

Deploying VMware PKS on VMware Cloud Foundation



vmware®

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2019 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

1	Overview	4
2	Prerequisites	6
	VMware Cloud Foundation - SDDC Manager	7
	vCenter Server	8
	Cluster	9
	ESXi Host	10
	NSX Manager	11
3	Deployment	19
	Assumptions	20
	Preparation	20
	Pivotal Cloud Foundry on vSphere Requirements	20
	vSphere Service Account Requirements	22
	Download Software	22
	DNS	22
	Prepare NSX-T	22
	Ops Manager	29
	Installing the Pivotal Container Service Ops Manager	29
	NSX Certificates	32
	Configure Bosh Director for vSphere	33
	Pivotal Container Service	37
	Harbor Registry	45
	Configure Harbor Registry	45
	Install PKS CLI Client	51
	Create a User for Create the Kubernetes clusters	51
	Set Up the BOSH Environment	52

Overview

This document describes the deployment of Pivotal Container Service in VMware Cloud Foundation.

VMware PKS is a container services solution to put Kubernetes in operation for multicloud enterprises and service providers. PKS simplifies the deployment and management of Kubernetes clusters with Day 1 and Day 2 operations support. PKS manages container deployment from the application layer all the way to the infrastructure layer according to the requirements for production-grade software. PKS supports high availability, autoscaling, health-checks and self-repairing of underlying virtual machines, and rolling upgrades for the Kubernetes clusters.

Intended Audience

The document is intended for the cloud administrators who want to provide the container environment to the organization.

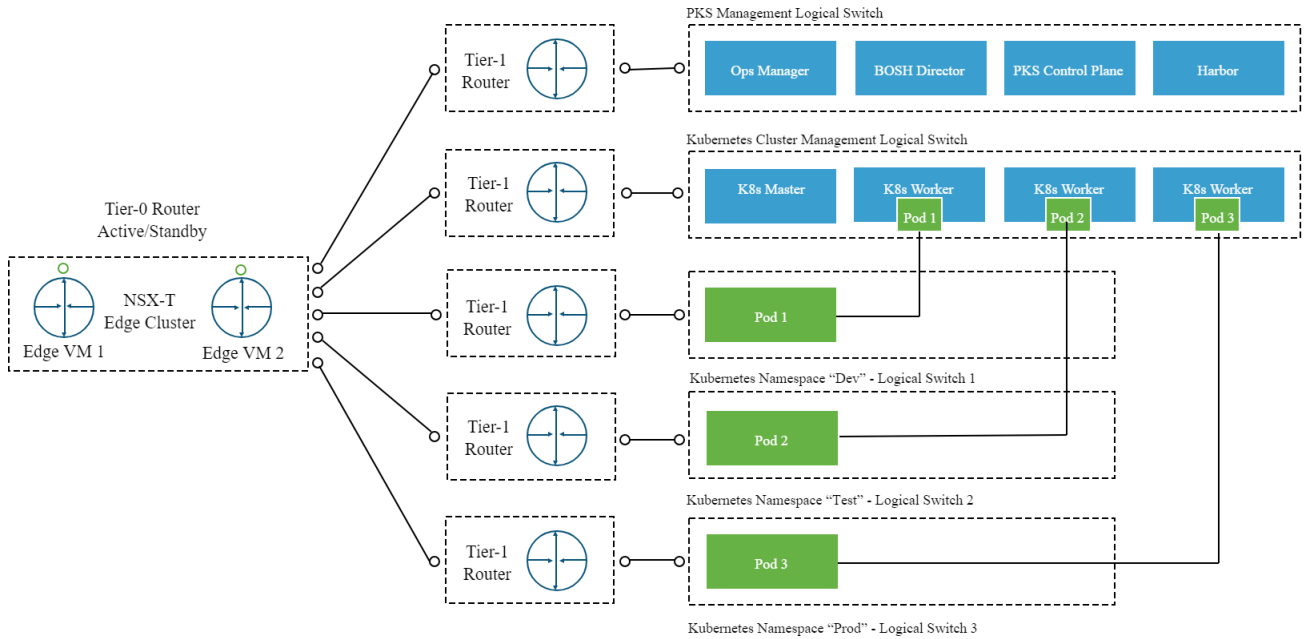
Naming Convention

The names and IP addresses in this document are used as examples.

PKS Network Topology

The No NAT topology with the NSX-T logical switches topology has the following features:

- The PKS control plane (Ops Manager, BOSH Director, and PKS VM) components use the corporate routable IP addresses.
- The Kubernetes cluster master and worker nodes use corporate non-routable IP addresses.
- The PKS control plane is deployed inside the NSX-T network. The PKS control plane components (VMs) use the corporate routable IP addresses.



For information on PKS with NSX-T, see https://docs.pivotal.io/pivotalcf/2-3/refarch/vsphere/vsphere_ref_arch.html#pks-nsxt.

PKS Availability

You can achieve the availability when deploying the Kubernetes cluster modes across the multiple availabilities zones (AZ). The AZ can be defined either with the additional cluster or with the additional resource pool in the same cluster.

Prerequisites

This document is based on the following assumptions:

- The VMware Cloud Foundation bring-up is successful.
- The VI workload domain with NSX-T is created and is used as a target workload domain for the PKS deployment.
- The shared NSX-T edges are deployed and the logical North-South routing is configured as per the [Deploy and Configure the NSX-T Instance for the Shared Edge and Compute Cluster](#) section in the *Deployment of VMware NSX-T for Workload Domains* guide.

Bill of Materials

Table 2-1.

Software Component	Version	Build
VMware Cloud Foundation Appliance	3.5	11215871
VMware vCenter Server on vCenter Server Appliance	6.7 U1	10244745
VMware vSphere (ESXi)	6.7 EP5	10764712
VMware NSX Data Center for vSphere	6.4.4	11197766
VMware NSX-T Data Center	2.3	10085361
Ops Manager	2.3	170
PKS	1.2.0	47
Harbor	2.3	GA

Known Issues

The NSX-T edge fails to establish a tunnel with the ESXi host on which it is running.

It is a known limitation in ESXi . When the FastPath (fp) interfaces of the edge node VM are mapped to the NSX-T-backed logical switches, the edge TEP cannot form a tunnel with the VTEP of the host that it has deployed.

This chapter includes the following topics:

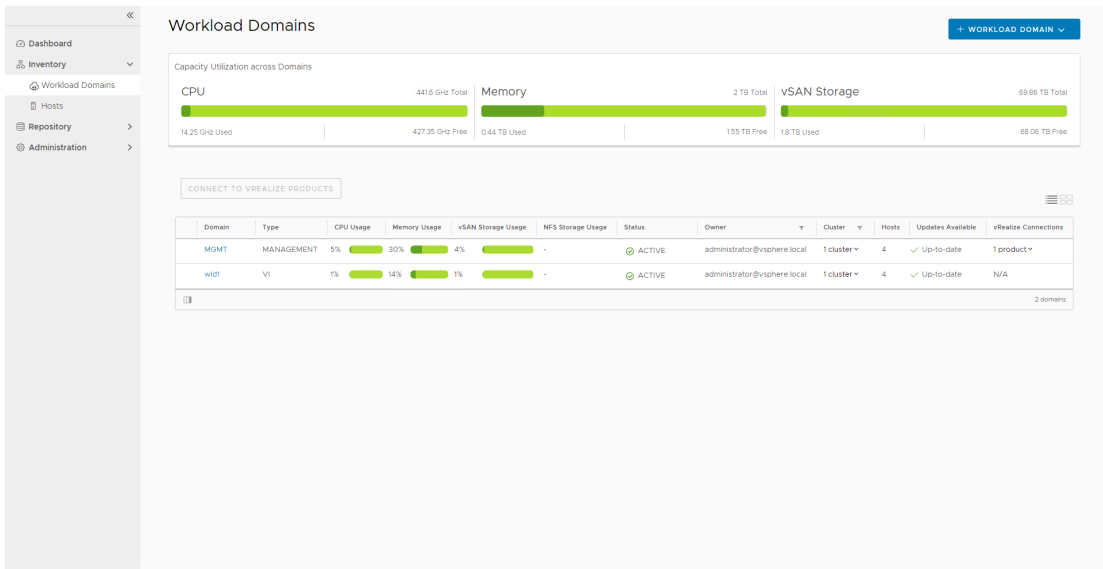
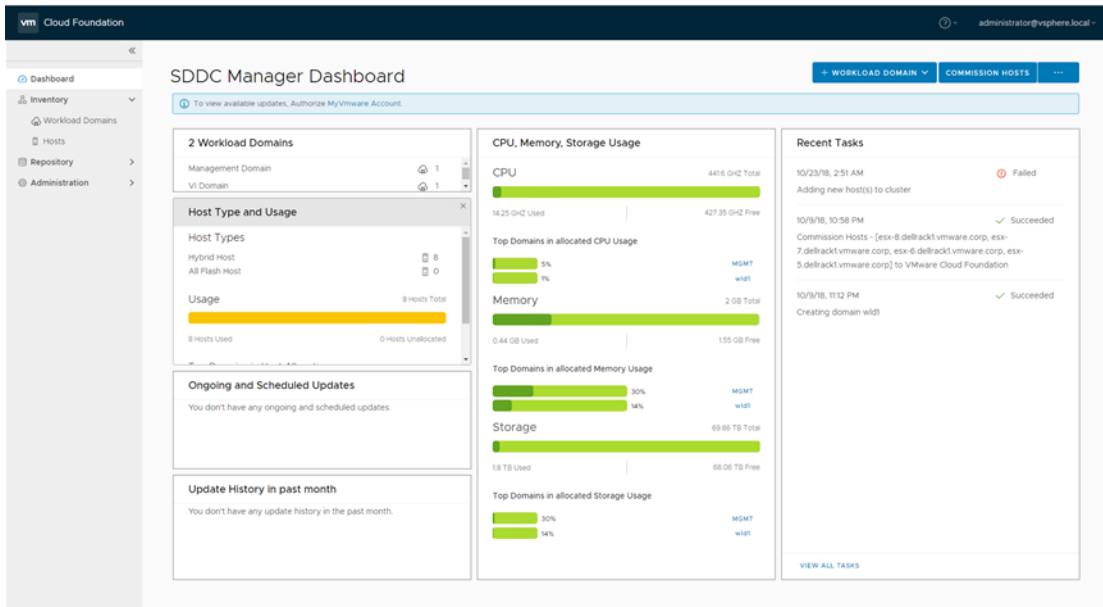
- [VMware Cloud Foundation - SDDC Manager](#)
- [vCenter Server](#)
- [NSX Manager](#)

VMware Cloud Foundation - SDDC Manager

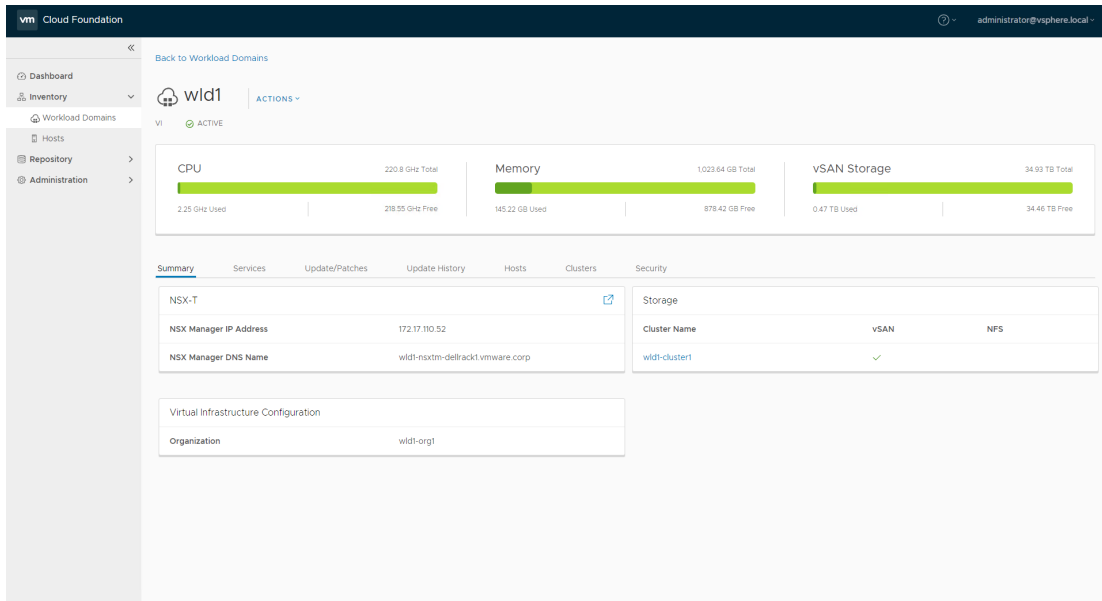
You should install and configure VMware Cloud Foundation.

Ensure the following prerequisites:

- Log in to VMware Cloud Foundation SDDC Manager to access the dashboard. Check the **Workload Domains** page to ensure that **Management** and **VI** workload are available.



- Under **Summary** on the **VI** workload details page, ensure NSX-T details shows up with the IP address of NSX-T Manager and FQDN.

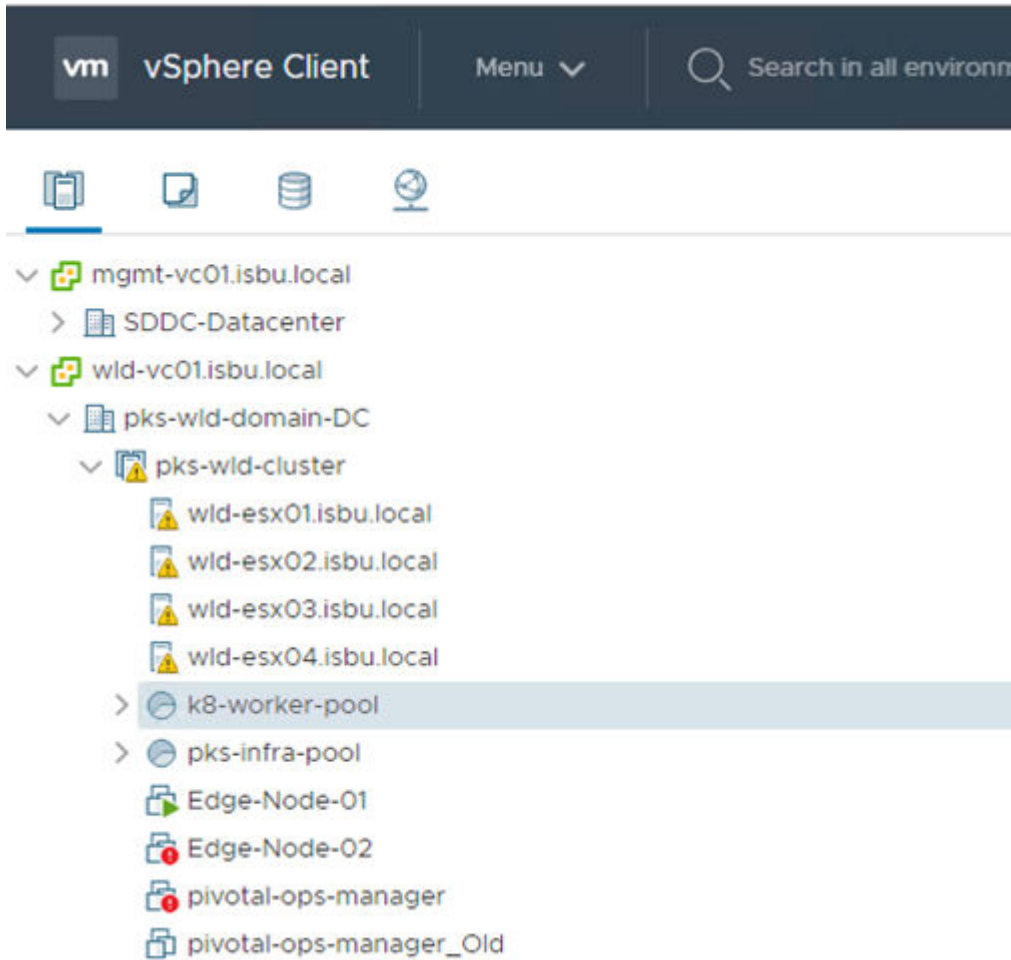


vCenter Server

Log in to the vCenter UI of the NSX-T enabled workload domain.

