

Deployment of VMware NSX-T Workload Domains with Multiple Availability Zones for VMware Cloud Foundation 3.8

Modified on 12 MAR 2020

VMware Validated Design

VMware Cloud Foundation 3.8

VMware Cloud Foundation 3.8.1



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 © 2019-2020 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

1 About Deployment of VMware NSX-T Workload Domains with Multiple Availability Zones for VMware Cloud Foundation 5

[Updated Information 6](#)

2 Specification of an NSX-T Workload Domain with Multiple Availability Zones 7

3 Example IP and DNS Configuration of an NSX-T Workload Domain with Multiple Availability Zones 9

4 Preparing for Deploying an NSX-T Workload Domain with Multiple Availability Zones 13

[Prepare the Virtual Infrastructure for an NSX-T Workload Domain with Multiple Availability Zones 13](#)

[System Prerequisites for Adding Availability Zone 2 to the NSX-T Workload Domain 14](#)

5 Configure the Virtual Infrastructure Components for a Second Availability Zone 16

[Detach the Host Transport Node Profile from the Cluster for an NSX-T Workload Domain 17](#)

[Add the ESXi Hosts to the Cluster for an NSX-T Workload Domain 18](#)

[Create Resource Pools for the Components in Availability Zone 2 in the Cluster for an NSX-T Workload Domain 19](#)

[Create a Virtual Machine Folder for the NSX-T Edge Appliances for an NSX-T Workload Domain 20](#)

[Create Port Groups for System Traffic in Availability Zone 2 on the vSphere Distributed Switch for an NSX-T Workload Domain 20](#)

[Connect the ESXi Hosts in Availability Zone 2 to the vSphere Distributed Switch for an NSX-T Workload Domain 21](#)

[Create the VMkernel Network Adapters for vSAN and vSphere vMotion in Availability Zone 2 for an NSX-T Workload Domain 22](#)

[Complete the Migration of the Physical Network Adapters to the vSphere Distributed Switch for Availability Zone 2 for an NSX-T Workload Domain 26](#)

[Enable SSH on the ESXi Hosts in Availability Zone 2 for an NSX-T Workload Domain 27](#)

[Create vSAN Disk Groups for the ESXi Hosts in Availability Zone 2 for an NSX-T Workload Domain 28](#)

[Add Static Routes for vSAN Traffic between the ESXi Host in Both Availability Zones for an NSX-T Workload Domain 28](#)

6 Deploy and Configure the vSAN Witness Host for an NSX-T Workload Domain 31

[Deploy the vSAN Witness Host in a Third Location for an NSX-T Workload Domain 32](#)

[Configure the Management Network on the vSAN Witness Host for an NSX-T Workload Domain 33](#)

[Add the vSAN Witness Appliance as a Standalone Host to the NSX-T Workload Domain 34](#)

[Configure NTP and SSH on the vSAN Witness Host for an NSX-T Workload Domain 35](#)

	Configure the Witness VMkernel Adapter on the vSAN Witness Host for an NSX-T Workload Domain	36
	Add Static Routes for the vSAN Witness Host for an NSX-T Workload Domain	37
7	Configure vSAN Stretched Cluster for an NSX-T Workload Domain	38
	Enable vSAN Stretched Cluster on the NSX-T Workload Domain Cluster	38
	Turn on the vSAN Performance Service for the Cluster for an NSX-T Workload Domain	39
	Update the vSAN Storage Policy for an NSX-T Workload Domain	40
	Update the vSphere Availability Settings of the Cluster for an NSX-T Workload Domain	41
8	Configure the NSX-T Instance for an NSX-T Workload Domain	43
	Create the Transport Zones for Uplink Traffic for an NSX-T Workload Domain	44
	Create Uplink Profiles for the ESXi Hosts in Availability Zone 2 for an NSX-T Workload Domain	45
	Create the NSX-T Segments for System, Uplink, and Overlay Traffic for an NSX-T Workload Domain	46
	Configure the ESXi Host Transport Nodes in Availability Zone 2 for an NSX-T Workload Domain	47
	Remove the ESXi Hosts for Availability Zone 2 from the vSphere Distributed Switch for an NSX-T Workload Domain	49
	Configure Dynamic Routing for an NSX-T Workload Domain	50
	Create an NSX-T Edge Cluster Profile for an NSX-T Workload Domain	51
	Deploy the NSX-T Edge Appliances for an NSX-T Workload Domain	52
	Join the NSX-T Edge Nodes to the Management Plane for an NSX-T Workload Domain	54
	Create an Anti-Affinity Rule for the NSX-T Edge Nodes for an NSX-T Workload Domain	54
	Create Host Groups and Rules for Both Availability Zones for an NSX-T Workload Domain	55
	Add the NSX-T Edge Nodes to the Transport Zones for an NSX-T Workload Domain	57
	Create the NSX-T Edge Cluster for an NSX-T Workload Domain	59
	Create and Configure the Tier-0 Gateway for an NSX-T Workload Domain	60
	Create and Configure the Tier-1 Gateway for an NSX-T Workload Domain	63
	Verify BGP Peering and Route Redistribution for an NSX-T Workload Domain	64
9	Register the vSAN Stretched Cluster in SDDC Manager for an NSX-T Workload Domain	66

About Deployment of VMware NSX-T Workload Domains with Multiple Availability Zones for VMware Cloud Foundation

1

Deployment of VMware NSX-T Workload Domains with Multiple Availability Zones for VMware Cloud Foundation provides step-by-step instructions for deploying a virtual infrastructure workload domain in VMware Cloud Foundation™ with VMware NSX-T Data Center™ for software-defined networking and with stretched clusters in VMware vSAN™ for multiple availability zones.

Intended Audience

The *Deployment of VMware NSX-T Workload Domains with Multiple Availability Zones for VMware Cloud Foundation* document is intended for cloud architects, infrastructure administrators, and cloud administrators who want to provide a highly available environment for provisioning tenant workloads on top of NSX-T and vSAN.

Required VMware Software

Deployment of VMware NSX-T Workload Domains with Multiple Availability Zones for VMware Cloud Foundation is compatible and validated with certain product versions.

- VMware Cloud Foundation 3.8 or VMware Cloud Foundation 3.8.1

Updated Information

This *Deployment of VMware NSX-T Workload Domains with Multiple Availability Zones for VMware Cloud Foundation* is updated with each release of VMware Cloud Foundation or when necessary.

This table provides the update history of the *Deployment of VMware NSX-T Workload Domains with Multiple Availability Zones for VMware Cloud Foundation*.

Revision	Description
03 SEP 2019	In VMware Cloud Foundation 3.8.1 or later, you can use the Supportability and Serviceability (SoS) utility without applying the patch in VMware Knowledge Base article 71278 . See Create the VMkernel Network Adapters for vSAN and vSphere vMotion in Availability Zone 2 for an NSX-T Workload Domain
19 AUG 2019	Initial release.

Specification of an NSX-T Workload Domain with Multiple Availability Zones

2

In VMware Cloud Foundation, you deploy the management domain and NSX-T workload domain of the SDDC with a configuration for supporting multiple availability zones. The use of vSAN stretched cluster also requires that you allocate a third location for deploying the vSAN witness appliance that is not local to the ESXi hosts on either side of the stretched cluster.

Table 2-1. Configuration of the Management Domain

Component	Value
Deployed management components	<ul style="list-style-type: none">■ ESXi hosts■ Platform Services Controller pair and Management vCenter Server■ NSX Data Center for vSphere<ul style="list-style-type: none">■ NSX Manager■ NSX Controller cluster■ vRealize Log Insight■ SDDC Manager
Minimum number of ESXi hosts	8 4 hosts per availability zone
Storage type	vSAN
Number of clusters	1 The cluster contains the management components of the SDDC.
Number of availability zones	2 <ul style="list-style-type: none">■ Availability Zone 1■ Availability Zone 2

Table 2-2. Configuration of the NSX-T Workload Domain

Component	Value
Deployed management components	<ul style="list-style-type: none"> ■ Management domain <ul style="list-style-type: none"> ■ ESXi hosts ■ vCenter Server for the workload domain ■ NSX-T Manager cluster ■ Workload domain <ul style="list-style-type: none"> ■ A pair of NSX-T Edge nodes in each availability zone in the workload domain
Minimum number of ESXi hosts	8 4 hosts per availability zone
Storage type	vSAN
Number of clusters	1 The hosts in the cluster are shared between the NSX-T Edge nodes and tenant workloads.
Number of availability zones	2 <ul style="list-style-type: none"> ■ Availability Zone 1 ■ Availability Zone 2 You add Availability Zone 2 after you deploy an NSX-T workload domain in VMware Cloud Foundation.

Table 2-3. Example Configuration of the Third Location for the vSAN Witness Host

Component	Value
Deployed management components	<ul style="list-style-type: none"> ■ ESXi host ■ Management vCenter Server ■ Virtual networking
Minimum number of ESXi Hosts	1

Example IP and DNS Configuration of an NSX-T Workload Domain with Multiple Availability Zones

3

You assign host names and IP address according to the VLAN setup and domain configuration for multiple availability zones in VMware Cloud Foundation. You can use an example configuration of IP subnets, VLAN IDs, and FQDNs in your environment.

Example IP Addresses and Host Names for the NSX-T Workload Domain

Verify that the static IP addresses and FQDNs for all components are allocated on the DNS server and are available for deployment.

Table 3-1. Example VLAN IDs and IP Subnets for Availability Zone 1 for the NSX-T Workload Domain

VLAN Function	VLAN ID	Subnet	Gateway
Management	1641 (Stretched Layer 2)	172.16.41.0/24	172.16.41.253
vSphere vMotion	1642	172.16.42.0/24	172.16.42.253
vSAN	1643	172.16.43.0/24	172.16.43.253
Host overlay	1644	172.16.44.0/24	172.16.44.253
Uplink01	1647	172.16.47.0/24	172.16.47.253
Uplink02	1648	172.16.48.0/24	172.16.48.253
Edge overlay	1649	172.16.49.0/24	172.16.49.253

Table 3-2. Example VLAN IDs and IP Subnets for Availability Zone 2 for the NSX-T Workload Domain

VLAN Function	VLAN ID	Subnet	Gateway
Management	1661	172.16.61.0/24	172.16.61.253
vSphere vMotion	1662	172.16.62.0/24	172.16.62.253
vSAN	1663	172.16.63.0/24	172.16.63.253
Host overlay	1664	172.16.64.0/24	172.16.64.253
Uplink01	1667	172.16.67.0/24	172.16.67.253

Table 3-2. Example VLAN IDs and IP Subnets for Availability Zone 2 for the NSX-T Workload Domain (continued)

VLAN Function	VLAN ID	Subnet	Gateway
Uplink02	1668	172.16.68.0/24	172.16.68.253
Edge overlay	1669	172.16.69.0/24	172.16.69.253

Table 3-3. Example FQDNs and IP Addresses for the ESXi Hosts of the NSX-T Workload Domain

Availability Zone	ESXi Host	Management IP Address	NTP Server
Availability Zone 1	sfo01w01esx01.sfo01.rainpole.local	172.16.41.101	ntp.sfo01.rainpole.local
	sfo01w01esx02.sfo01.rainpole.local	172.16.41.102	ntp.sfo01.rainpole.local
	sfo01w01esx03.sfo01.rainpole.local	172.16.41.103	ntp.sfo01.rainpole.local
	sfo01w01esx04.sfo01.rainpole.local	172.16.41.104	ntp.sfo01.rainpole.local
Availability Zone 2	sfo02w01esx01.sfo01.rainpole.local	172.16.61.101	ntp.sfo01.rainpole.local
	sfo02w01esx02.sfo01.rainpole.local	172.16.61.102	ntp.sfo01.rainpole.local
	sfo02w01esx03.sfo01.rainpole.local	172.16.61.103	ntp.sfo01.rainpole.local
	sfo02w01esx04.sfo01.rainpole.local	172.16.61.104	ntp.sfo01.rainpole.local

Table 3-4. Example FQDNs and IP Addresses for NSX-T Manager in the Management Domain

Role	FQDN	IP Address
NSX-T Manager instances	sfo01wnsx01a.sfo01.rainpole.local	172.16.11.82
	sfo01wnsx01b.sfo01.rainpole.local	172.16.11.83
	sfo01wnsx01c.sfo01.rainpole.local	172.16.11.84
NSX-T Manager cluster	sfo01wnsx01.sfo01.rainpole.local	172.16.11.81 (VIP)

Table 3-5. Example FQDNs and IP Addresses for the NSX-T Edge Nodes

Availability Zone	Role	FQDN	IP Address
Availability Zone 1	Edge Node 01	sfo01wesg01.sfo01.rainpole.local	172.16.41.21 (Management)
			172.16.49.21 (Overlay)
			172.16.47.2 (Uplink 1)
			172.16.48.2 (Uplink 2)
	Edge Node 02	sfo01wesg02.sfo01.rainpole.local	172.16.41.22 (Management)

Table 3-5. Example FQDNs and IP Addresses for the NSX-T Edge Nodes (continued)

Availability Zone	Role	FQDN	IP Address
Availability Zone 2			172.16.49.21 (Overlay)
			172.16.47.3 (Uplink 1)
			172.16.48.3 (Uplink 2)
	Edge Node 01	sfo02wesg01.sfo01.rainpole.local	172.16.61.21 (Management)
			172.16.69.21 (Overlay)
			172.16.67.2 (Uplink 1)
			172.16.68.2 (Uplink 2)
	Edge Node 02	sfo02wesg02.sfo01.rainpole.local	172.16.61.22 (Management)
			172.16.69.22 (Overlay)
			172.16.67.3 (Uplink 1)
			172.16.68.3 (Uplink 2)
Subnet mask		-	255.255.255.0
DNS		-	■ 172.16.11.5 ■ 172.16.11.4
NTP		ntp.sfo01.rainpole.local	■ 172.16.11.251 ■ 172.16.11.252

Example Host Names of in the Management Domain

Component	FQDN	IP Address
Management vCenter Server	sfo01m01vc01.sfo01.rainpole.local	172.16.11.62
vCenter Server for the NSX-T workload domain	sfo01w01vc01.sfo01.rainpole.local	172.16.11.64
SDDC Manager	sfo01sddcm01.sfo01.rainpole.local	172.16.11.65

Third Location Components

Component	FQDN	IP Address
vCenter Server in the third location	lax01m01vc01.lax01.rainpole.local	172.17.11.62
vSAN witness host	sfo03w01vsanw01.sfo01.rainpole.local	■ 172.17.11.203 (Management) ■ 172.17.13.203 (vSAN)

Example Naming Convention for Host and Object Names

In your environment, you can use an example naming convention where host and object names indicate the location and scope of the management components.

