# Installing and Configuring vSphere with Tanzu

Update 2

VMware vSphere 8.0

VMware vCenter 8.0

VMware ESXi 8.0

vmware by Broadcom

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

https://docs.vmware.com/

**VMware by Broadcom**
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

# Contents

# *Installing and Configuring vSphere with Tanzu*

*Installing and Configuring vSphere with Tanzu* provides information about configuring and managing vSphere with Tanzu by using the vSphere Client.

*Installing and Configuring vSphere with Tanzu* provides instructions for enabling vSphere with Tanzu on existing vSphere clusters, creating and managing namespaces. This information also provides guidelines about establishing a session with the Kubernetes control plane through kubectl.

## Intended Audience

*Installing and Configuring vSphere with Tanzu* is intended for vSphere administrators who want to enable vSphere with Tanzu in vSphere, configure and provide namespaces to DevOps teams. vSphere administrators who want to use vSphere with Tanzu should have basic knowledge of containers and Kubernetes.

# Updated Information

This *Installing and Configuring vSphere with Tanzu* is updated with each release of the product or when necessary.

This table provides the update history of the *Installing and Configuring vSphere with Tanzu* documentation.

| Revision | Description |
|---|---|
| 18 MAR 2024 | Updated the Replace the VIP Certificate to Securely Connect to the Supervisor API Endpoint topic with a note about importing the entire certificate chain. |
| 29 FEB 2024 | ■ Added steps for creating a custom cloud during the initial configuration of the Controller. See Configure the Controller.<br>■ Added steps for selecting the cloud while deploying the Supervisor. See Deploy a Three-Zone Supervisor with the VDS Networking Stack and Deploy a One-Zone Supervisor with the VDS Networking Stack.<br>■ Added steps for configuring FQDN login with the Supervisor. See Deploy a Three-Zone Supervisor with the VDS Networking Stack, Deploy a One-Zone Supervisor with the VDS Networking Stack, and Change the Management Network Settings on a Supervisor.<br>■ Added a step to create NSX overlay segment. See Create a Tier-0 Uplink Segment and Overlay Segment. |
| 24 JAN 2024 | ■ Updated Register the NSX Advanced Load Balancer Controller with NSX Manager with a note about DNS and NTP settings.<br>■ Added content for steps to perform if the Supervisor deployment does not complete and NSX Advanced Load Balancer configuration is not applied when a private Certificate Authority (CA) signed certificate is provided. See NSX Advanced Load Balancer Configuration Is Not Applied. |
| 23 DEC 2023 | ■ Added content for changing the load balancer settings on a Supervisor configured with VDS networking. See #unique_11.<br>■ Updated the content for changing the workload networking settings of Supervisor configured with VDS networking. See Change the Workload Network Settings on a Supervisor Configured with VDS Networking. |
| 13 DEC 2023 | Added a reference for preparing ESXi hosts as transport nodes. See VDS Required for Host Transport Node Traffic. |
| 21 NOV 2023 | Updated the documentation to indicate that Multi NSX is not supported on the Supervisor Cluster. See Add a Compute Manager. |
| 29 SEP 2023 | ■ Updates on Configuring HTTP Proxy Settings in vSphere with Tanzu for use with Tanzu Mission Control.<br>■ Updated requirements for customizing the HAProxy load balancer. See Customize the HAProxy Load Balancer.<br>. |
| 21 SEP 2023 | Updated the networking section with information for installing and configuring NSX Advanced Load Balancer with NSX. See Install and Configure NSX and NSX Advanced Load Balancer. |

| Revision | Description |
|---|---|
| 30 JUN 2023 | Added the Supervisor control plane sizes in the Supervisor installation topics and Change the Control Plane Size of a Supervisor . |
| 23 JUN 2023 | Updated a link to create and edit content libraries. See Import the NSX Advanced Load Balancer OVA to a Local Content Library. |
| 15 JUN 2023 | Added a note that you can only use and HTTP proxy to register a .Supervisor with Tanzu Mission Control. See Configuring HTTP Proxy Settings in vSphere with Tanzu for use with Tanzu Mission Control. |
| 15 MAY 2023 | Added a note that you should not enable consumption domain in the storage policies used for a Supervisor or a namespace in a one-zone Supervisor. See Chapter 2 Create Storage Policies for vSphere with Tanzu. |
| 12 MAY 2023 | Added a note that if you have upgraded your vSphere with Tanzu environment from a vSphere version earlier than 8.0 and want to use vSphere Zones, you must create a new three-zone Supervisor. See Chapter 5 Deploy a Three-Zone Supervisor . |
| 26 APR 2023 | Moved Configuring and Managing vSphere Namespaces to *vSphere with Tanzu Services and Workloads*. |
| 18 APR 2023 | Updated the Install and Configure the NSX Advanced Load Balancer section to include support for NSX Advanced Load Balancer version 22.1.3. |

# vSphere with Tanzu Installation and Configuration Workflow

1

Review the workflows for turning vSphere clusters to a platform for running Kubernetes workloads on vSphere.

## Workflow for Deploying a Supervisor with VDS Networking and NSX Advanced Load Balancer

As a vSphere administrator, you can deploy a Supervisor with the networking stack based on VDS networking with the NSX Advanced Load Balancer.

Figure 1-1. Workflow for Deploying a Supervisor with NSX Advanced Load Balancer

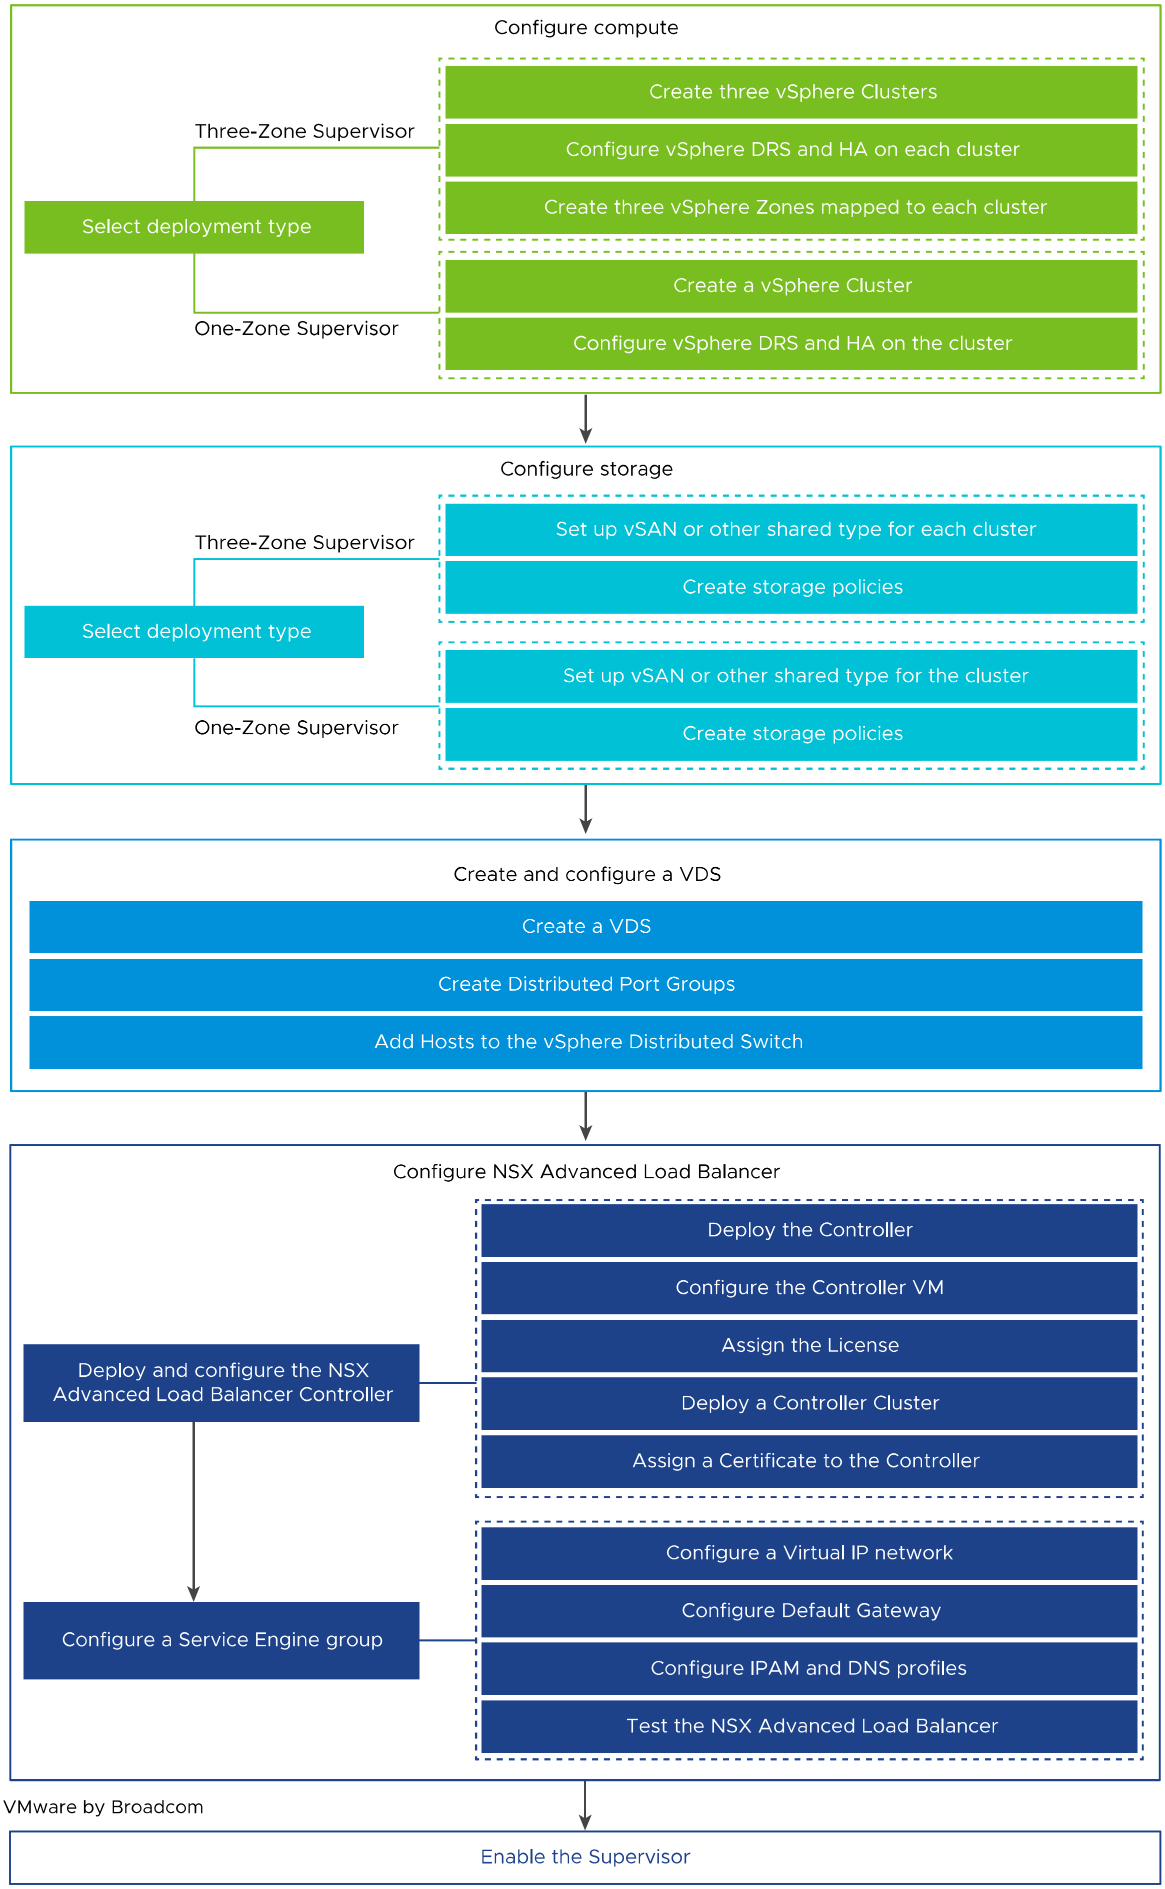
Configure compute
Three-Zone Supervisor
Select deployment type
One-Zone Supervisor
Create three vSphere Clusters
Configure vSphere DRS and HA on each cluster
Create three vSphere Zones mapped to each cluster
Create a vSphere Cluster
Configure vSphere DRS and HA on the cluster
Configure storage
Three-Zone Supervisor
Select deployment type
One-Zone Supervisor
Set up vSAN or other shared type for each cluster
Create storage policies
Set up vSAN or other shared type for the cluster
Create storage policies
Create and configure a VDS
Create a VDS
Create Distributed Port Groups
Add Hosts to the vSphere Distributed Switch
Configure NSX Advanced Load Balancer
Deploy and configure the NSX Advanced Load Balancer Controller
Deploy the Controller
Configure the Controller VM
Assign the License
Deploy a Controller Cluster
Assign a Certificate to the Controller
Configure a Service Engine group
Configure a Virtual IP network
Configure Default Gateway
Configure IPAM and DNS profiles
Test the NSX Advanced Load Balancer
Enable the Supervisor

# Supervisor with NSX Networking and NSX Advanced Load Balancer Controller Workflow

As a vSphere administrator, you can deploy a Supervisor with the NSX networking stack and the NSX Advanced Load Balancer Controller.