The NSX-T feature in VMware Cloud Foundation supports a single cluster to simulate the multiple availability zones in vCenter. There must be a minimum of two resource pools in the cluster:

- `pkc-infra-pool`: It is used for the PKS VMs.
- `k8-worker-pool`: It is used for the Kubernetes work nodes.



Cluster

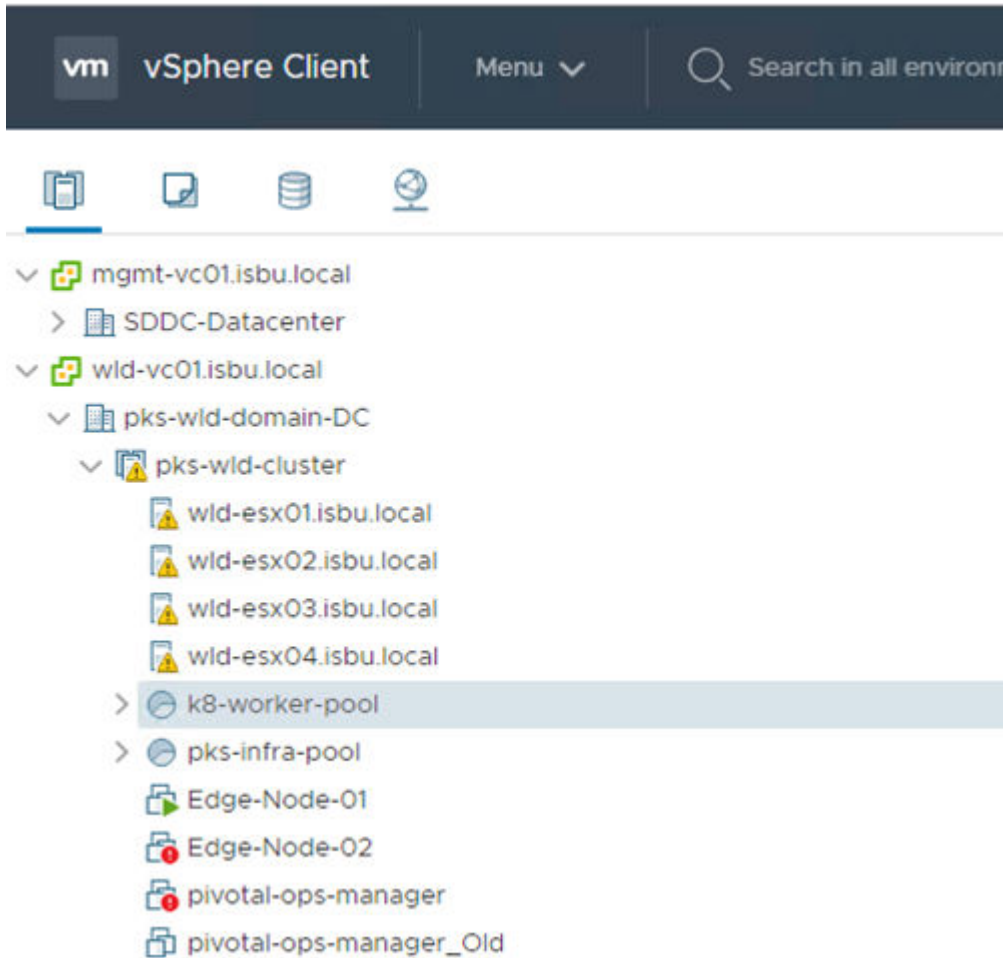
As the implementation of the current NSX-T automation in VMware Cloud Foundation supports a single cluster, the vCenter cluster is used as a source of resource pools to simulate multiple availability zones.

Resource Pools

There should be two resource pools in the cluster:

- `pks-infra-pool`: It is used for the PKS VMs.

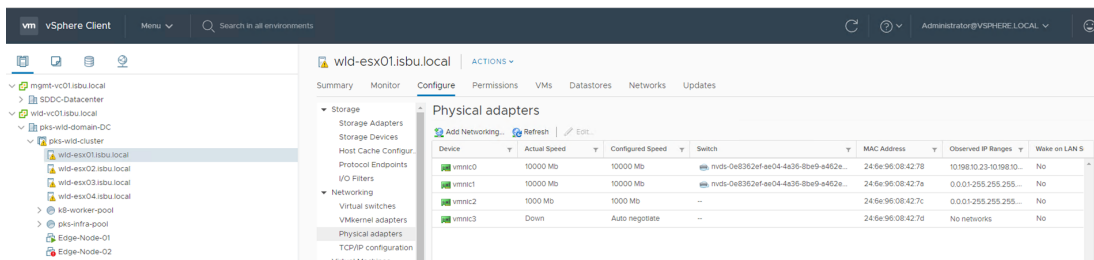
- `k8-worker-pool`: It is used for the Kubernetes work nodes




ESXi Host

Ensure the following prerequisites:

- The ESXi hosts are included in the cluster of the VI workload domains for NSX-T. Ensure this cluster has at least two physical network adapters assigned to N-VDS.



- All VMkernel adapters should be connected to the logical switches on N-VDS.



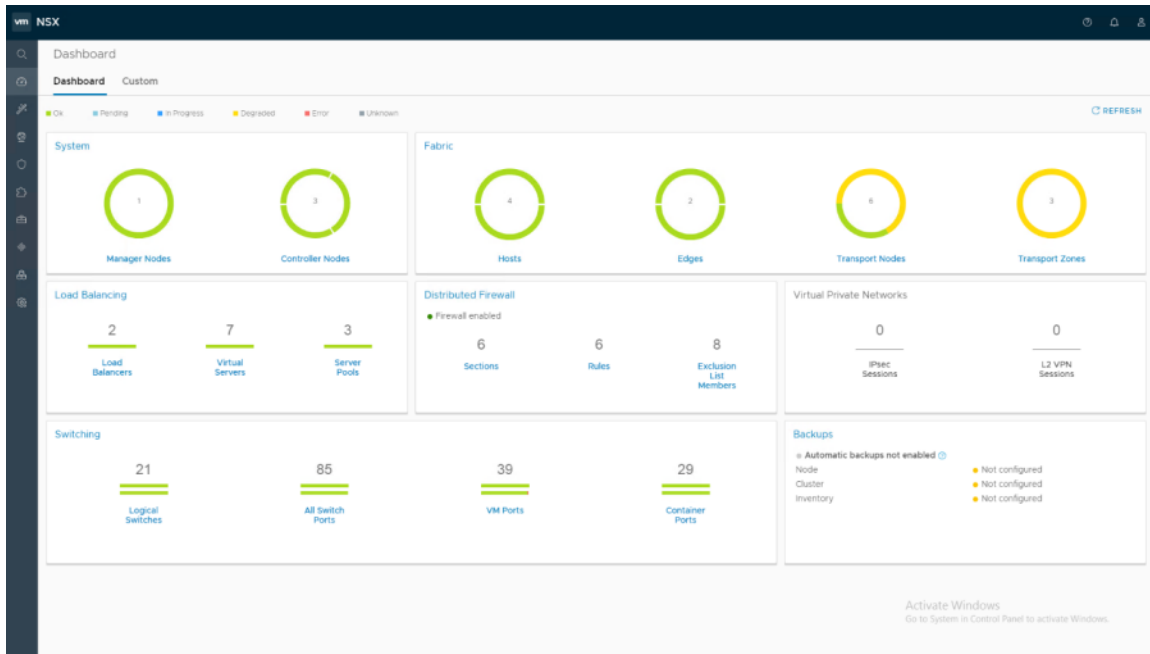
NSX Manager

Ensure the following prerequisites are met:

- Open the UI of NSX-T Manager in a browser.

`https://<FQDN or IP of the NSX-T manager>`

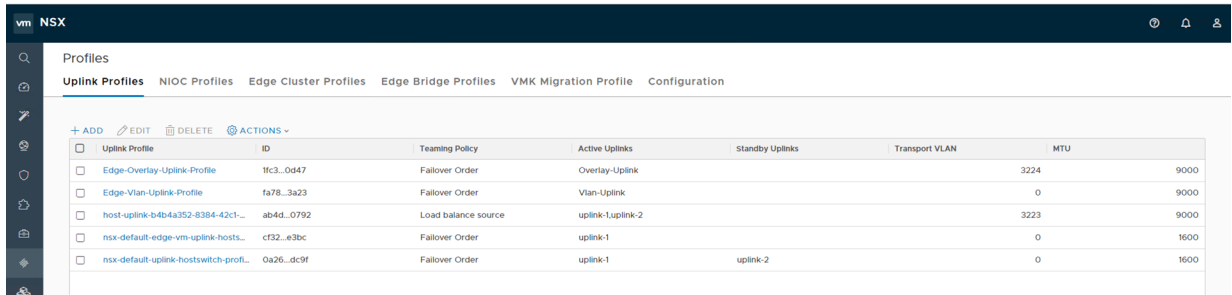
- Log in as the admin user and use the password created during the creation of workload domain in the NSX-T parameters section.
- Ensure that the dashboard does not display any errors. On the NSX-T Manager dashboard, navigate to **Fabric > Profiles**.



- Verify that there are uplink profiles with the corresponding teaming policy, VLANs, and MTU.

Table 2-2.


Uplink Profile	Teaming Policy	Active Uplinks	VLAN	MTU
Hosts-Overlay	Load balance source	Uplink1, Uplink2	<Host TEP vlan>	9000
Edge-Vlan	Failover Order	Uplink1	<Vlan uplink>	9000
Edge-Overlay	Failover Order	Uplink2	<Edge TEP vlan>	9000



- There must be three transport zones as follows:

Table 2-3.

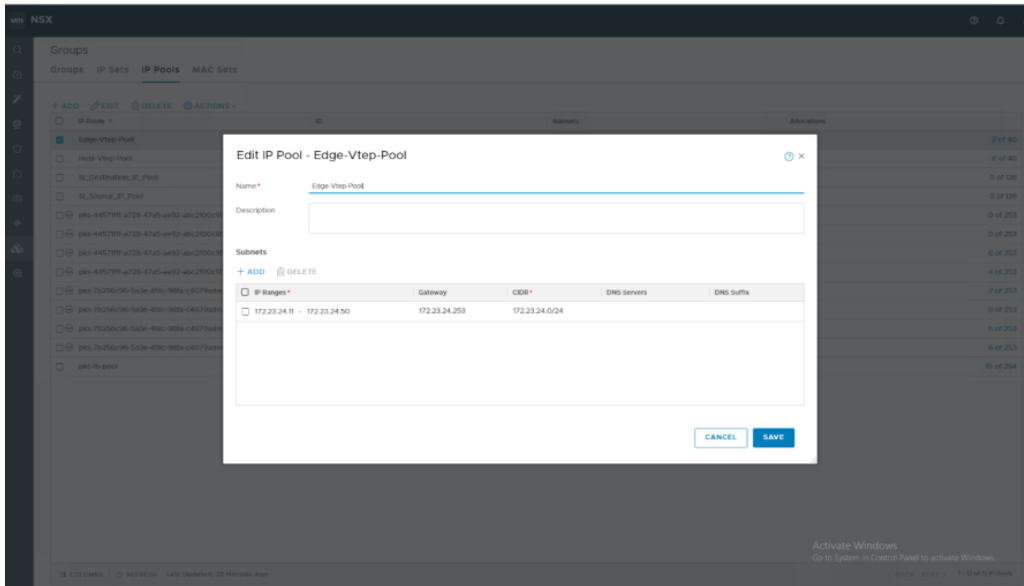
Transport Zone	Traffic Type	N-VDS Name
Vlan-tz-<ID>	van	Edge-Vlan-Nvds
Overlay-tz-<ID>	overlay	nvds-<ID>
Edge-Vlan-tz	vlan	nvds-<ID>



- You need an IP pool from where PKS to assign IP addresses for VIP of the load balancers for Kubernetes Pod services.

Table 2-4.


Name	IP Ranges	Gateway	CIDR	DNS Servers	DNS Suffix
pks-lb-pool	192.168.2.1-192.168.2.254		192.168.2.0/24		



- The network requirements are as follows:
 - a To connect the NSX-T edges to the logical switches and to route the South-North traffic, several logical switches have to be used.

Table 2-5.

Name	Traffic Type	Transport Zone
Edge-External-LS	<External VLAN>	Edge-Vlan-tz
Edge-Mgmt-LS	<VLAN to mgmt network>	Edge-Vlan-tz
Edge-VM-Trunk-LS	<VLAN 0 - 4094>	vlan-tz-<ID>



- b Verify that all the ESXi hosts from the VI-NSX-T workload domain are prepared for NSX-T and display

NSX Activated in the **Deployment Status** column.

Host	ID	IP Addresses	OS Type	OS Version	Deployment Status	NSX Version	Controller Connectivity	Manager Connectivity	Transport Node (TN)
Install/Uninstall NSX actions cannot be performed on cluster member hosts if Auto-install NSX configuration is enabled.									
CONFIGURE CLUSTER + INSTALL NSX UNINSTALL NSX ACTIONS									
View All Hosts									
pkc-wld-cluster (4)	MoRef ID: domain-c7				Auto-Install NSX Enabled				Auto-Create TN Enabled
wld-esx04.isbu.local	3387...4565	172.25.15.24, ...	ESXi	6.7.0	NSX Installed	2.3.0.0.0.10085378	Up	Up	wld-esx04.isbu.local
172.25.15.21	fa1b...81c4	172.25.15.21, L...	ESXi	6.7.0	NSX Installed	2.3.0.0.0.10085378	Up	Up	wld-esx01.isbu.local
172.25.15.22	d877...Offa	172.25.15.22, ...	ESXi	6.7.0	NSX Installed	2.3.0.0.0.10085378	Up	Up	wld-esx02.isbu.local
172.25.15.23	a218...696b	172.25.15.23, ...	ESXi	6.7.0	NSX Installed	2.3.0.0.0.10085378	Up	Up	wld-esx03.isbu.local

- c Ensure that the edges (for redundancy) are connected to NSX-T Manager.

Edge	ID	Deployment Type	Management IP	Host	Deployment Status	Controller Connecti	Manager Connectivi	Transport Node	Edge Cluster	Logical Routers
Edge-Node-01	bc35...c2e6	Virtual Machine	172.25.15.30	172.25.15.22	Node Ready	Up	Up	Edge-TN-01	PKS-Edge-Cluster	4
Edge-Node-02	9af7...37de	Virtual Machine	172.25.15.31	172.25.15.21	Node Ready	Up	Up	Edge-TN-02	PKS-Edge-Cluster	4

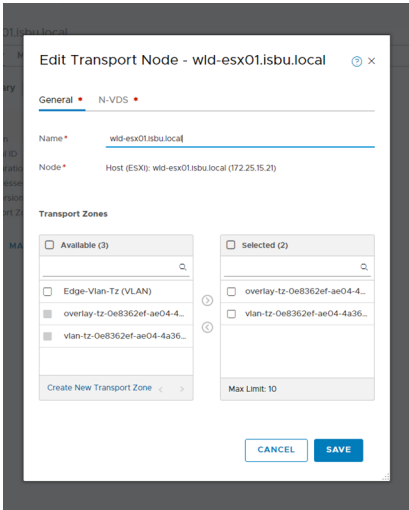
- d All the ESXi hosts and the edges must be configured as the transport nodes.