Table 3-6. Example Naming Convention for NSX- T Workload Domains

Component Type	
■ Host names	<location><az-number><workload-domain><workload-domain-number>
■ Virtual machine names	<mgmt-component><mgmt-component-id>
■ Clusters	<location><az-number>-<workload-domain><workload-domain-number>-<object-name><object-number>-<sub-object>
■ Virtual switches and port groups	
■ Datastores	
■ Data centers	<location><az-number>-<workload-domain><workload-domain-number><object-type>-<object-name><object-number>
■ Resource pools	
■ Virtual machine folders	

The <location> segment represents an instance of VMware Cloud Foundation at a specific location, also called region. For example, you can use city airport codes, such as SFO or LAX.

- Components with object names that begin with *sfo01* are part of the SFO region and Availability Zone 1.
- Components with object names that begin with *sfo02* are part the SFO region and Availability Zone 2.
- Components with object names that begin with *lax01* are part the LAX region and Availability Zone 1.
- Components with object names that begin with *sfo* without a number are part of the SFO region and both availability zones.
- Components with object names that begin with *sfo01w01* are part of the SFO region and the first workload domain.
- The vSAN witness node is deployed in a third availability zone and its host name starts with the *sfo03* prefix.

Preparing for Deploying an NSX-T Workload Domain with Multiple Availability Zones

4

Before you start the deployment of an NSX-T virtual infrastructure domain in VMware Cloud Foundation with two availability zones, your environment must meet target prerequisites and be in a specific starting state. Prepare the SDDC by deploying and configuring the necessary infrastructure and management components.

Planning the Deployment of an NSX-T Workload Domain with Multiple Availability Zones

- Because of intermittent network loss during migration of the host VMkernel network adapters between virtual switches, perform the configuration during a maintenance window. Schedule a maintenance window that is suitable for your organization and tenants.
- Allocate time in your maintenance window to run operational verification tests.
- Plan for any potential impact to tenant workloads while setting up the vSAN stretched cluster in the NSX-T workload domain.

Procedure

1 [Prepare the Virtual Infrastructure for an NSX-T Workload Domain with Multiple Availability Zones](#)

The configuration of the management domain and NSX-T workload domain in your VMware Cloud Foundation environment must support extending the vSAN datastore across two availability zones. Adding a second availability zone to the NSX-T workload domain also requires ESXi hosts for the zone and a location for the vSAN witness appliance.

2 [System Prerequisites for Adding Availability Zone 2 to the NSX-T Workload Domain](#)

Verify that your environment satisfies the following prerequisites for adding an availability zone to the NSX-T workload domain in a VMware Cloud Foundation environment.

Prepare the Virtual Infrastructure for an NSX-T Workload Domain with Multiple Availability Zones

The configuration of the management domain and NSX-T workload domain in your VMware Cloud Foundation environment must support extending the vSAN datastore across two availability zones.

Adding a second availability zone to the NSX-T workload domain also requires ESXi hosts for the zone and a location for the vSAN witness appliance.

- 1 Deploy the management domain in a configuration with multiple availability zones.
 - a Deploy the VMware Cloud Foundation stack in the management domain of your SDDC by using VMware Cloud Builder. See *VMware Cloud Foundation Planning and Preparation Guide* and *VMware Cloud Foundation Architecture and Deployment Guide*.
 - b Configure the management domain with a second availability zone by using a vSAN stretched cluster. See the [Stretching Clusters](#) section of the *VMware Cloud Foundation Operations and Administration Guide* documentation.

- 2 By using SDDC Manager, deploy an NSX-T workload domain with a configuration that supports the subsequent addition of a second availability zone.

SDDC Manager supports automatic deployment of workload domains with a single availability zone. See [Working with the Management Domain and VI Workload Domains](#) in *VMware Cloud Foundation Operations and Administration Guide*.

- 3 In SDDC Manager, create a network pool for the vSphere vMotion and vSAN subnets in Availability Zone 2.
- 4 Install the ESXi Hosts for the second availability zone.

See [Installing ESXi Software on Cloud Foundation Servers](#) in *VMware Cloud Foundation Operations and Administration Guide*.

Commissioning the hosts for the deployed workload domain is not required because you must configure IP subnets for the second availability zone that are different from the subnets in the initial availability zone.

- 5 Provide a third location for the vSAN witness appliance.

The vSAN witness appliance must have access to the SDDC management network.

For information about the configuration of the environment, see [Chapter 2 Specification of an NSX-T Workload Domain with Multiple Availability Zones](#) and [Chapter 3 Example IP and DNS Configuration of an NSX-T Workload Domain with Multiple Availability Zones](#).

System Prerequisites for Adding Availability Zone 2 to the NSX-T Workload Domain

Verify that your environment satisfies the following prerequisites for adding an availability zone to the NSX-T workload domain in a VMware Cloud Foundation environment.

Prerequisite Type	Value
Network connectivity	<ul style="list-style-type: none"> ■ Verify the two management networks for the NSX-T workload domain are stretched across the two availability zones. ■ Verify that all other network communication across the availability zones is over routed Layer 3 connections. ■ Verify that jumbo frames are enabled across the two availability zones. ■ Verify that the round trip latency between the availability zones is less than 5 ms.
Software features	<ul style="list-style-type: none"> ■ Verify that the Management vCenter Server is operational. ■ Verify that the management cluster has vSphere DRS and vSphere HA enabled. ■ Verify that you have the Postman REST client installed.
Installation packages	<ul style="list-style-type: none"> ■ Download the .ova file for the vSAN witness appliance 6.7 Update 3 to the machine where you use the vSphere Client from My VMware. ■ Download the .ova file for the NSX-T 2.4.2 Edge appliance for VMware ESXi from My VMware. ■ For VMware Cloud Foundation 3.8, download the vcf-stretch-cluster-patch.zip patch for the SoS utility from VMware Knowledge Base article 71278 and update the SoS utility on the SDDC Manager appliance by following the instructions in the VMware Knowledge Base article. <p>The SoS utility in VMware Cloud Foundation 3.8.1 supports deploying NSX-T workload domains with multiple availability zones.</p>
Access to the data center	Provide a host virtual machine or a physical server to connect to the SDDC and store software downloads. The host must have access to the Internet and to the ESXi management network.
Network pools	Verify that you have a network pool with IP addresses in the vSphere vMotion and vSAN subnets for Availability Zone 2.

Configure the Virtual Infrastructure Components for a Second Availability Zone

5

VMware Cloud Foundation deploys the virtual infrastructure for the NSX-T workload domain according to a pre-defined configuration. For vSAN stretched cluster and NSX-T support, you must modify the cluster configuration and virtual networking in the domain.

Procedure

1 Detach the Host Transport Node Profile from the Cluster for an NSX-T Workload Domain

Because the hosts in the two availability zones use different VLANs and subnets, you cannot use the transport node profile created by VMware Cloud Foundation for both availability zones.

2 Add the ESXi Hosts to the Cluster for an NSX-T Workload Domain

After you deploy the NSX-T workload domain, you add the ESXi hosts for Availability Zone 2 to the cluster of the workload domain. When you turn on the stretched cluster features, you place these hosts in the secondary fault domain of the cluster.

3 Create Resource Pools for the Components in Availability Zone 2 in the Cluster for an NSX-T Workload Domain

You create a resource pools for NSX-T Edge node pair and for tenant workloads in the cluster of the NSX-T workload domain in VMware Cloud Foundation.

4 Create a Virtual Machine Folder for the NSX-T Edge Appliances for an NSX-T Workload Domain

In the vCenter Server for the workload domain, create a virtual machine folder for the NSX-T Edge appliances so that you can manage the appliances together.

5 Create Port Groups for System Traffic in Availability Zone 2 on the vSphere Distributed Switch for an NSX-T Workload Domain

In VMware Cloud Foundation, after you add the ESXi hosts for Availability Zone 2 to the workload domain cluster, you create additional port groups on the vSphere Distributed Switch for the management, vSphere vMotion, and vSAN traffic. This switch handles traffic until you migrate the hosts to the N-VDS instance for the cluster.

6 Connect the ESXi Hosts in Availability Zone 2 to the vSphere Distributed Switch for an NSX-T Workload Domain

In VMware Cloud Foundation, you add all ESXi hosts in Availability Zone 2 to the vSphere Distributed Switch for the workload domain so that the switch can start handling management traffic until you migrate all hosts to the N-VDS instance.

7 Create the VMkernel Network Adapters for vSAN and vSphere vMotion in Availability Zone 2 for an NSX-T Workload Domain

In VMware Cloud Foundation, you create and configure the VMkernel network adapters for vSAN and vSphere vMotion traffic for the hosts in Availability Zone 2.

8 Complete the Migration of the Physical Network Adapters to the vSphere Distributed Switch for Availability Zone 2 for an NSX-T Workload Domain

In VMware Cloud Foundation, after you add all hosts in Availability Zone 2 to the vSphere Distributed Switch and configure the VMkernel adapters for vSphere vMotion and vSAN traffic, migrate the last physical network adapters to complete the networking configuration.

9 Enable SSH on the ESXi Hosts in Availability Zone 2 for an NSX-T Workload Domain

In VMware Cloud Foundation, you use SSH for configuring access to the ESXi hosts in Availability Zone 2 and for configuring static routes on the ESXi hosts for the vSAN stretched cluster.

10 Create vSAN Disk Groups for the ESXi Hosts in Availability Zone 2 for an NSX-T Workload Domain

In VMware Cloud Foundation, to extend the storage space of the vSAN datastore, create vSAN disk groups on each ESXi host in Availability Zone 2.

11 Add Static Routes for vSAN Traffic between the ESXi Host in Both Availability Zones for an NSX-T Workload Domain

In VMware Cloud Foundation, because vSphere uses a single default gateway, all routed traffic attempts to reach its destination through this gateway. Configure static routes on the ESXi hosts in both availability zones to route vSAN traffic using a dedicated gateway.

Detach the Host Transport Node Profile from the Cluster for an NSX-T Workload Domain

Because the hosts in the two availability zones use different VLANs and subnets, you cannot use the transport node profile created by VMware Cloud Foundation for both availability zones.

Procedure

- 1 In a Web browser, log in to the user interface of NSX-T Manager.

Setting	Value
URL	https://sfo01wnsx01.sfo01.rainpole.local
User name	admin
Password	nsx_admin_password

- 2 On the main navigation bar, click **System**.
- 3 In the navigation pane, select **Fabric > Nodes**.
- 4 On the **Host Transport Nodes** tab, from the **Managed by** drop-down menu, select **sfo01w01vc01**.
- 5 Select the sfo01-w01-shared01 cluster.

- 6 Click **Actions** and click **Detach TN Profile**.

Add the ESXi Hosts to the Cluster for an NSX-T Workload Domain

After you deploy the NSX-T workload domain, you add the ESXi hosts for Availability Zone 2 to the cluster of the workload domain. When you turn on the stretched cluster features, you place these hosts in the secondary fault domain of the cluster.

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 Right-click the **sfo01-w01-shared01** cluster, and select **Add hosts**.