Figure 1-2. Workflow for deploying a Supervisor with NSX networking and NSX Advanced Load Balancer Controller

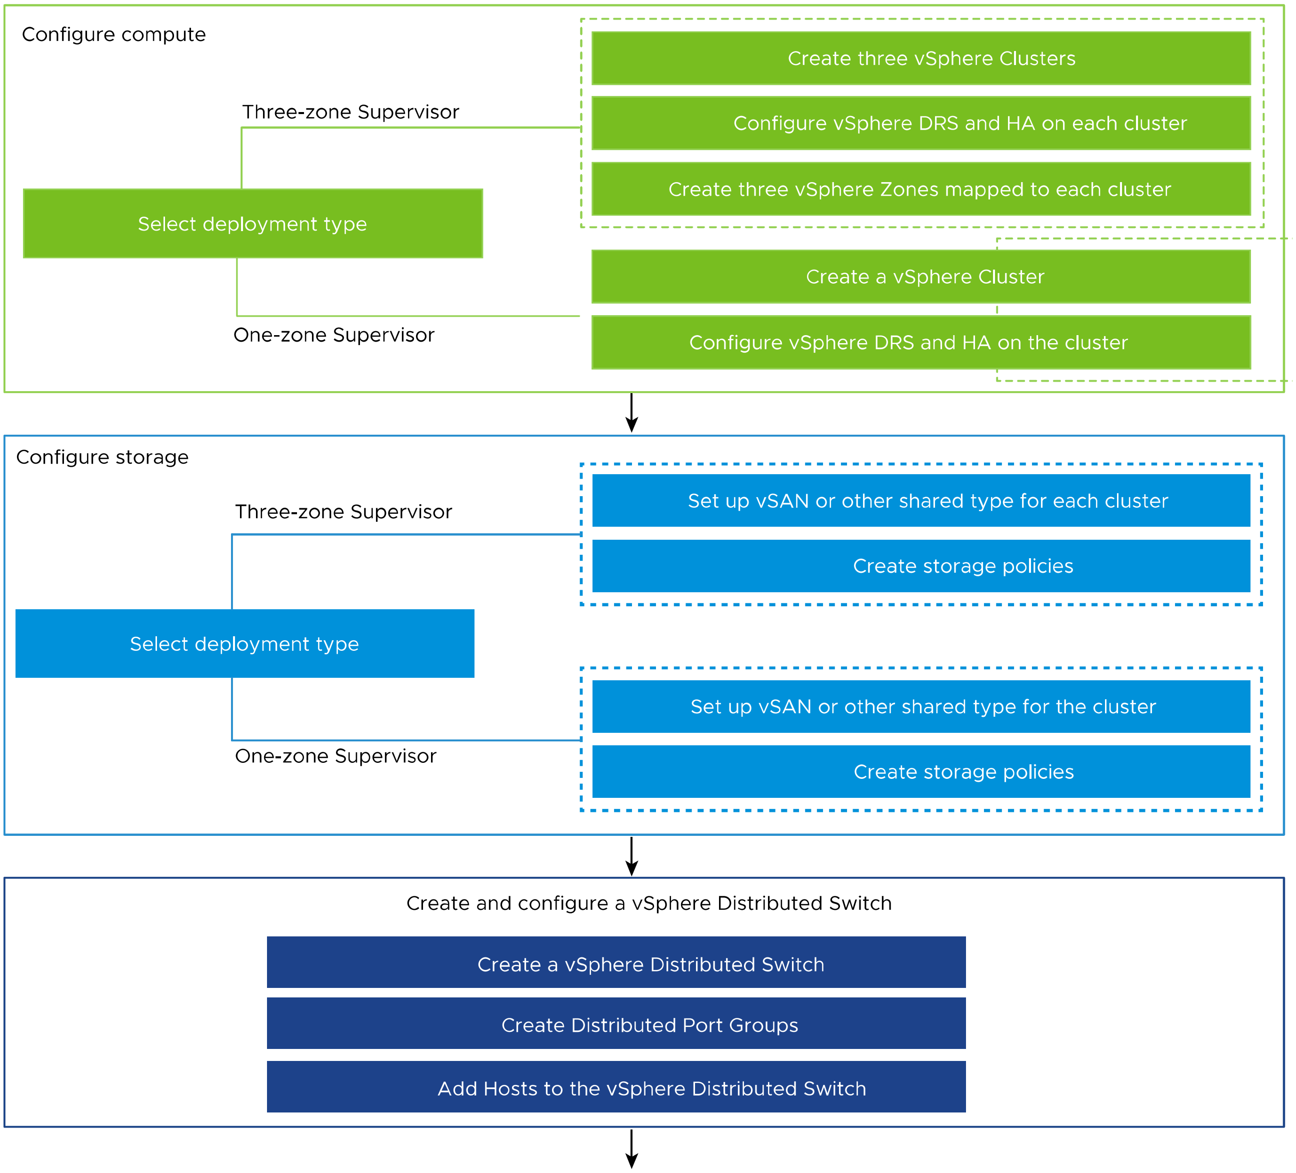
Configure compute
Three-zone Supervisor
Select deployment type
One-zone Supervisor
Create three vSphere Clusters
Configure vSphere DRS and HA on each cluster
Create three vSphere Zones mapped to each cluster
Create a vSphere Cluster
Configure vSphere DRS and HA on the cluster
Configure storage
Three-zone Supervisor
Select deployment type
One-zone Supervisor
Set up vSAN or other shared type for each cluster
Create storage policies
Set up vSAN or other shared type for the cluster
Create storage policies
Create and configure a vSphere Distributed Switch
Create a vSphere Distributed Switch
Create Distributed Port Groups
Add Hosts to the vSphere Distributed Switch

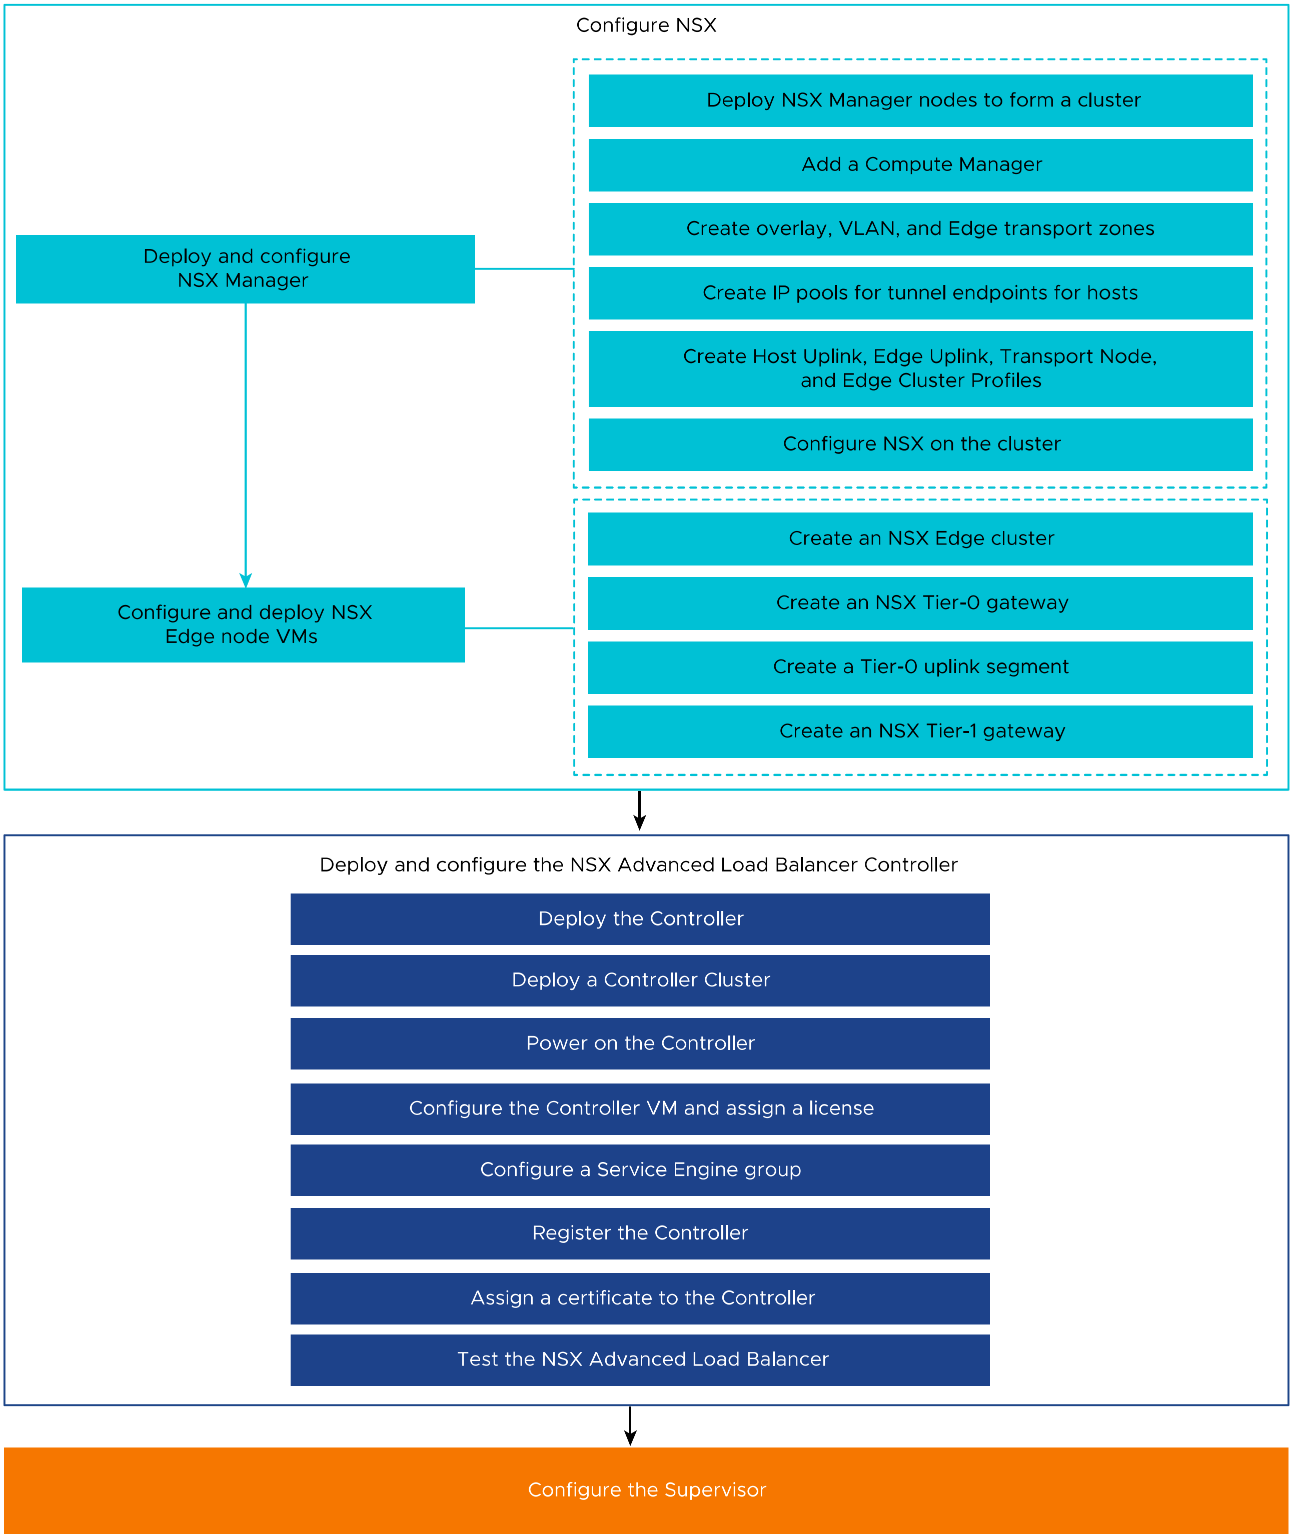
Configure NSX
Deploy and configure NSX Manager
Deploy NSX Manager nodes to form a cluster
Add a Compute Manager
Create overlay, VLAN, and Edge transport zones
Create IP pools for tunnel endpoints for hosts
Create Host Uplink, Edge Uplink, Transport Node, and Edge Cluster Profiles
Configure NSX on the cluster
Configure and deploy NSX Edge node VMs
Create an NSX Edge cluster
Create an NSX Tier-0 gateway
Create a Tier-0 uplink segment
Create an NSX Tier-1 gateway
Deploy and configure the NSX Advanced Load Balancer Controller
Deploy the Controller
Deploy a Controller Cluster
Power on the Controller
Configure the Controller VM and assign a license
Configure a Service Engine group
Register the Controller
Assign a certificate to the Controller
Test the NSX Advanced Load Balancer
Configure the Supervisor

## Workflow for Deploying a Supervisor with NSX Networking

As a vSphere administrator, you can deploy a Supervisor with the networking stack basedNSX.

