between the ESXi hosts and FlashArray for vVols. Establishing this connection at the FlashArray level is a simple, but is a required step before the VMware Cloud Foundation Workload Domain can be successfully deployed.

As mentioned earlier, it is expected that the ESXi hosts to be used in the Workload Domain have been added to the FlashArray and their WWNs have been associated with that host object. Another required step is to create a Host Group and add those hosts to it.

From there, select the **Host Group (1)**, click on the radio button under **Connected Volumes** and click on the **Connect...** button (2).




Figure 56. Creating a host group.

Click the checkbox next to **pure-protocol-endpoint (1)** and then click **Connect (2)** to complete the operation.






Figure 57. Connecting the host group.

Inspect the Host Group to confirm that the pure-protocol-endpoint has been successfully connected to it.




Figure 58. Ensuring that the pure-protocol-endpoint has been successfully connected.

Alternatively, the Purity CLI can be used to connect the PE to a Host Group via the following command, which also confirms PE connectivity to each host within the host group:

```
pureuser@FlashArray> purevol connect --hgroup VCF-Sisters pure-protocol-endpoint
```

Name	Host Group	Host	LUN
pure-protocol-endpoint	VCF-Sisters	sn1-m5-ch1-05	254
pure-protocol-endpoint	VCF-Sisters	sn1-m5-ch1-06	254
pure-protocol-endpoint	VCF-Sisters	sn1-m5-ch1-07	254



Now that the PE has been connected to the host group, we can move forward with commissioning the ESXi hosts into SDDC Manager for use.

Commission ESXi Hosts to SDDC Manager

With all preparatory work completed for our vVols-based Workload Domain, we can now proceed to commission the hosts into SDDC Manager.

To get started, click on **Hosts** under **Inventory**, and then select the **Commission Hosts** button on the top-right.




Figure 59. Selecting commission hosts.

The below figure shows the various fields are and how to populate them.

	FQDN	Network Pool	IP Address	Confirm FingerPrint	Validation Status
No hosts added					
0 hosts					

Figure 60. Host addition and validation fields.

How to populate the fields:

1. **ESXi Host FQDN:** Enter in the ESXi Host name you wish to commission.
2. **Storage Type:** Select the vVol option.
3. **vVol Storage Protocol Type:** Pick FC from the drop-down menu of available protocols.



4. **Network Pool Name:** Select a **Network Pool** that can be used with the Workload Domain ESXi hosts. Only a vMotion network is required for FC-based vVols.
5. Provide the **root username** for the ESXi host.
6. Provide the **root password** for the ESXi host.
7. Click the **Add** button to save the hosts entries and then repeat the above process for as many additional hosts as you will be adding. Note that a JSON template can be populated and imported to speed up this process by importing ESXi hosts in batches.

Once all hosts have been added, the following actions are needed to finish the commissioning process.

Commission Hosts

Host Addition and Validation

Storage Type: ☐ VSAN ☐ NFS ☐ VMFS on FC ☐ vVol

vVol Storage Protocol Type:

Network Pool Name:

User Name:

Password:

ADD

Hosts Added

Click on Confirm FingerPrint button ☒ the below grid to enable or disable to **validate** hosts before proceeding to commission

☒ Host Validated Successfully

REMOVE

☐	FQDN	Network Pool	IP Address	☒ Confirm FingerPrint	Validation Status
☐	sn1-m5-ch1-07.puretec.purestorage.com	vMotion Only	10.21.143.36	☒	Valid
☐	sn1-m5-ch1-06.puretec.purestorage.com	vMotion Only	10.21.143.35	☒	Valid
☐	sn1-m5-ch1-05.puretec.purestorage.com	vMotion Only	10.21.143.34	☒	Valid

VALIDATE ALL

CANCEL **NEXT**

Figure 61. Finishing the commissioning process.

How to populate the fields:

1. Select **all** or **individual hosts** you wish to validate.
2. Click on **Confirm FingerPrint**.
3. Click on **Validate All** to precheck the hosts to confirm that they are ready for use in SDDC Manager and Workload Domains.
4. Once validation has completed successfully click on **Next** to proceed.

Confirm that the hosts are as expected (making certain that Storage Type is **VVOL**) before clicking the **Commission** button to add them to SDDC Manager inventory.






Figure 62. Commissioning the hosts.