The **Add hosts** wizard opens.

- 3 Add the ESXi hosts for Availability Zone 2 to the cluster of the workload domain.

- a On the **Add hosts** page, select **Use same credentials for all hosts**, configure the settings, and click **Next**.

ESXi Host FQDN	User name	Password
sfo02w01esx01.sfo01.rainpole.local	root	esxi_root_password
sfo02w01esx02.sfo01.rainpole.local		
sfo02w01esx03.sfo01.rainpole.local		
sfo02w01esx04.sfo01.rainpole.local		

- b In the **Security Alert** dialog box, select all ESXi hosts, and, to confirm the validity of the host certificates, click **OK**.

A trusted connection between vCenter Server and the ESXi hosts is established using the host certificates for SSL handshake.

- c On the **Host summary** page, review the host information and click **Next**.

- d On the **Ready to complete** page, click **Finish**.

- 4 On the **sfo01-w01-shared01** cluster page, click the **Hosts** tab.

- 5 Select all ESXi hosts, right-click and select **Maintenance mode > Exit maintenance mode**.

- 6 Select all ESXi hosts, and from the **Actions** drop-down menu, select **Assign License**.

- 7 On the **Assign license** dialog box, select the ESXi license from the inventory, and click **OK**.

Create Resource Pools for the Components in Availability Zone 2 in the Cluster for an NSX-T Workload Domain

You create a resource pools for NSX-T Edge node pair and for tenant workloads in the cluster of the NSX-T workload domain in VMware Cloud Foundation.

Allocate more CPU and memory resources to the resource pool for the NSX-T Edge appliances. The edge nodes provide dynamic routing to the tenant workloads in the workload domain. During contention, the edge nodes must receive enough resources so that network traffic is not interrupted.

Table 5-1. Resource Pools for the NSX-T Edge Appliances and Tenant Workloads for Availability Zone 2

Setting	Value for the NSX-T Edge Appliance Pool	Value for the Tenant Workload Pool
Name	sfo01-w01rp-sddc-edge	sfo01-w01rp-user-vm
CPU-Shares	High	Normal
CPU-Reservation	0	0
CPU-Reservation Type: Expandable	Selected	Selected
CPU-Limit	Unlimited	Unlimited
Memory-Shares	Normal	Normal
Memory-Reservation	128 GB	0
Memory-Reservation type: Expandable	Selected	Selected
Memory-Limit	Unlimited	Unlimited

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Hosts and clusters** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree and expand the sfo01-m01dcsfo01-w01dc data center.
- 3 Right-click the **sfo01-w01-shared01** cluster and select **New Resource Pool**.
- 4 In the **New Resource Pool** dialog box, configure the settings for the sfo01-w01rp-sddc-edge pool and click **OK**.
- 5 Repeat the step to create the sfo01-w01rp-user-vm resource pool.

Create a Virtual Machine Folder for the NSX-T Edge Appliances for an NSX-T Workload Domain

In the vCenter Server for the workload domain, create a virtual machine folder for the NSX-T Edge appliances so that you can manage the appliances together.

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **VMs and templates** inventory, expand **sfo01w01vc01.sfo01.rainpole.local**.
- 3 Right-click the **sfo01-w01dc** data center and select **New folder > New VM and template folder**.
- 4 In the **New Folder** dialog box, enter **sfo01-w01fd-nsx** and click **OK**.

Create Port Groups for System Traffic in Availability Zone 2 on the vSphere Distributed Switch for an NSX-T Workload Domain

In VMware Cloud Foundation, after you add the ESXi hosts for Availability Zone 2 to the workload domain cluster, you create additional port groups on the vSphere Distributed Switch for the management, vSphere vMotion, and vSAN traffic. This switch handles traffic until you migrate the hosts to the N-VDS instance for the cluster.

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Networking** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree and expand the **sfo01-w01dc** data center.
- 3 Right-click the **sfo01-w01-vds01** distributed switch, and select **Distributed Port Group > New Distributed Port Group**.

- 4 Create the port groups for system traffic for the ESXi hosts in Availability Zone 2 on the sfo01-w01-vds01 distributed switch and click **Next**.

Port Group Name	Port Binding	Port Allocation	Number of Ports	VLAN Type	VLAN ID
sfo02-w01-vds01-management	Ephemeral	Elastic	8	VLAN	1661
sfo02-w01-vds01-vmotion	Ephemeral	Elastic	8	VLAN	1662
sfo02-w01-vds01-vsan	Ephemeral	Elastic	8	VLAN	1663

- 5 On the **Ready to complete** page, review the configuration and click **Finish**.

- 6 Repeat this procedure to create the remaining port groups.

Connect the ESXi Hosts in Availability Zone 2 to the vSphere Distributed Switch for an NSX-T Workload Domain

In VMware Cloud Foundation, you add all ESXi hosts in Availability Zone 2 to the vSphere Distributed Switch for the workload domain so that the switch can start handling management traffic until you migrate all hosts to the N-VDS instance.

You configure the physical adapters for all ESXi hosts in Availability Zone 2.

Host	Physical Adapter	Uplink
sfo02w01esx01.sfo01.rainpole.local	vmnic1	Uplink2
sfo02w01esx02.sfo01.rainpole.local	vmnic1	
sfo02w01esx03.sfo01.rainpole.local	vmnic1	
sfo02w01esx04.sfo01.rainpole.local	vmnic1	

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Networking** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree.

- 3 Right-click **sfo01-w01-vds01**, and select **Add and manage hosts**.

The **Add and manage hosts** wizard opens.

- 4 On the **Select task** page, select **Add hosts** and click **Next**.

- 5 On the **Select hosts** page, click **New hosts**.
- 6 In the **Select new hosts** dialog box, select all hosts from Availability Zone 2 and click **OK**.
- 7 On the **Select hosts** page, click **Next**.
- 8 Add the second physical network adapter of the hosts to the switch.
 - a On the **Manage physical network adapters** page, select **vmnic1** and click **Assign uplink**.
 - b In the **Select an uplink** dialog box, select **Uplink2**.
 - c Select **Apply this uplink assignment to the rest of the hosts**, click **OK**.
 - d On the **Manage physical network adapters** page, click **Next**
- 9 Configure the VMkernel network adapters.
 - a On the **Manage VMkernel network adapters** page, click **vmk0**, click **Assign port group**.
 - b Select **sfo01-w01-vds01-management** and click **OK**.
 - c Select **Apply this port group assignment to the rest of the hosts** and click **OK**.
 - d On the **Manage VMkernel network adapters** page, click **Next**.
- 10 On the **Migrate VM Networking** page, click **Next**.
- 11 On the **Analyze impact** page, click **Next**.
- 12 On the **Ready to complete** page, review the configuration and click **Finish**.

Create the VMkernel Network Adapters for vSAN and vSphere vMotion in Availability Zone 2 for an NSX-T Workload Domain

In VMware Cloud Foundation, you create and configure the VMkernel network adapters for vSAN and vSphere vMotion traffic for the hosts in Availability Zone 2.

You repeat this procedure to create VMkernel network adapters on all ESXi hosts in Availability Zone 2.

Host	IPv4 Address for sfo02-w01-vds01-vmotion	IPv4 Address for sfo02-w01-vds01-vsan	VMkernel Gateway
sfo02w01esx01.sfo01.rainpole.local	172.16.62.101	172.16.63.101	172.16.62.253
sfo02w01esx02.sfo01.rainpole.local	172.16.62.102	172.16.63.102	
sfo02w01esx03.sfo01.rainpole.local	172.16.62.103	172.16.63.103	
sfo02w01esx04.sfo01.rainpole.local	172.16.62.104	172.16.63.104	

Prerequisites

For VMware Cloud Foundation 3.8, download the `vcf-stretch-cluster-patch.zip` patch for the SoS utility from VMware Knowledge Base article [71278](#) and update the SoS utility on the SDDC Manager appliance by following the instructions in the VMware Knowledge Base article.

The SoS utility in VMware Cloud Foundation 3.8.1 supports deploying NSX-T workload domains with multiple availability zones.

Procedure

- 1 Reserve the IP addresses for the VMkernel adapters for vSAN and vSphere vMotion traffic in the workload domain.
 - a Log in to the SDDC Manager appliance by using a Secure Shell (SSH) client.

Setting	Value
FQDN	sfo01sddcm01.sfo01.rainpole.local
User name	vcf
Password	<i>sddc_manager_vcf_password</i>

- b On the SDDC Manager appliance, run the `su` command to switch to the **root** account.

- c Run the following command for the host IP addresses for vSphere vMotion traffic and note down the reserved IP addresses from the command output .

```
./opt/vmware/sddc-support/sos --reserve-ips --network-pool sfo02w01-network-pool --network-type vmotion --ip-count 4
```

The command reserves IP addresses according to the configuration of the network pool for the workload domain. For example:

- 172.16.62.101
- 172.16.62.102
- 172.16.62.103
- 172.16.62.104

- d Run the following command for the host IP addresses for vSAN traffic and note down the reserved IP addresses from the command output.

```
./opt/vmware/sddc-support/sos --reserve-ips --network-pool sfo02w01-network-pool --network-type vsan --ip-count 4
```

The command reserves IP addresses according to the configuration of the network pool for the workload domain. For example:

- 172.16.63.101
- 172.16.63.102
- 172.16.63.103
- 172.16.63.104

- 2 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 3 Configure the VMkernel network adapters.

- a In the **Hosts and clusters** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree.
- b Select the **sfo02w01esx01.sfo01.rainpole.local** host and click the **Configure** tab.
- c In the **Networking** section, click **VMkernel adapters**.
- d On the **VMkernel adapters** page, click the **Add networking** button.

The **Add networking** wizard opens.

- e On the **Select connection type** page, select **VMkernel network adapter**, and click **Next**.

- f On the **Select target device** page, click **Select an existing network**, select the **sfo02-w01-vds01-vmotion** port group, and click **Next**.
- g On the **Port properties** page, configure the settings and click **Next**.

Setting	Value for sfo02-w01-vds01-vmotion	Value for sfo02-w01-vds01-vsan
IP settings	IPv4	IPv4
MTU	Get MTU from switch	Get MTU from switch
TCP/IP stack	vMotion	Default
Enabled services	-	vSAN

- h On the **IPv4 settings** page, configure the following settings and click **Next**.
Assign the IP addresses you have reserved in [Step 1](#).

Setting	Value for sfo02-w01-vds01-vmotion	Value for sfo02-w01-vds01-vsan
Use static IPv4 settings	Selected	Selected
IPv4 address	172.16.62.101	172.16.63.101
Subnet mask	255.255.255.0	255.255.255.0

- i On the **Ready to complete** page, click **Finish**.
 - j Repeat the step to add the VMkernel network adapter for vSAN traffic.
- 4 Configure the vMotion TCP/IP stack.
 - a In the **Networking** section, click **TCP/IP configuration**
 - b Select **vMotion** and click the **Edit** icon.
 - c On the **Routing** page, set **VMkernel gateway** to **172.16.62.253** and click **OK**.
 - d Repeat this step to configure the vMotion TCP/IP stack for the remaining ESXi hosts in Availability Zone 2.
 - 5 Repeat [Step 3](#) to [Step 4](#) to add the VMkernel network adapters on the other ESXi hosts.
 - 6 Update the inventory in SDDC Manager.
 - a On the SDDC Manager appliance, create a `update.inventory.json` file in the `/tmp` directory by using the sample in the `/tmp/vcf-stretch-cluster-patch` directory.
 - b Run the following command.

```
./opt/vmware/sddc-support/sos --update-domain-inventory /tmp/update.inventory.json
```

Complete the Migration of the Physical Network Adapters to the vSphere Distributed Switch for Availability Zone 2 for an NSX-T Workload Domain

In VMware Cloud Foundation, after you add all hosts in Availability Zone 2 to the vSphere Distributed Switch and configure the VMkernel adapters for vSphere vMotion and vSAN traffic, migrate the last physical network adapters to complete the networking configuration.

You repeat this procedure for all ESXi hosts in Availability Zone 2.

Host	Physical Adapter	Uplink
sfo02w01esx01.sfo01.rainpole.local	vmnic0	Uplink1
sfo02w01esx02.sfo01.rainpole.local	vmnic0	
sfo02w01esx03.sfo01.rainpole.local	vmnic0	
sfo02w01esx04.sfo01.rainpole.local	vmnic0	

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Networking** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree and expand the **sfo01-w01dc** data center.
- 3 Right-click **sfo01-w01-vds01**, and select **Add and manage hosts**.
The **Add and manage hosts** wizard opens.
- 4 On the **Select task** page, select **Attached hosts** and click **Next**.
- 5 On the **Select hosts** page, click **New hosts**.
- 6 On the **Select hosts** page, click **Attached hosts**.
- 7 On the **Select member hosts** page, select all ESXi hosts from Availability Zone 2, and click **OK**.
- 8 Migrate the last physical adapter from the standard switch to the sfo01-w01-vds01 distributed switch.
 - a On the **Manage physical adapters** page, select **vmnic0**, and click **Assign uplink**.
 - b In the **Select an Uplink** dialog box, select **Uplink1**, select **Apply this uplink assignment to the rest of the hosts**, and click **OK**.
 - c On the **Manage physical network adapters** page, click **Next**.
- 9 On the **Manage VMkernel adapter** page, click **Next**.