Figure 1-3. Workflow for deploying a Supervisor with NSX as a networking stack

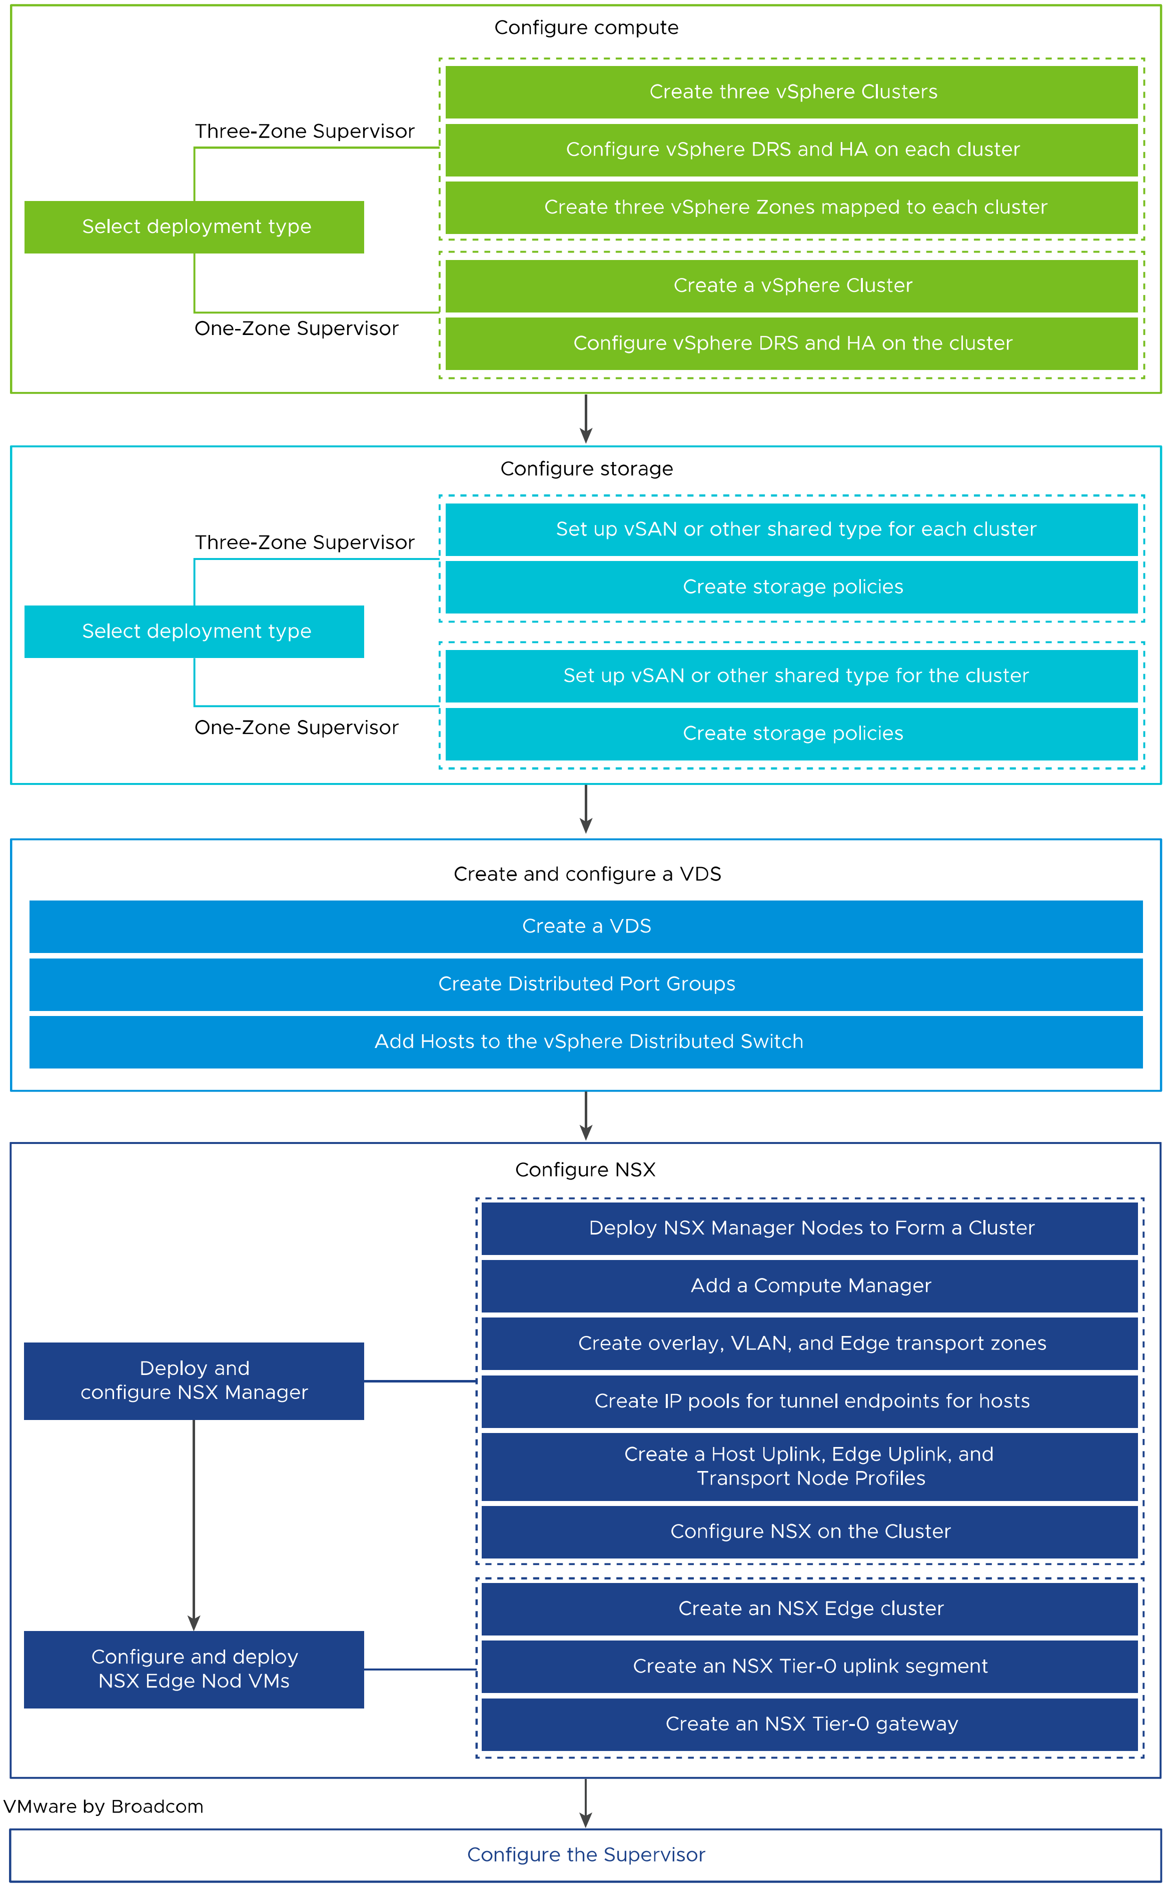
Configure compute
Three-Zone Supervisor
Select deployment type
One-Zone Supervisor
Create three vSphere Clusters
Configure vSphere DRS and HA on each cluster
Create three vSphere Zones mapped to each cluster
Create a vSphere Cluster
Configure vSphere DRS and HA on the cluster
Configure storage
Three-Zone Supervisor
Select deployment type
One-Zone Supervisor
Set up vSAN or other shared type for each cluster
Create storage policies
Set up vSAN or other shared type for the cluster
Create storage policies
Create and configure a VDS
Create a VDS
Create Distributed Port Groups
Add Hosts to the vSphere Distributed Switch
Configure NSX
Deploy and configure NSX Manager
Deploy NSX Manager Nodes to Form a Cluster
Add a Compute Manager
Create overlay, VLAN, and Edge transport zones
Create IP pools for tunnel endpoints for hosts
Create a Host Uplink, Edge Uplink, and Transport Node Profiles
Configure NSX on the Cluster
Configure and deploy NSX Edge Nod VMs
Create an NSX Edge cluster
Create an NSX Tier-0 uplink segment
Create an NSX Tier-0 gateway
Configure the Supervisor

# Workflow for Deploying a Supervisor with VDS Networking and HAProxy Load Balancer

As a vSphere administrator, you can deploy a Supervisor with the networking stack based on VDS and the HAProxy load balancer.

Figure 1-4. Workflow for deploying Supervisor with VDS networking and HAProxy

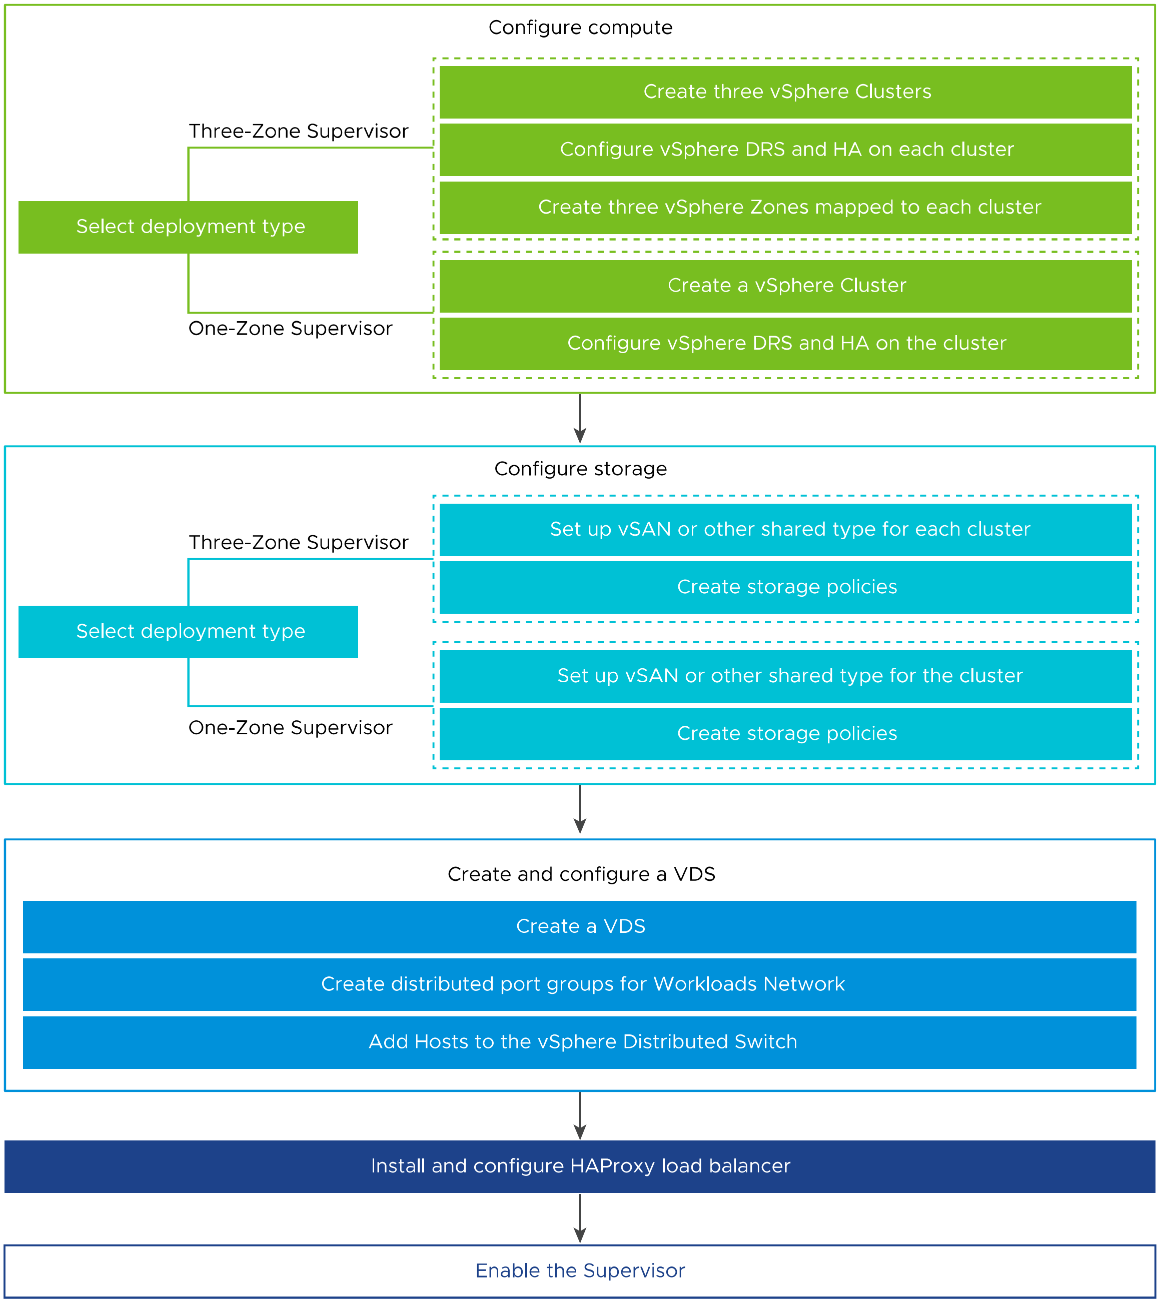
Configure compute
Three-Zone Supervisor
Select deployment type
One-Zone Supervisor
Create three vSphere Clusters
Configure vSphere DRS and HA on each cluster
Create three vSphere Zones mapped to each cluster
Create a vSphere Cluster
Configure vSphere DRS and HA on the cluster
Configure storage
Three-Zone Supervisor
Select deployment type
One-Zone Supervisor
Set up vSAN or other shared type for each cluster
Create storage policies
Set up vSAN or other shared type for the cluster
Create storage policies
Create and configure a VDS
Create a VDS
Create distributed port groups for Workloads Network
Add Hosts to the vSphere Distributed Switch
Install and configure HAProxy load balancer
Enable the Supervisor

Read the following topics next:

- Prerequisites for Configuring vSphere with Tanzu on vSphere Clusters

# Prerequisites for Configuring vSphere with Tanzu on vSphere Clusters

Verify prerequisites for enabling vSphere with Tanzu in your vSphere environment. To run container-based workloads natively on vSphere, as a vSphere administrator you enable vSphere clusters as Supervisors. A Supervisor has a Kubernetes layer that allows you to run Kubernetes workloads on vSphere by deploying vSphere Pods, provision Tanzu Kubernetes clusters, and VMs.

## Create and Configure vSphere Clusters

A Supervisor can run on either one or three vSphere clusters associated with vSphere Zones. Each vSphere Zone maps to one vSphere cluster, and you can deploy a Supervisor on either one or three zones. A three-zone Supervisor provides greater amount of resources for running your Kubernetes workloads and has high-availability at a vSphere cluster level that protects your workloads against cluster failure. A one-zone Supervisor has host-level high availably provided by vSphere HA and utilizes the resources of only one cluster for running your Kubernetes workloads.

**Note** Once you deploy a Supervisor on one vSphere Zone, you cannot expand the Supervisor to a three-zone deployment.

Each vSphere cluster where you intend to deploy a Supervisor must meet the following requirements:

- Create and configure a vSphere cluster with at least two ESXi hosts. If you are using vSAN, the cluster must have at least three hosts or four for optimal performance. See Creating and Configuring Clusters.
- Configure the cluster with shared storage such as vSAN. Shared storage is required for vSphere HA, DRS, and for storing persistent container volumes. See Creating a vSAN Cluster.
- Enable the cluster with vSphere HA. See Creating and Using vSphere HA Clusters.
- Enable the cluster with vSphere DRS in fully-automated mode. See Creating a DRS Cluster.
- Verify that your user account has the **Modify cluster-wide configuration** on the vSphere cluster so that you can deploy the Supervisor.
- To deploy a three-zone Supervisor, create three vSphere Zones, see Chapter 3 Create vSphere Zones for a Multi-Zone Supervisor Deployment.
- If you want to use vSphere Lifecycle Manager images with the Supervisor, switch the vSphere clusters where you want to activate **Workload Management** to use vSphere Lifecycle Manager images before activating **Workload Management**. You can manage the lifecycle of a Supervisor with either vSphere Lifecycle Manager baselines or vSphere Lifecycle

Manager images. However, you cannot convert a Supervisor that uses vSphere Lifecycle Manager baselines to a Supervisor that uses vSphere Lifecycle Manager images. Therefore, switching the vSphere clusters to using vSphere Lifecycle Manager images before you activate **Workload Management** is required.

## Create Storage Policies

Before deploying a Supervisor, you must create storage policies that determine the datastore placement of the Supervisor control plane VMs. If the Supervisor supports vSphere Pods, you also need storage policies for containers and images. You can create storage policies associated with different levels of storage services.

See Chapter 2 Create Storage Policies for vSphere with Tanzu.

## Choose and Configure the Networking Stack

To deploy a Supervisor, you must configure the networking stack to use with it. You have two options: NSX or vSphere Distributed Switch (vDS) networking with a load balancer. You can configure the NSX Advanced Load Balancer or the HAProxy load balancer.

To use NSX networking for the Supervisor:

- Review the system requirements and topologies for NSX networking. See Requirements for Enabling a Three-Zone Supervisor with NSX and Requirements for Setting Up a Single-Cluster Supervisor with NSX in *vSphere with Tanzu Concepts and Planning*.
- Install and configure NSX for vSphere with Tanzu. See Install and Configure NSX for vSphere with Tanzu.

To use vDS networking with the NSX Advanced Load Balancer for the Supervisor:

- Review the NSX Advanced Load Balancer requirements. See Requirements for a Three-Zone Supervisor with NSX Advanced Load Balancer and Requirements for Enabling a Single Cluster Supervisor with NSX Advanced Load Balancer in *vSphere with Tanzu Concepts and Planning*.
- Create a vSphere Distributed Switch (vDS) and add all ESXi hosts from the cluster to the vDS and create port groups for Workload Networks. See Create a vSphere Distributed Switch for a Supervisor for Use with NSX Advanced Load Balancer.
- Deploy and configure the NSX Advanced Load Balancer. See Deploy the NSX Advanced Load Balancer Controller.

**Note** vSphere with Tanzu supports the NSX Advanced Load Balancer with vSphere 7 U2 and later.

To use vDS networking with HAProxy load balancing for the Supervisor:

- Review the system requirements and network topologies for vSphere networking with HAProxy load balancer. See Requirements for Enabling a Three-Zone Supervisor with HA Proxy Load Balancer and Requirements for Enabling a Single-Cluster Supervisor with VDS Networking and HAProxy Load Balancer *vSphere with Tanzu Concepts and Planning*.