Transport Node	ID	N-VDS	Configuration State	Status	IP Addresses	Fabric Node Type	Transport Zones	NSX Version
Edge-TN-01	bc35...c2e6	2	Success	Degraded	172.25.15.30	Edge - Virtual Machine	Edge-Vlan-Tz overlay-tz-0e8362ef-a...	2.3.0.0.0.10085448
Edge-TN-02	9af7...37de	2	Success	Degraded	172.25.15.31	Edge - Virtual Machine	Edge-Vlan-Tz overlay-tz-0e8362ef-a...	2.3.0.0.0.10085448
wld-esx01.isbu.local	fa1b...81c4	1	Success	Degraded	172.25.15.21	Host - ESXi 6.7.0	overlay-tz-0e8362ef-a... vlan-tz-0e8362ef-ae04...	2.3.0.0.0.10085378
wld-esx02.isbu.local	d877...Offa	1	Success	Degraded	172.25.15.22	Host - ESXi 6.7.0	overlay-tz-0e8362ef-a... vlan-tz-0e8362ef-ae04...	2.3.0.0.0.10085378
wld-esx03.isbu.local	a218...696b	1	Success	Up	172.25.15.23	Host - ESXi 6.7.0	overlay-tz-0e8362ef-a... vlan-tz-0e8362ef-ae04...	2.3.0.0.0.10085378
wld-esx04.isbu.local	3387...4565	1	Success	Up	172.25.15.24, 172.23.21.24, 172.23.15.24	Host - ESXi 6.7.0	overlay-tz-0e8362ef-a... vlan-tz-0e8362ef-ae04...	2.3.0.0.0.10085378

- e Each host transport node has to be a part of the following transport zones:

- Overlay-tz-<ID>

■ Vlan-tz-<ID>

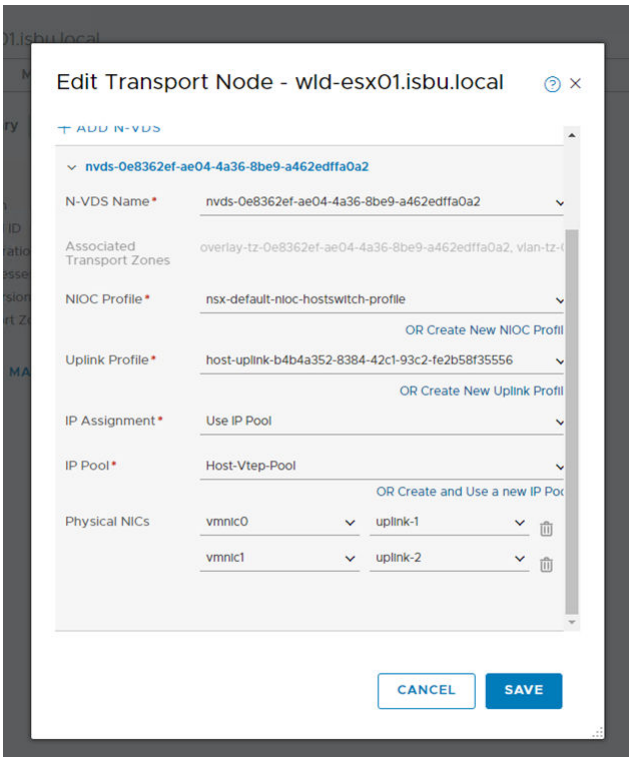


f You have to configure a single N-VDS on each host transport node:

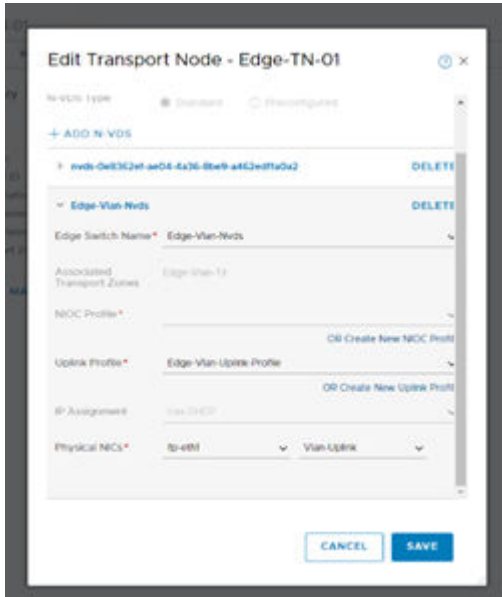
Table 2-6.

N-VDS Switch	Uplink Profile	IP Assignment	Physical NICs
nvds-<ID>	host-uplink-<ID>	DHCP	vmnic0 → uplink1 vmnic1 → uplink2

Note If you do not use DHCP to assign the IP addresses to the TEP adapters, then you can use the IP pools.



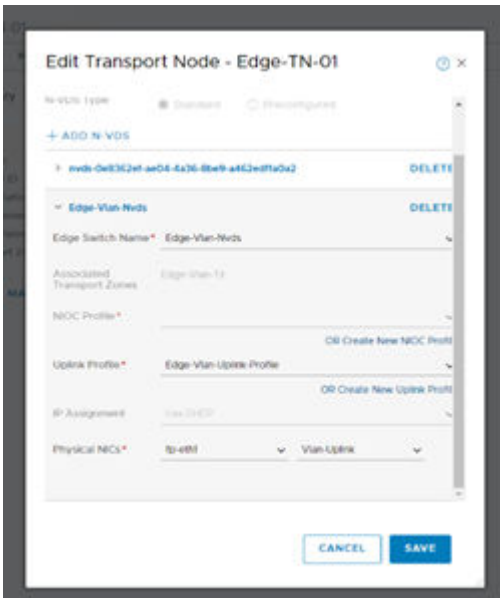
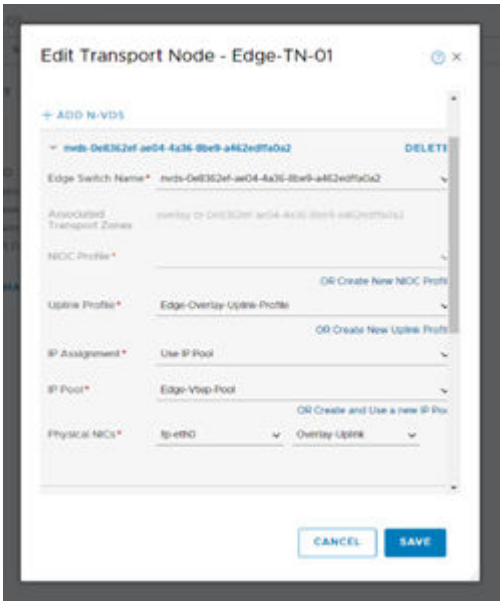
- g Each edge transport node has to be a part of the following transport zones:
 - Edge-vlan-tz
 - Overlay-tz-<ID>



- h You need to have two N-VDS switches as follows:

Table 2-7.

N-VDS Switch	Uplink Profile	IP Assignment	Physical NICs
nvds-<ID>	Edge-Overlay-Uplink-Profile	IP Pool or Static IP	fp-eth0 → uplink1
edge-vlan-nvds	Edge-Vlan-Uplink-profile		fp-eth1 → vlan-uplink



- i Both the edges need to be part of the edge cluster.

Edit Edge Cluster - PKS-Edge-Cluster

Name*

Description

Edge Cluster Profile

Transport Nodes

Member Type

Available (0)

Selected (2)

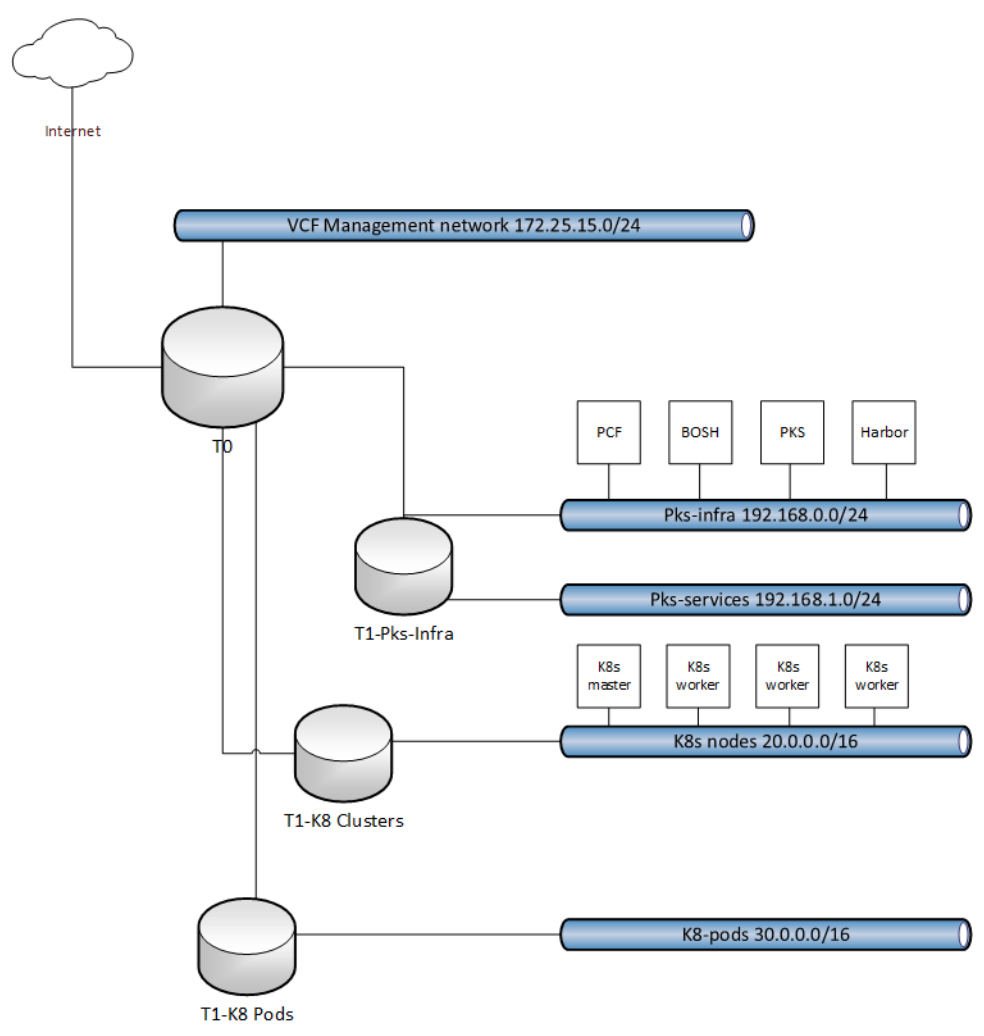
Edge-TN-01

Edge-TN-02

CANCEL SAVE

Deployment

PKS Example Network Topology



The networks used in the deployment example are as follows:

Table 3-1.

N	Description	CIDR	Range	NSX Object	vCenter Object	PKS Object
1	Pks-infra (OpsMgr, Bosh, PKS, Harbor)	192.168.0.0/24	192.168.0.22-192.168.0.254	pks-infra-ls	pks-infra-ls	pks-infra-ntk
2	Pks-service (Upgrade purposes)	192.168.1.0/24	192.168.1.22-192.168.1.254	pks-services-ls	pks-services-ls	pks-services-ntk
3	Kubernetes service load balance	192.168.2.0/24	192.168.2.1-192.168.2.254	pks-lb-pool		
4	Kubernetes Nodes	20.0.0.0/16	20.0.0.1-20.0.255.254	Node-Block		
5	Kubernetes Pods	30.0.0.0/16	30.0.0.1-30.0.255.254	Pod-Block		

This chapter includes the following topics:

- [Assumptions](#)
- [Preparation](#)
- [Ops Manager](#)
- [Harbor Registry](#)
- [Install PKS CLI Client](#)
- [Create a User for Create the Kubernetes clusters](#)
- [Set Up the BOSH Environment](#)

Assumptions

The following assumptions are made in this document:

- We are going to use internal authentication for OpsManager. In this model, the components of PKS will store the credentials in OpsManager.
- We use self-signed certificates for OpsManager, Bosh, PKS, and Harbor.

Preparation

Pivotal Cloud Foundry on vSphere Requirements

Minimum Resource Requirements for Pivotal Cloud Foundry Deployment with Pivotal Application Service

The following are the minimum resource requirements for maintaining a Pivotal Cloud Foundry deployment with Ops Manager and Pivotal Application Service on vSphere:

- vSphere v6.7, v6.5, v6.0, or v5.5
- 2 TB of recommended disk space
- 120 GB memory
- One public IP address for Pivotal Application Service and other public address for Ops Manager
- 80 vCPU cores
- Overall CPU of 28 GHz
- vSphere Standard and further editions
- Ops Manager must have the HTTPS access to vCenter and ESX hosts on TCP port 443.
- Configure the vSphere cluster as follows:
 - If you enable vSphere Distributed Resource Scheduler (DRS) for the cluster, you must set the Automation level to `Partially automated` or `Fully automated`. If you set the Automation level to `Manual`, the BOSH automated installation will fail with a `power_on_vm` error when BOSH attempts to create virtual machines (VMs).
 - Disable hardware virtualization if your vSphere hosts do not support VT-X/EPT. If you are unsure whether the VM hosts support VT-x/EPT, you should disable this setting. If you leave this setting enabled and the VM hosts do not support VT-x/EPT, then each VM requires manual intervention in vCenter to continue powering on without the Intel virtualized VT-x/EPT. Refer the vCenter help topic at [Configuring Virtual Machines > Setting Virtual Processors and Memory > Set Advanced Processor Options](#) for more information.

Note Refer [PCF on vSphere Requirements](#) for more information.

Configure Firewall for Pivotal Cloud Foundry

Ops Manager and Pivotal Application Service require the following open TCP ports:

Table 3-2.

Ports	Description
25555	Routes from Ops Manager to the BOSH Director.
443	Routes to HA Proxy or, if configured, your own load balancer.
80	Routes to HA Proxy or, if configured, your own load balancer.
8844	Routes from Ops Manager to BOSH CredHub.

Table 3-2. (Continued)

Ports	Description
2222	Necessary for using Application SSH.
25595	Routes from the Traffic Controller to the BOSH Director to enable sending BOSH health metrics to the Firehose.

The UDP port 123 must be open if you want to use an external NTP server.

vSphere Service Account Requirements

<https://docs.pivotal.io/pivotalcf/2-3/customizing/vsphere-service-account.html>

Download Software

On a jump station with access to the management network, download the following:

Table 3-3.

Product	Download Link
Pivotal Cloud Foundry Ops Manager for vSphere - 2.3-build.170	https://network.pivotal.io/products/ops-manager#/releases/200482
Pivotal Container Service 1.2.0-build.47	https://network.pivotal.io/products/pivotal-container-service#/releases/191865
Harbor Container Registry for Pivotal Cloud Foundry 1.6.0	https://network.pivotal.io/products/harbor-container-registry#/releases/190421

DNS

The FQDNs are used in the example:

Table 3-4.