- 10 On the **Migrate VM networking** page, click **Next**.
- 11 On the **Ready to Complete** page, click **Finish**.
- 12 Remove the standard switch on the host.
 - a In the **Hosts and clusters** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree and expand the **sfo01-w01dc** data center.
 - b Select the **sfo02w01esx01.sfo01.rainpole.local** host and click the **Configure** tab.
 - c In the **Networking** section, click **Virtual switches**.
 - d Expand the **vSwitch0** section and click **Remove**.
 - e In the **Remove standard switch** dialog box, click **Yes**.
 - f Repeat this step to remove the standard switch for the remaining ESXi hosts in Availability Zone 2.

Enable SSH on the ESXi Hosts in Availability Zone 2 for an NSX-T Workload Domain

In VMware Cloud Foundation, you use SSH for configuring access to the ESXi hosts in Availability Zone 2 and for configuring static routes on the ESXi hosts for the vSAN stretched cluster.

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 Enable SSH.
 - a In the **Hosts and clusters** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree and expand the **sfo01-w01dc** data center.
 - b Expand the **sfo01-w01-shared01** cluster and select the **sfo01w01esx01.sfo01.rainpole.local** host.
 - c On the **Configure** tab, select **System > Services**.
 - d Select **SSH** and click the **Start** button.
 - e Click the **Edit startup policy** button, select **Start and stop with host**, and click **OK**.
- 3 Disable the SSH warning banner.
 - a On the **Configure** tab for the host, under **System**, select **Advanced system settings**.
 - b Click **Edit**.

- c In the **Filter** text box, enter **ssh**.
 - d Change the value of `UserVars.SuppressShellWarning` to **1** and click **OK**.
- 4 Repeat this procedure for the remaining ESXi hosts in Availability Zone 2.

Create vSAN Disk Groups for the ESXi Hosts in Availability Zone 2 for an NSX-T Workload Domain

In VMware Cloud Foundation, to extend the storage space of the vSAN datastore, create vSAN disk groups on each ESXi host in Availability Zone 2.

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Hosts and clusters** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree and expand the **sfo01-w01dc** data center.
- 3 Select the **sfo01-w01-shared01** cluster and click the **Configure** tab.
- 4 In the **vSAN** section, click **Disk management**.
- 5 Select **sfo02w01esx01.sfo01.rainpole.local** and click the **Create a new disk group** button.
- 6 In the **Create disk group** dialog box, select a flash disk for the **cache tier**, two hard disk drives for the **capacity Type**, and click **OK**.
- 7 Repeat this procedure for the remaining ESXi hosts in Availability Zone 2.
 - sfo02w01esx02.sfo01.rainpole.local
 - sfo02w01esx03.sfo01.rainpole.local
 - sfo02w01esx04.sfo01.rainpole.local

Add Static Routes for vSAN Traffic between the ESXi Host in Both Availability Zones for an NSX-T Workload Domain

In VMware Cloud Foundation, because vSphere uses a single default gateway, all routed traffic attempts to reach its destination through this gateway. Configure static routes on the ESXi hosts in both availability zones to route vSAN traffic using a dedicated gateway.

Table 5-2. ESXi Hosts of the NSX-T Workload Domain

Availability Zone	ESXi Host
Availability Zone 1	sfo01w01esx01.sfo01.rainpole.local
	sfo01w01esx02.sfo01.rainpole.local
	sfo01w01esx03.sfo01.rainpole.local
	sfo01w01esx04.sfo01.rainpole.local
Availability Zone 2	sfo02w01esx01.sfo01.rainpole.local
	sfo02w01esx02.sfo01.rainpole.local
	sfo02w01esx03.sfo01.rainpole.local
	sfo02w01esx04.sfo01.rainpole.local

Procedure

- 1 Configure static routes on the first ESXi host in Availability Zone 1.

- a Log in to the sfo01w01esx01 ESXi host by using a Secure Shell (SSH) client.

Setting	Value
FQDN	sfo01w01esx01.sfo01.rainpole.local
User name	root
Password	<i>esxi_root_user_password</i>

- b Add static route to the vSAN network in Availability Zone 2 by running this command.

```
esxcli network ip route ipv4 add -n 172.16.63.0/24 -g 172.16.43.253
```

- c Add static route to the network of the vSAN witness host by running this command.

```
esxcli network ip route ipv4 add -n 172.17.13.0/24 -g 172.16.43.253
```

- 2 Repeat [Step 1](#) on the other hosts in Availability Zone 1.

3 Configure static routes on the first ESXi host of Availability Zone 2.

- a Log in to the sfo02w01esx01.sfo01.rainpole.local ESXi host by using a Secure Shell (SSH) client.

Setting	Value
FQDN	sfo02w01esx01.sfo01.rainpole.local
User name	root
Password	<i>esxi_root_user_password</i>

- b Add static route to the vSAN network in Availability Zone 1 by running this command.

```
esxcli network ip route ipv4 add -n 172.16.43.0/24 -g 172.16.63.253
```

- c Add static route to the vSAN witness host witness network by running this command.

```
esxcli network ip route ipv4 add -n 172.17.13.0/24 -g 172.16.63.253
```

4 Repeat [Step 3](#) on the remaining ESXi hosts in Availability Zone 2.

Deploy and Configure the vSAN Witness Host for an NSX-T Workload Domain

6

vSAN stretched cluster requires a witness host deployed in a third location, different from the location of both availability zones.

In VMware Cloud Foundation, you deploy the vSAN witness appliance instead of using a dedicated physical ESXi host as a witness host. Unlike a general purpose ESXi host, the witness appliance does not run virtual machines and its only purpose is to serve as a vSAN witness.

Procedure

1 [Deploy the vSAN Witness Host in a Third Location for an NSX-T Workload Domain](#)

Start the deployment of multiple availability zones by deploying the vSAN witness host as a virtual appliance in a third location, such as an instance of VMware Cloud Foundation in another region.

2 [Configure the Management Network on the vSAN Witness Host for an NSX-T Workload Domain](#)

After the initial boot, use the ESXi Direct Console User Interface (DCUI) to configure the management network on the appliance of the vSAN witness host in VMware Cloud Foundation.

3 [Add the vSAN Witness Appliance as a Standalone Host to the NSX-T Workload Domain](#)

To use the witness appliance in the third location as a witness host for the vSAN stretched cluster in the NSX-T workload domain in VMware Cloud Foundation, you add the witness host as a standalone host in the vCenter Server instance for the workload domain.

4 [Configure NTP and SSH on the vSAN Witness Host for an NSX-T Workload Domain](#)

Prevent time synchronization problems by configuring the NTP service on the vSAN witness host in VMware Cloud Foundation. Enable SSH to add static routes to the networks of Availability Zone 1 and Availability Zone 2.

5 [Configure the Witness VMkernel Adapter on the vSAN Witness Host for an NSX-T Workload Domain](#)

To enable the vSAN data communication to both availability zones, configure the witness network on the vSAN witness host in VMware Cloud Foundation.

6 [Add Static Routes for the vSAN Witness Host for an NSX-T Workload Domain](#)

In VMware Cloud Foundation, configure routing for vSAN witness traffic so that the vSAN witness host can exchange system data with the management components for the workload domain.

Deploy the vSAN Witness Host in a Third Location for an NSX-T Workload Domain

Start the deployment of multiple availability zones by deploying the vSAN witness host as a virtual appliance in a third location, such as an instance of VMware Cloud Foundation in another region.

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Settings	Value
URL	https://lax01m01vc01.lax01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Hosts and clusters** inventory, expand the **lax01m01vc01.lax01.rainpole.local** tree and expand the **lax01-m01dc** data center.
- 3 Deploy the vSAN witness host virtual machine.
 - a Right-click the **lax01-m01-mgmt01** cluster and click **Deploy OVF template**.
The **Deploy OVF template** wizard opens.
 - b On the **Select an OVF template** page, select **Local file** and click **Choose files**.
 - c Select the vSAN witness host .ova file on your local file system, and click **Next**.
 - d On the **Select name and folder** page, configure the settings, and click **Next**.

Setting	Value
Virtual machine name	sfo03w01vsanw01
Virtual machine location	lax01-m01fd-mgmt

- e On the **Select a compute resource** page, select **lax01-m01-mgmt01**, and click **Next**.
- f On the **Review details** page, click **Next**.
- g On the **License agreements** page, select **I accept all license agreements** and click **Next**.
- h On the **Configuration** page, select **Medium**.
- i On the **Select storage** page, configure the settings, and click **Next**.

Setting	Value
Select virtual disk format	Thin provision
VM storage policy	vSAN Default Storage Policy
Datastore	lax01-m01-vsan01

- j On the **Select networks** page, configure the settings, and click **Next**.

Setting	Value
Witness network	lax01-m01-vds01-vsan
Management network	lax01-m01-vds01-management

- k On the **Customize template** page, configure the settings, and click **Next**.

Setting	Value
Root password	<i>vsan_witness_root_password</i>
Confirm root password	<i>vsan_witness_root_password</i>

- l On the **Ready to Complete** page, review your entries and click **Finish**.

- 4 In the **Hosts and clusters** inventory, expand the **lax01m01vc01.lax01.rainpole.local** tree and expand the **lax01-m01dc** data center.
- 5 Right-click the **sfo03w01vsanw01** virtual machine and select **Power > Power on**.

Configure the Management Network on the vSAN Witness Host for an NSX-T Workload Domain

After the initial boot, use the ESXi Direct Console User Interface (DCUI) to configure the management network on the appliance of the vSAN witness host in VMware Cloud Foundation.

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Settings	Value
URL	https://lax01m01vc01.lax01.rainpole.local/ui
User name	administrator@vsphere.local
Password	<i>vsphere_admin_password</i>

- 2 In the **Hosts and clusters** inventory, expand the **lax01m01vc01.lax01.rainpole.local** tree and expand the **lax01-m01dc** data center.
- 3 Open the Direct Console User Interface (DCUI) on the vSAN witness appliance.
 - a Right-click the **sfo03w01vsanw01** virtual machine, and click **Open remote console**.
 - b To enter the DCUI, press F2.
 - c Log in by using the following credentials.

Setting	Value
User name	root
Password	<i>vsan_witness_root_password</i>

4 Configure the management network.

- a Select **Configure management network** and press Enter.
- b Select **IPv4 configuration** and press Enter.
- c Configure the IPv4 settings, and press Enter.

Setting	Value
Set static IPv4 address and network configuration	Selected
IPv4 Address	172.17.11.203
Subnet Mask	255.255.255.0
Default Gateway	172.17.11.253

- d Select **DNS Configuration** and press Enter.
- e Configure the DNS settings and press Enter.

Setting	Value
Use the following DNS Server address and hostname	Selected
Primary DNS Server	172.17.11.5
Alternate DNS Server	172.17.11.4
Hostname	sfo03w01vsanw01.sfo01.rainpole.local

- f Select **Custom DNS Suffixes** and press Enter.
- g Verify that the list contains no suffixes and press Enter.

5 Press Escape to exit and press Y to confirm the changes.

6 Close the remote console to the vSAN witness appliance.

Add the vSAN Witness Appliance as a Standalone Host to the NSX-T Workload Domain

To use the witness appliance in the third location as a witness host for the vSAN stretched cluster in the NSX-T workload domain in VMware Cloud Foundation, you add the witness host as a standalone host in the vCenter Server instance for the workload domain.

Procedure

1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Hosts and clusters** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree.
- 3 Right-click the **sfo01-w01dc** data center and select **Add host**.

The **Add host** wizard appears.

- 4 On the **Name and location** page, configure the settings and click **Next**.

Setting	Value
Host name or IP address	sfo03w01vsanw01.sfo01.rainpole.local
Location	sfo01-w01dc

- 5 On the **Connection settings** page, enter the following credentials, and click **Next**.

Setting	Value
User name	root
Password	<i>vsan_witness_root_password</i>

- 6 In the **Security alert** dialog box, click **Yes**.
- 7 On the **Host summary** page, review the host configuration and click **Next**.
- 8 On the **Assign license** page, use the default license and click **Next**.
- 9 On the **Lockdown mode** page, click **Next**.
- 10 On the **VM location** page, click **Next**.
- 11 On the **Ready to complete** page, click **Finish**.

Configure NTP and SSH on the vSAN Witness Host for an NSX-T Workload Domain

Prevent time synchronization problems by configuring the NTP service on the vSAN witness host in VMware Cloud Foundation. Enable SSH to add static routes to the networks of Availability Zone 1 and Availability Zone 2.