- Create a vSphere Distributed Switch (VDS) and add all ESXi hosts from the cluster to the vDS and create port groups for Workload Networks. See Create a vSphere Distributed Switch for a Supervisor for Use with HAProxy Load Balancer.
- Install and configure the HAProxy load balancer instance that is routable to the vDS that is connected to the hosts from the vSphere clusters where you deploy the Supervisor. The HAProxy load balancer supports the network connectivity to workloads from client networks and to load balance traffic between Tanzu Kubernetes clusters. See Install and Configure the HAProxy Load Balancer.

**Note** vSphere with Tanzu supports the HAProxy load balancer with vSphere 7 U1 and later.

# Create Storage Policies for vSphere with Tanzu

2

Before you enable vSphere with Tanzu, create storage policies to be used in the Supervisor and namespaces. The policies represent datastores and manage storage placement of such components and objects as Supervisor control plane VMs, vSphere Podephemeral disks, and container images. You might also need policies for storage placement of persistent volumes and VM content libraries. If you use Tanzu Kubernetes clusters, the storage policies also dictate how the Tanzu Kubernetes cluster nodes are deployed.

Depending on your vSphere storage environment and the needs of DevOps, you can create several storage policies for different classes of storage. For example, if your vSphere storage environment has three classes of datastores, Bronze, Silver, and Gold, you can create storage policies for all datastore types.

When you enable a Supervisor and set up namespaces, you can assign different storage policies to be used by various objects, components, and workloads.

**Note** Storage policies that you create for a Supervisor or for a namespace in a one-zone Supervisor do not need to be topology aware. Do not enable consumption domain for those policies.

Storage policies that you create for a namespace in a three-zone Supervisor must be topology aware and have the consumption domain enabled in Step 4b. The three-zone namespace prevents you from assigning storage policies that are not topology aware.

The following example creates the storage policy for the datastore tagged as Gold.

### Prerequisites

- Be familiar with information about storage policies in vSphere with Tanzu, see About Storage Policies in *vSphere with Tanzu Concepts and Planning*.
- If you use vSAN Data Persistence platform for persistent storage and need to create custom storage policies for vSAN Direct or vSAN SNA datastores, see Creating Custom Storage Policies for vSAN Data Persistence Platform in *vSphere with Tanzu Services and Workloads*.
- If you need to create topology aware storage policies to use for persistent storage in a three-zone Supervisor, be familiar with the guidelines in Using Persistent Storage on a Three-Zone Supervisor in *vSphere with Tanzu Services and Workloads*.

- Make sure that the datastore you reference in the storage policy is shared between all ESXi hosts in the cluster. Any shared datastores in your environment are supported, including VMFS, NFS, vSAN, or vVols.
- Required privileges: **VM storage policies. Update** and **VM storage policies. View**.

Procedure

1 Add tags to the datastore.

 a Right-click the datastore you want to tag and select **Tags and Custom Attributes > Assign Tag**.

 b Click **Add Tag** and specify the tag's properties.

| Property | Description |
|---|---|
| **Name** | Specify the name of the datastore tag, for example, `Gold`. |
| **Description** | Add the description of the tag. For example, `Datastore for Kubernetes objects`. |
| **Category** | Select an existing category or create a new category. For example, `Storage for Kubernetes`. |

2 In the vSphere Client, open the **Create VM Storage Policy** wizard.

 a Click **Menu > Policies and Profiles**.

 b Under **Policies and Profiles**, click **VM Storage Policies**.

 c Click **Create VM Storage Policy**.

3 Enter the policy name and description.

| Option | Action |
|---|---|
| **vCenter Server** | Select the vCenter Server instance. |
| **Name** | Enter the name of the storage policy, for example, `goldsp`.<br>Note When vSphere with Tanzu converts storage policies that you assign to namespaces into Kubernetes storage classes, it changes all upper case letters into lower case and replaces spaces with dashes (-). To avoid confusion, use lower case and no spaces in the VM storage policy names. |
| **Description** | Enter the description of the storage policy. |

4 On the **Policy structure** page, select the following options and click **Next**.

 a Under **Datastore specific rules**, enable tag-based placement rules.

 b To create a topology aware policy, under **Storage topology**, select **Enable consumption domain**.

 This step is necessary only if you are creating topology aware policies to be used for persistent storage on a namespace in a three-zone Supervisor.

5 On the **Tag based placement** page, create the tag rules.

Select the options using the following example.

| Option | Description |
|---|---|
| **Tag category** | From the drop-down menu, select the tag's category, such as **Storage for Kubernetes**. |
| **Usage option** | Select **Use storage tagged with**. |
| **Tags** | Click **Browse Tags**, and select the datastore tag, for example, **Gold**. |

6 If you enabled **Storage topology**, on the **Consumption domain** page, specify the storage topology type.

| Option | Description |
|---|---|
| **Zonal** | Datastore is shared across all hosts in a single zone. |

7 On the **Storage compatibility** page, review the list of datastores that match this policy.

In this example, only the datastore that is tagged as Gold is displayed.

8 On the **Review and finish** page, review the storage policy settings and click **Finish**.

### Results

The new storage policy for the datastore tagged as Gold appears on the list of existing storage policies.

### What to do next

After creating storage policies, a vSphere administrator can perform the following tasks:

- Assign the storage policies to the Supervisor. The storage policies configured on the Supervisor ensure that the control plane VMs, pod ephemeral disks, and container images are placed on the datastores that the policies represent.
- Assign the storage policies to the vSphere Namespace. Storage policies visible to the namespace determine which datastores the namespace can access and use for persistent volumes. The storage policies appear as matching Kubernetes storage classes in the namespace. They are also propagated to the Tanzu Kubernetes cluster on this namespace. DevOps engineers can use the storage classes in their persistent volume claim specifications. See Create and Configure a vSphere Namespace.

# Create vSphere Zones for a Multi-Zone Supervisor Deployment

3

Check out how to create vSphere Zones that you can use to provide cluster-level high availability to your Kubernetes workloads running on a Supervisor. To provide cluster-level high availability to your Kubernetes workloads, you deploy the Supervisor on three vSphere Zones. Each vSphere Zone is mapped to one vSphere cluster that has a minimum of 2 hosts.

### Prerequisites

- Create three vSphere clusters with at least 3 hosts in each zone. For storage with vSAN, the cluster must have 4 hosts.
- Configure storage with vSAN or other shared storage solution for each cluster.
- Enable vSphere HA and vSphere DRS on Fully Automate or Partially Automate mode.
- Configure networking with NSX or vSphere Distributed Switch (vDS) networking for the clusters.

### Procedure

1. In the vSphere Client, navigate to vCenter Server.
2. Select **Configure** and select **vSphere Zones**.
3. Click **Add New vSphere Zone**.
4. Name the zone, for example `zone1` and add an optional description.
5. Select a vSphere cluster to add to the zone and click **Finish**.
6. Repeat the steps to create three vSphere Zones.

### What to do next

- - Configure a networking stack to use with the Supervisor. See Chapter 4 Networking for vSphere with Tanzu
  - Activate the Supervisor on the three vSphere Zones that you created. See Chapter 5 Deploy a Three-Zone Supervisor .

  If you need to do any changes to a vSphere Zone, you can do them before you deploy the Supervisor on it.

# Managing vSphere Zones

If you need to make changes to a vSphere Zone, you have do that before you deploy a Supervisor on the zone. You can change the cluster that is associated with it, or delete the zone. Deleting a vSphere Zone removes its associated cluster and then deletes the zone from vCenter Server.

## Removing a Cluster from a vSphere Zone

To remove a cluster from a vSphere Zone, click the three dots (…) on the zone card and select **Remove Cluster**. The cluster is removed from the zone and you can add a different one.

**Note** You cannot remove a cluster from a vSphere Zone when there is a Supervisor already enabled on that zone.

## Deleting a vSphere Zone

To delete a vSphere Zone, click the three dots (…) on the zone card and select **Delete Zone**.

**Note** You cannot delete a vSphere Zone when there is a Supervisor already enabled on that zone.

# Networking for vSphere with Tanzu

4

A Supervisor can either use the vSphere networking stack or VMware NSX® to provide connectivity to Kubernetes control plane VMs, services, and workloads. The networking used for Tanzu Kubernetes clusters provisioned by the Tanzu Kubernetes Grid is a combination of the fabric that underlies the vSphere with Tanzu infrastructure and open-source software that provides networking for cluster pods, services, and ingress.

Read the following topics next:

## Supervisor Networking

In a vSphere with Tanzu environment, a Supervisor can either use a vSphere networking stack or NSX to provide connectivity to Supervisor control plane VMs, services, and workloads.

When a Supervisor is configured with the vSphere networking stack, all hosts from the Supervisor are connected to a vDS that provides connectivity to workloads and Supervisor control plane VMs. A Supervisor that uses the vSphere networking stack requires a load balancer on the vCenter Server management network to provide connectivity to DevOps users and external services.

A Supervisor that is configured with NSX, uses the software-based networks of the solution and an NSX Edge load balancer or the NSX Advanced Load Balancer to provide connectivity to external services and DevOps users. You can configure the NSX Advanced Load Balancer on NSX if your environment meets the following conditions:

- NSX version is 4.1.1 or later.
- The NSX Advanced Load Balancer version is 22.1.4 or later with the Enterprise license.
- The NSX Advanced Load Balancer Controller you plan to configure is registered on NSX.
- An NSX load balancer is not already configured on the Supervisor.

## Supervisor Networking with VDS

In a Supervisor that is backed by VDS as the networking stack, all hosts from the vSphere clusters backing the Supervisor must be connected to the same VDS. The Supervisor uses distributed port groups as workload networks for Kubernetes workloads and control plane traffic. You assign workload networks to namespaces in the Supervisor.

Depending on the topology that you implement for the Supervisor, you can use one or more distributed port groups as workload networks. The network that provides connectivity to the Supervisor control plane VMs is called Primary workload network. You can assign this network to all the namespaces on the Supervisor, or you can use different networks for each namespace. The Tanzu Kubernetes Grid clusters connect to the Workload Network that is assigned to the namespace where the clusters reside.

A Supervisor that is backed by a VDS uses a load balancer for providing connectivity to DevOps users and external services. You can use the NSX Advanced Load Balancer or the HAProxy load balancer.

For more information, see Install and Configuring NSX Advanced Load Balancer and Install and Configure HAProxy Load Balancer.

In a single-cluster Supervisor setup, the Supervisor is backed by only one vSphere cluster. All hosts from the cluster must be connected to a VDS.

Figure 4-1. Single-cluster Supervisor networking with VDS

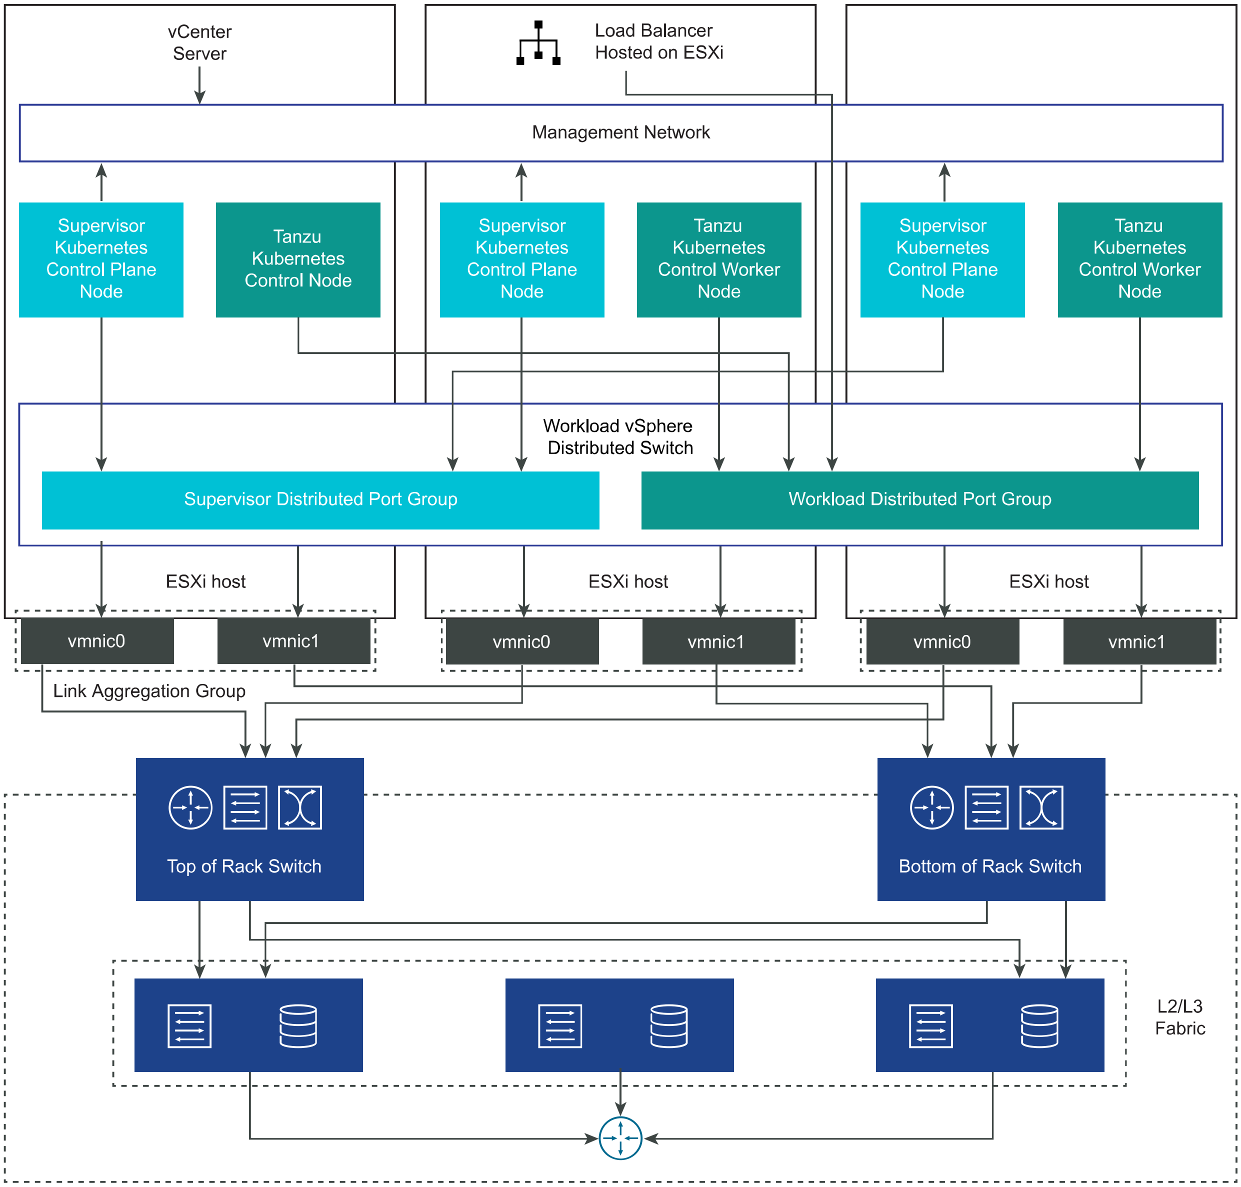
vCenter Server
Load Balancer Hosted on ESXi
Management Network
Supervisor Kubernetes Control Plane Node
Tanzu Kubernetes Control Node
Supervisor Kubernetes Control Plane Node
Tanzu Kubernetes Control Worker Node
Supervisor Kubernetes Control Plane Node
Tanzu Kubernetes Control Worker Node
Workload vSphere Distributed Switch
Supervisor Distributed Port Group
Workload Distributed Port Group
ESXi host
ESXi host
ESXi host
vmnic0
vmnic1
vmnic0
vmnic1
vmnic0
vmnic1
Link Aggregation Group
Top of Rack Switch
Bottom of Rack Switch
L2/L3 Fabric

In a three-zone Supervisor, you deploy the Supervisor on three vSphere zones, each mapped to a vSphere cluster. All hosts from these vSphere clusters must be connected to the same VDS. All physical servers must be connected to a L2 device. Workload networks that you configure to namespace span across all three vSphere zones.

