VMware Cloud on AWS Networking and Security

01 June 2021

SDDC Version 1.14

VMware Cloud on AWS

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

https://docs.vmware.com/

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

Contents

About VMware Cloud on AWS Networking and Security 5

1 NSX-T Networking Concepts 6

Features Supported with NSX-T 11

2 Configuring VMware Cloud on AWS Networking and Security Using NSX-T 13

Assign NSX Service Roles to Organization Members 14

Configure AWS Direct Connect Between Your SDDC and On-Premises Data Center 15

Set Up an AWS Direct Connect Connection 16

Configure Direct Connect to a Private Virtual Interface for SDDC Management and Compute Network Traffic 17

Configure Direct Connect to a Public Virtual Interface for Access to AWS Services 21

Specify the Direct Connect MTU 21

Configure a VPN Connection Between Your SDDC and On-Premises Data Center 22

Create a Route-Based VPN 23

Create a Policy-Based VPN 26

Configure a Layer 2 VPN and Extended Network Segment 29

View VPN Tunnel Status and Statistics 33

IPsec VPN Settings Reference 33

Configure Management Gateway Networking and Security 35

Set vCenter Server FQDN Resolution Address 35

Set HCX FQDN Resolution Address 35

Add or Modify Management Gateway Firewall Rules 36

Configure Compute Gateway Networking and Security 39

Create or Modify a Network Segment 40

Add or Modify Compute Gateway Firewall Rules 46

Add or Modify Distributed Firewall Rules 49

Configure DNS Services 54

View Routes Learned and Advertised over VMware Transit Connect 55

Configure a Multi-Edge SDDC With Traffic Groups 56

Working With Inventory Groups 60

Add a Management Group 60

Add or Modify a Compute Group 61

Add a Custom Service 62

View Virtual Machine Inventory 63

Managing Workload Connections 64

Attach a VM to or Detach a Workload VM from a Compute Network Segment 64

Request or Release a Public IP Address 65

Create or Modify NAT Rules 66

Creating Firewall Rules to Manage Traffic Between the Compute and Management Networks 67

3 Configure Monitoring and Troubleshooting Features 69

Configure IPFIX 69

Configure Port Mirroring 70

View Connected VPC Information 72

About VMware Cloud on AWS Networking and Security

The *VMware Cloud on AWS Networking and Security* guide provides information about configuring NSX-T networking and security for VMware Cloud on AWS.

Intended Audience

This information is intended for anyone who wants to use VMware Cloud on AWS to create an SDDC that has the networking and security infrastructure necessary to migrate workloads off premises and run them securely in the cloud. It was written for readers who have used vSphere in an on-premises environment and are familiar with the fundamentals of IP networking using NSX-T or another networking solution. In-depth knowledge of vSphere or Amazon Web Services is not required.

For a detailed discussion of how VMware Cloud on AWS uses NSX-T networking, see the VMware Press eBook VMware Cloud on AWS: NSX Networking and Security.

NSX-T Networking Concepts

1

VMware Cloud on AWS uses NSX-T to create and manage internal SDDC networks and provide endpoints for VPN connections from your on-premises network infrastructure.

SDDC Network Topology

When you create an SDDC, it includes a Management Network. Single-host trial SDDCs also include a small Compute Network. You specify the Management Network CIDR block when you create the SDDC. It cannot be changed after the SDDC has been created. See Deploy an SDDC from the VMC Console for details. The Management Network has two subnets:

Appliance Subnet

This subnet is used by the vCenter, NSX, and HCX appliances in the SDDC. When you add appliance-based services such as SRM to the SDDC, they also connect to this subnet.

Infrastructure Subnet

This subnet is used by the ESXi hosts in the SDDC.

The Compute Network includes an arbitrary number of logical segments for your workload VMs. See VMware Configuration Maximums for current limits on logical segments. In a Single Host SDDC starter configuration, we create a compute network with a single routed segment. In SDDC configurations that have more hosts, you'll have to create compute network segments to meet your needs. See VMware Configuration Maximums for applicable limits.

An SDDC network has two notional tiers:

- Tier 0 handles north-south traffic (traffic leaving or entering the SDDC, or between the Management and Compute gateways).
- Tier 1 handles east-west traffic (traffic between routed network segments within the SDDC).