With the ESXi hosts now available in SDDC Manager inventory, we can now use them to build a vVols-based Workload Domain.

Create vVols-based Workload Domain

All of the previous steps come together in this section when we create our Workload Domain as this next procedure will showcase.

To get started, select **Workload Domains** from under the **Inventory** menu item, and then click on **+ Workload Domain** and **VI - Workload Domains**.





Figure 63. Adding workload domains.

In the first window that spawns, select **vVol** from the available Storage Selections and then click on **Begin**.





Storage Selection ?

Select the type of storage you would like to use for this Workload Domain.

☐ vSAN
 Configure vSAN based workload domain.

☐ NFS
 Configure NFS based workload domain.

☐ VMFS on FC
 Configure Fibre Channel based workload domain.


☒ vVol
 Configure vVol based workload domain.

CANCEL
 BEGIN

Annotations: A red box highlights the vVol option with a red arrow and the number 1. Another red box highlights the BEGIN button with a red arrow and the number 2.

Figure 64. Selecting vVol as storage.

Provide a descriptive **Virtual Infrastructure** and **Organization Name**. For our vVols-based Workload Domain we will later be using Workload Management, so we select the option to **Enable vSphere Lifecycle Manager Baselines** rather than using vSphere Lifecycle Manager and then click on **Next**.



VI Configuration

1 Name

2 Cluster

3 Compute

4 Networking

5 vVol Storage

6 Host Selection

7 License

8 Object Names

9 Review

Name ?

Virtual Infrastructure Name ? WLD1

Organization Name ? WLD1

vSphere Lifecycle Manager and Kubernetes - Workload Management Readiness ?

Select the vSphere Lifecycle Manager for this workload domain. The selected vSphere Lifecycle Manager determines whether Kubernetes - Workload Management can be deployed. This cannot be changed once the workload domain has been deployed.

☒ Enable vSphere Lifecycle Manager Baselines

- Kubernetes - Workload Management can be deployed.
- Firmware upgrades not supported natively.

☐ Enable vSphere Lifecycle Manager Images

- Kubernetes - Workload Management cannot be deployed.
- Firmware upgrades using cluster images can be performed.

CANCEL
 NEXT

Annotations: A red box highlights the WLD1 input field with a red arrow and the number 1. Another red box highlights the 'Enable vSphere Lifecycle Manager Baselines' radio button with a red arrow and the number 2. A third red box highlights the NEXT button with a red arrow and the number 3.

Figure 65. Adding a descriptive Virtual Infrastructure and Organization Name.

Provide a **Cluster Name** for the Workload Domain and click on **Next**.



VI Configuration

1 Name

2 Cluster

3 Compute

4 Networking

5 vVol Storage

6 Host Selection

7 License

8 Object Names

9 Review

Cluster

Enter the details for the first cluster that will be created as default in this new workload domain.

Cluster Name ① Sisters

Cluster Image ① vLCM is disabled for this workload domain.

CANCEL BACK NEXT

Figure 66. Adding a cluster name..

Input the **vCenter FQDN**. This should have already been added to your DNS server which can be confirmed when the IP address, subnet mask and gateway all auto-populate when the FQDN is correctly resolved. Provide the **vCenter Root Password** then click **Next**.

VI Configuration

1 Name

2 Cluster

3 Compute

4 Networking

5 vVol Storage

6 Host Selection

7 License

8 Object Names

9 Review

Compute

vCenter

vCenter FQDN ① vcsa-vcf-workload1.puretec.purestorage.com

vCenter IP Address ① 10.21.143.140

vCenter Subnet Mask ① 255.255.255.0

vCenter Default Gateway ① 10.21.143.1

vCenter Root Password ①

Confirm vCenter Root Password

CANCEL BACK NEXT

Figure 67. Inputting the vCenter FQDN.