Figure 4-2. Three-zone Supervisor networking with VDS

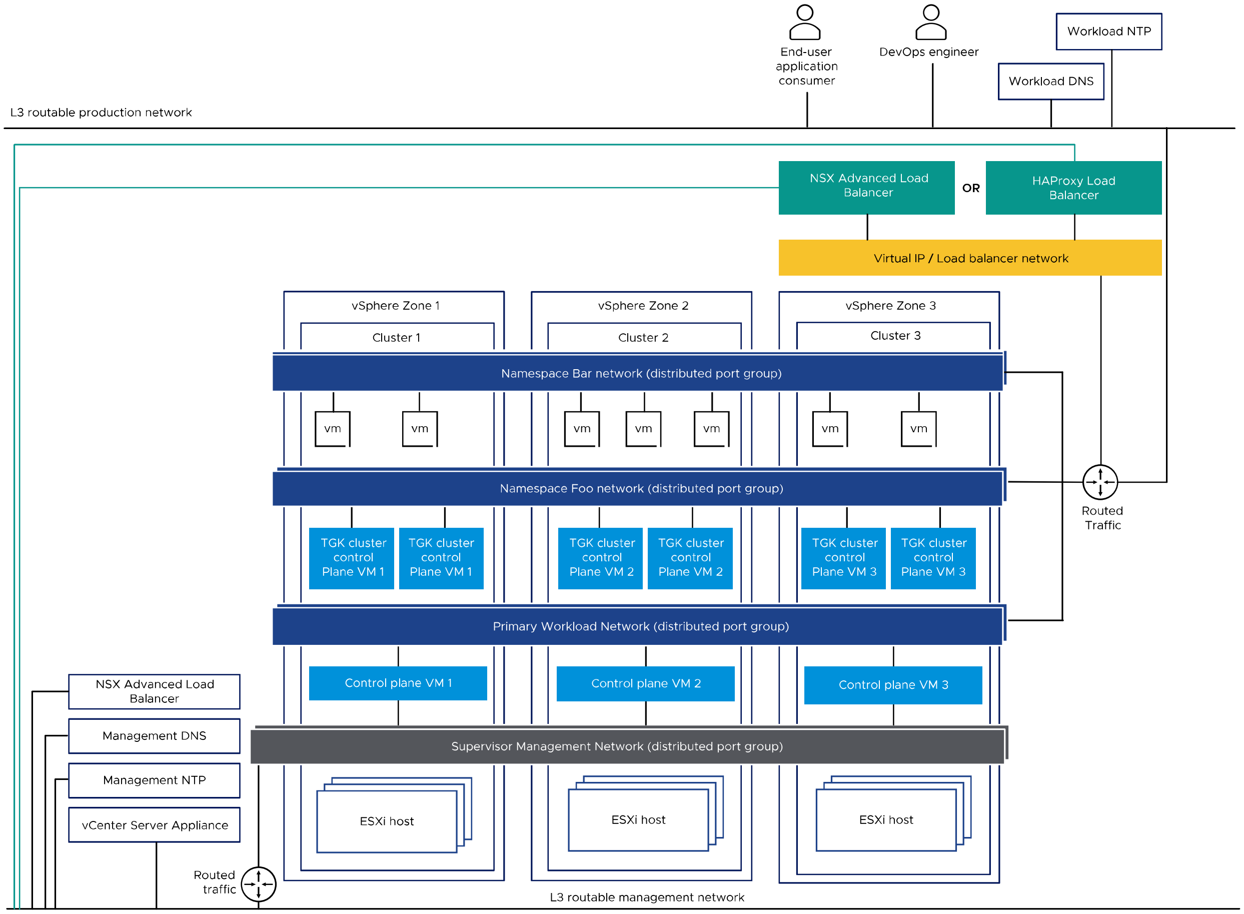
End-user application consumer
DevOps engineer
Workload NTP
Workload DNS
L3 routable production network
NSX Advanced Load Balancer
OR
HAProxy Load Balancer
Virtual IP / Load balancer network
vSphere Zone 1
vSphere Zone 2
vSphere Zone 3
Cluster 1
Cluster 2
Cluster 3
Namespace Bar network (distributed port group)
vm
vm
vm
vm
vm
vm
vm
Namespace Foo network (distributed port group)
Routed Traffic
TGK cluster control Plane VM 1
TGK cluster control Plane VM 1
TGK cluster control Plane VM 2
TGK cluster control Plane VM 2
TGK cluster control Plane VM 3
TGK cluster control Plane VM 3
Primary Workload Network (distributed port group)
NSX Advanced Load Balancer
Control plane VM 1
Control plane VM 2
Control plane VM 3
Management DNS
Supervisor Management Network (distributed port group)
Management NTP
vCenter Server Appliance
ESXi host
ESXi host
ESXi host
Routed traffic
L3 routable management network

## Supervisor Networking with NSX

NSX provides network connectivity to the objects inside the Supervisor and external networks. Connectivity to the ESXi hosts comprising the cluster is handled by the standard vSphere networks.

You can also configure the Supervisor networking manually by using an existing NSX deployment or by deploying a new instance of NSX.

For more information, see Install and Configure NSX for vSphere with Tanzu.

Figure 4-3. Supervisor networking with NSX

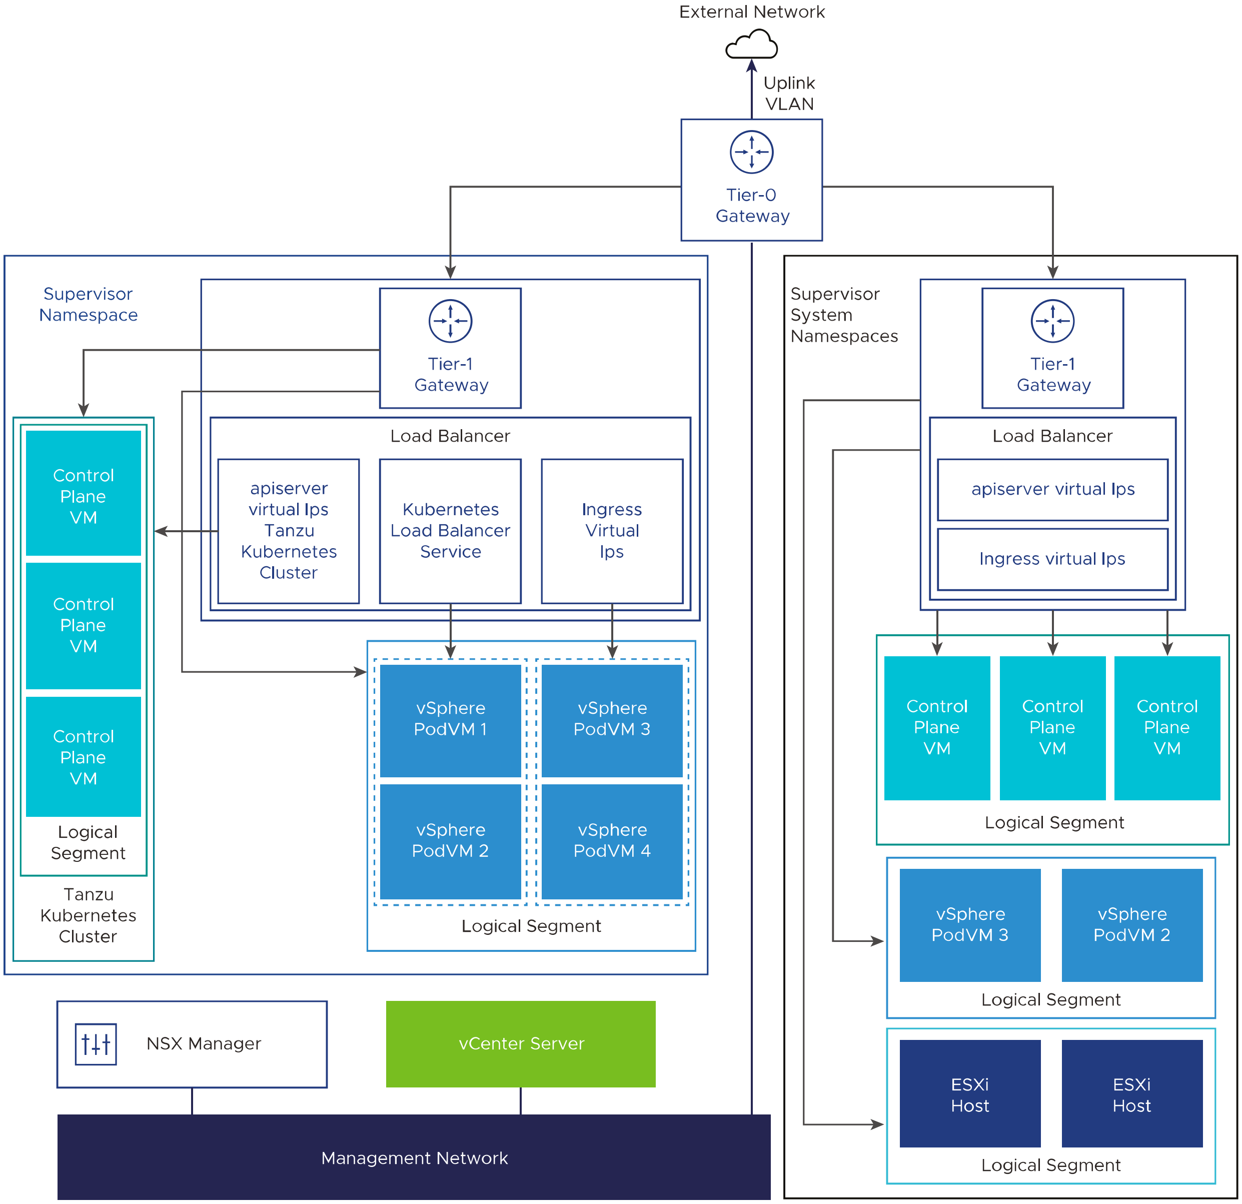

- NSX Container Plugin (NCP) provides integration between NSX and Kubernetes. The main component of NCP runs in a container and communicates with NSX Manager and with the Kubernetes control plane. NCP monitors changes to containers and other resources and manages networking resources such as logical ports, segments, routers, and security groups for the containers by calling the NSX API.

  The NCP creates one shared tier-1 gateway for system namespaces and a tier-1 gateway and load balancer for each namespace, by default. The tier-1 gateway is connected to the tier-0 gateway and a default segment.

System namespaces are namespaces that are used by the core components that are integral to functioning of theSupervisor and Tanzu Kubernetes Grid clusters. The shared network resources that include the tier-1 gateway, load balancer, and SNAT IP are grouped in a system namespace.

- NSX Edge provides connectivity from external networks to Supervisor objects. The NSX Edge cluster has a load balancer that provides a redundancy to the Kubernetes API servers residing on the Supervisor control plane VMs and any application that must be published and be accessible from outside the Supervisor.
- A tier-0 gateway is associated with the NSX Edge cluster to provide routing to the external network. The uplink interface uses either the dynamic routing protocol, BGP, or static routing.
- Each vSphere Namespace has a separate network and set of networking resources shared by applications inside the namespace such as, tier-1 gateway, load balancer service, and SNAT IP address.
- Workloads running in vSphere Pods, regular VMs, or Tanzu Kubernetes Grid clusters, that are in the same namespace, share a same SNAT IP for North-South connectivity.
- Workloads running in vSphere Pods or Tanzu Kubernetes Grid clusters will have the same isolation rule that is implemented by the default firewall.
- A separate SNAT IP is not required for each Kubernetes namespace. East west connectivity between namespaces will be no SNAT.
- The segments for each namespace reside on the VDS functioning in Standard mode that is associated with the NSX Edge cluster. The segment provides an overlay network to the Supervisor.
- Supervisors have separate segments within the shared tier-1 gateway. For each Tanzu Kubernetes Grid cluster, segments are defined within the tier-1 gateway of the namespace.
- The Spherelet processes on each ESXi hosts communicate with vCenter Server through an interface on the Management Network.

In a three-zone Supervisor configured with NSX as the networking stack, all hosts from all three vSphere clusters mapped to the zones must use be connected to the same VDS and participate in the same NSX Overlay Transport Zone. All hosts must be connected to the same L2 physical device.