Component	FQDN	IP Address
ops-manager	ops-manager.isbu.local	
pks	pks-api.isbu.local	192.168.0.23
harbor	pks-harbor.isbu.local	192.168.0.24
pks-cluster-01	pks-cluster-01.isbu.local	192.168.2.1

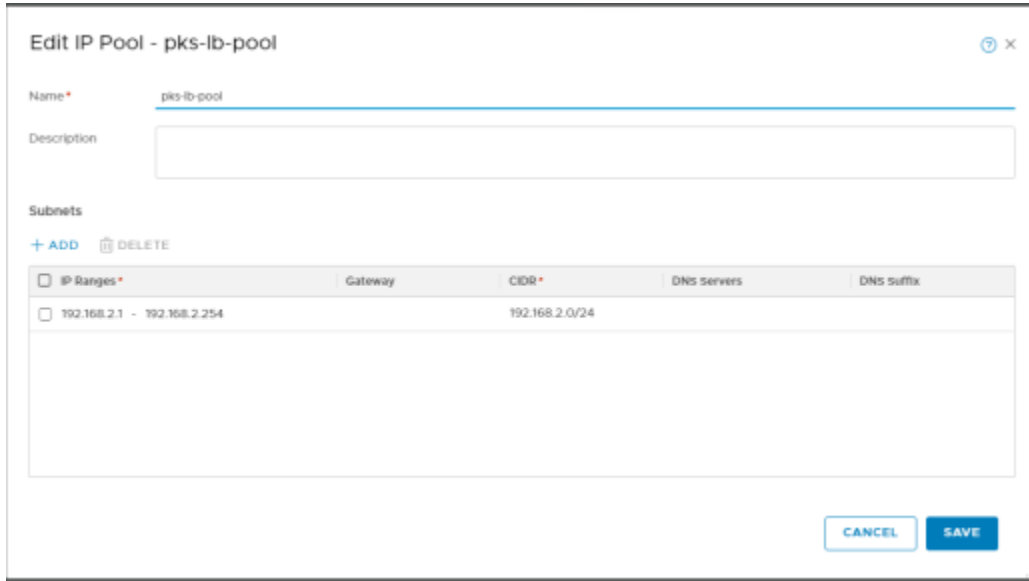
Prepare NSX-T

Open the UI of NSX-T Manager. Perform the following tasks to prepare NSX-T:

- Inventory: You will require an IP pool from where PKS assigns the IP addresses for VIP of the load balancers for Kubernetes Pod services. Add a new IP pool.

Table 3-5.

Name	IP Ranges	Gateway	CIDR	DNS Servers	DNS Suffix
pks-lb-pool	192.168.2.1-192.168.2.254		192.168.2.0/24		



- Network:
 - IPAM: You require two IP blocks in the IPAM for PKS to retrieve the IP addresses for Kubernetes Nodes and Pod. Add the two new IP blocks.

Table 3-6.

Name	CIDR
Node-Block	20.0.0.0/16
Pod-Block	30.0.0.0/16

Edit IP Block - Node-Block ? ×

Name*

Description

CIDR*

CANCEL
SAVE

Edit IP Block - Pod-Block ? ×

Name*

Description

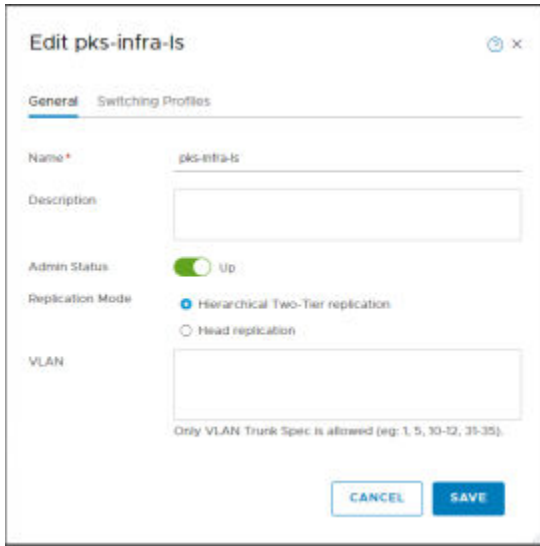
CIDR*

CANCEL
SAVE

- Switching: Create two new logical switches.

Table 3-7.

Logical Switch	Transport Zone
pks-infra-ls	overlay-tz-<ID>
pks-services-ls	overlay-tz-<ID>

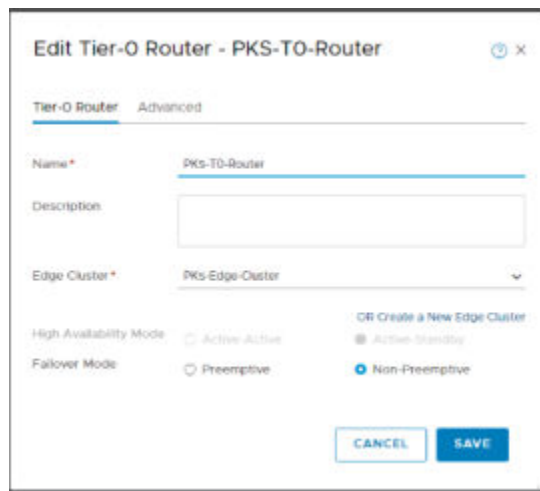



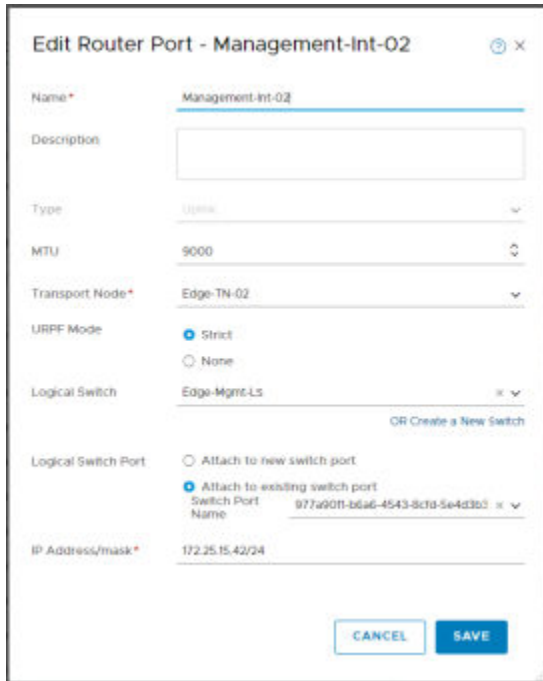
■ Routers:

- 1 Tier 0 router: Configure T0 router as Active-Standby. The details of the logical router ports are as follows:

Table 3-8.

Port	Type	MTU	Transport Node	Logical Switch	IP Address/Mask
External-Int	Uplink	9000	Edge-TN-01	Edge-External-LS	10.198.10.117/24
Management-Int-01	Uplink	9000	Edge-TN-01	Edge-Mgmt-LS	172.25.15.41/24
Management-Int-01	Uplink	9000	Edge-TN-01	Edge-Mgmt-LS	172.25.15.41/24





Configure dynamic or static routes for the external access. In the example in this document, we use static routing.

- 2 Tier 1 Router: PKS-Infra-T1-Router is used to connect the pks-infra-ls and pks-services-ls logical switches that are used to deploy PCF, Bosh, PKS, and Harbor.

Table 3-9.

Name	Tier 0 Router	Edge Cluster	Failover Mode	Edge Cluster Members
PKS-Infra-T1-Router	PKS-T0-Router	PKS-Edge-Cluster	Non-Preemptive	Edge-TN-01 Edge-TN-02

New Tier-1 Router ? ×

Tier-1 Router **Advanced**

Name *

Description

Tier-0 Router × ▾

Edge Cluster × ▾

Failover Mode Preemptive Non-Preemptive

Edge Cluster Members ⓘ × × × ▾

The logical router ports are PKS-Infra-Int and PKS-Services-Int.

Edit Router Port - PKS-Infra-Int

Name*

Description

Type

URPF Mode Strict None

Logical Switch OR Create a New Switch

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

Switch Port Name

IP Address/mask*

Relay Service

Edit Router Port - PKS-services-Int

Name*

Description

Type

URPF Mode Strict None

Logical Switch OR Create a New Switch

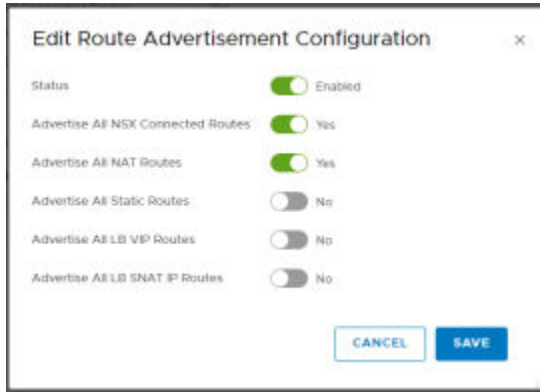
Logical Switch Port Attach to new switch port Attach to existing switch port

Switch Port Name

IP Address/mask*

Relay Service

Enable the route advertisement for the connected and all NAT routes.



Ops Manager

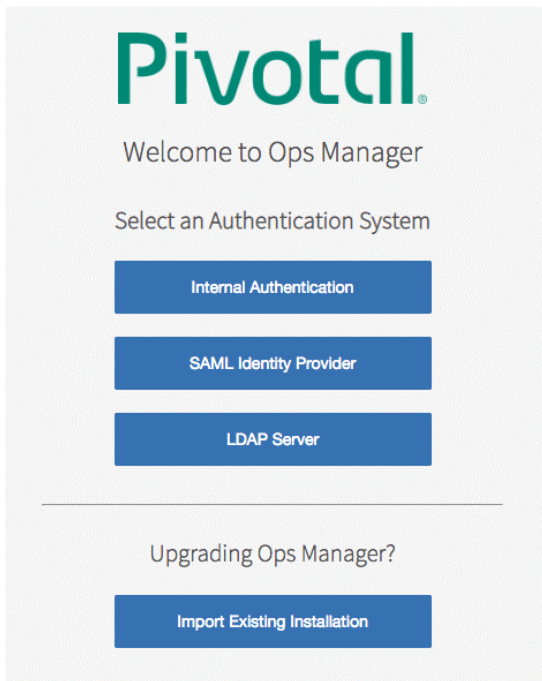
For the official Pivotal documentation, refer <https://docs.pivotal.io/pivotalcf/2-3/customizing/deploying-vm.html> .

Installing the Pivotal Container Service Ops Manager

Procedure

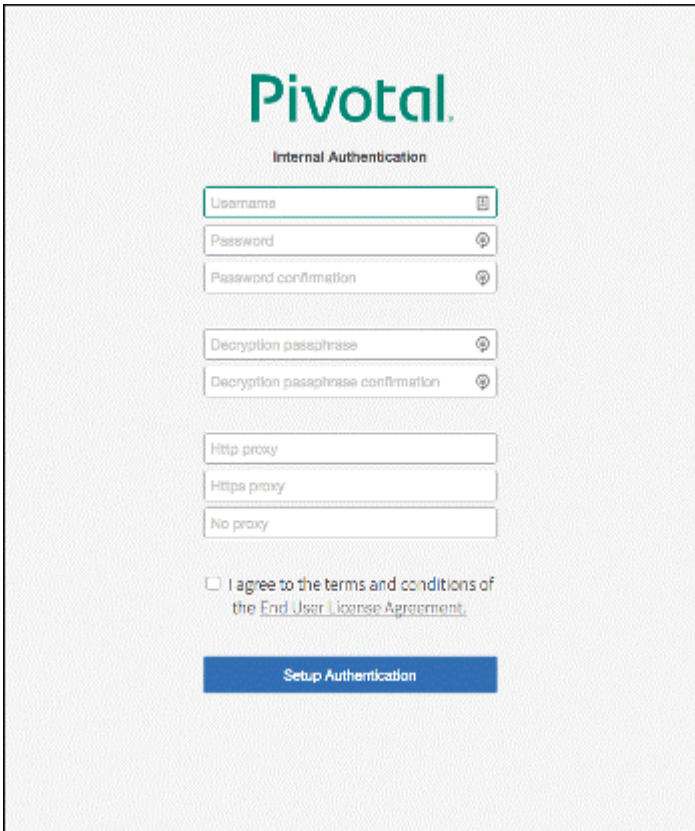
- 1 Download [Pivotal Cloud Foundry \(PCF\) Ops Manager](#) for vSphere from [Pivotal Network](#).
- 2 Use the vSphere client to deploy the OVA.
 - a Select `pks-infra-ls` from the drop-down menu to connect the network adapter.
 - b Provide the IP address, network mask, gateway. Provide DNS and NTP corresponding to the `pks-infra-ls` network.
 - c Create a DNS entry for the IP address that you used for Ops Manager.

- 3 Access <https://<FQDN of OpsManager>> or <https://<IP of the Ops Manager>> in a browser.



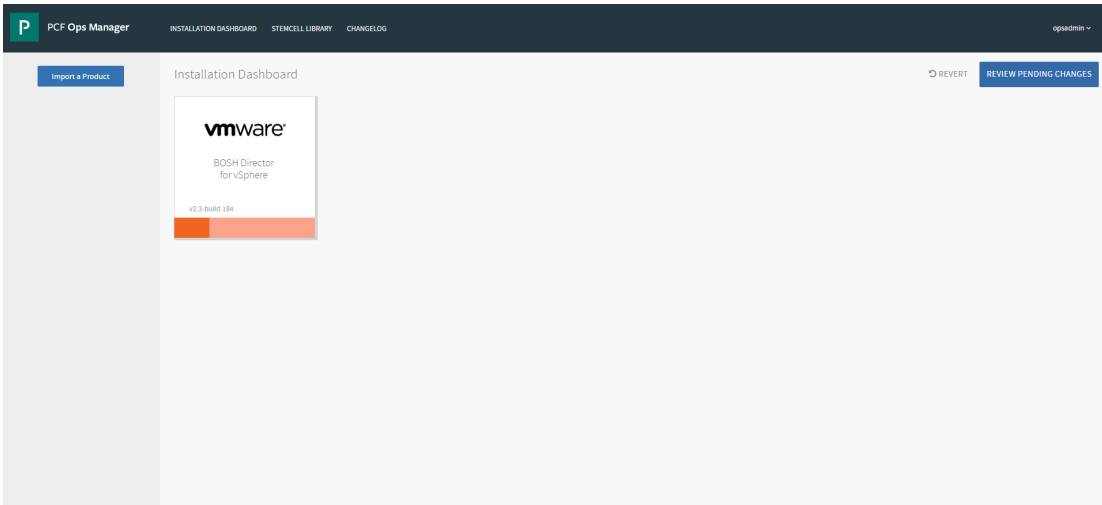
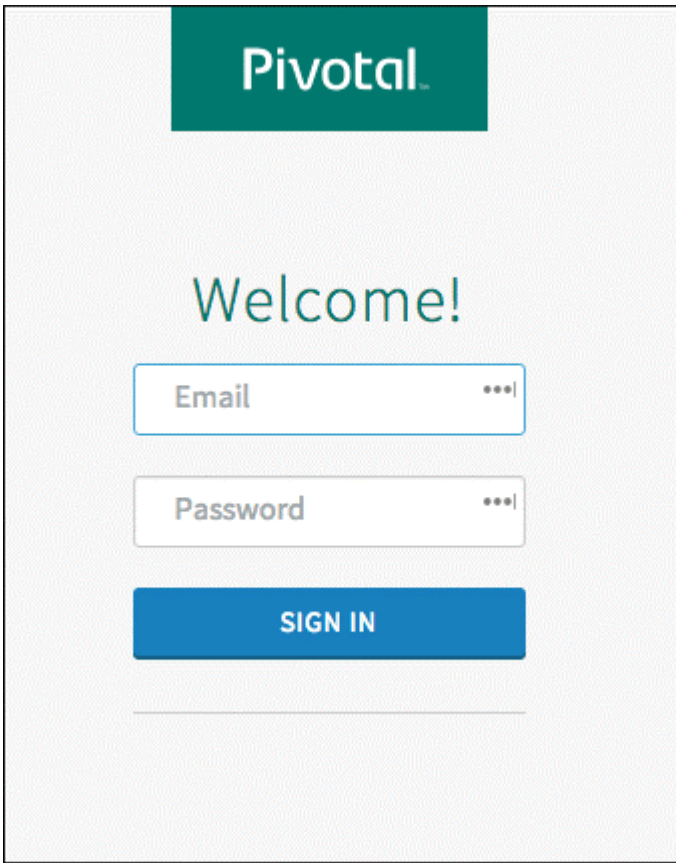
- 4 Select **Internal Authentication**.