NSX-T deployment parameters are provided in more detail below:



VI Configuration

- 1 Name
- 2 Cluster
- 3 Compute
- 4 Networking**
- 5 vVol Storage
- 6 Host Selection
- 7 License
- 8 Object Names
- 9 Review

Networking

Overlay Networking

VLAN ID ①: 2143

NSX-T Manager

Cluster FQDN ②: vcsa-nsx-t-1.puretec.purestorage.com

Cluster IP ②: 10.21.143.130

FQDN 1 ②: vcsa-nsx-t-1.puretec.purestorage.com

IP Address 1 ②: 10.21.143.131

FQDN 2 ②: vcsa-nsx-t-2.puretec.purestorage.com

IP Address 2 ②: 10.21.143.132

FQDN 3 ②: vcsa-nsx-t-3.puretec.purestorage.com

IP Address 3 ②: 10.21.143.133

Admin Password ③: [Redacted]

Confirm Admin Password: [Redacted]

CANCEL BACK NEXT ④

Figure 68. NSX-T deployment parameters.

1. **Host Overlay (TEP) VLAN** needs to be provided. This VLAN should have an available DHCP scope that the ESXi hosts can grab an IP address from. This is a critical piece for Workload Management/vSphere with Kubernetes to function properly and should also be routable to the Edge TEP network on a separate VLAN.
2. Similar to vCenter, all **NSX-T component FQDNs** should be added to DNS and when the FQDNs are added the IP addresses associated with them should be automatically resolved.
3. Provide a strong **Admin Password** for NSX-T.
4. Click on **Next** to proceed.

The vVol Storage section allows us to specify what array and protocol we wish to associate the Workload Domain vVol datastore to.



The screenshot shows the 'vVol Storage' configuration screen. On the left is a sidebar with a list of steps: 1 Name, 2 Cluster, 3 Compute, 4 Networking, 5 vVol Storage (highlighted), 6 Host Selection, 7 License, 8 Object Names, and 9 Review. The main area is titled 'vVol Storage' and contains the following fields:

- Select VASA Protocol Type:** A dropdown menu with 'FC' selected. Callout 1 points to this field.
- Select VASA Provider:** A dropdown menu with 'sn1-x70-b10-2t-ct0' selected. Callout 2 points to this field.
- Select Storage Container:** A dropdown menu with 'Vvol container' selected. Callout 3 points to this field.
- Select VASA User:** A dropdown menu with 'purevvols' selected. Callout 4 points to this field.
- Datastore Name:** A text input field containing 'vvols-FC-WLD1'. Callout 5 points to this field.

At the bottom right, there are three buttons: 'CANCEL', 'BACK', and 'NEXT'. Callout 6 points to the 'NEXT' button.

Figure 69. Specifying the array and protocol.

1. Select the **FC** protocol from the drop-down list.
2. Select the name of the VASA Provider we entered in the first section of this KB article.
3. Pick **Vvol container** as the **Storage Container**. Note that this is the only container option currently supported.
4. Pick the **FlashArray user account** added during VASA registration.
5. Provide a descriptive **Datastore Name** for the vVol datastore that will be deployed with the Workload Domain.
6. Click the **Next** button to proceed.

Hosts that match the vVol storage protocol (FC in our example) will be shown as available to be used with the Workload Domain. Select at minimum three hosts and then click **Next** to proceed.



VI Configuration

- 1 Name
- 2 Cluster
- 3 Compute
- 4 Networking
- 5 vVol Storage
- 6 Host Selection
- 7 License
- 8 Object Names
- 9 Review

Host Selection

As a best practice, VMware recommends deploying ESXi hosts with similar or identical configurations across all cluster members, including similar or identical storage configurations. The minimum configuration required for vVol is 3 hosts. For more detail, please check product documentation.

⚠ Add VI only supports hosts that have physical NICs 0 and 1, please ensure these are connected and active, as these will be used to connect to DVS from UI. Use API to select hosts with other physical NIC configurations.

Selected resources: 72 Cores, 895.2 GB Memory, 0 GB Storage

☐ Show only selected hosts

RESET FILTER CLEAR SELECTION

☒	FOON	Network Pool	Memory	Dirty Host
☒	sn1-m5-ch1-06.puretec.purestorage.com	vMotion Only	255.73 GB	
☒	sn1-m5-ch1-05.puretec.purestorage.com	vMotion Only	383.73 GB	
☒	sn1-m5-ch1-07.puretec.purestorage.com	vMotion Only	255.73 GB	

3

CANCEL BACK NEXT

Figure 70. Selecting hosts (minimum of three required).

Pick the licenses you wish to use for vSphere and NSX-T (page redacted to not show license info). Review the object names to be used with the Workload Domain deployment and click **Next**.