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	<i>vsphere_admin_password</i>

- 2 In the **Hosts and clusters** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** and expand the **sfo01-w01dc** data center.
- 3 Select the **sfo03w01vsanw01.sfo01.rainpole.local** host and click the **Configure** tab.

- 4 Configure the NTP client on the vSAN witness host.
 - a Under the **System** section, click **Time configuration** and click the **Edit** button.
The **Edit time configuration** dialog box opens.
 - b Configure the settings and click **OK**.

Setting	Value
Use Network Time Protocol (enable NTP client)	Selected
NTP Servers	ntp.lax01.rainpole.local, ntp.sfo01.rainpole.local
Start NTP Service	Selected
NTP Service Startup Policy	Start and stop with host

- 5 Enable the SSH service on the vSAN witness host.
 - a Under the **System** section, select **Services** and select **SSH**.
 - b Click the **Edit startup policy** button.
 - c On the **Edit startup policy** page, select **Start and stop with host** and click **OK**.

Configure the Witness VMkernel Adapter on the vSAN Witness Host for an NSX-T Workload Domain

To enable the vSAN data communication to both availability zones, configure the witness network on the vSAN witness host in VMware Cloud Foundation.

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Hosts and clusters** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree and expand the **sfo01-w01dc** data center.
- 3 Select the **sfo03w01vsanw01.sfo01.rainpole.local** host, and click the **Configure** tab.
- 4 Under the **Networking** section, select **VMkernel adapters**.
- 5 Select the **vmk1** VMkernel adapter and click **Edit**.
The network label of the adapter is witnessPg.

6 In the **vmk1 - edit settings** dialog box, click **IPv4 settings**, configure the settings, and click **OK**.

Setting	Value
Use static IPv4 settings.	Selected
IPv4 address	172.17.13.203
Subnet mask	255.255.255.0
Override default gateway for this adapter	Deselected

Add Static Routes for the vSAN Witness Host for an NSX-T Workload Domain

In VMware Cloud Foundation, configure routing for vSAN witness traffic so that the vSAN witness host can exchange system data with the management components for the workload domain.

Procedure

1 Log in to the appliance of the vSAN witness host by using a Secure Shell (SSH) client.

Setting	Value
FQDN	sfo03w01vsanw01.sfo01.rainpole.local
User name	root
Password	<i>vsan_witness_root_password</i>

2 Add a static route to the vSAN network in Availability Zone 1 by running this command.

```
esxcli network ip route ipv4 add -n 172.16.43.0/24 -g 172.17.13.253
```

3 Add a static route to the vSAN network in Availability Zone 2 by running this command.

```
esxcli network ip route ipv4 add -n 172.16.63.0/24 -g 172.17.13.253
```

Configure vSAN Stretched Cluster for an NSX-T Workload Domain

7

In VMware Cloud Foundation, after preparing all the ESXi hosts and the network for multiple availability zones, you turn on and configure vSAN stretched cluster in the NSX-T workload domain.

Procedure

1 [Enable vSAN Stretched Cluster on the NSX-T Workload Domain Cluster](#)

In VMware Cloud Foundation, after you configure Availability Zone 2 and deploy the vSAN witness appliance in a third location, turn on the vSAN stretched cluster functionality, and set fault domains and the vSAN witness host.

2 [Turn on the vSAN Performance Service for the Cluster for an NSX-T Workload Domain](#)

In VMware Cloud Foundation, to monitor the performance of the vSAN stretched cluster, hosts, disks, and VMs, turn on the vSAN performance service.

3 [Update the vSAN Storage Policy for an NSX-T Workload Domain](#)

In VMware Cloud Foundation, to tolerate an availability zone failure, update the vSAN storage policy of the domain cluster for a secondary level of failures.

4 [Update the vSphere Availability Settings of the Cluster for an NSX-T Workload Domain](#)

In VMware Cloud Foundation, after you configure the vSAN stretched cluster, you change the admission control policy to reserve 50% of the capacity to support recovery if an availability zone failure occurs.

Enable vSAN Stretched Cluster on the NSX-T Workload Domain Cluster

In VMware Cloud Foundation, after you configure Availability Zone 2 and deploy the vSAN witness appliance in a third location, turn on the vSAN stretched cluster functionality, and set fault domains and the vSAN witness host.

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Hosts and clusters** inventory, expand **sfo01w01vc01.sfo01.rainpole.local**.

- 3 Select the **sfo01-w01-shared01** cluster and click the **Configure** tab.

- 4 In the **vSAN** section, select **Fault Domains**.

- 5 In the **Stretched cluster** card, click **Configure**.

The **Configure stretched cluster** wizard opens.

- 6 On the **Configure fault domains** page, configure the settings and click **Next**.

Setting	Value
Preferred domain	Preferred
Secondary domain	Secondary
Preferred domain members	-
Secondary domain members	<ul style="list-style-type: none"> ■ sfo01w01esx01.sfo01.rainpole.local ■ sfo01w01esx02.sfo01.rainpole.local ■ sfo01w01esx03.sfo01.rainpole.local ■ sfo01w01esx04.sfo01.rainpole.local

- 7 On the **Select witness host** page, expand the **sfo01-w01dc** center in the **sfo01w01vc01.sfo01.rainpole.local** tree, select the **sfo03w01vsanw01.sfo01.rainpole.local** host, and click **Next**.

- 8 On the **Claim disks for witness host** page, select the disks for each tier, click **Next**, and click **Finish**.

Setting	Value
Local VMware Disk (mpx.vmhba1:C0:T2:L0)	Cache tier
Local VMware Disk (mpx.vmhba1:C0:T1:L0)	Capacity tier

Turn on the vSAN Performance Service for the Cluster for an NSX-T Workload Domain

In VMware Cloud Foundation, to monitor the performance of the vSAN stretched cluster, hosts, disks, and VMs, turn on the vSAN performance service.

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Hosts and clusters** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree and expand the **sfo01-w01dc** data center.
- 3 Select the **sfo01-w01-shared01** cluster and click the **Configure** tab.
- 4 Under the **vSAN** section, select **Services**.
- 5 In the **Performance service** section, click **Edit**.
- 6 In the **vSAN Performance Service Settings** dialog box, configure the settings and click **Apply**.

Setting	Value
Enable vSAN Performance Service	On
Storage policy	vSAN Default Storage Policy
Verbose mode	Deselected
Network diagnostic mode	Deselected

Update the vSAN Storage Policy for an NSX-T Workload Domain

In VMware Cloud Foundation, to tolerate an availability zone failure, update the vSAN storage policy of the domain cluster for a secondary level of failures.

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 Set the default storage policy for the domain vSAN datastore to the vSAN default storage policy.
 - a In the **Storage** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree and expand the **sfo01-w01dc** data center.
 - b Select the **sfo01-w01-shared01-vsan01** datastore and click the **Configure** tab.

- c Select **General** and next to **Default Storage Policy** click **Edit**.
 - d Select **vSAN Default Storage Policy** and click **OK**.
- 3 Turn on the dual-site support in the vSAN default storage policy.
- a Select **Menu > Policies and Profiles**.
 - b In the **Policies and Profiles** pane, select **VM Storage Policies**.
 - c Select the **vSAN Default Storage Policy** entry for the sfo01w01vc01.sfo01.rainpole.local vCenter Server instance and click **Edit Settings**.
 - d On the **Name and description** page, click **Next**.
 - e On the **vSAN** page, click the **Availability** tab, configure the settings, and click **Next**.

Setting	Value
Site disaster tolerance	Dual site mirroring (stretched cluster)
Failures to tolerate	1 failure - RAID-1 (Mirroring)

- f On the **Storage compatibility** page, click **Next**.
- g On the **Review and finish** page, click **Finish**.

Update the vSphere Availability Settings of the Cluster for an NSX-T Workload Domain

In VMware Cloud Foundation, after you configure the vSAN stretched cluster, you change the admission control policy to reserve 50% of the capacity to support recovery if an availability zone failure occurs.

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Hosts and clusters** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree and expand the **sfo01-w01dc** data center.
- 3 Select the **sfo01-w01-shared01** cluster, and click the **Configure** tab.
- 4 In the **Services** section, select **vSphere Availability**, and, in the **vSphere HA** section, click **Edit**.
- 5 Click the **Admission control** tab and set **Host failures cluster tolerates** to 4.

6 Click the **Advanced options** tab, click **Add**, configure the settings, and click **OK**.

Setting	Value
das.isolationaddress0	172.16.43.253
das.isolationaddress1	172.16.63.253
das.usedefaultisolationaddress	false

Configure the NSX-T Instance for an NSX-T Workload Domain

8

NSX-T Manager implements both the management and central control planes in an NSX-T system. For dynamic routing between the tenant workloads in the domain, in VMware Cloud Foundation, you deploy a pair of NSX-T Edge nodes.

NSX-T Manager also provides the user interface and REST APIs for creating, configuring, and monitoring NSX-T components in a workload domain, such as segments, gateways, and security policies.

For high availability of the management and control planes, VMware Cloud Foundation deploys a cluster of three NSX-T Manager nodes.

Procedure

1 [Create the Transport Zones for Uplink Traffic for an NSX-T Workload Domain](#)

In VMware Cloud Foundation, after you deploy the NSX-T Manager cluster, create the transport zones for uplink traffic for the NSX-T Edge nodes in the workload domain. SDDC Manager creates the transport zones for management and overlay traffic when you add an NSX-T workload domain to your SDDC.

2 [Create Uplink Profiles for the ESXi Hosts in Availability Zone 2 for an NSX-T Workload Domain](#)

Uplink profiles define the policies for the links from ESXi hosts to NSX-T segments or from NSX-T Edge nodes to top of rack switches. Because you add the hosts for Availability Zone 2 manually to the workload domain in VMware Cloud Foundation, for integration as transport nodes in the NSX-T configuration, you must provide an uplink profiles for these hosts and NSX-T Edge node.

3 [Create the NSX-T Segments for System, Uplink, and Overlay Traffic for an NSX-T Workload Domain](#)

Create the segments to connect nodes that send VLAN and overlay traffic in the NSX-T workload domain in VMware Cloud Foundation.

4 [Configure the ESXi Host Transport Nodes in Availability Zone 2 for an NSX-T Workload Domain](#)

When you configure the hosts in Availability Zone 2 as NSX-T transport nodes, NSX-T Manager installs the NSX-T kernel modules on the hosts as VIB files. Because you add manually the hosts in Availability Zone 2,

5 Remove the ESXi Hosts for Availability Zone 2 from the vSphere Distributed Switch for an NSX-T Workload Domain

After you configure the ESXi hosts in Availability Zone 2 as transport nodes, the NSX-T infrastructure starts handling the system and virtual machine traffic to the hosts. You can remove the hosts from the vSphere Distributed Switch for the availability zone.

6 Configure Dynamic Routing for an NSX-T Workload Domain

IN VMware Cloud Foundation, to support the communication between tenant workloads by using application virtual networks in NSX-T and to connect tenant workloads to the external network, configure dynamic routing in the initial cluster of the NSX-T workload domain.

Create the Transport Zones for Uplink Traffic for an NSX-T Workload Domain

In VMware Cloud Foundation, after you deploy the NSX-T Manager cluster, create the transport zones for uplink traffic for the NSX-T Edge nodes in the workload domain. SDDC Manager creates the transport zones for management and overlay traffic when you add an NSX-T workload domain to your SDDC.

Table 8-1. Transport Zones for the NSX-T Edge Nodes in a Workload Domain

Name	N-VDS Name	N-VDS Mode	Traffic Type
sfo01-w-uplink01	sfo01-w-uplink01	Standard	VLAN
sfo01-w-uplink02	sfo01-w-uplink02	Standard	VLAN

Procedure

- 1 In a Web browser, log in to the user interface of NSX-T Manager.

Setting	Value
URL	https://sfo01wnsx01.sfo01.rainpole.local
User name	admin
Password	<i>nsx_admin_password</i>

- 2 On the main navigation bar, click **System**.
- 3 Navigate to **Fabric > Transport zones** and click **Add**.
- 4 On the **New transport zone** page, configure the settings for the first transport zone and click **Add**.

Setting	Value
Name	sfo01-w-uplink01
N-VDS Name	sfo01-w-uplink01
N-VDS Mode	Standard
Traffic Type	VLAN

- 5 Repeat the previous step to create the remaining transport zones.

Create Uplink Profiles for the ESXi Hosts in Availability Zone 2 for an NSX-T Workload Domain

Uplink profiles define the policies for the links from ESXi hosts to NSX-T segments or from NSX-T Edge nodes to top of rack switches. Because you add the hosts for Availability Zone 2 manually to the workload domain in VMware Cloud Foundation, for integration as transport nodes in the NSX-T configuration, you must provide an uplink profiles for these hosts and NSX-T Edge node.