Figure 4-4. Three-Zone Supervisor networking with NSX

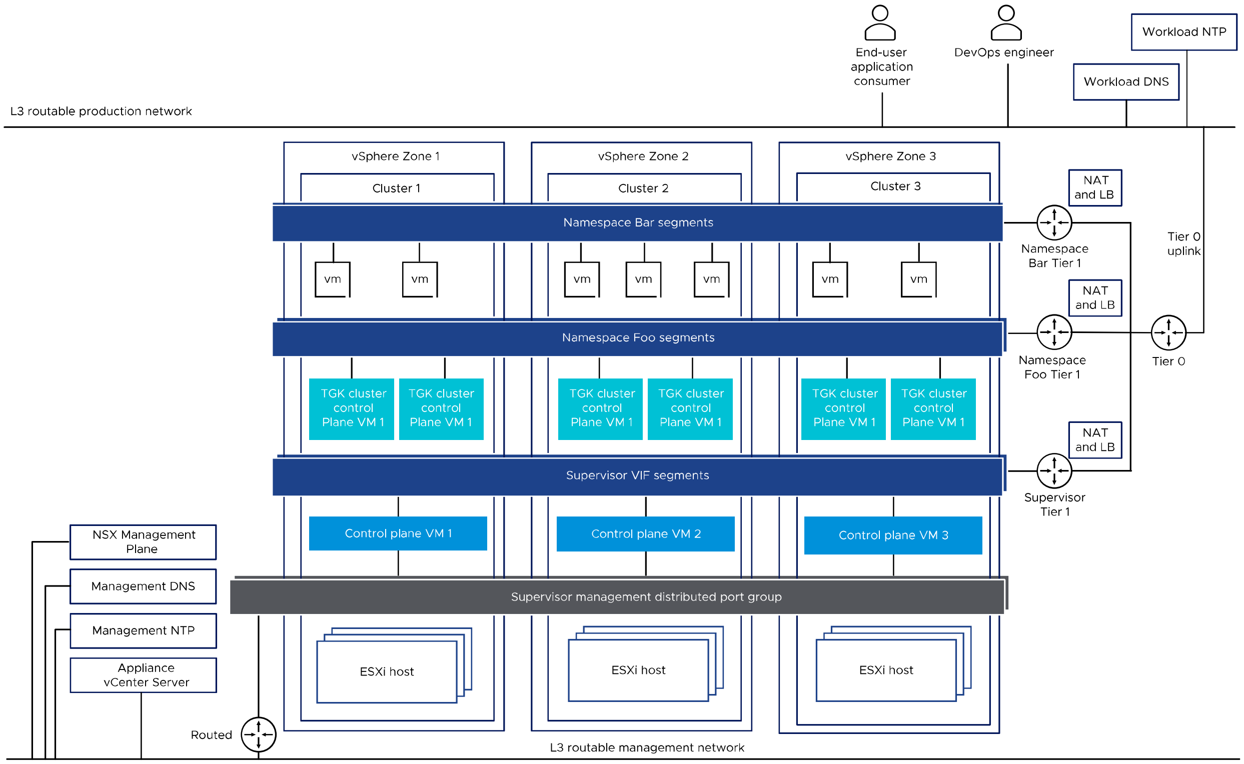

## Supervisor networking with NSX and NSX Advanced Load Balancer

NSX provides network connectivity to the objects inside the Supervisor and external networks. A Supervisor that is configured with NSX can use the NSX Edge or the NSX Advanced Load Balancer.

The components of the NSX Advanced Load Balancer include the NSX Advanced Load Balancer Controller cluster, Service Engines (data plane) VMs and the Avi Kubernetes Operator (AKO).

The NSX Advanced Load Balancer Controller interacts with the vCenter Server to automate the load balancing for the Tanzu Kubernetes Grid clusters. It is responsible for provisioning service engines, coordinating resources across service engines, and aggregating service engine metrics and logging. The Controller provides a Web interface, command-line interface, and API for user operation and programmatic integration.After you deploy and configure the Controller VM, you can deploy a Controller Cluster to set up the control plane cluster for HA.

The Service Engine, is the data plane virtual machine. A Service Engine runs one or more virtual services. A Service Engine is managed by the NSX Advanced Load Balancer Controller. The Controller provisions Service Engines to host virtual services.

The Service Engine has two types of network interfaces:

- The first network interface, `vnic0` of the VM, connects to the Management Network where it can connect to the NSX Advanced Load Balancer Controller.
- The remaining interfaces, `vnic1 - 8`, connect to the Data Network where virtual services run.

The Service Engine interfaces automatically connect to correct vDS port groups. Each Service Engine can support up to 1000 virtual services.

A virtual service provides layer 4 and layer 7 load balancing services for Tanzu Kubernetes Grid cluster workloads. A virtual service is configured with one virtual IP and multiple ports. When a virtual service is deployed, the Controller automatically selects an ESX server, spins up a Service Engine, and connects it to the correct networks (port groups).

The first Service Engine is created only after the first virtual service is configured. Any subsequent virtual services that are configured use the existing Service Engine.

Each virtual server exposes a layer 4 load balancer with a distinct IP address of type load balancer for a Tanzu Kubernetes Grid cluster. The IP address assigned to each virtual server is selected from the IP address block given to the Controller when you configure it.

The Avi Kubernetes operator (AKO) watches Kubernetes resources and communicates with the NSX Advanced Load Balancer Controller to request the corresponding load balancing resources. The Avi Kubernetes Operator is installed on the Supervisors as part of the enablement process.

For more information, see Install and Configure NSX and NSX Advanced Load Balancer.

Figure 4-5. Supervisor networking with NSX and NSX Advanced Load Balancer Controller

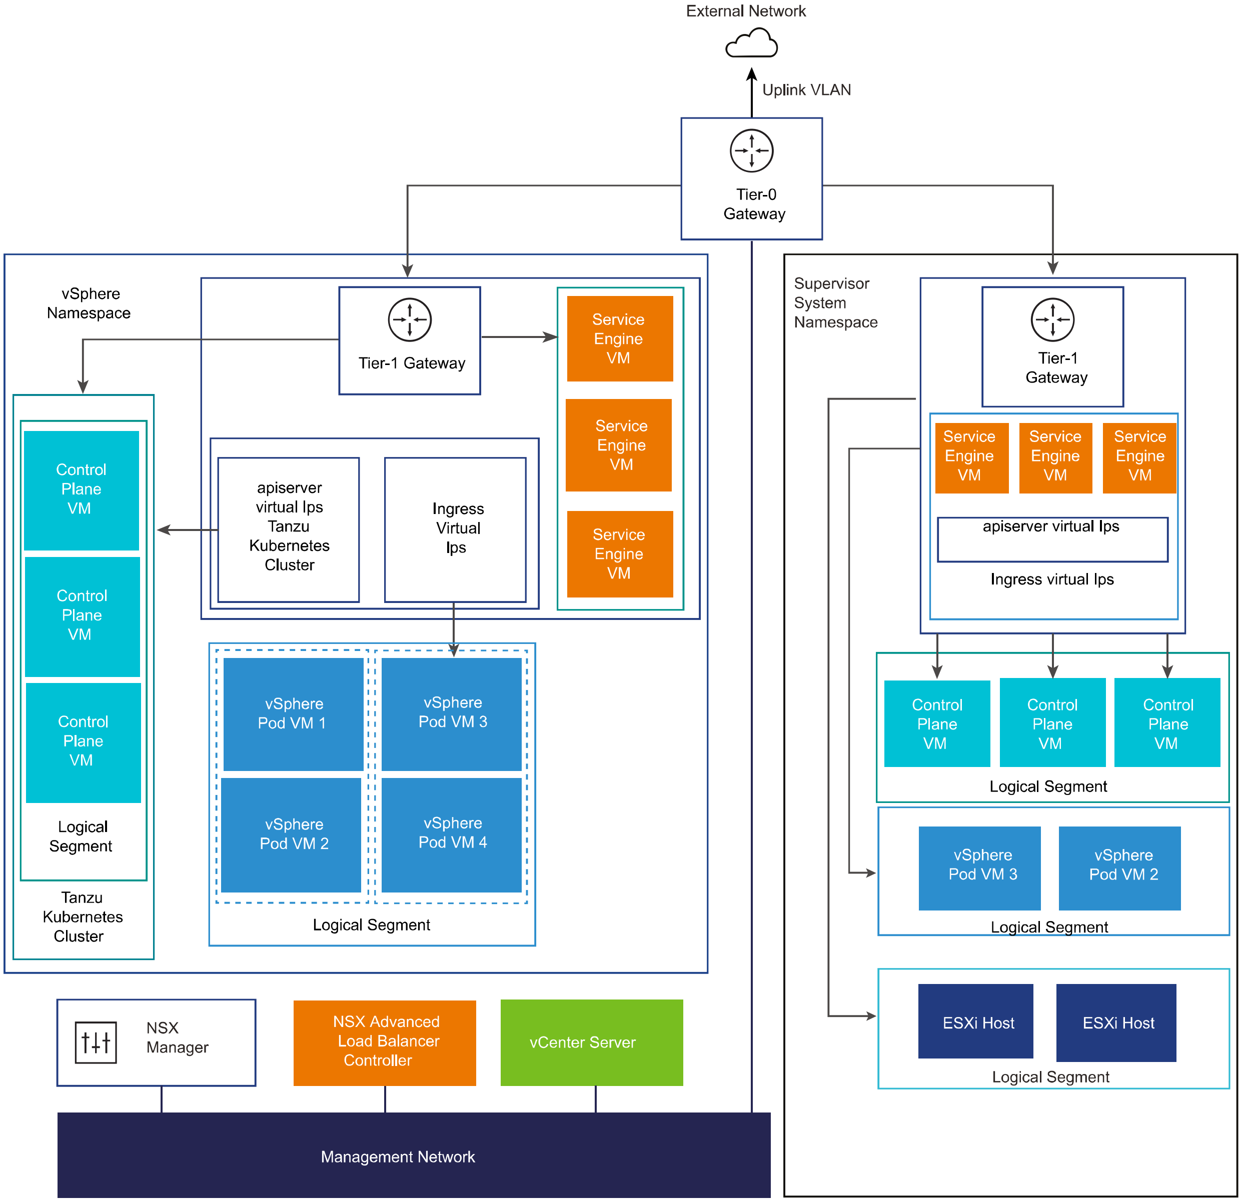

Figure 4-6. Three-zone Supervisor networking with NSX and NSX Advanced Load Balancer Controller

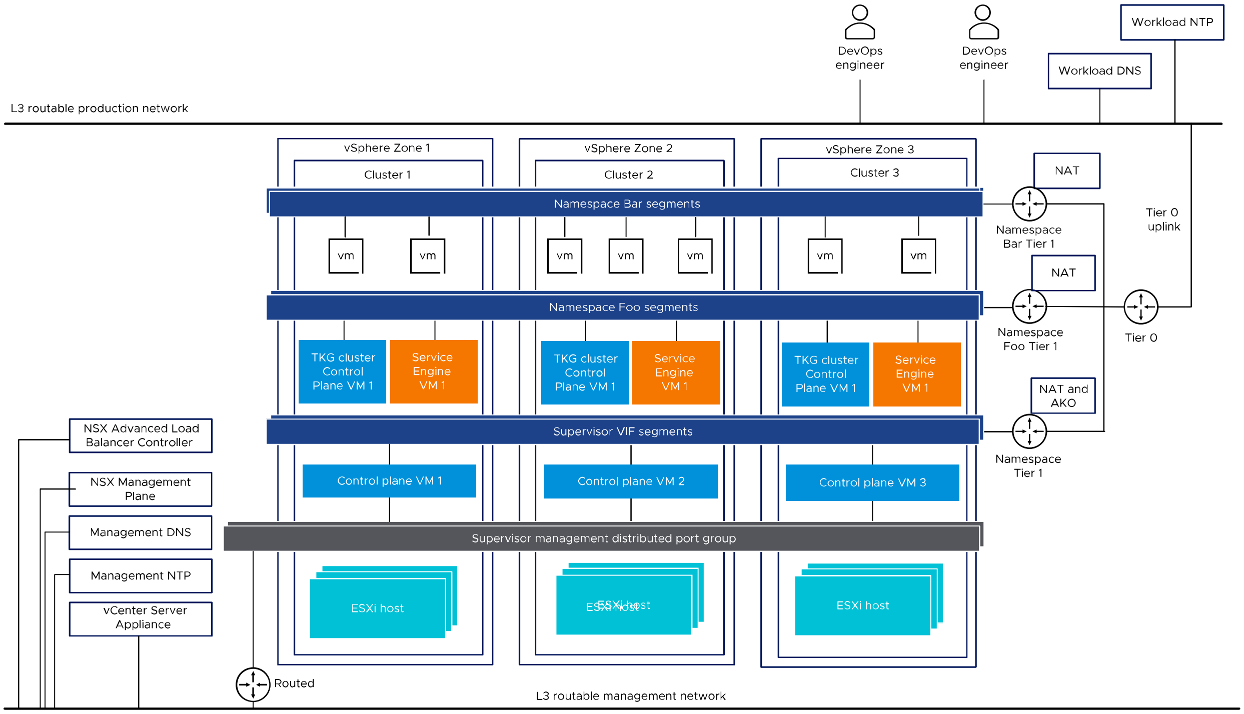

**Important** When you configure the NSX Advanced Load Balancer Controller in an NSX deployment, keep in mind the following considerations:

- You cannot deploy the NSX Advanced Load Balancer Controller in an vCenter Server Enhanced Linked Mode deployment. You can only deploy the NSX Advanced Load Balancer Controller in a single vCenter Server deployment. If more than one vCenter Server is linked, only one of them can be used while configuring the NSX Advanced Load Balancer Controller.
- You cannot configure the NSX Advanced Load Balancer Controller in a multi-tier tier-0 topology. If the NSX environment is set up with a multi-tier tier-0 topology, you can only use one tier-0 gateway while configuring the NSX Advanced Load Balancer Controller.

## Networking Configuration Methods with NSX

Supervisor uses an opinionated networking configuration. Two methods exist to configure the Supervisor networking with NSX that result in deploying the same networking model for a once-zone Supervisor:

- The simplest way to configure the Supervisor networking is by using the VMware Cloud Foundation SDDC Manager. For more information, see the VMware Cloud Foundation SDDC Manager documentation. For information, see VMware Cloud Foundation Administration Guide.

- You can also configure the Supervisor networking manually by using an existing NSX deployment or by deploying a new instance of NSX. See Install and Configure NSX for vSphere with Tanzu for more information.

# Install and Configure NSX for vSphere with Tanzu

vSphere with Tanzu requires specific networking configuration to enable connectivity to the Supervisors, vSphere Namespaces, and all objects that run inside the namespaces, such as vSphere Pods, VMs, and Tanzu Kubernetes clusters. As a vSphere administrator, install and configure NSX for vSphere with Tanzu.