- 5 When redirected to the Internal Authentication page, you must complete the following steps:
 - a Enter **Username**, and **Password confirmation** to create an admin user.
 - b Enter **Decryption passphrase** and the **Decryption passphrase confirmation**. This passphrase encrypts the Ops Manager datastore and is not recoverable.
 - c Click **Setup Authentication**.



The screenshot shows the Pivotal Internal Authentication page. At the top, the Pivotal logo is displayed in green. Below the logo, the text "Internal Authentication" is centered. The page contains several input fields: "Username" (with a help icon), "Password" (with a strength indicator), "Password confirmation" (with a strength indicator), "Decryption passphrase" (with a strength indicator), and "Decryption passphrase confirmation" (with a strength indicator). Below these are three radio button options for proxy settings: "Http proxy", "Https proxy", and "No proxy". A checkbox is present for "I agree to the terms and conditions of the End User License Agreement," with a link to the agreement. At the bottom, there is a blue button labeled "Setup Authentication".

- 6 Log in to Ops Manager with the admin user name and password you created in the previous step.



NSX Certificates

You need to prepare and replace the NSX certificates.

For information on NSX Manager Superuser Principal Identity, refer <https://docs.pivotal.io/runtimes/pks/1-2/generate-nsx-pi-cert.html>.

For information on generating and registering the NSX Manager certificate, refer <https://docs.pivotal.io/runtimes/pks/1-2/generate-nsx-ca-cert.html>.

Configure Bosh Director for vSphere

The **vCenter Config** page consists of the following information:

Procedure

- 1 Enter the required information on the **vCenter Config** page. This page has the following information.

Table 3-10.

Name	Description
Name	A name you provide for your vCenter configuration.
vCenter Host	The hostname of the vCenter that manages ESXi/vSphere.
vCenter Username	A vCenter username with create and delete privileges for virtual machines (VMs) and folders.
vCenter Password	The password for the vCenter user specified above
Datacenter Name	The name of the datacenter as it appears in vCenter
Virtual Disk Type	The Virtual Disk Type to provision for all VMs
Ephemeral Datastore Names (comma delimited)	The names of the datastores that store ephemeral VM disks deployed by Ops Manager
Persistent Datastore Names (comma delimited)	The names of the datastores that store persistent VM disks deployed by Ops Manager.
VM Folder	The vSphere datacenter folder where Ops Manager places VMs. It defaults to <code>pcf_vms</code> .
Template Folder	The vSphere datacenter folder where Ops Manager places the VMs. This defaults to <code>pcf_templates</code> .
Disk Path Folder	The vSphere datastore folder where Ops Manager creates attached disk images. This defaults to <code>pcf_disk</code> .

BOSH Director for vSphere

Settings | Status | Credentials

vCenter Config Add vCenter Config

- Director Config
- Create Availability Zones
- Create Networks
- Assign AZs and Networks
- Security
- Syslog
- Resource Config

Name:

vCenter Host: The hostname of the vCenter that manages ESXi / vSphere

vCenter Username:

vCenter Password:
[Change](#)

Datacenter Name:

Virtual Disk Type:

Ephemeral Datastore Names (comma delimited):

NOTE: Removing an Ephemeral Datastore after an initial deploy can result in a system outage and/or data loss.

Persistent Datastore Names (comma delimited):

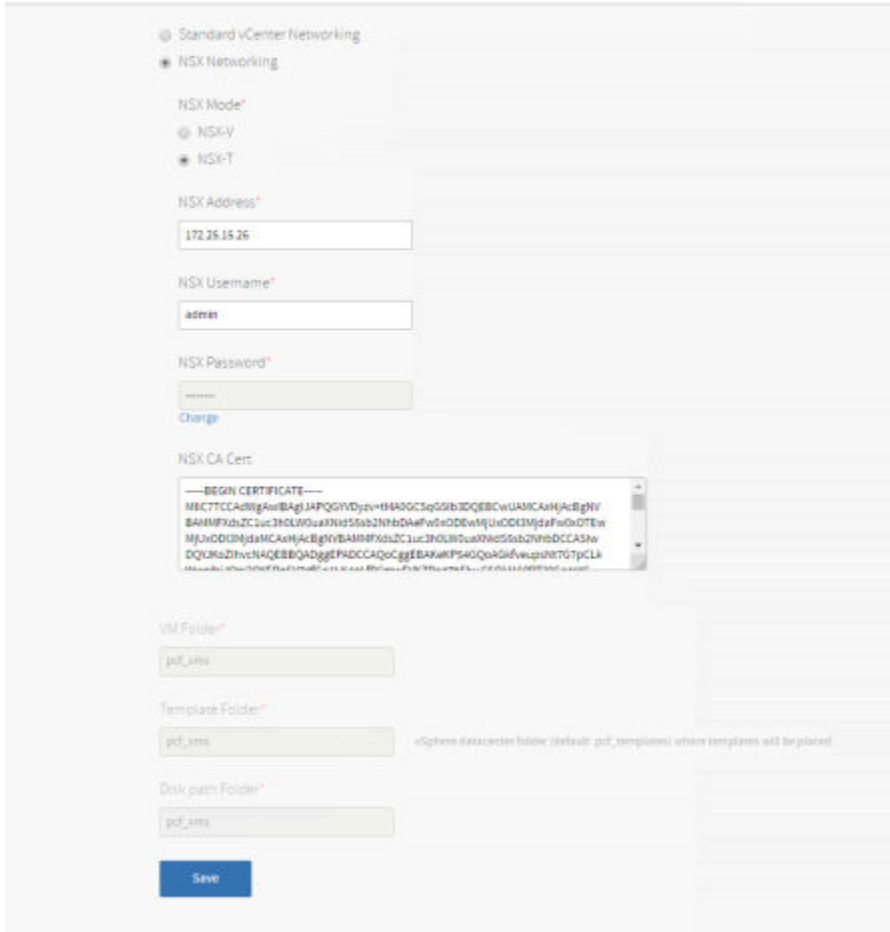
NOTE: Removing a Persistent Datastore after an initial deploy can result in a system outage and/or data loss.

Standard vCenter Networking

NSX Networking

Activate Windows
Go to System in Control Panel to activate Windows.

- 2 Select **NSX Networking** to enable VMware NSX Network Virtualization.
 - a Enter the NSX-T manager IP address.
 - b Enter the NSX-T manager user name and password.
 - c In **NSX CA Cert**, paste the certificate generated in the Generating and Registering the NSX Manager Certificate.



- 3 For the director configuration, enter the IP address of the NTP server. Leave the defaults as is.
- 4 Add two availability zones in the same cluster in the workload with different resource pools (these should be created before.)

Table 3-11.

Availability Zone	IaaS Configuration	Clusters	Resource Pool
pks-infra	default	pks-wld-cluster	pks-infra-pool
pks-services	default	pks-wld-cluster	k8-worker-pool

- a Click **Add**.
- b Enter a unique name for the availability zone.

- c Select a vCenter config name from the IaaS Configuration drop-down menu to associate your availability zone with a vCenter.
- d Enter the name of an existing vCenter Cluster to use as an availability zone.
- e Enter the name of a Resource Pool in the vCenter cluster that you specified above. The jobs running in this availability zone share the CPU and memory resources defined by the pool.

5 Create the two networks.

Table 3-12.

Network	vSphere Network Name	CIDR	Reserved IP Ranges	DNS	Gateway	Availability Zones
pkc-infra-ntk	pkc-infra-ls	192.168.0.0/24	192.168.0.1-19 2.168.0.21	172.25.15.45	192.168.0.1	pkc-infra pkc-worker
pkc-services-ntk	pkc-services-ls	192.168.1.0/24	192.168.1.1-19 2.168.1.21	172.25.15.45	192.168.1.1	pkc-infra pkc-worker

- a Select **Enable ICMP checks** to enable ICMP on your networks. Ops Manager uses ICMP checks to confirm that the components within your network are reachable.
- b Use the following steps to create one or more Ops Manager networks (use the values from the table above):
 - 1 Click **Add Network**.
 - 2 Enter a unique name for the network.
 - 3 Click **Add Subnet** to create one or more subnets for the network.
 - 4 Enter the full path and **vSphere Network Name** as it displays in vCenter. For example, enter YOUR-DIRECTORY-NAME/YOUR-NETWORK-NAME. If your vSphere network name contains a forward slash character, replace the forward slash with the URL-encoded forward slash character %2f.
 - 5 For CIDR, enter a valid CIDR block in which to deploy VMs. For example, enter 192.0.2.0/24.
 - 6 For Reserved IP Ranges, enter any IP addresses from the CIDR that you want to blacklist from the installation. Ops Manager will not deploy VMs to any address in this range.
 - 7 Enter your DNS IP addresses.
 - 8 Select the availability zones to use with the network.
- 6 To assign the availability zones and the networks, use the drop-down menu to select a Singleton Availability Zone . The BOSH Director installs in this Availability Zone such as pkc-infra. Select a Network for your BOSH Director such as pkc-infra-ntk.

- 7 To configure security, perform the following tasks:
 - a In **Trusted Certificates**, enter your custom certificate authority (CA) certificates to insert in your organization certificate trust chain. This feature enables all BOSH-deployed components in your deployment to trust the custom root certificates.
 - b Choose **Generate passwords** or **Use default BOSH password**. Pivotal recommends that you use the **Generate passwords** option for greater security.
- 8 Do not configure the syslog for now.
- 9 Leave the defaults as is for **Resource Config**.
- 10 Click **Review Pending Changes** and **Apply Changes** on the right navigation bar on the Installation Dashboard.

Pivotal Container Service

The general guidance from Pivotal is available at <https://docs.pivotal.io/runtimes/pks/1-2/installing-nsx-t.html#next-steps>.

Install Pivotal Container Service

Before installing Pivotal Container Service (PKS), you must have deployed and configured Ops Manager.

To install PKS, perform the following steps:

Prerequisites

If you use an instance of Ops Manager that you configured previously to install other run time instances, confirm the following settings before you install PKS:

- 1 Navigate to Ops Manager.
- 2 Open the **Director Config** pane.
- 3 Select **Enable Post Deploy Scripts**.
- 4 Clear the Disable BOSH DNS server for troubleshooting purposes.
- 5 Click the **Installation Dashboard** link to return to the Installation Dashboard.
- 6 Click **Apply Changes**.


Procedure

- 1 Download the product file from [Pivotal Network](#).
- 2 Navigate to <https://YOUR-OPS-MANAGER-FQDN/> in a browser to log in to the Ops Manager Installation Dashboard.
- 3 Click **Import a Product** to upload the product file.
- 4 Under Pivotal Container Service in the left column, click the plus sign to add this product to your staging area.

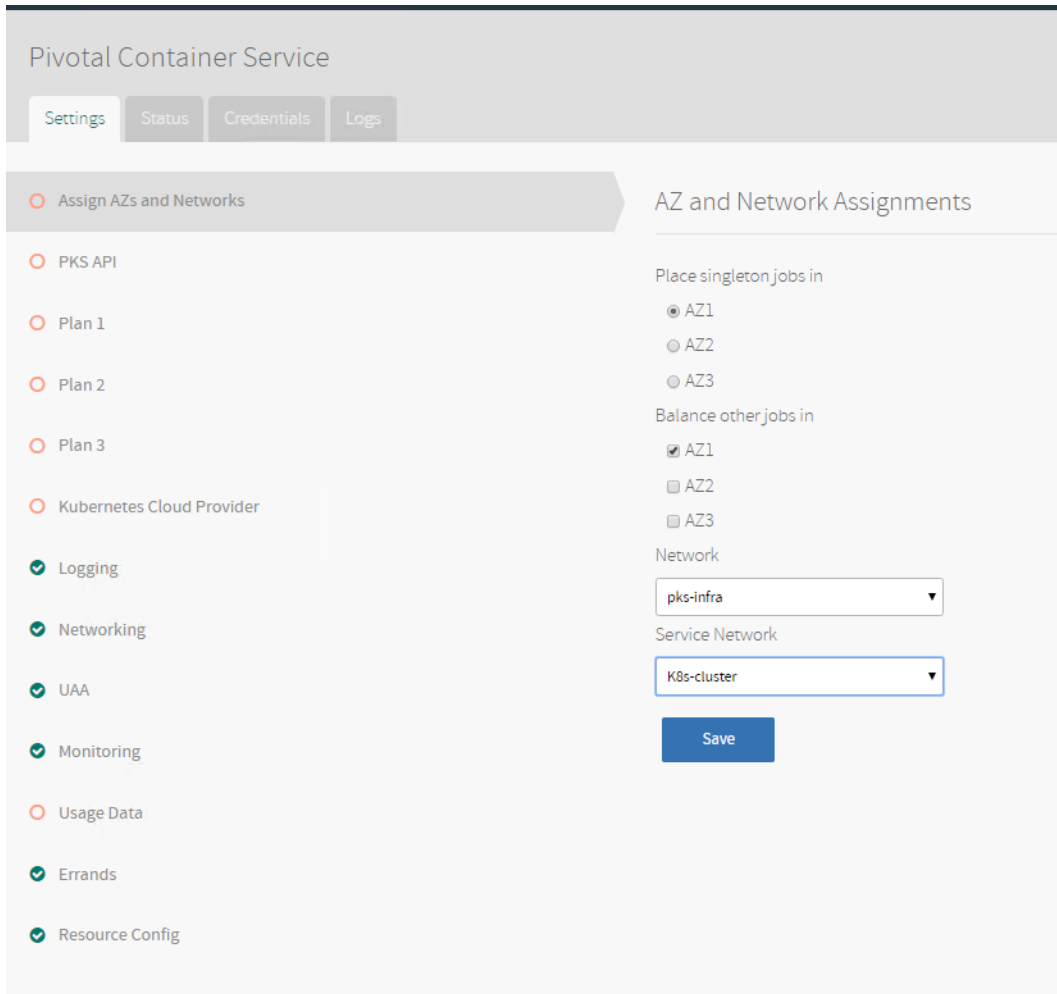
Configure Pivotal Container Service

Procedure

- 1 Click the orange Pivotal Container Service (PKS) tile to start the configuration process.