VI Configuration

- 1 Name
- 2 Cluster
- 3 Compute
- 4 Networking
- 5 vVol Storage
- 6 Host Selection
- 7 License
- 8 Object Names
- 9 Review

Object Names

Virtual Infrastructure Name WLD1

Cluster Name Sisters

vCenter Name vcsa-vcf-workload1

Your input above will be used as a pre-fix to generate vSphere Object Names.

Object Names	Description	Generated Name
resource.vds	vSphere Distributed Switch	WLD1-vcsa-vcf-workload1-Sisters-vds01
resource.portgroup.management	Distributed Port Group for Management Traffic	WLD1-vcsa-vcf-workload1-Sisters-vds01-management
resource.portgroup.vmotion	Distributed Port Group for vMotion Traffic	WLD1-vcsa-vcf-workload1-Sisters-vds01-vmotion
resource.portgroup.fc	resource-name-desc-distributed-port-group-fc-traffic	WLD1-vcsa-vcf-workload1-Sisters-vds01-fc

CANCEL BACK NEXT

Figure 71. Reviewing the object names.



Review the overall Workload Domain deployment specification and then click **Finish** to kick off the deployment process.




Figure 72. Reviewing the overall Workload Domain deployment specifications.

Typical deployments can take around an hour to complete. Once the Workload Domain has been built, we can see it within SDDC Manager and that it is indeed using the vVol storage option as shown below:




Figure 73. Host selection.



Pick the licenses you wish to use for vSphere and NSX-T (page not shown to protect license details).

Review the object names to be used with the Workload Domain deployment and click **Next**.

VI Configuration

- 1 Name
- 2 Cluster
- 3 Compute
- 4 Networking
- 5 vVol Storage
- 6 Host Selection
- 7 License
- 8 Object Names**
- 9 Review

Object Names

Virtual Infrastructure Name: WLD1
Cluster Name: Sisters
vCenter Name: vcsa-vcf-workload1

Your input above will be used as a pre-fix to generate vSphere Object Names.

Object Names	Description	Generated Name
resource.vds	vSphere Distributed Switch	WLD1-vcsa-vcf-workload1-Sisters-vds01
resource.portgroup.management	Distributed Port Group for Management Traffic	WLD1-vcsa-vcf-workload1-Sisters-vds01-management
resource.portgroup.vmotion	Distributed Port Group for vMotion Traffic	WLD1-vcsa-vcf-workload1-Sisters-vds01-vmotion
resource.portgroup.fc	resource-name-desc-distributed-port-group-fc-traffic	WLD1-vcsa-vcf-workload1-Sisters-vds01-fc

CANCEL BACK **NEXT**

Figure 74. Reviewing the object names to be used with the Workload Domain deployment

Review the overall Workload Domain deployment specification and then click **Finish** to kick off the deployment process.

VI Configuration

- 1 Name
- 2 Cluster
- 3 Compute
- 4 Networking
- 5 vVol Storage
- 6 Host Selection
- 7 License
- 8 Object Names
- 9 Review**

Review

sn1-m5-ch1-06.puretec.purestorage.com
Network Pool Name: vMotion Only
Server: Cisco Systems Inc UCSB-B200-M4
Memory Size : 255.73 GB
Processor: 24 x Intel(R) Xeon(R) CPU E5-2670 v3 @ 2.30GHz

sn1-m5-ch1-05.puretec.purestorage.com
Network Pool Name: vMotion Only
Server: Cisco Systems Inc UCSB-B200-M4
Memory Size : 383.73 GB
Processor: 24 x Intel(R) Xeon(R) CPU E5-2670 v3 @ 2.30GHz

sn1-m5-ch1-07.puretec.purestorage.com
Network Pool Name: vMotion Only
Server: Cisco Systems Inc UCSB-B200-M4
Memory Size : 255.73 GB
Processor: 24 x Intel(R) Xeon(R) CPU E5-2670 v3 @ 2.30GHz

vVol Storage

VASA Protocol Type: FC
VASA Provider: sn1-x70-b10-21-ct0
Storage Container: Vvol container
VASA User: purevvols
Datastore Name: vvols-FC-WLD1

licenses

CANCEL BACK **FINISH**

Figure 75. Reviewing the overall Workload Domain deployment specification



Typical deployments can take around an hour to complete. Once the Workload Domain has been built, we can see it within SDDC Manager and that it is indeed using the vVol storage option as shown here:




Figure 76. Confirming the vVol storage option is being used.