Procedure

- 1 In a Web browser, log in to the user interface of NSX-T Manager.

Setting	Value
URL	https://sfo01wnsx01.sfo01.rainpole.local
User name	admin
Password	nsx_admin_password

- 2 On the main navigation bar, click **System**.
- 3 In the navigation pane, select **Fabric > Profiles**.
- 4 To define policies for the links in Availability Zone 2 between the ESXi hosts and segments and between NSX-T Edge nodes and top of rack switches, create uplink profiles.
 - a On the **Profiles** page, click the **Uplink profiles** tab and click **Add**.
 - b On the **New uplink profile** page, enter the following values and click **Add**.

Name	Teaming - Teaming Policy	Teaming - Active Uplinks	Transport VLAN	MTU
sfo02-esxi-w01-uplink-profile	Load Balance Source	uplink-1,uplink-2	1664	9000
sfo02-w-overlay-profile	Failover Order	uplink-1	1669	9000
sfo02-w-uplink01-profile	Failover Order	uplink-1	1667	9000
sfo02-w-uplink02-profile	Failover Order	uplink-1	1668	9000

- 5 In the **sfo02-esxi-w01-uplink-profile** profile, create two teaming policies for ECMP uplinks.
 - a On the **Uplink profiles** tab, select the **sfo02-esxi-w01-uplink-profile** profile and click **Edit**.
 - b In the **Edit uplink profile** dialog box, under **Teamings**, click the **Add** button, enter the following information, and click **Save**.

Name	Teaming Policy	Active Uplinks
Uplink01	Failover Order	uplink-1
Uplink02	Failover Order	uplink-2

6 Include the new teaming policies in the transport zone.

- a Navigate to **Fabric > Transport zones**, select **sfo-esxi-vlan**, and click **Edit**.
- b In the **Edit Transport Zone** dialog box, add **Uplink01** and **Uplink02** to **Uplink teaming policy names**, and click **Save**.

Create the NSX-T Segments for System, Uplink, and Overlay Traffic for an NSX-T Workload Domain

Create the segments to connect nodes that send VLAN and overlay traffic in the NSX-T workload domain in VMware Cloud Foundation.

You create each segment and assign each uplink segments an uplink port.

Table 8-2. NSX-T Logical Networks in the NSX-T Workload Domain

Segment Name	Uplink & Type	Transport Zone	VLAN
sfo01-w-nvds01-uplink01	None	sfo-esxi-vlan	0-4094
sfo01-w-nvds01-uplink02	None	sfo-esxi-vlan	0-4094
sfo01-w-uplink01	None	sfo-w-uplink01	1647
sfo01-w-uplink02	None	sfo-w-uplink02	1648
sfo02-w-nvds01-management	None	sfo-esxi-vlan	1661
sfo02-w-nvds01-vmotion	None	sfo-esxi-vlan	1662
sfo02-w-nvds01-vsan	None	sfo-esxi-vlan	1663
sfo02-w-nvds01-uplink01	None	sfo-esxi-vlan	0-4094
sfo02-w-nvds01-uplink02	None	sfo-esxi-vlan	0-4094
sfo02-w-uplink01	None	sfo-w-uplink01	1667
sfo02-w-uplink02	None	sfo-w-uplink02	1668
sfo-w-overlay	None	sfo-esxi-vlan	0-4094

Procedure

1 In a Web browser, log in to the user interface of NSX-T Manager.

Setting	Value
URL	https://sfo01wnsx01.sfo01.rainpole.local
User name	admin
Password	nsx_admin_password

2 On the main navigation bar, click **Networking**.

3 In the navigation pane, select **Segments**.

4 On the **Segments** tab, click **Add Segment**.

5 Configure the values for the sfo02-w-nvds01-management segment and click **Save**.

Setting	Value
Name	sfo02-w-nvds01-management
Uplink & Type	None
Transport Zone	sfo01-esxi-vlan
VLAN	1661

6 Repeat this step to create the remaining segments.

7 Change the teaming policy for the uplink segments.

- On the main navigation bar, click **Advanced Networking and Security**.
- Under **Networking**, click **Switching**.
- Select the uplink segment, click **Edit** and configure the following uplink teaming policy.

Uplink Segment	Uplink teaming policy name
sfo01-w-nvds01-uplink01	Uplink01
sfo01-w-nvds01-uplink02	Uplink02
sfo02-w-nvds01-uplink01	Uplink01
sfo02-w-nvds01-uplink02	Uplink02

Configure the ESXi Host Transport Nodes in Availability Zone 2 for an NSX-T Workload Domain

When you configure the hosts in Availability Zone 2 as NSX-T transport nodes, NSX-T Manager installs the NSX-T kernel modules on the hosts as VIB files. Because you add manually the hosts in Availability Zone 2,

The NSX-T kernel modules provide services such as distributed routing and distributed firewall. You configure the same settings on all ESXi hosts in Availability Zone 2.

Table 8-3. Configuration of the ESXi Hosts in Availability Zone 2

Setting	Value
Transport Zone	<ul style="list-style-type: none"> ■ sfo-esxi-vlan ■ sfo-w-overlay
N-VDS Name	sfo-w-nvds01
NIOC Profile	sfo-w-nioc-profile
Uplink Profile	sfo02-esxi-w01-uplink-profile
LLDP Profile	LLDP [Send Packet Enabled]
IP Assignment	Use DHCP

Table 8-3. Configuration of the ESXi Hosts in Availability Zone 2 (continued)

Setting	Value
Physical NICs	<ul style="list-style-type: none"> ■ vmnic0 - uplink-1 ■ vmnic1 - uplink-2
vmk0	sfo02-w-nvds01-management
vmk1	sfo02-w-nvds01-vmotion
vmk2	sfo02-w-nvds01-vsan

Procedure

- 1 In a Web browser, log in to the user interface of NSX-T Manager.

Setting	Value
URL	https://sfo01wnsx01.sfo01.rainpole.local
User name	admin
Password	<i>nsx_admin_password</i>

- 2 On the main navigation bar, click **System**.
- 3 In the navigation pane, select **Fabric > Nodes**.
- 4 On the **Host transport nodes** tab, from the **Managed by** drop-down menu, select **sfo01w01vc01**.
- 5 Expand the **sfo01-w01-shared01** cluster.
- 6 Select the **sfo01w01esx01.sfo01.rainpole.local** ESXi host and click **Configure NSX**.
 - a On the **Host details** page, click **Next**.
 - b On the **Configure NSX** page, configure the settings and click **Next**.

Setting	Value
Transport Zone	<ul style="list-style-type: none"> ■ sfo-esxi-vlan ■ sfo-w-overlay
N-VDS Name	sfo-w-nvds01
NIOC Profile	sfo-w-nioc-profile
Uplink Profile	sfo02-esxi-w01-uplink-profile
LLDP Profile	LLDP [Send Packet Enabled]
IP Assignment	Use DHCP
Physical NICs	<ul style="list-style-type: none"> ■ vmnic0 - uplink-1 ■ vmnic1 - uplink-2

- c On the **Network mappings to install** page, click **Add mapping**, configure the settings, and click **Add**.

VMK	Segment
vmk0	sfo02-w-nvds01-management
vmk1	sfo02-w-nvds01-vmotion
vmk2	sfo02-w-nvds01-vsan

- d Click **Finish**.

- 7 Repeat this procedure to configure the remaining ESXi hosts in Availability Zone 2.

- sfo01w01esx02.sfo01.rainpole.local
- sfo01w01esx03.sfo01.rainpole.local
- sfo01w01esx04.sfo01.rainpole.local

Remove the ESXi Hosts for Availability Zone 2 from the vSphere Distributed Switch for an NSX-T Workload Domain

After you configure the ESXi hosts in Availability Zone 2 as transport nodes, the NSX-T infrastructure starts handling the system and virtual machine traffic to the hosts. You can remove the hosts from the vSphere Distributed Switch for the availability zone.

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Networking** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree and expand the **sfo01-w01dc** data center.
- 3 Right-click the **sfo01-w01-vds01** vSphere Distributed Switch and select **Add and manage hosts**.
- 4 In the **sfo01-w01-vds01 - Add and manage hosts** wizard, select **Remove hosts** and click **Next**.
- 5 Click **Attached hosts**, select all hosts in Availability Zone 2, click **OK**, and click **Next**.
- 6 On the **Ready to complete** page, click **Finish**.

Configure Dynamic Routing for an NSX-T Workload Domain

IN VMware Cloud Foundation, to support the communication between tenant workloads by using application virtual networks in NSX-T and to connect tenant workloads to the external network, configure dynamic routing in the initial cluster of the NSX-T workload domain.

Routing occurs in both the North-South and East-West directions.

- North-South traffic leaving or entering the workload domain, for example, a virtual machine on an overlay network communicating with an end-user device on the corporate network.
- East-West traffic remains in the workload domain, for example, two virtual machines on the same or different segments communicating with each other.

Procedure

1 [Create an NSX-T Edge Cluster Profile for an NSX-T Workload Domain](#)

In VMware Cloud Foundation, for availability of the routing services and connectivity to the external network, you create a multi-node cluster of NSX-T Edge nodes. To define a common configuration for NSX-T Edge nodes, you create an edge cluster profile.

2 [Deploy the NSX-T Edge Appliances for an NSX-T Workload Domain](#)

In VMware Cloud Foundation, to provide tenant workloads with routing, services, and connectivity to networks that are external to the workload domain, deploy a pair of NSX-T Edge nodes in each availability zone.

3 [Join the NSX-T Edge Nodes to the Management Plane for an NSX-T Workload Domain](#)

In VMware Cloud Foundation, after you deploy the NSX-T Edge appliances in the NSX-T workload domain cluster, to connect them to the NSX-T Manager cluster, join them to the management plane.

4 [Create an Anti-Affinity Rule for the NSX-T Edge Nodes for an NSX-T Workload Domain](#)

To ensure that the two NSX-T Edge appliances run on different ESXi hosts in VMware Cloud Foundation, create a vSphere DRS VM-host anti-affinity rule. If a failure occurs on one of the hosts, the appliance on the other host continues providing routing services.

5 [Create Host Groups and Rules for Both Availability Zones for an NSX-T Workload Domain](#)

To ensure that all virtual machines that in an availability zone run on ESXi hosts in the same zone in VMware Cloud Foundation, you create and configure host and virtual machine groups rules in vSphere DRS .

6 [Add the NSX-T Edge Nodes to the Transport Zones for an NSX-T Workload Domain](#)

After you deploy the NSX-T Edge nodes and join them to the management plane in VMware Cloud Foundation, add them to the transport zones for uplink and overlay traffic, and configure the N-VDS switches on each edge node.

7 Create the NSX-T Edge Cluster for an NSX-T Workload Domain

Adding multiple NSX-T Edge nodes to a cluster increases the availability of networking services. An NSX-T Edge cluster is necessary to support the Tier-0 and Tier-1 gateways in the workload domain in VMware Cloud Foundation.

8 Create and Configure the Tier-0 Gateway for an NSX-T Workload Domain

The Tier-0 gateway in the NSX-T Edge cluster provides a gateway service between the logical and physical network in VMware Cloud Foundation. The NSX-T Edge cluster can back multiple Tier-0 gateways.

9 Create and Configure the Tier-1 Gateway for an NSX-T Workload Domain

To redistribute routes to the Tier-0 gateway and to provide routing between tenant workloads in VMware Cloud Foundation, create and configure the Tier-1 gateway.

10 Verify BGP Peering and Route Redistribution for an NSX-T Workload Domain

The Tier-0 gateway must establish a connection to each of the upstream Layer 3 devices in its availability zone before BGP updates can be exchanged. Verify that the NSX-T Edge nodes are successfully peering and that BGP routing is established in VMware Cloud Foundation.

Create an NSX-T Edge Cluster Profile for an NSX-T Workload Domain

In VMware Cloud Foundation, for availability of the routing services and connectivity to the external network, you create a multi-node cluster of NSX-T Edge nodes. To define a common configuration for NSX-T Edge nodes, you create an edge cluster profile.

Procedure

- 1 In a Web browser, log in to the user interface of NSX-T Manager.

Setting	Value
URL	https://sfo01wnsx01.sfo01.rainpole.local
User name	admin
Password	nsx_admin_password

- 2 On the main navigation bar, click **System**.
- 3 In the navigation pane, select **Fabric > Profiles**.
- 4 On the **Edge cluster profiles** tab, click **Add**.
- 5 On the **New Edge cluster profile** page, configure the settings and click **Add**.

Setting	Value
Name	sfo-w-edge-cluster01-profile
BFD Probe	1000

Setting	Value
BFD Allowed Hops	255
BFD Declare Dead Multiple	3

Deploy the NSX-T Edge Appliances for an NSX-T Workload Domain

In VMware Cloud Foundation, to provide tenant workloads with routing, services, and connectivity to networks that are external to the workload domain, deploy a pair of NSX-T Edge nodes in each availability zone.