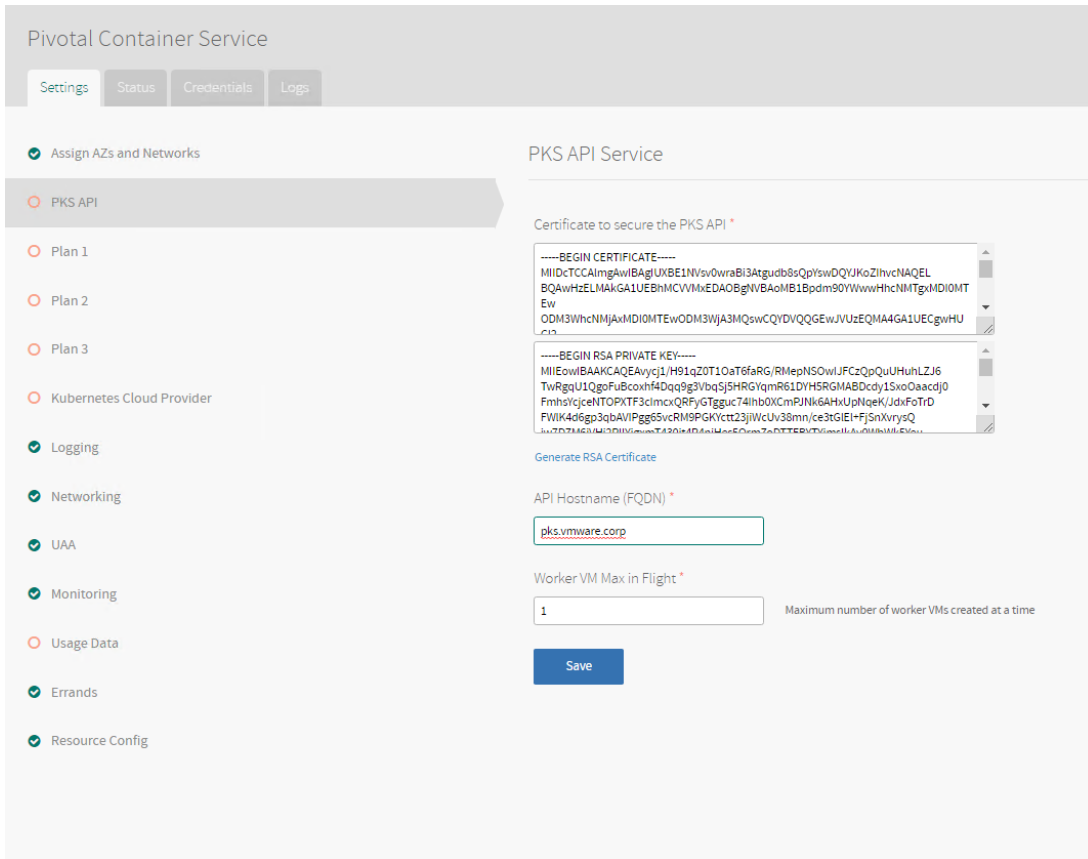
- 2 Click **Assign AZs and Networks** and select the availability zone (AZ) where you want to deploy the PKS API VM as a singleton job.



- a Under **Network**, select **PKS Management Network** linked to the `pks-infra-ls` NSX-T logical switch you created in the **Create Networks Page** step. This provides network placement for the PKS API VM.
- b Under **Service Network**, your selection depends on whether you are upgrading from a previous PKS version or installing an original PKS deployment.
 - If you are deploying PKS with NSX-T for the first time, the **Service Network** field does not apply because PKS instructs NSX-T to create a service network on-demand each time a new Kubernetes cluster is requested. However, the PKS tile requires you to make a selection. Therefore, select the same network you specified in the Network.
 - If you are upgrading from a previous PKS version, select **Service Network** linked to the `pks-service-ls` NSX-T logical switch that is created by PKS during installation. The service network provides network placement for the existing on-demand Kubernetes cluster service instances created by the PKS broker.

Click **Save**.

3 To configure the PKS API service:



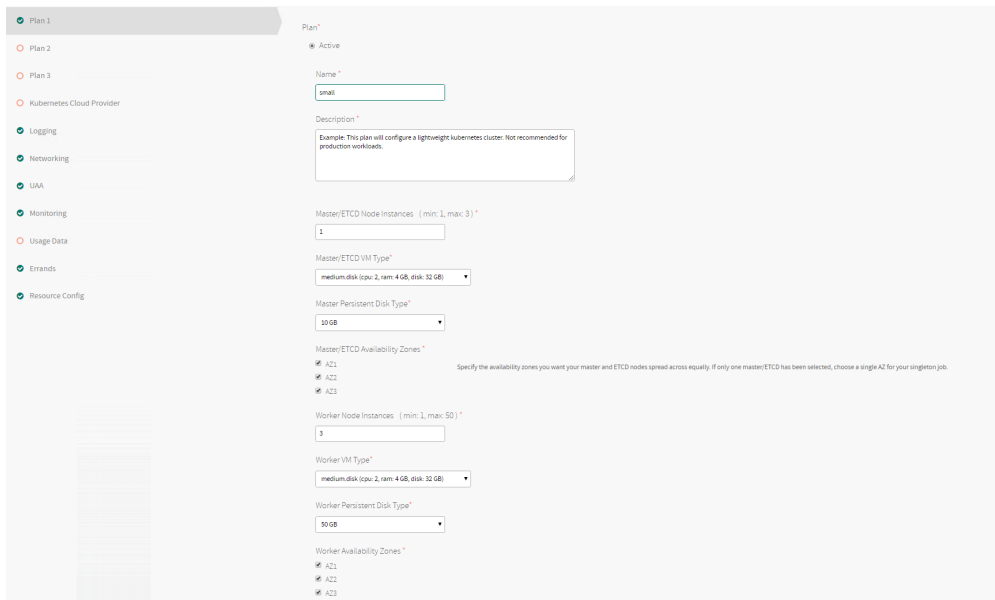
- a Click **PKS API**.
- b Under **Certificate to secure the PKS API**, provide your own certificate and private key pair. The certificate you enter here should cover the domain that routes to the PKS API VM with TLS termination on the ingress.
- c (Optional) If you do not have a certificate and private key pair, you can have Ops Manager generate one for you. Perform the following steps:
 - 1 Select **Generate RSA Certificate**.
 - 2 Enter the wildcard domain for your API hostname. For example, if your PKS API domain is `api.pks.example.com`, then enter `*.pks.example.com`.
 - 3 Click **Generate**.
- d Under **API Hostname (FQDN)**, enter a Fully Qualified Domain Name (FQDN) to access the PKS API. For example, `api.pks.example.com`.
- e Click **Save**.

4 To activate a plan, perform the following tasks:

- a Click the Plan 1, Plan 2, or Plan 3 tab.

Table 3-13.

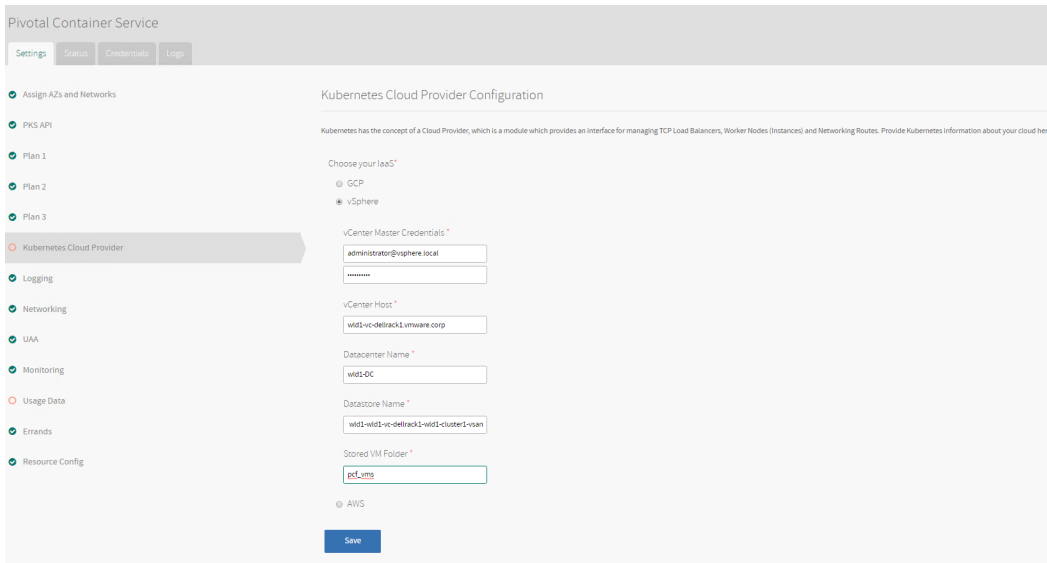
Plan	State	Master/ETCD Node Instances	Master/ETCD Availability Zones	Worker Node Instances	Worker Availability Zones
Plan 1	Active	1	pkc-infra	3	pkc-worker
Plan 2	Active	3	pkc-infra	5	pkc-worker
Plan 3	Active	3	pkc-infra	7	pkc-worker



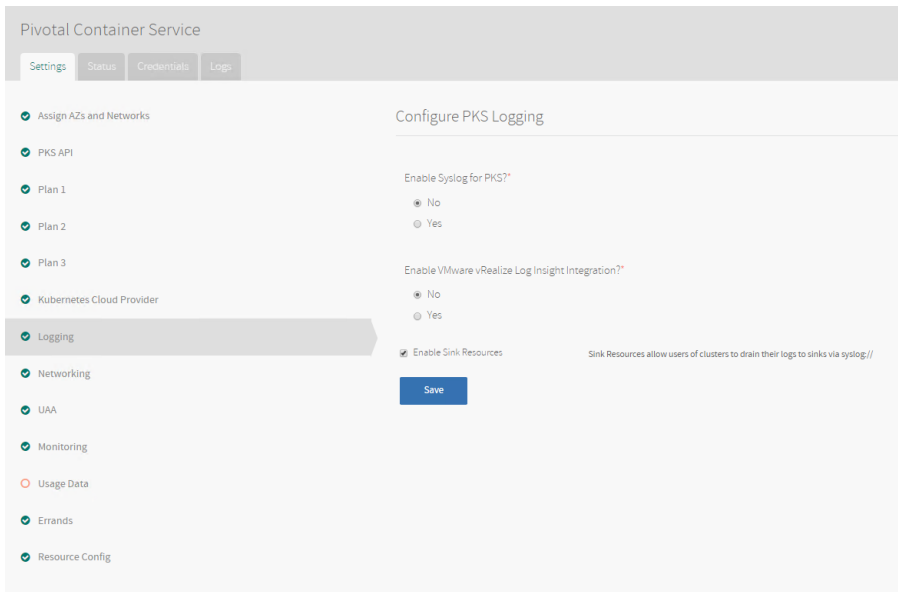
For Plan1

- 1 Under **Master/ETCD Node Instances**, select the default number of Kubernetes master/etcd nodes to provision for each cluster. Select **1**.
- 2 Under **Master/ETCD Availability Zones**, select one or more availability zones for the Kubernetes clusters deployed by PKS. If you select more than one AZ, PKS deploys the master VM in the first AZ and the worker VMs across the remaining AZs. Select **Master/ETCD Availability Zones**. For example, pkc-infra.
- 3 Under Worker Node Instances, select the default number of Kubernetes worker nodes to provision for each cluster. For high availability, create clusters with a minimum of three worker nodes, or two per AZ if you intend to use persistent volumes. For example, if you deploy across three AZs, you should have six worker nodes. Select **3**.
- 4 Select **Worker Availability Zones**. For example, pkc-worker.

- 5 For Kubernetes Cloud Provider settings, select **vSphere** under **Choose your IaaS**. Ensure the values in the following procedure match those in the **vCenter Config** section of the **Ops Manager** tile.



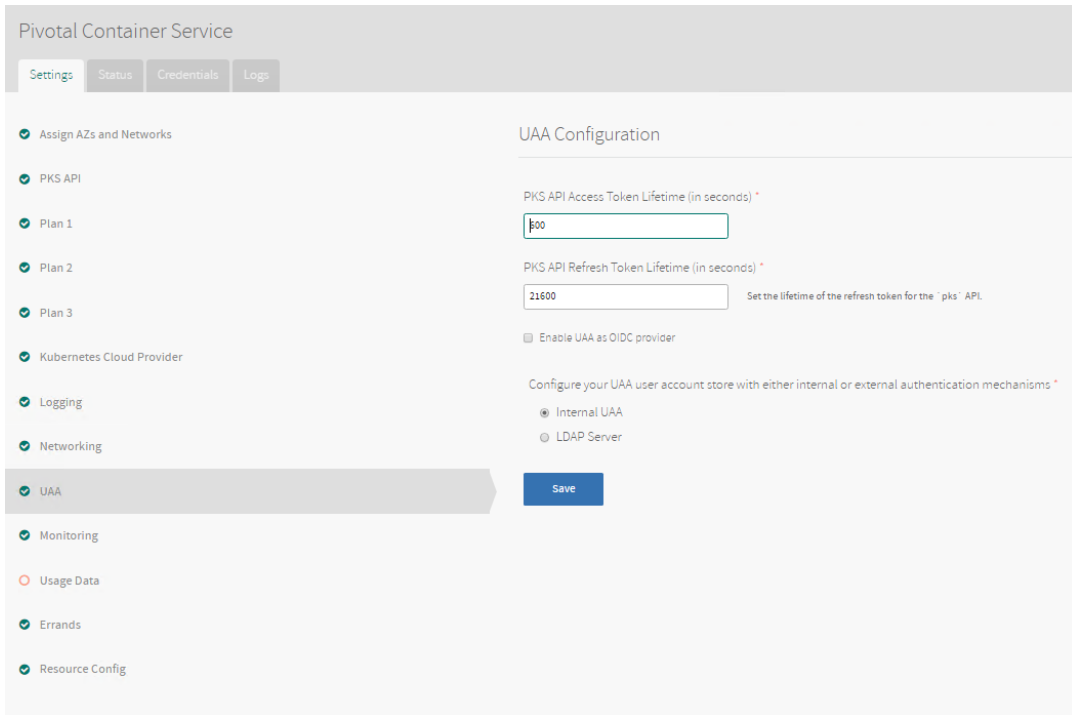
- 6 The **Configure PKS Logging** step is optional.



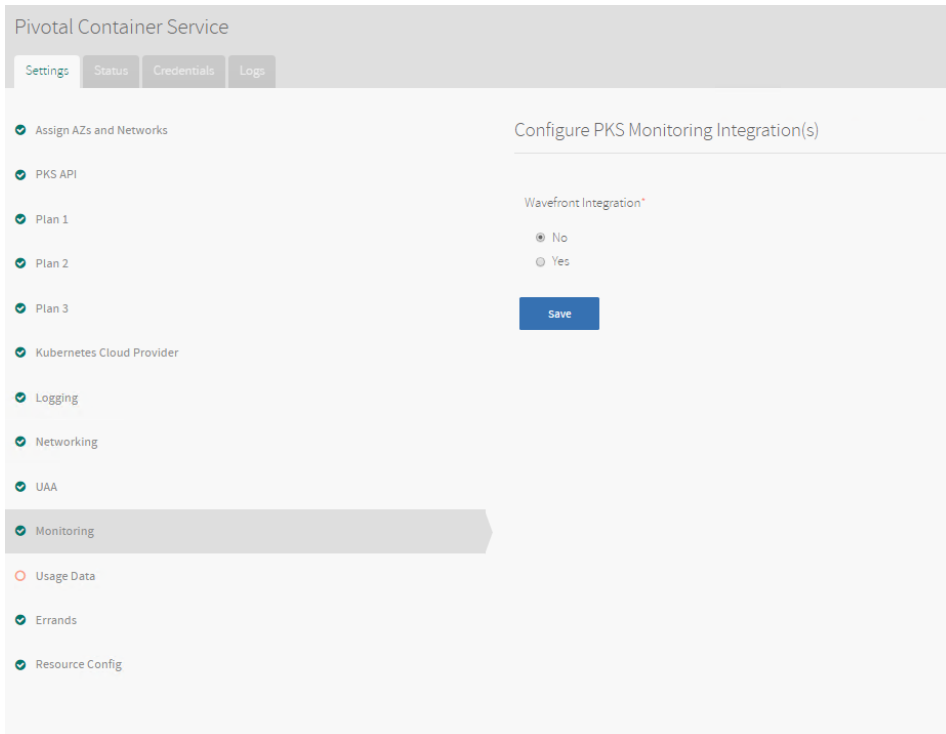
- 7 For networking configuration, perform the following tasks:
 - a Under Container Networking Interface, select NSX-T.
 - b For NSX Manager hostname, enter the hostname or IP address of your NSX Manager.
 - c For NSX Manager Super User Principal Identify Certificate, copy and paste the contents and private key of the Principal Identity certificate you created in [Generating and Registering the NSX Manager Superuser Principal Identity Certificate and Key](#).