Upon logging in to the Workload Domain vCenter instance, we can see that the correct FlashArray has been added as a **Storage Provider**.

However, for failover and other performance considerations, it is required to add the second FlashArray controller as a Storage Provider as well. We will outline that procedure in the next section.




Figure 77. Adding the second FlashArray controller as a Storage Provider

Complete VASA Registration with Pure Storage vSphere Plugin

The first step is to install the Pure Storage vSphere Plugin. There are multiple ways to install the plugin which can be found [here](#).

NOTE: It is not required to install the Pure Storage Plugin for the vSphere Client, but it is generally recommended. There are many other ways to register the VASA provider besides the plugin, including manually in the vSphere Client, PowerShell, or vRealize Orchestrator. Find more information [here](#).

The first step is to register an array against the plugin. Click on **Menu** and then **Pure Storage**.






Figure 78. Registering an array against the plugin.

Click on the **+ Add** button to register an array.




Figure 79. Registering an array.

Arrays can be added individually or via Pure1 if you have authenticated against it. The procedure for authenticating and then importing arrays via Pure1 is available [here](#).

To add a single array, we break out each field below.



Figure 80. Adding a single array.

To add an array:

1. Click on **Add a Single Array**.
2. Provide a descriptive Array Name.
3. Give the FlashArray IP address (management VIP or ideally the correlating FQDN is recommended).
4. Provide the pureuser username and password.
5. Click on **Submit** to add the array.

Next, select the array we just added and select the **Register Storage Provider** button.




Figure 81. Registering Storage Provider.

To register the storage provider, we recommend providing the FlashArray username and password created specifically for VASA use and then select **Workload Domain vCenter**. Select **Register** to complete the process.

Register Storage Provider



Registering the storage provider requires a valid username and password.

Username

Password

☐ Select the server(s) to register with:

- ☒ vcsa-vcf-workload1.puretec.purestorage.com
- ☐ vcsa-vcf1.puretec.purestorage.com

Figure 82. Registering the storage provider.

Returning to the storage providers registered against our Workload Domain vCenter instance, we can see that both FlashArray controllers have been added. The Workload Domain is now ready for production use.

Storage Providers

Storage Provider/Storage Sys...	Status	Active/Standby	Priority	URL	Last Rescan Time	VASA API Ver
sn1-x70-610-21-ct0	Online	...	...	https://10.21.143.118084/version.xml	09/18/2020, 15:43	3.0
sn1-x70-610-21 (2/2 online)	Active	...	200			
sn1-x70-610-21-ct1	Online	...	...	https://10.21.143.128084/version.xml	...	3.0
sn1-x70-610-21 (2/2 online)	Standby	...	200			

Figure 83. Checking that both FlashArray controllers have been added.



Pure Storage Support

- [Pure Storage Support](#)
- [VMware Platform Guide](#)
- [VMware Cloud Foundation How-to](#)
- [VMware Cloud Foundation Quick Reference](#)
- [VMware Cloud Foundation Video Guide](#)
- [VMware Cloud Foundation on FlashStack Deployment Guide](#)
- [FlashStack SmartConfig](#)

Related Contacts

- [VMware Team at Pure Storage](#)
- [National Technical Partner Managers](#)

©2020 Pure Storage, the Pure P Logo, and the marks on the Pure Trademark List at <https://www.purestorage.com/legal/productenduserinfo.html> are trademarks of Pure Storage, Inc. Other names are trademarks of their respective owners. Use of Pure Storage Products and Programs are covered by End User Agreements, IP, and other terms, available at: <https://www.purestorage.com/legal/productenduserinfo.html> and <https://www.purestorage.com/patents>

The Pure Storage products and programs described in this documentation are distributed under a license agreement restricting the use, copying, distribution, and decompilation/reverse engineering of the products. No part of this documentation may be reproduced in any form by any means without prior written authorization from Pure Storage, Inc. and its licensors, if any. Pure Storage may make improvements and/or changes in the Pure Storage products and/or the programs described in this documentation at any time without notice.

THIS DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PURE STORAGE SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

Pure Storage, Inc.
650 Castro Street, #400
Mountain View, CA 94041

purestorage.com

800.379.PURE