Figure 4-7. Workflow for Configuring a Supervisor with NSX

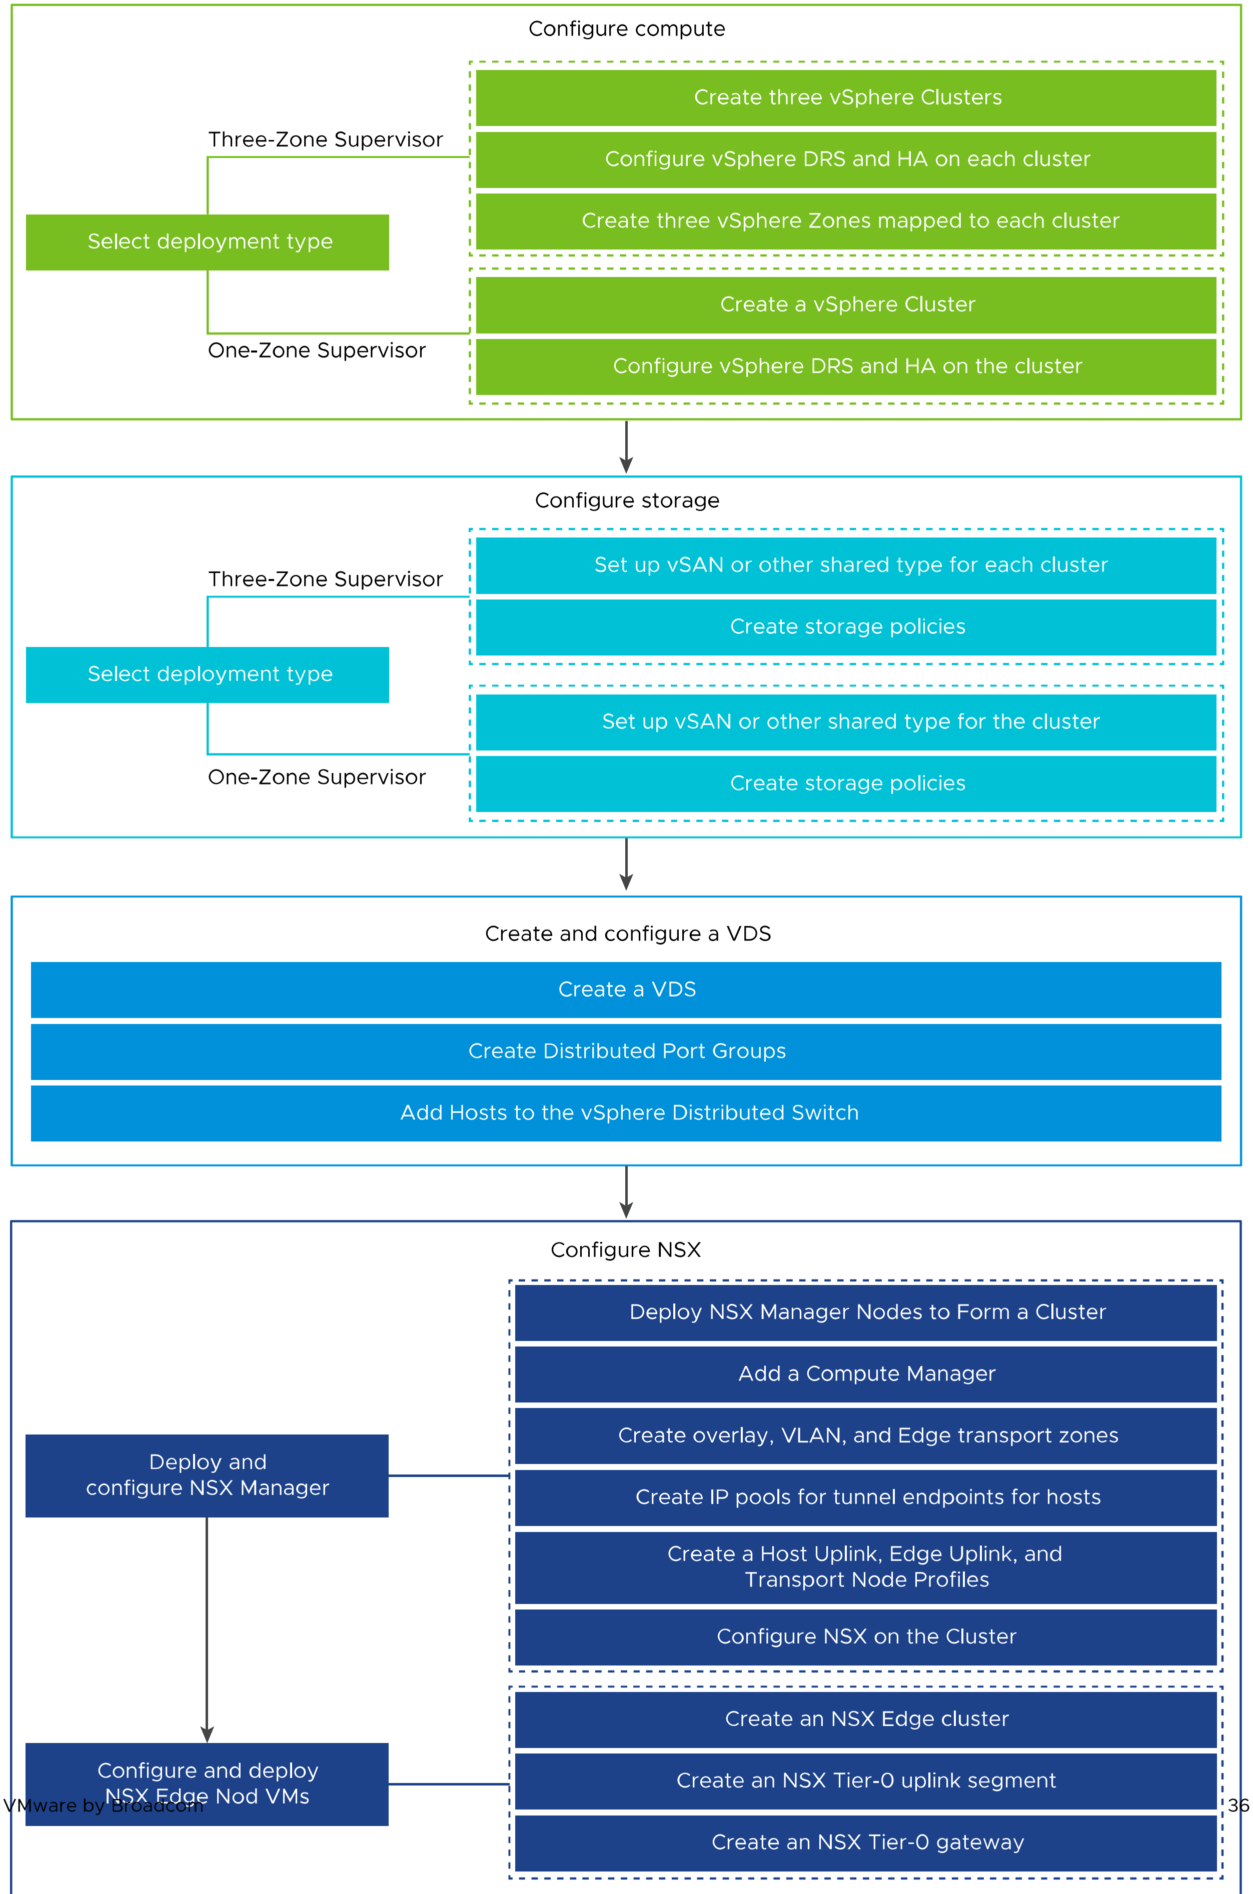
Configure compute
Three-Zone Supervisor
Select deployment type
One-Zone Supervisor
Create three vSphere Clusters
Configure vSphere DRS and HA on each cluster
Create three vSphere Zones mapped to each cluster
Create a vSphere Cluster
Configure vSphere DRS and HA on the cluster
Configure storage
Three-Zone Supervisor
Select deployment type
One-Zone Supervisor
Set up vSAN or other shared type for each cluster
Create storage policies
Set up vSAN or other shared type for the cluster
Create storage policies
Create and configure a VDS
Create a VDS
Create Distributed Port Groups
Add Hosts to the vSphere Distributed Switch
Configure NSX
Deploy and configure NSX Manager
Deploy NSX Manager Nodes to Form a Cluster
Add a Compute Manager
Create overlay, VLAN, and Edge transport zones
Create IP pools for tunnel endpoints for hosts
Create a Host Uplink, Edge Uplink, and Transport Node Profiles
Configure NSX on the Cluster
Configure and deploy NSX Edge Nod VMs
Create an NSX Edge cluster
Create an NSX Tier-0 uplink segment
Create an NSX Tier-0 gateway

This section describes how to configure the Supervisor networking by deploying a new NSX instance, but the procedures are applicable against an existing NSX deployment as well. This section also provides background to understand what VMware Cloud Foundation SDDC Manager is doing when it sets up the Supervisor workload domain.

Prerequisites

- Verify that your environment meets the system requirements for configuring a vSphere cluster as a Supervisor. For information about requirements, see Requirements for a Zonal Supervisor with NSX and Requirements for Cluster Supervisor Deployment with NSX in *vSphere with Tanzu Concepts and Planning*.
- Assign a Tanzu edition license to the Supervisor.
- Create storage policies for the placement of control plane VMs, pod ephemeral disks, and container images.
- Configure shared storage for the cluster. Shared storage is required for vSphere DRS, HA, and storing persistent volumes of containers.
- Verify that DRS and HA is enabled on the vSphere cluster, and DRS is in the fully automated mode.
- Verify that you have the **Modify cluster-wide configuration** privilege on the cluster.

Procedure

1 Create and Configure a vSphere Distributed Switch

  To handle the networking configuration for all hosts in the Supervisor, create a vSphere Distributed Switch, create distributed port groups, and associate hosts with the switch.

2 Deploy and Configure NSX Manager

  You can use the vSphere Client to deploy the NSX Manager to the vSphere cluster and use it with vSphere with Tanzu.

3 Create Transport Zones

  Transport zones indicate which hosts and VMs can use a particular network. A transport zone can span one or more host clusters.

4 Configure and Deploy an NSX Edge Transport Node

  You can add an NSX Edge virtual machine (VM) to the NSX fabric and proceed to configure it as an NSX Edge transport node VM.

## Create and Configure a vSphere Distributed Switch

To handle the networking configuration for all hosts in the Supervisor, create a vSphere Distributed Switch, create distributed port groups, and associate hosts with the switch.

Procedure

1 In the vSphere Client, navigate to a data center.

2 In the navigator, right-click the data center and select **Distributed Switch > New Distributed Switch**.

3 Enter a name for the new distributed switch.

For example, `DSwitch`.

4 In **Select Version**, enter a version for the distributed switch.

Select **8.0**.

5 In **Configure Settings**, enter the number of uplinks ports.

Enter a value of `2`.

6 Review the settings and click **Finish**.

7 Right-click the distributed switch you created and select **Settings > Edit settings**.

8 On the **Advanced** tab, enter a value of more than 1700 as the MTU (Bytes) value and click **OK**.

The MTU size must be 1700 or greater on any network that carries overlay traffic.

For example, `9000`.

NSX takes the global default MTU value 1700.

## Create Distributed Port Groups

Create distributed port groups for each NSX Edge node uplink, Edge node TEP, management network, and shared storage.

The default port group and the default uplinks are created when you create the vSphere Distributed Switch. You must create the management port group, vSAN port group. Edge TEP port group, and the NSX Edge uplink port group.

### Prerequisites

Verify that you have created a vSphere Distributed Switch.

### Procedure

1 In the vSphere Client, navigate to a data center.

2 In the navigator, right-click the distributed switch and select **Distributed Port Group > New Distributed Port Group**.

3 Create a port group for the NSX Edge uplink.

For example, `DPortGroup-EDGE-UPLINK`.

4 Configure **VLAN Type** as VLAN Trunking.

5 Accept the default VLAN trunk range **(0-4094)**.

6 Click **Next** and then click **Finish**.

7 Right-click the distributed switch and from the **Actions** menu, select **Distributed Port Group > Manage Distributed Port Groups**.

8 Select **Teaming and failover** and click **Next**.

9 Configure active and standby uplinks.

For example, active uplink is `Uplink1` and standby uplink is `Uplink2`.

10 Click **OK** to complete the configuration of the port group.

11 Repeat steps from 2 through 10 to create port groups for the Edge node TEP, management network, and shared storage.

For example, create the following port groups:

| Port group | Name | VLAN Type |
|---|---|---|
| Edge node TEP | `DPortGroup-EDGE-TEP` | Configure **VLAN Type** as VLAN Trunking.<br>Configure the active uplink as `Uplink2` and the standby uplink as `Uplink1`.<br>**Note** The VLAN used for the Edge nodes TEP must be different than the VLAN used for ESXi TEP. |
| Management | `DPortGroup-MGMT` | Configure **VLAN Type** as VLAN and enter the VLAN ID of the management network. For example, `1060`. |
| Shared storage or VSAN | `DPortGroup-VSAN` | Configure **VLAN Type** as VLAN and enter the VLAN ID. For example, `3082`. |

12 Create port groups for the following components:

- **vSphere vMotion**. This port group is required for Supervisor updates. Configure the default port group for vMotion.
- **VM traffic**. Configure the default port group to handle VM traffic.

## Add Hosts to the vSphere Distributed Switch

To manage the networking of your environment by using the vSphere Distributed Switch, you must associate hosts from the Supervisor with the switch. Connect the physical NICs, VMkernel adapters, and virtual machine network adapters of the hosts to the distributed switch.

### Prerequisites

- Verify that enough uplinks are available on the distributed switch to assign to the physical NICs that you want to connect to the switch.
- Verify that at least one distributed port group is available on the distributed switch.

- Verify that the distributed port group has active uplinks configured in its teaming and failover policy.

Procedure

1. In the vSphere Client, select **Networking** and navigate to the distributed switch.
2. From the **Actions** menu, select **Add and Manage Hosts**.
3. On the **Select task** page, select **Add hosts**, and click **Next**.
4. On the **Select hosts** page, click **New hosts**, select the hosts in your data center, click **OK**, and then click **Next**.
5. On the **Manage physical adapters** page, configure physical NICs on the distributed switch.

   a. From the **On other switches/unclaimed** list, select a physical NIC.

      If you select physical NICs that are already connected to other switches, they are migrated to the current distributed switch.

   b. Click **Assign uplink**.

   c. Select an uplink.

   d. To assign the uplink to all the hosts in the cluster, select **Apply this uplink assignment to the rest of the hosts**.

   e. Click **OK**.

   For example, assign `Uplink 1` to `vmnic0` and `Uplink 2` to `vmnic1`.
6. Click **Next**.
7. On the **Manage VMkernel adapters** page, configure VMkernel adapters.

   a. Select a VMkernel adapter and click **Assign port group**.

   b. Select a distributed port group.

      For example, **DPortGroup**.

   c. To apply the port group to all hosts in the cluster, select **Apply this port group assignment to the rest of the hosts**.

   d. Click **OK**.
8. Click **Next**.
9. (Optional) On the **Migrate VM networking** page, select the **Migrate virtual machine networking** check box to configure virtual machine networking.

   a. To connect all network adapters of a virtual machine to a distributed port group, select the virtual machine, or select an individual network adapter to connect only that adapter.

   b. Click **Assign port group**.

c Select a distributed port group from the list and click **OK**.

d Click **Next**.

### What to do next

Deploy and configure NSX Manager. See Deploy and Configure NSX Manager

## Deploy and Configure NSX Manager

You can use the vSphere Client to deploy the NSX Manager to the vSphere cluster and use it with vSphere with Tanzu.

To deploy the NSX Manager using the OVA file, perform the steps in this procedure.

For information about deploying the NSX Manager through the user interface or CLI, see the *NSX Installation Guide*.

### Prerequisites

- Verify that your environment meets the networking requirements. For information about requirements, see Requirements for a Three-Zone Supervisor with NSX Advanced Load Balancer and Requirements for Enabling a Single Cluster Supervisor with NSX Advanced Load Balancer in *vSphere with Tanzu Concepts and Planning*.
- Verify that the required ports are open. For information about port and protocols, see the *NSX Installation Guide*.

### Procedure

1 Locate the NSX OVA file on the VMware download portal.

  Either copy the download URL or download the OVA file.