- d (Optional) For NSX Manager CA Cert, copy and paste the contents of the NSX Manager CA certificate you created in [Generating and Registering the NSX Manager Certificate](#). Use this certificate and key to connect to the NSX Manager.
- e The Disable SSL certificate verification checkbox is not selected by default. In order to disable TLS verification, select the check box. You can disable TLS verification if you did not enter a CA certificate, or if your CA certificate is self-signed.
- f If you are using a NAT deployment topology, leave the **NAT mode** check box selected. If you are using a No-NAT topology, clear this checkbox. For more information, see [Deployment Topologies](#).
- g Enter the following values:
 - Pods IP Block ID: Enter the UUID (located in the NSX-T Manager) of the IP block to be used for Kubernetes pods. PKS allocates IP addresses for the pods when they are created in Kubernetes. Each time a namespace is created in Kubernetes, a subnet from this IP block is allocated.
 - Nodes IP Block ID: Enter the UUID (located in the NSX-T Manager) of the IP block to be used for Kubernetes nodes. PKS allocates IP addresses for the nodes when they are created in Kubernetes. The node networks are created on a separate IP address space from the pod networks.
 - For T0 Router ID, enter the t0-pks T0 router UUID. Locate this value in the NSX-T Manager UI router overview.
 - For Floating IP Pool ID, enter the ip-pool-vips ID that you created for load balancer VIPs. PKS uses the floating IP pool to allocate IP addresses to the load balancers created for each of the clusters. The load balancer routes the API requests to the master nodes and the data plane.
 - For Nodes DNS, enter one or more Domain Name Servers used by the Kubernetes nodes.
 - For vSphere Cluster Names, enter the vSphere cluster of the workload where you deploy Kubernetes clusters.

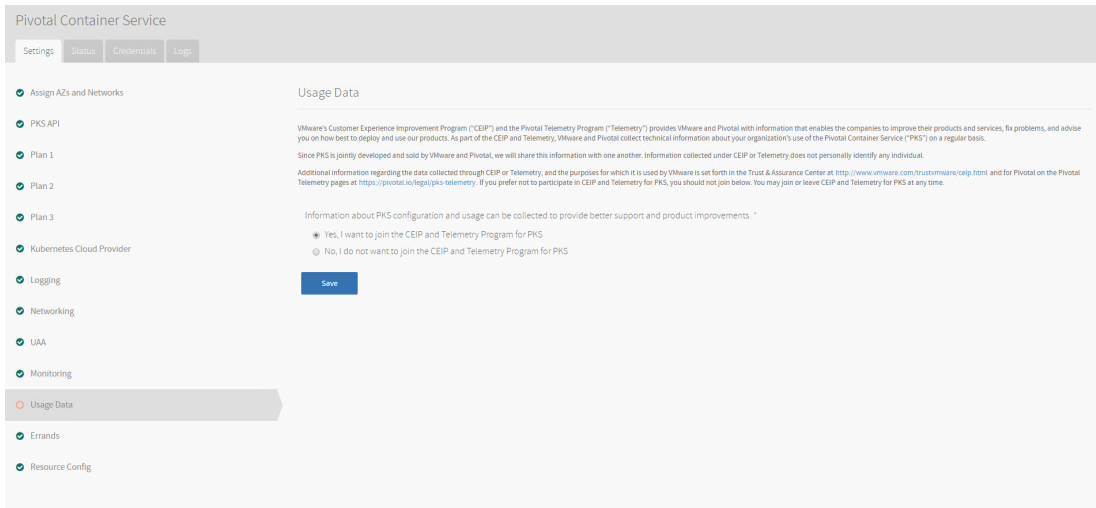
8 For UAA settings, leave the default values as is and select **Internal UAA**.



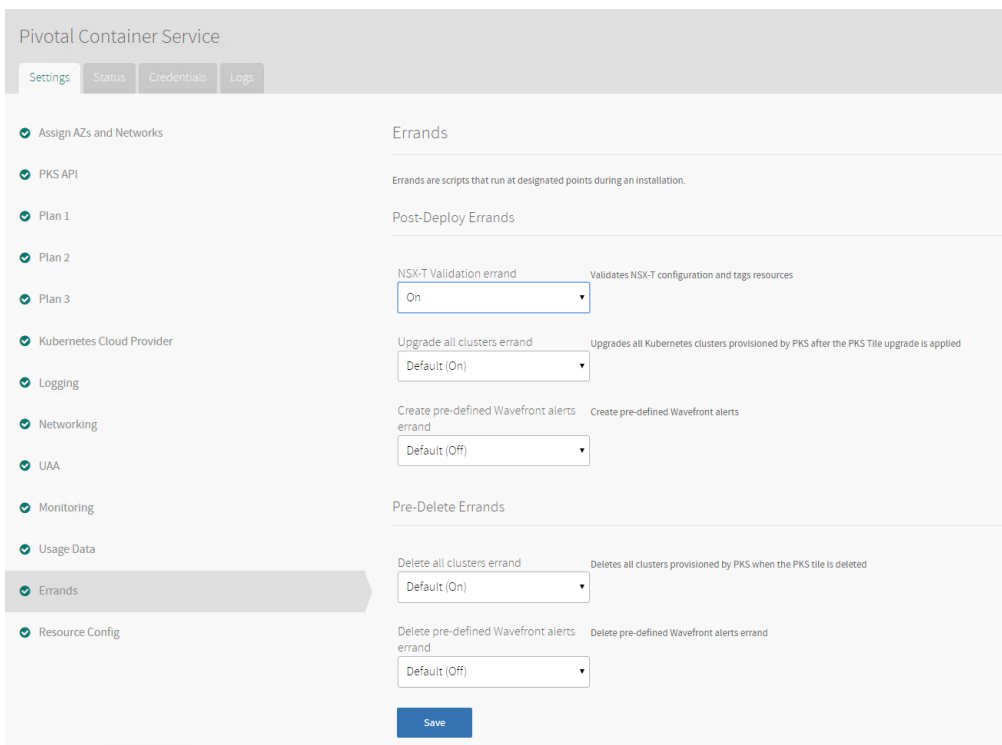
9 For the monitoring integration settings, select **No** for the example in this document. It implies that Wavefront is not integrated.



10 For the usage data, you can choose whether to join CEIP and Telemetry Program for PKS.



11 To configure when post-deploy and pre-delete errands for PKS are run, make a selection in the drop-down menu next to the errand. For example, select **NSX-T Validation errand: ON**.



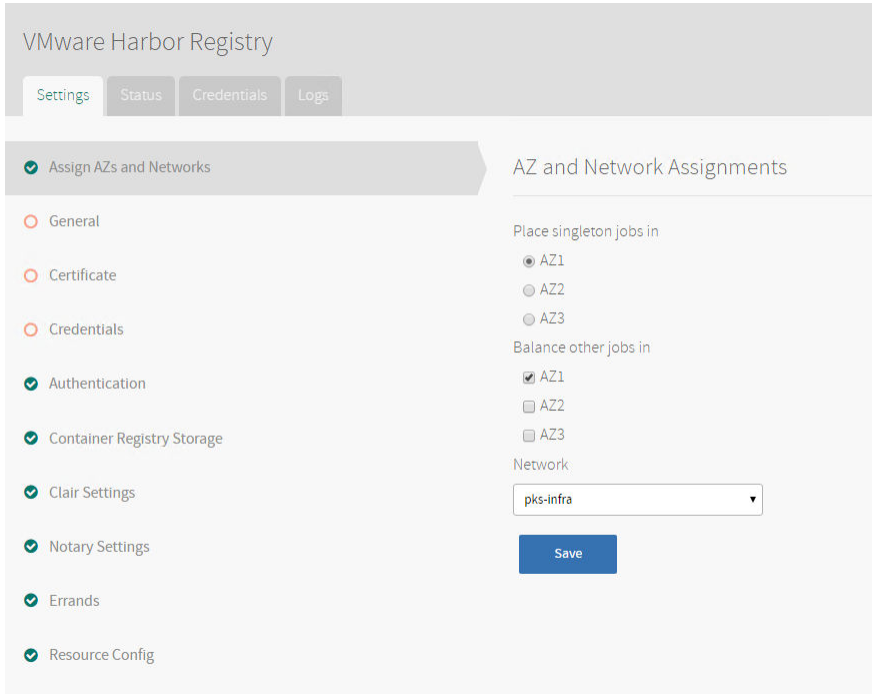
12 For resource configuration, leave default values as is.

Harbor Registry

Configure Harbor Registry

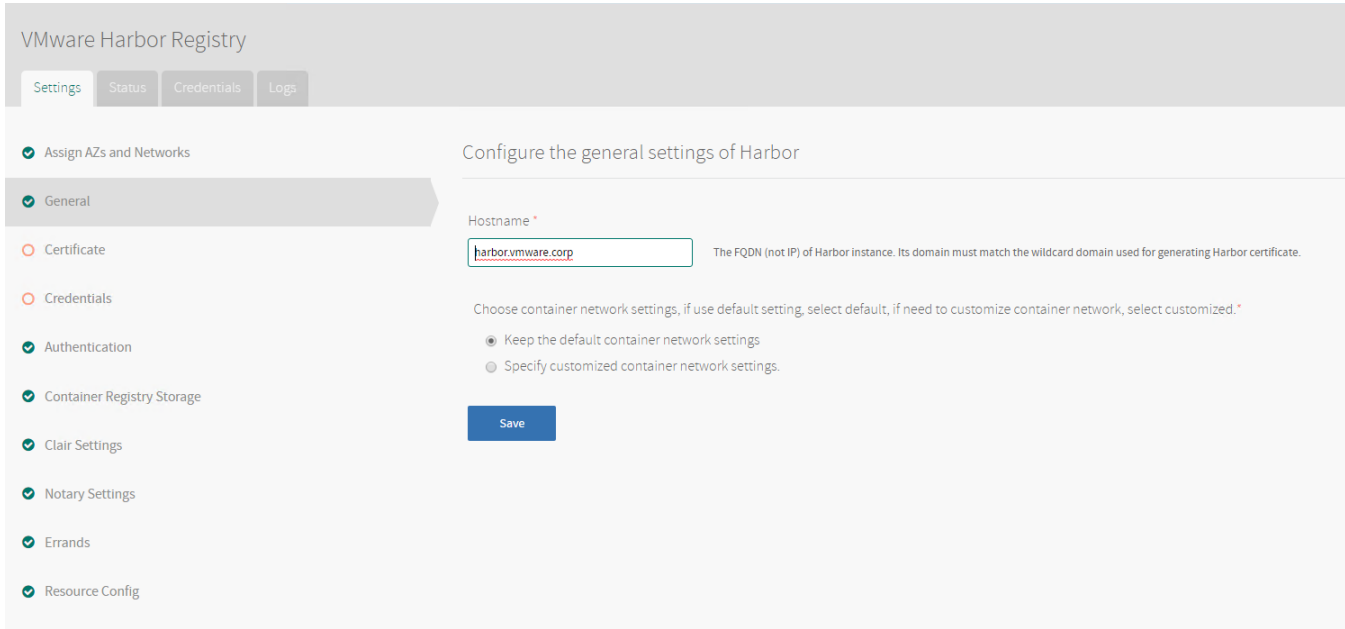
Procedure

- 1 To assign the availability zones and networks:
 - a Select the availability zone under **place singleton jobs in**. Harbor is a singleton job and will be placed on this network. Select **pks-infra**.
 - b Select the availability zone where to balance other jobs in . For PKS, this is the same availability zone as the singleton. Select **pks-infra**.



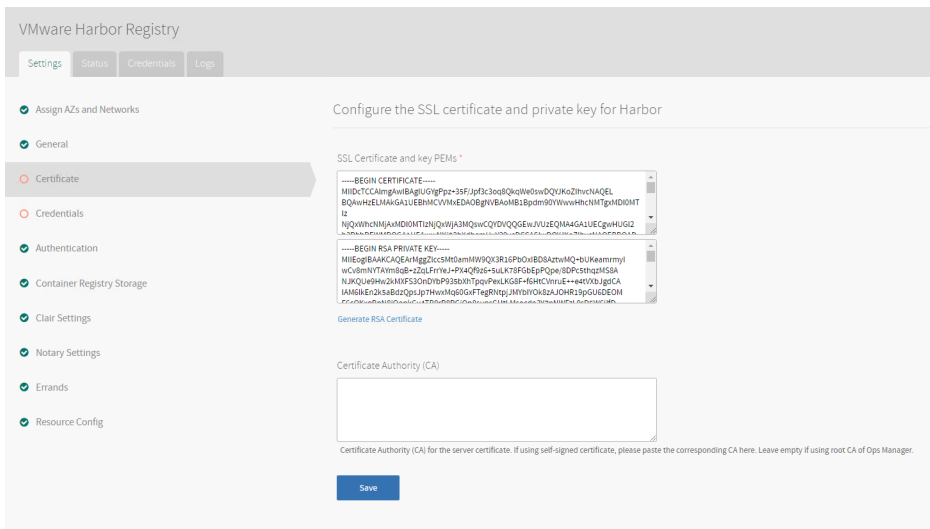
- 2 To configure the general settings of Harbor, enter a Hostname (FDQN) to access the Harbor administration UI and the registry service. For example, `pks-harbor.isbu.local`. The host name must include a domain and must be able to resolve to the IP address of the Harbor instance VM by an external DNS server. The Kubernetes worker nodes can resolve the Harbor FQDN through the

local BOSH DNS server. You must provide a Harbor FQDN that is resolvable by an external DNS server so that the Docker clients external to Kubernetes worker nodes can resolve the Harbor FQDN. When Harbor is successfully deployed, you need to update the Harbor external DNS record with the IP address of the Harbor VM.