You repeat this procedure to deploy four NSX-T Edge nodes.

Table 8-4. NSX-T Edge Nodes

Setting	Value for sfo01wesg01	Value for sfo01wesg02	Value for sfo02wesg01	Value for sfo02wesg02
Network 0	sfo01-w-nvds01-management	sfo01-w-nvds01-management	sfo02-w-nvds01-management	sfo02-w-nvds01-management
Network 1	sfo-w-overlay	sfo-w-overlay	sfo-w-overlay	sfo-w-overlay
Network 2	sfo01-w-nvds01-uplink01	sfo01-w-nvds01-uplink01	sfo02-w-nvds01-uplink01	sfo02-w-nvds01-uplink01
Network 3	sfo01-w-nvds01-uplink02	sfo01-w-nvds01-uplink02	sfo02-w-nvds01-uplink02	sfo02-w-nvds01-uplink02
Management IP address	172.16.41.21	172.16.41.22	172.16.61.21	172.16.61.22
Default gateway	172.16.41.253	172.16.41.253	172.16.61.253	172.16.61.253

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Hosts and Clusters** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree and expand the **sfo01-w01dc** data center.
- 3 Expand the **sfo01-w01-shared01** cluster.
- 4 Right-click the **sfo01-w01rp-sddc-edge** resource pool and select **Deploy OVF template**.
- 5 On the **Deploy OVF template** page, navigate to the .ova file of the NSX-T Edge appliance and click **Next**.

6 On the **Select name and location** page, configure the settings and click **Next**.

Setting	Value
Name	sfo01wesg01
Folder or data center	sfo01-w01fd-nsx

7 On the **Select a resource** page, select the **sfo01-w01rp-sddc-edge** resource pool and click **Next**.

8 On the **Select storage** page, select **sfo-w01-vsan** and click **Next**.

9 On the **Select networks** page, configure the settings and click **Next**.

Source Network	Destination Network
Network 3	sfo01-w-nvds01-uplink02
Network 2	sfo01-w-nvds01-uplink01
Network 1	sfo-w-overlay
Network 0	sfo01-w-nvds01-management

10 On the **Customize template** page, configure the settings, and click **Next**.

Setting	Value
System Root User Password / Confirm Password	<i>nsx_edge_root_password</i>
CLI "admin" User Password / Confirm Password	<i>nsx_edge_admin_password</i>
CLI "audit" User Password / Confirm Password	<i>nsx_edge_audit_password</i>
Hostname	sfo01wesg01.sfo01.rainpole.local
Default IPv4 Gateway	172.16.41.253
Management Network IPv4 Address	172.16.41.21
Management Network Netmask	255.255.255.0
DNS Server List	<ul style="list-style-type: none"> ■ 172.16.11.5 ■ 172.16.11.4
Domain Search List	sfo01.rainpole.local
NTP Server List	ntp.sfo01.rainpole.local
Enable SSH	Selected
Allow root SSH login	Deselected

11 On the **Ready to complete** page, click **Finish**.

12 After the deployment finishes, power on the NSX-T Edge appliance.

- a In the **VMs and Templates** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree.
- b Expand the sfo01-w01fd-nsx folder, right-click the **sfo01wesg01** virtual machine, and select **Power > Power On**.

13 Repeat this procedure to deploy the other NSX-T Edge appliances.

Join the NSX-T Edge Nodes to the Management Plane for an NSX-T Workload Domain

In VMware Cloud Foundation, after you deploy the NSX-T Edge appliances in the NSX-T workload domain cluster, to connect them to the NSX-T Manager cluster, join them to the management plane.

Table 8-5. NSX Edge Nodes

Setting	Value for sfo01wesg01	Value for sfo01wesg02	Value for sfo02wesg01	Value for sfo02wesg02
Name	sfo01wesg01	sfo01wesg02	sfo02wesg01	sfo02wesg02
Primary IP Address	172.16.41.21	172.16.41.22	172.16.61.21	172.16.61.22

Procedure

- 1 Log in to the first NSX-T Manager node by using Secure Shell (SSH) client.

Setting	Value
FQDN	sfo01wnsx01a.sfo01.rainpole.local
User name	admin
Password	<i>nsx_admin_password</i>

- 2 Retrieve the thumbprint ID of the certificate for the NSX-T Manager cluster by running and copying the output from the following command.

```
get certificate cluster thumbprint
```

- 3 Log in to the first NSX-T Edge node by using Secure Shell (SSH) client.

Setting	Value
FQDN	sfo01wesg01.sfo01.rainpole.local
User name	admin
Password	<i>edge_admin_password</i>

- 4 Join the NSX-T Edge node to the management plane by running the following command.

```
join management-plane sfo01wnsx01.sfo01.rainpole.local thumbprint thumbprint_id username admin
```

- 5 Enter the password for the **admin** account.
- 6 Repeat [Step 3](#) to [Step 5](#) for the remaining NSX-T Edge appliances.

Create an Anti-Affinity Rule for the NSX-T Edge Nodes for an NSX-T Workload Domain

To ensure that the two NSX-T Edge appliances run on different ESXi hosts in VMware Cloud Foundation, create a vSphere DRS VM-host anti-affinity rule. If a failure occurs on one of the hosts, the appliance on the other host continues providing routing services.

Because you allocate an NSX-T Edge pair to each availability zone, you configure a rule for each zone. You repeat this procedure twice to create the anti-affinity rules for both availability zones.

Setting	Value for Availability Zone 1	Value for Availability Zone 2
Name	anti-affinity-rule-az1-ecmpedges	anti-affinity-rule-az2-ecmpedges
Enable rule	Selected	Selected
Type	Separate Virtual Machines	Separate Virtual Machines
Members	<ul style="list-style-type: none"> ■ sfo01wesg01 ■ sfo01wesg02 	<ul style="list-style-type: none"> ■ sfo02wesg01 ■ sfo02wesg02

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Hosts and Clusters** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree and expand the **sfo01-w01dc** tree.
- 3 Select the **sfo01-w01-shared01** cluster and click the **Configure** tab.
- 4 In the **Configuration** section, select **VM/Host rules** and click **Add**.
- 5 In the **sfo01-w01-shared01- Create VM/Host rule** dialog box, configure the settings for Availability Zone 1 and click **Add**.

Setting	Value
Name	anti-affinity-rule-az1-ecmpedges
Enable rule	Selected
Type	Separate Virtual Machines

- 6 In the **Add rule member** dialog box, select **sfo01wesg01** and **sfo01wesg02**, and click **OK**.
- 7 In the **sfo01-w01-shared01- Create VM/Host rule** dialog box, click **OK**.
- 8 Repeat this procedure to create the anti-affinity rule for Availability Zone 2.

Create Host Groups and Rules for Both Availability Zones for an NSX-T Workload Domain

To ensure that all virtual machines that in an availability zone run on ESXi hosts in the same zone in VMware Cloud Foundation, you create and configure host and virtual machine groups rules in vSphere DRS .

Procedure

- 1 In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://sfo01m01vc01.sfo01.rainpole.local/ui
User name	administrator@vsphere.local
Password	vsphere_admin_password

- 2 In the **Hosts and Clusters** inventory, expand the **sfo01w01vc01.sfo01.rainpole.local** tree.

- 3 Select the **sfo01-w01-shared01** cluster and click the **Configure** tab.

- 4 Create host groups, each containing the ESXi hosts in the availability zone.

- a In the **Configuration** section, click **VM/Host groups** and click **Add**.
- b In the **Create VM/Host group** dialog box, configure the settings and click **OK**.

Setting	Value for Availability Zone 1	Value for Availability Zone 2
Name	availability-zone-1-hosts	availability-zone-2-hosts
Type	Host group	Host group
Members	<ul style="list-style-type: none"> ■ sfo01w01esx01.sfo01.rainpole.local ■ sfo01w01esx02.sfo01.rainpole.local ■ sfo01w01esx03.sfo01.rainpole.local ■ sfo01w01esx04.sfo01.rainpole.local 	<ul style="list-style-type: none"> ■ sfo02w01esx01.sfo01.rainpole.local ■ sfo02w01esx02.sfo01.rainpole.local ■ sfo02w01esx03.sfo01.rainpole.local ■ sfo02w01esx04.sfo01.rainpole.local

- c Repeat this step to create the host group for the other availability zone.

- 5 Create virtual machine groups, each containing the two edge nodes in the availability zone.

- a In the **Configuration** section, click **VM/Host groups** and click **Add**.
- b In the **Create VM/Host group** dialog box, configure the settings and click **OK**.

Setting	Value for Availability Zone 1	Value for Availability Zone 2
Name	availability-zone-1-vms	availability-zone-2-vms
Type	VM group	VM group
Members	<ul style="list-style-type: none"> ■ sfo01wesg01 ■ sfo01wesg02 	<ul style="list-style-type: none"> ■ sfo02wesg01 ■ sfo02wesg02

- c Repeat this step to create the virtual machine group for the other availability zone

- 6 Create rules for running virtual machines on hosts in their initial availability zone.

- a In the **Configuration** section, click **VM/Host rules**.
- b On the **VM/Host rules** page, click the **Add** button.

- c In the **Create VM/Host rule** dialog box, enter the settings, and click **OK**.

Setting	Value for Availability Group 1	Value for Availability Group 2
Name	hostgroup-availability-zone-1	hostgroup-availability-zone-2
Enable rule	Selected	Selected
Type	Virtual Machines to Hosts	Virtual Machines to Hosts
VM group	availability-zone-1-vms	availability-zone-2-vms
Action	Should run on hosts in group.	Should run on hosts in group.
Host group	availability-zone-1-hosts	availability-zone-2-hosts

- d Repeat this step to create the rule for the other availability zone.

Add the NSX-T Edge Nodes to the Transport Zones for an NSX-T Workload Domain

After you deploy the NSX-T Edge nodes and join them to the management plane in VMware Cloud Foundation, add them to the transport zones for uplink and overlay traffic, and configure the N-VDS switches on each edge node.

You repeat this procedure to configure all four NSX-T edge transport nodes.

Setting	Value for sfo01wesg01	Value for sfo01wesg02	Value for sfo02wesg01	Value for sfo02wesg02
Management IP address	172.16.41.21	172.16.41.22	172.16.61.21	172.16.61.22
Transport zones	<ul style="list-style-type: none"> ■ sfo-w-overlay ■ sfo-w-uplink01 ■ sfo-w-uplink02 	<ul style="list-style-type: none"> ■ sfo-w-overlay ■ sfo-w-uplink01 ■ sfo-w-uplink02 	<ul style="list-style-type: none"> ■ sfo-w-overlay ■ sfo-w-uplink01 ■ sfo-w-uplink02 	<ul style="list-style-type: none"> ■ sfo-w-overlay ■ sfo-w-uplink01 ■ sfo-w-uplink02
N-VDS	3	3	3	3

Procedure

- 1 In a Web browser, log in to the user interface of NSX-T Manager.

Setting	Value
URL	https://sfo01wnsx01.sfo01.rainpole.local
User name	admin
Password	<i>nsx_admin_password</i>

- 2 On the main navigation bar, click **System**.
- 3 In the navigation pane, select **Fabric > Nodes**.
- 4 Click the **Edge Transport Nodes** tab.
- 5 Select the **sfo01wesg01** edge node and click **Configure NSX**.
- 6 On the **Edit transport node - sfo01wesg01** dialog box, click the **General** tab.

7 In the **Transport zones** section, move the transport zones to the **Selected** list and click **Add**.

Setting	Value for sfo01wesg01	Value for sfo01wesg02	Value for sfo02wesg01	Value for sfo02wesg02
Transport Zones	sfo-w-uplink01(VLAN)	sfo-w-uplink01(VLAN)	sfo-w-uplink01(VLAN)	sfo-w-uplink01(VLAN)
	sfo-w-uplink02(VLAN)	sfo-w-uplink02(VLAN)	sfo-w-uplink02(VLAN)	sfo-w-uplink02(VLAN)
	sfo-w-overlay(Overlay)	sfo-w-overlay(Overlay)	sfo-w-overlay(Overlay)	sfo-w-overlay(Overlay)

8 In the **Edit transport node - sfo01wesg01** dialog box, click the **N-VDS** tab.

9 In the **New node switch**, configure the settings.

Setting	Value for sfo01wesg01	Value for sfo01wesg02	Value for sfo02wesg01	Value for sfo02wesg02
Edge Switch Name	sfo-w-nvds01	sfo-w-nvds01	sfo-w-nvds01	sfo-w-nvds01
Uplink Profile	sfo01-w-overlay-profile	sfo01-w-overlay-profile	sfo02-w-overlay-profile	sfo02-w-overlay-profile
IP Assignment	Use Static IP List	Use Static IP List	Use Static IP List	Use Static IP List
Static IP List	172.16.49.21	172.16.49.22	172.16.69.21	172.16.69.22
Gateway	172.16.49.253	172.16.49.253	172.16.69.253	172.16.69.253
Subnet Mask	255.255.255.0	255.255.255.0	255.255.255.0	255.255.255.0
Virtual NICs: fp-eth0	uplink-1	uplink-1	uplink-1	uplink-1

10 Click **Add N-VDS**, in the **New node switch** section configure the settings for the Uplink-1 transport nodes.

Setting	Value for sfo01wesg01	Value for sfo01wesg02	Value for sfo02wesg01	Value for sfo02wesg02
Edge Switch Name	sfo01-w-nvds01-uplink01	sfo01-w-nvds01-uplink01	sfo02-w-nvds01-uplink01	sfo02-w-nvds01-uplink01
Uplink Profile	sfo01-w-uplink01-profile	sfo01-w-uplink01-profile	sfo02-w-uplink01-profile	sfo02-w-uplink01-profile
IP Assignment	-	-	-	-
Virtual NICs: fp-eth1	uplink-1	uplink-1	uplink-1	uplink-1

- 11 Click **Add N-VDS**, in the **New node switch section** configure the settings for the Uplink-2 transport nodes, and click **Save**.