2 Right-click and select **Deploy OVF template** to start the installation wizard.

3 In the **Select an OVF template** tab, enter the download OVA URL or navigate to the OVA file.

4 In the **Select a name and folder** tab, enter a name for the NSX Manager virtual machine (VM).

5 In the **Select a compute resource** tab, select the vSphere cluster on which to deploy the NSX Manager.

6 Click **Next** to review details.

7 In the **Configuration tab**, select the NSX deployment size.

  The recommended minimum deployment size is Medium.

8 In the **Select storage** tab, select the shared storage for deployment.

9 Enable thin provisioning by selecting the **Thin Provision** in **Select virtual disk format**.

  The virtual disks are thick provisioned by default.

10 In the **Select networks** tab, select the management port group or destination network for the NSX Manager in **Destination Network**.

   For example, `DPortGroup-MGMT`.

11 In the **Customize template** tab, enter the system root, CLI admin, and audit passwords for the NSX Manager. Your passwords must comply with the password strength restrictions.

   - At least 12 characters.
   - At least one lower-case letter.
   - At least one upper-case letter.
   - At least one digit.
   - At least one special character.
   - At least five different characters.
   - Default password complexity rules are enforced by the Linux PAM module.

12 Enter the default IPv4 gateway, management network IPv4, management network netmask, DNS server, domain search list, and NTP IP address.

13 Enable SSH and allow root SSH login to the NSX Manager command line.

   By default, the SSH options are disabled for security reasons.

14 Verify that your custom OVF template specification is accurate, and click **Finish** to initiate the installation.

15 After the NSX Manager boots, log in to the CLI as admin and run the `get interface eth0` command to verify that the IP address was applied as expected.

16 Enter the `get services` command to verify that all the services are running.

## Deploy NSX Manager Nodes to Form a Cluster

An NSX Manager cluster provides high availability. You can deploy NSX Manager nodes using the user interface only on ESXi hosts managed by vCenter Server. To create an NSX Manager cluster, deploy two additional nodes to form a cluster of three nodes total. When you deploy a new node from the UI, the node connects to the first deployed node to form a cluster. All the repository details and the password of the first deployed node are synchronized with the newly deployed node.

### Prerequisites

- Verify that an NSX Manager node is installed.
- Verify that a compute manager is configured.
- Verify that the required ports are open.
- Verify that a datastore is configured on the ESXi host.

- Verify that you have the IP address and gateway, DNS server IP addresses, domain search list, and the NTP server IP address for the NSX Manager to use.
- Verify that you have a target VM port group network. Place the NSX appliances on a management VM network.

### Procedure

1 From a browser, log in with admin privileges to the NSX Manager at https://<manager-ip-address>.

2 To deploy an appliance, select **System > Appliances > Add NSX Appliance**.

3 Enter the appliance details.

| Option | Description |
|---|---|
| **Hostname** | Enter the host name or FQDN to use for the node. |
| **Management IP/Netmask** | Enter an IP address to be assigned to the node. |
| **Management Gateway** | Enter a gateway IP address to be used by the node. |
| **DNS servers** | Enter the list of DNS server IP addresses to be used by the node. |
| **NTP server** | Enter the list of NTP server IP addresses |
| **Node Size** | Select **Medium (6 vCPU, 24 GB RAM, 300 GB storage)** form factor from the options. |

4 Enter the appliance configuration details

| Option | Description |
|---|---|
| **Compute Manager** | Select the vCenter Server that you configured as compute manager. |
| **Compute Cluster** | Select the cluster that the node must join. |
| **Datastore** | Select a datastore for the node files. |
| **Virtual Disk Format** | Select **Thin Provision** format. |
| **Network** | Click **Select Network** to select the management network for the node. |

5 Enter the access and credentials details.

| Option | Description |
|---|---|
| **Enable SSH** | Toggle the button to allow SSH login to the new node. |
| **Enable Root Access** | Toggle the button to allow root access to the new node. |

| Option | Description |
|---|---|
| **System Root Credentials** | Set and confirm the root password for tne new node.<br>Your password must comply with the password strength restrictions.<br>■ At least 12 characters.<br>■ At least one lower-case letter.<br>■ At least one upper-case letter.<br>■ At least one digit.<br>■ At least one special character.<br>■ At least five different characters.<br>■ Default password complexity rules are enforced by the Linux PAM module. |
| **Admin CLI Credentials and Audit CLI Credentials** | Select the **Same as root password** check box to use the same password that you configured for root, or deselect the check box and set a different password. |

6 Click **Install Appliance**.

The new node is deployed. You can track the deployment process in the **System> > Appliances** page. Do not add additional nodes until the installation is finished and the cluster is stable.

7 Wait for the deployment, cluster formation, and repository synchronization to finish.

The joining and cluster stabilizing process might take from 10 to 15 minutes. Verify that the status for every cluster service group is `UP` before making any other cluster changes.

8 After the node boots, log in to the CLI as admin and run the `get interface eth0` command to verify that the IP address was applied as expected.

9 If your cluster has only two nodes, add another appliance. Select **System > Appliances > Add NSX Appliance** and repeat the configuration steps.

## Add a License

Add a license using the NSX Manager.

### Prerequisites

Obtain an NSX Advanced or higher license.

### Procedure

1 Log in to the NSX Manager.

2 Select **System > Licenses > Add**.

3 Enter the license key.

4 Click **Add**.

## Add a Compute Manager

A compute manager is an application that manages resources such as hosts and virtual machines. Configure the vCenter Server that is associated with the NSX as a compute manager in the NSX Manager.

For more information see the *NSX Administration Guide*.

### Procedure

1 Log in to the NSX Manager.

2 Select **System > Fabric > Compute Managers > Add**

3 Enter the compute manager details.

| Option | Description |
|---|---|
| **Name and Description** | Enter the name and description of the vCenter Server. |
| **Type** | The default type is VMware vCenter. |
| **Multi NSX** | Leave this option unselected.<br>Multi NSX option allows you to register the same vCenter Server with multiple NSX Managers. This option is not supported on Supervisor and vSphere Lifecycle Manager clusters. |
| **FQDN or IP Address** | Enter the FQDN or the IP address of the vCenter Server. |
| **HTTPS Port of Reverse Proxy** | The default port is 443. If you use another port, verify that the port is open on all the NSX Manager appliances.<br>Set the reverse proxy port to register the compute manager in NSX. |
| **User name and Password** | Enter the vCenter Server login credentials. |
| **SHA-256 Thumbprint** | Enter the vCenter Server SHA-256 thumbprint algorithm value. |

You can leave the defaults for the other settings.

If you left the thumbprint value blank, you are prompted to accept the server provided thumbprint. After you accept the thumbprint, it takes a few seconds for NSX to discover and register the vCenter resources.

4 Select **Enable Trust** to allow vCenter Server to communicate with NSX.

5 If you did not provide a thumbprint value for NSX Manager, the system identifies the thumbprint and displays it.

6 Click **Add** to accept the thumbprint.

### Results

After some time, the compute manager is registered with vCenter Server and the connection status changes to `Up`. If the FQDN/PNID of vCenter Server changes, you must re-register it with the NSX Manager. For more information, see Register vCenter Server with NSX Manager.

**Note** After the vCenter Server is successfully registered, do not power off and delete the NSX Manager VM without deleting the compute manager first. Otherwise, when you deploy a new NSX Manager, you will not be able to register the same vCenter Server again. You will get an error stating that the vCenter Server is already registered with another NSX Manager.

You can click the compute manager name to view the details, edit the compute manager, or to manage tags that apply to the compute manager.

## Create Transport Zones

Transport zones indicate which hosts and VMs can use a particular network. A transport zone can span one or more host clusters.

As a vSphere administrator, you use the default transport zones or create the following ones:

- An overlay transport zone that is used by the Supervisor Control Plane VMs.
- A VLAN transport zone for the NSX Edge nodes to use for uplinks to the physical network.

### Procedure

1 Log in to the NSX Manager.

2 Select **System > Fabric > Transport Zones > Add**.

3 Enter a name for the transport zone and optionally a description.

4 Select a traffic type.

   You can select **Overlay** or **VLAN**.

   The following transport zones exist by default:

   - A VLAN transport zone with name `nsx-vlan-transportzone`.
   - An overlay transport zone with name `nsx-overlay-transportzone`.

5 (Optional) Enter one or more uplink teaming policy names.

   The segments attached to the transport zones use these named teaming policies. If the segments do not find a matching named teaming policy, then the default uplink teaming policy is used.

### Results

The new transport zone appears on the **Transport Zones** page.

## Create an IP Pool for Host Tunnel Endpoint IP Addresses

Create IP pools for the ESXi host tunnel endpoints (TEPs). TEPs are the source and destination IP addresses used in the external IP header to identify the ESXi hosts that originate and end the NSX encapsulation of overlay frames. You can use DHCP or manually configured IP pools for TEP IP addresses.

### Procedure

1 Log in to the NSX Manager.

2 Select **Networking > IP Address Pools > Add IP Address Pool**.

3 Enter the following IP pool details.

| Option | Description |
|---|---|
| **Name and Description** | Enter the IP pool name and optional description.<br>For example, `ESXI-TEP-IP-POOL`. |
| **IP Ranges** | Enter the IP allocation range.<br>For example, `192.23.213.158 - 192.23.213.160` |
| **Gateway** | Enter the gateway IP address.<br>For example, `192.23.213.253`. |
| **CIDR** | Enter the network address in a CIDR notation.<br>For example, `192.23.213.0/24`. |

4 Click **Add** and **Apply**.

### Results

Verify that the TEP IP pools you created are listed in the **IP Pool** page.

## Create an IP Pool for Edge Nodes

Create IP pools for the Edge Nodes. The TEP addresses are not required to be routable. You can use any IP addressing scheme that enables the Edge TEP to talk to the Host TEP.

### Procedure

1 Log in to the NSX Manager.

2 Select **Networking > IP Address Pools > Add IP Address Pool**.

3 Enter the following IP pool details.

| Option | Description |
|---|---|
| **Name and Description** | Enter the IP pool name and optional description.<br>For example, `EDGE-TEP-IP-POOL`. |
| **IP Ranges** | Enter the IP allocation range.<br>For example, `192.23.213.1 - 192.23.213.10`. |

| Option | Description |
|---|---|
| **Gateway** | Enter the gateway IP address.<br>For example, `192.23.213.253`. |
| **CIDR** | Enter the network address in a CIDR notation.<br>For example, `192.23.213.0/24`. |

4 Click **Add** and **Apply**.

### Results

Verify that the IP pools you created are listed in the **IP Pool** page.

## Create a Host Uplink Profile

A host uplink profile defines policies for the uplinks from the ESXi hosts to NSX segments.

### Procedure

1 Log in to the NSX Manager.

2 Select **System** > **Fabric** > **Profiles** > **Uplink Profiles** > **Add**.

3 Enter an uplink profile name, and optionally, an uplink profile description.

For example, `ESXI-UPLINK-PROFILE`.

4 In the **Teaming** section, click **Add** to add a naming teaming policy, and configure a **Failover Order** policy.

A list of active uplinks is specified, and each interface on the transport node is pinned to one active uplink. This configuration allows use of several active uplinks at the same time.

5 Configure active and standby uplinks.

For example, configure `uplink-1` as the active uplink and `uplink-2` as the standby uplink.

6 Enter a transport VLAN value.

The transport VLAN set in the uplink profile tags overlay traffic and the VLAN ID is used by the tunnel endpoint (TEP).

For example, `1060`.

7 Enter the MTU value.

The default value for uplink profile MTU is 1600.

**Note** The value must be at least 1600 but not higher than the MTU value on the physical switches and the vSphere Distributed Switch.

## Create an Edge Uplink Profile

Create an uplink profile with the failover order teaming policy with one active uplink for edge virtual machine overlay traffic.

### Procedure

1. Log in to the NSX Manager.
2. Select **System** > **Fabric** > **Profiles** > **Uplink Profiles** > **Add**.
3. Enter an uplink profile name, and optionally, add an uplink profile description.

   For example, `EDGE-UPLINK-PROFILE`.
4. In the **Teaming** section, click **Add** to add a naming teaming policy, and configure a **Failover** policy.

   A list of active uplinks is listed, and each interface on the transport node is pinned to one active uplink. This configuration allows use of several active uplinks at the same time.
5. Configure an active uplinks.

   For example, configure `uplink-1` as active uplink.
6. View the uplinks in the **Uplink Profile** page.

## Create a Transport Node Profile

A transport node profile defines how NSX is installed and configured on the hosts in a particular cluster the profile is attached to.

### Prerequisites

Verify that you have created an overlay transport zone.

### Procedure

1. Log in to the NSX Manager.
2. Select **System** > **Fabric** > **Profiles** > **Transport Node Profiles** > **Add**.
3. Enter a name for the transport node profile and optionally a description.

   For example, `HOST-TRANSPORT-NODE-PROFILE`.
4. In the **New Node Switch** section, Select **Type** as `VDS`.
5. Select **Mode** as `Standard`.
6. Select the vCenter Server and the Distributed Switch names from the list.

   For example, `DSwitch`
7. Select the overlay transport zone created previously.

   For example, `NSX-OVERLAY-TRANSPORTZONE`.
8. Select the host uplink profile created previously.

   For example, `ESXI-UPLINK-PROFILE`.
9. Select **Use IP Pool** from the **IP Assignment** list.

10 Select the host TEP pool created previously.

   For example, `ESXI-TEP-IP-POOL`.

11 In the **Teaming Policy Switch Mapping**, click the edit icon and map the uplinks defined in the NSX uplink profile with the vSphere Distributed Switch uplinks.

   For example, map `uplink-1 (active)` to `Uplink 1` and `uplink-2 (standby)` to `Uplink 2`.

12 Click **Add**.

13 Verify that the profile that you created in listed in the **Transport Node Profiles** page.

### Configure NSX on the Cluster

To install NSX and prepare the overlay TEPs, apply the transport node profile to the vSphere cluster.

#### Prerequisites

Verify that you have created a transport node profile.

#### Procedure

1 Log in to the NSX Manager.

2 Select **System > Fabric > Nodes > Host Transport Nodes**.

3 From the **Managed By** drop-down menu, select an existing vCenter Server.

   The page lists the available vSphere clusters.

4 Select the compute cluster on which you want to configure NSX.

5 Click **Configure NSX**.

6 Select the transport node profile created previously and click **Apply**.

   For example, `HOST-TRANSPORT-NODE-PROFILE`.

7 From the **Host Transport Node** page, verify that the NSX configuration state is `Success` and NSX Manager connectivity status of hosts in the cluster is `Up`.

#### Results

The transport node profile created previously is applied to the vSphere cluster to install NSX and prepare the overlay TEPs.

## Configure and Deploy an NSX Edge Transport Node

You can add an NSX Edge virtual machine (VM) to the NSX fabric and proceed to configure it as an NSX Edge transport node VM.

#### Prerequisites

Verify that you have created transport zones, edge uplink profile, and edge TEP IP pool.