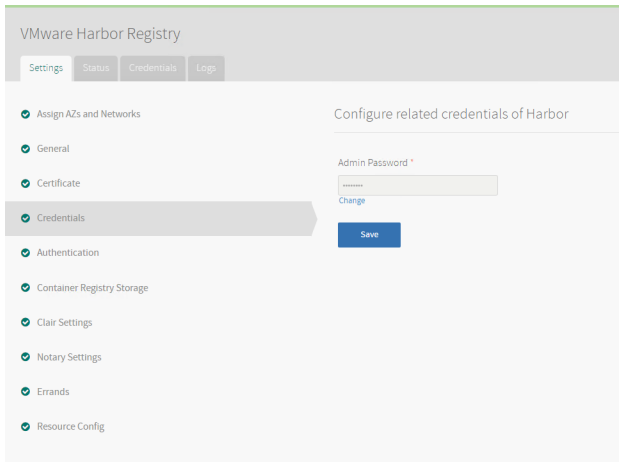


3 To use a certificate that Ops Manager generates automatically, perform the following steps:

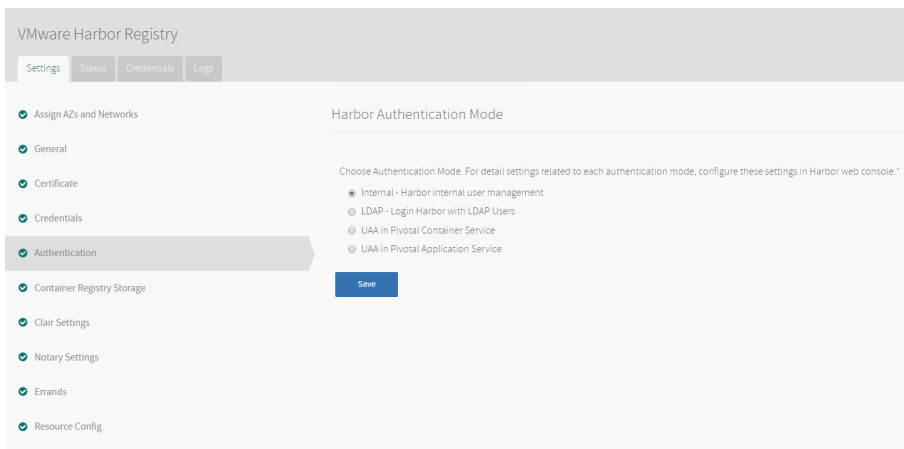
- a Click **Generate RSA Certificate**.
- b Enter a domain name wildcard in the **Generate RSA Certificate** text box. The domain name wildcard must match the DNS resolvable domain name that you used when you specified the hostname for Harbor. For example, if you set the Harbor hostname to [harbor.isbu.local](#), enter `*.isbu.local` in the **Generate RSA Certificate** text box.
- c Click **Generate**.



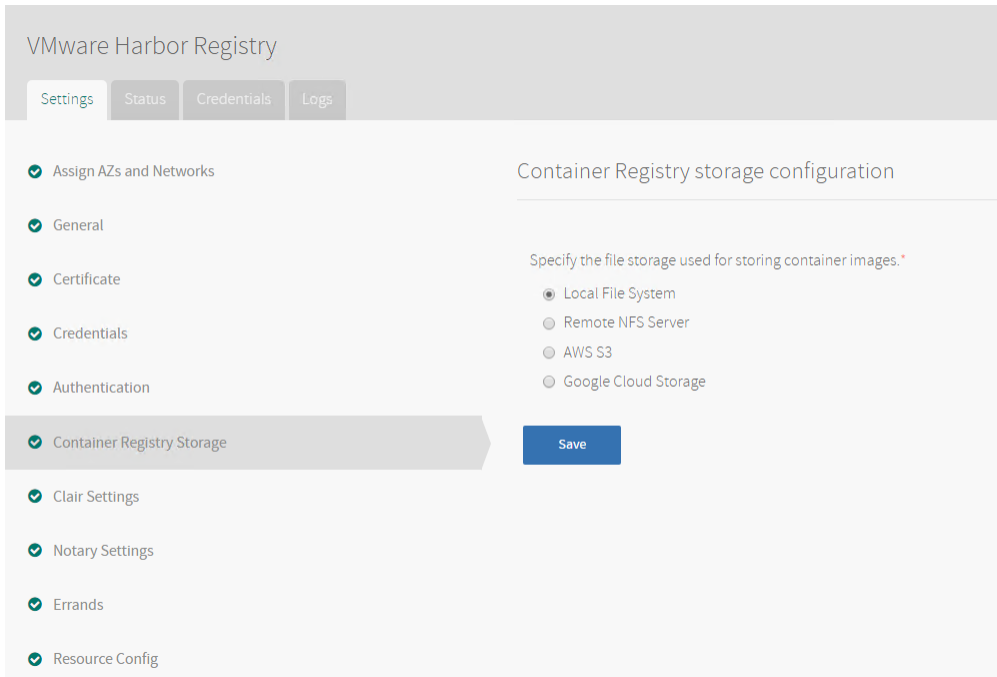
- 4 To configure the related credentials of Harbor, set a password for the Harbor system administrator account. You cannot change the admin password in Ops Manager after you set it. Use the Harbor interface to make the subsequent changes to the password after deployment.



- 5 Select the authentication mode. For example, **UAA** in Pivotal Container Service.

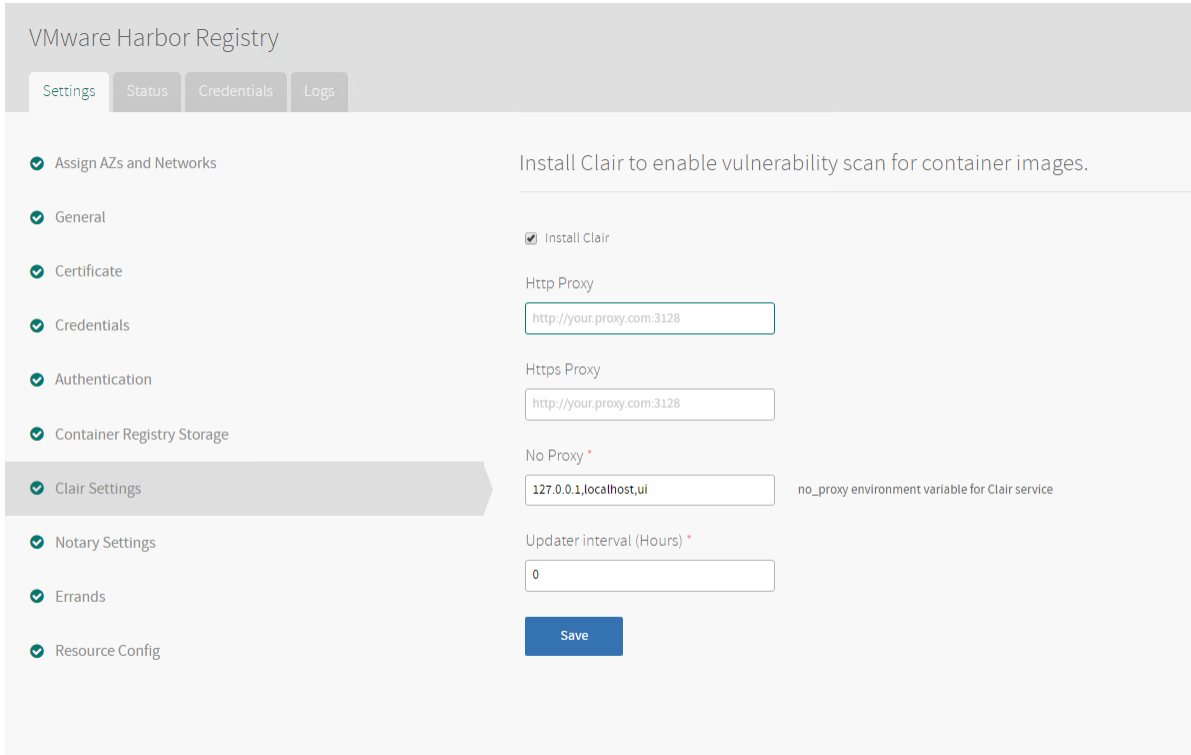


- 6 Select the desired storage for container images. For example, **Local File System** which is the default value.

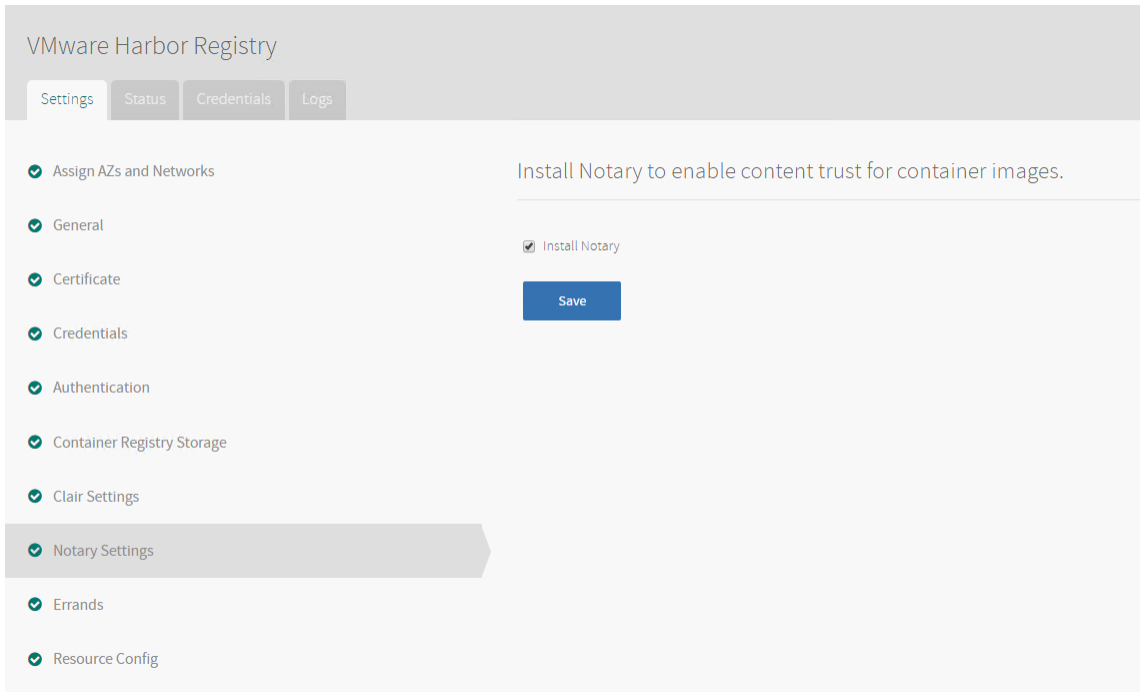


7 For the Clair settings, perform the following tasks:

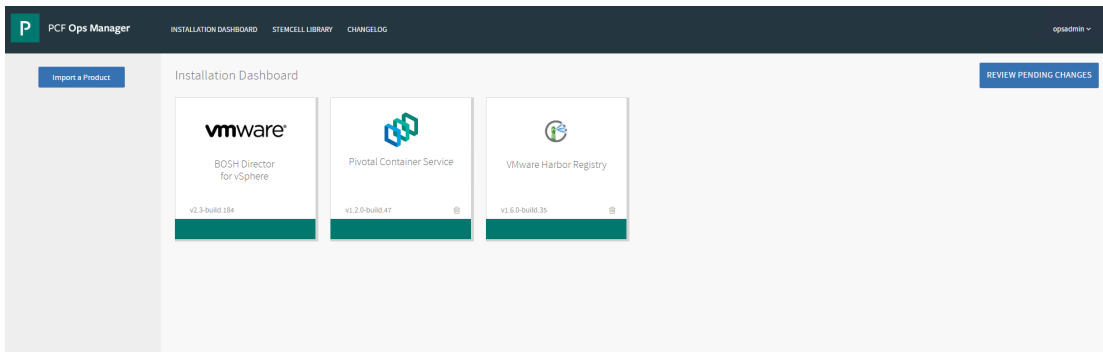
- a To enable the container image vulnerability scanning, ensure Install Clair is selected. If you deselect this option, Clair is not installed.
- b In the **No Proxy** text box, specify the endpoints that will not have proxies. This text box is required if Clair is installed. The default values, 127.0.0.1, localhost, and ui, are populated for you.
- c Let the value of **Updater interval (Hours)**



- For the Notary settings, by default, the **Install Notary** option is selected. Deselect this option if you do not want to install Notary.



- For **Errands** and **Resource Config**, leave the default values as is.
- To deploy the Harbor registry, click **Apply Changes** in the Ops manager Installation dashboard.



Note For information on integrating Harbor registry with PKS, refer <https://docs.pivotal.io/partners/vmware-harbor/integrating-pks.html>.

Install PKS CLI Client

For information on installing the PKS CLI client, refer <https://docs.pivotal.io/runtimes/pks/1-2/installing-pks-cli.html>.

Create a User for Create the Kubernetes clusters

To create the Kubernetes clusters, you have to create a user.

Procedure

- 1 Perform the following steps to retrieve the secret password:
 - a In a web browser, navigate to the fully qualified domain name of Ops Manager and click the **Pivotal Container Service** tile.
 - b Click **Credentials**.
 - c To view the secret, click **Link to Credential** next to `pkcs_uaa_management_admin_client`. The client username is `admin`.
- 2 On the command line, run the following commands:

```
uaac target https://pkcs-api.isbu.local:8443 --skip-ssl-validation
uaac token client get admin -s <SECRET_PWD>
uaac user add pkcs-admin --emails pkcs-admin@isbu.local -p VMware1!
uaac member add pkcs.clusters.admin pkcs-admin
```

Set Up the BOSH Environment

To set the BOSH environment for your PKS deployment, follow the steps below:

Procedure

- 1 Run the following on CLI:

```
om --target https://192.168.0.21 -u pkcs-admin -p 'VMware123!' -k curl -p
/api/v0/certificate_authorities -s | jq -r '.certificate_authorities |
select(map(.active == true))[0] | .cert_pem' > /root/opsmanager.pem
om --target https://192.168.0.21 -u pkcs-admin -p 'VMware123!' -k curl -p
/api/v0/deployed/director/credentials/bosh2_commandline_credentials -s |
jq -r '.credential'
```

- 2 Retrieve the list of the environment variables.

```
BOSH_CLIENT=ops_manager BOSH_CLIENT_SECRET=hpUV_A6SUQxjYazHtmzr4KgLb-
LOTUCb BOSH_CA_CERT=/var/tempest/workspaces/default/root_ca_certificate
BOSH_ENVIRONMENT=172.22.33.101 bosh
```

- 3 Replace the `CA_CERT` string with `/root/opsmanager.pem` and export the variables.

```
export BOSH_CLIENT=ops_manager
export BOSH_CLIENT_SECRET=hpUV_A6SUQxjYazHtmzr4KgLb-LOTUCb
export BOSH_CA_CERT=/root/opsmanager.pem
export BOSH_ENVIRONMENT=172.22.33.101
```

- 4 When redirected to the **Internal Authentication** page, perform the following steps:
 - a Enter the user name, password, and password confirmation to create an admin user.
 - b Enter a decryption passphrase and the decryption passphrase confirmation. This passphrase encrypts the Ops Manager datastore and is not recoverable.
 - c If you are using an HTTP proxy or HTTPS proxy, follow the instructions in the [Configuring Proxy Settings](#) for the BOSH CPI topic.
 - d Read the End-User License Agreement, and select the check box to accept the terms.
 - e Click **Setup Authentication**.
 - f Run the following on CLI:

```
om --target https://192.168.0.21 -u pks-admin -p 'VMware123!' -k curl -p  
/api/v0/certificate_authorities -s | jq -r '.certificate_authorities |  
select(map(.active == true))[0] | .cert_pem' > /root/opsmanager.pem  
om --target https://192.168.0.21 -u pks-admin -p 'VMware123!' -k curl -p  
/api/v0/deployed/director/credentials/bosh2_commandline_credentials -s |  
jq -r '.credential'
```