Setting	Value for sfo01wesg01	Value for sfo01wesg02	Value for sfo02wesg01	Value for sfo02wesg02
Edge Switch Name	sfo01-w-nvds01-uplink02	sfo01-w-nvds01-uplink02	sfo02-w-nvds01-uplink02	sfo02-w-nvds01-uplink02
Uplink Profile	sfo01-w-uplink02-profile	sfo01-w-uplink02-profile	sfo02-w-uplink02-profile	sfo02-w-uplink02-profile
IP Assignment	-	-	-	-
Virtual NICs: fp-eth2	uplink-2	uplink-2	uplink-2	uplink-2

- 12 Repeat this procedure to configure the other NSX-T Edge nodes.

Create the NSX-T Edge Cluster for an NSX-T Workload Domain

Adding multiple NSX-T Edge nodes to a cluster increases the availability of networking services. An NSX-T Edge cluster is necessary to support the Tier-0 and Tier-1 gateways in the workload domain in VMware Cloud Foundation.

Procedure

- 1 In a Web browser, log in to the user interface of NSX-T Manager.

Setting	Value
URL	https://sfo01wnsx01.sfo01.rainpole.local
User name	admin
Password	nsx_admin_password

- 2 On the main navigation bar, click **System**.
- 3 In the navigation pane, select **Fabric > Nodes**.
- 4 On the **Edge clusters** tab, click **Add**.
- 5 In the **Add edge cluster** dialog box, configure the settings and click **Add**.

Setting	Value
Name	sfo-w-edge-cluster01
Edge cluster profile	sfo-w-edge-cluster01-profile
Member type	Edge Node
Members	<ul style="list-style-type: none"> ■ sfo01wesg01.sfo01.rainpole.local ■ sfo01wesg02.sfo01.rainpole.local ■ sfo02wesg01.sfo01.rainpole.local ■ sfo02wesg02.sfo01.rainpole.local

Create and Configure the Tier-0 Gateway for an NSX-T Workload Domain

The Tier-0 gateway in the NSX-T Edge cluster provides a gateway service between the logical and physical network in VMware Cloud Foundation. The NSX-T Edge cluster can back multiple Tier-0 gateways.

Procedure

- 1 In a Web browser, log in to the user interface of NSX-T Manager.

Setting	Value
URL	https://sfo01wnsx01.sfo01.rainpole.local
User name	admin
Password	nsx_admin_password

- 2 On the main navigation bar, click **Networking**.
- 3 In the navigation pane, click **Tier-0 gateways**.
- 4 Create the Tier-0 gateway.
 - a On the **Tier-0 gateways** page, click **Add Tier-0 gateway**.
 - b Configure the settings and click **Save**.

Setting	Value
Name	sfo01-w-tier0-01
High Availability Mode	Active-Active
Edge Cluster	sfo-w-edge-cluster01

- 5 Configure route re-distribution.
 - a Expand **Route Re-Distribution** and click **Set**.
 - b On the **Set route re-distribution** dialog box, select all **Tier-0 subnets** and **Advertised Tier-1 subnet** sources, and click **Apply**.
 - c On the **Add Tier-0 gateway** page, in the **Route re-distribution** section, click **Save**.

6 Add the uplink interfaces to the NSX-T Edge nodes.

- a Expand **Interfaces** and click **Set**.
- b In the **Set interfaces** dialog box, click **Add interface**, configure the settings, and click **Save**.

Name	Type	IP Address / Mask	Connected to (Segment)	Edge Node	MTU
sfo01wesg01-Uplink01	External	172.16.47.2/24	sfo01-w-uplink01	sfo01wesg01	9000
sfo01wesg01-Uplink02	External	172.16.48.2/24	sfo01-w-uplink02	sfo01wesg01	9000
sfo01wesg02-Uplink01	External	172.16.47.3/24	sfo01-w-uplink01	sfo01wesg02	9000
sfo01wesg02-Uplink02	External	172.16.48.3/24	sfo01-w-uplink02	sfo01wesg02	9000
sfo02wesg01-Uplink01	External	172.16.67.2/24	sfo02-w-uplink01	sfo02wesg01	9000
sfo02wesg01-Uplink02	External	172.16.68.2/24	sfo02-w-uplink02	sfo02wesg01	9000
sfo02wesg02-Uplink01	External	172.16.67.3/24	sfo02-w-uplink01	sfo02wesg02	9000
sfo02wesg02-Uplink02	External	172.16.68.3/24	sfo02-w-uplink02	sfo02wesg02	9000

- c Repeat this step for the remaining interfaces and click **Close**.
- d On the **Add Tier-0 gateway** page, in the **Interfaces** section, click **Save**.

7 Create an IP prefix list.

- a Expand **Routing**.
- b In the **IP prefix list** section, click **Set**.
- c In the **Set IP prefix list** dialog box, click **Add IP prefix list**.
- d Enter **Any-Prefix** as the **Name** and click **Set**.
- e In the **Set prefixes** dialog box, click **Add Prefix**, configure the settings, click **Add** and click **Apply**.

Setting	Value
Network	any
Action	Permit

- f In the **Set IP prefix list** dialog box, click **Save** and click **Close**.

8 Create a route map.

- a Expand **Routing**.
- b In the **Route maps** section, click **Set**.
- c In the **Set route maps** dialog box, click **Add route map**.
- d Enter **az2-route-map** for **Name**.
- e In the **Match criteria** column, click **Set**.
- f On the **Set match criteria** dialog box, click **Add match criteria**, configure the settings, click **Add** and click **Apply**.

Setting	Value
Type	IP Prefix
Members	Any-Prefix
AS path prepend	65000 65000

- g In the **Set route maps** dialog box, click **Save** and click **Close**.

9 Configure BGP.

- a Expand **BGP**, configure the settings, and click **Save**.

Setting	Value
Local AS	65000
BGP	On
Graceful Restart	Off
Inter SR iBGP	On
ECMP	On
Multipath Relax	On

- b Click **Set** for **BGP neighbors**.

- c In the **Set BGP neighbors** dialog box, click **Add BGP neighbor**, configure the settings for the first neighbour, and click **Save**.

IP Address	BFD	Remote AS	Hold Down Time	Keep Alive Time	Password	Out Filter	In Filter
172.16.47.1	Disabled	65001	12	4	<i>bgp_password</i>	-	-
172.16.48.1	Disabled	65001	12	4	<i>bgp_password</i>	-	-
172.16.67.1	Disabled	65002	12	4	<i>bgp_password</i>	az2-route-map	az2-route-map
172.16.68.1	Disabled	65002	12	4	<i>bgp_password</i>	az2-route-map	az2-route-map

Note Enable BFD if the network supports and is configured for BFD.

- d Repeat for the other neighbor, click **Save** and click **Close**.
- e On the **Add Tier-0 gateway** page, in the **BGP** section, click **Close editing**.
- 10** Generate a BGP summary for the Tier-0 gateway.

- a On the main navigation bar, click **Advanced networking & security**.
- b In the navigation pane, click **Routers** and select **sfo-w-tier0-01**.
- c From the **Actions** drop-down menu, select **Generate BGP summary**.
- d Verify that each transport node has five established connections: one to each neighbor in its availability zone and one to each of the other three transport nodes.

Create and Configure the Tier-1 Gateway for an NSX-T Workload Domain

To redistribute routes to the Tier-0 gateway and to provide routing between tenant workloads in VMware Cloud Foundation, create and configure the Tier-1 gateway.

Tier-1 gateways have downlink ports to connect to NSX-T segments and uplink ports to connect to NSX-T Tier-0 gateways.

Procedure

- 1** In a Web browser, log in to the user interface of NSX-T Manager.

Setting	Value
URL	<code>https://sfo01wnsx01.sfo01.rainpole.local</code>
User name	<code>admin</code>
Password	<code>nsx_admin_password</code>

- 2** On the main navigation bar, click **Networking**.

3 In the navigation pane, click **Tier-1 gateways**.

4 Create the Tier-1 gateway.

- a On the **Tier-1 gateways** page, click **Add Tier-1 Gateway**.
- b Enter the values.

Setting	Value
Name	sfo-w-tier1-01
Linked Tier-0 Gateway	sfo-w-tier0-01
Failover	Non Preemptive
Edge Cluster	sfo-w-edge-cluster01

- c In the **Edges** section, click **Set**.
- d In the **Select edges** dialog box, click **Add edge**.
- e Add the **sfo01wesg01** and sfo02wesg01 edge nodes.
- f Ensure that sfo01wesg01 is the first edge in the sequence and click **Apply**.
- g On the **Tier-1 gateways** page, click **Save**.

5 Expand **Route Advertisement**, enable all route advertisement types, and click **Save**.

6 On the **Tier-1 gateways** page, click **Close editing**.

Verify BGP Peering and Route Redistribution for an NSX-T Workload Domain

The Tier-0 gateway must establish a connection to each of the upstream Layer 3 devices in its availability zone before BGP updates can be exchanged. Verify that the NSX-T Edge nodes are successfully peering and that BGP routing is established in VMware Cloud Foundation.

Table 8-6. NSX-T Edge Nodes

Availability Zone	NSX-T Edge Node
Availability Zone 1	sfo01wesg01
	sfo01wesg02
Availability Zone 2	sfo02wesg01
	sfo02wesg02

Procedure

- 1 Log in to sfo01wesg01 NSX-T Edge node by using a Secure Shell (SSH) client.

Setting	Value
FQDN	sfo01wesg01
User name	admin
Password	<i>nsx_edge_admin_password</i>

- 2 Get information about the Tier-0 and Tier-1 service routers and distributed router by running the command.

```
get logical-router
```

The output of the command might contains the following configuration.

UUID	VRF	LR-ID	Name	Type	Ports
<i>sample_uuid</i>	0	0	-	TUNNEL	3
<i>sample_uuid</i>	1	5	SR-tier0-01	SERVICE_ROUTE R_TIER0	6
<i>sample_uuid</i>	2	2	DR-tier1-01	DISTRIBUTED_R OUTER_TIER1	5
<i>sample_uuid</i>	3	3	DR-tier0-01	DISTRIBUTED_R OUTER_TIER0	4
<i>sample_uuid</i>	4	11	SR-tier1-01	SERVICE_ROUTE R_TIER1	5

- 3 By using the VRF value for SERVICE_ROUTER_TIER0 connect to the service router for Tier 0.

```
vrf 1
```

The prompt changes to sfo01wesg01(tier0_sr)>. All commands are associated with this object.

- 4 Verify the BGP connections to the neighbors of the service router for Tier 0.

```
get bgp neighbor summary
```

The BGP State for each neighbor in the edge node availability zone, and each edge node, appears as Established.

- 5 Verify that you are receiving routes by using BGP and that multiple routes to BGP-learned networks exist.

```
get route
```

The routing table contains routes beginning with b, learned via BGP.

- 6 Repeat this procedure for the other NSX-T Edge nodes.

Register the vSAN Stretched Cluster in SDDC Manager for an NSX-T Workload Domain

9

After you configure manually an NSX-T workload domain with a second availability zone in VMware Cloud Foundation by using a vSAN stretched cluster, register the cluster as stretched in the SDDC Manager database by using the SoS utility.

Procedure

- 1 Log in to the SDDC Manager appliance by using a Secure Shell (SSH) client.

Setting	Value
FQDN	sfo01sddcm01.sfo01.rainpole.local
User name	vcf
Password	<i>sddc_manager_vcf_password</i>

- 2 On the SDDC Manager appliance, run the `su` command to switch to the **root** account.
- 3 Run this command .

```
/opt/vmware/sddc-support/sos --enable-stretch-cluster-flag sfo01-w01-shared01 --domain-name sfo
```