Figure 1-1. SDDC Network Topology

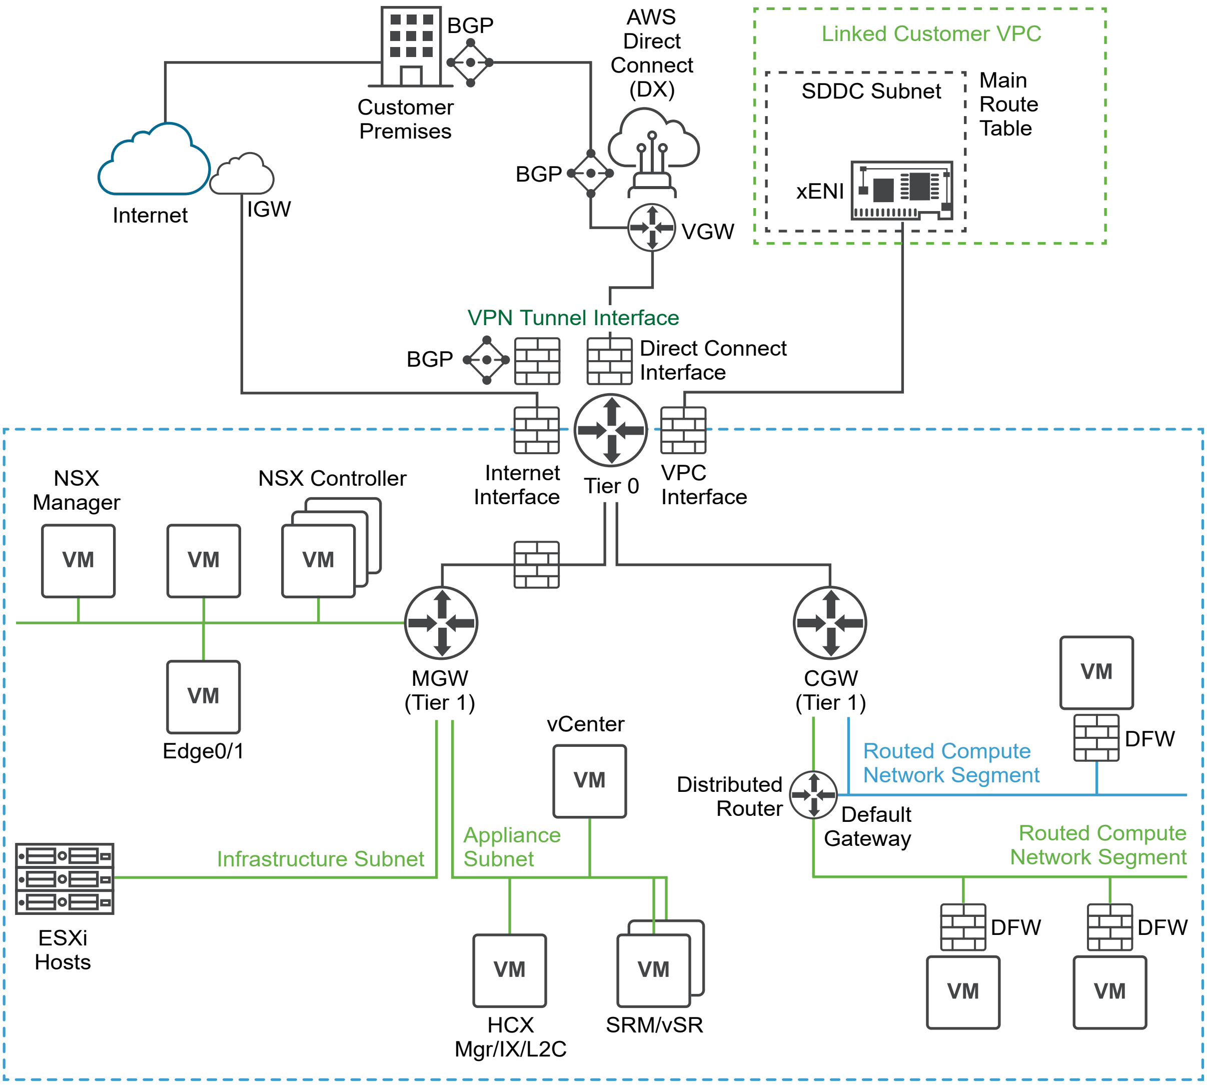

NSX Edge Appliance

The default NSX Edge Appliance is implemented as a pair of VMs that run in active/standby mode. This appliance provides the platform on which the default Tier 0 and Tier 1 routers run, along with IPsec VPN connections and their BGP routing machinery. All north-south traffic goes through the Tier 0 router. To avoid sending east-west traffic through the Edge Appliance, a component of each Tier 1 router runs on every ESXi host that handles routing for destinations within the SDDC.

If you need additional bandwidth for the subset of this traffic routed to SDDC group members, a Direct Connect Gateway attached to an SDDC group, HCX Service Mesh, or to the connected VPC, you can reconfigure your SDDC to be Multi-Edge by creating traffic groups, each of which creates an additional T0 router. See Configure a Multi-Edge SDDC With Traffic Groups for details.

Note VPN traffic, as well as DX traffic to a private VIF must pass through on the default T0 and cannot be routed to a non-default traffic group. In addition, because NAT rules always run on the default T0 router, additional T0 routers cannot handle traffic affected by SNAT or DNAT rules. This includes traffic to and from the SDDC's native Internet connection. It also includes traffic to the Amazon S3 service, which uses a NAT rule and must go through the default T0.

Management Gateway (MGW)

The MGW is a Tier 1 router that handles routing and firewalling for the vCenter Server and other management appliances running in the SDDC. Management gateway firewall rules run on the MGW and control access to management VMs. In the default configuration, these rules block all inbound traffic to the management network (see Add or Modify Management Gateway Firewall Rules).

Compute Gateway (CGW)

The CGW is a Tier 1 router that handles network traffic for workload VMs connected to routed compute network segments. Compute gateway firewall rules, along with NAT rules, run on the Tier 0 router. In the default configuration, these rules block all traffic to and from compute network segments (see Configure Compute Gateway Networking and Security).

Routing Between Your SDDC and the Connected VPC

Important Any VPC Subnets on which AWS services or instances communicate with the SDDC must be associated with the main route table of the connected VPC. Use of a custom route table or replacement of the main route table is not supported.

When you create an SDDC, we pre-allocate 17 AWS Elastic Network Interfaces (ENIs) in the selected VPC owned by the AWS account you specify at SDDC creation. We assign each of these ENIs an IP address from the subnet you specify at SDDC creation, then attach each of the hosts in the SDDC cluster `Cluster-1` to one of these ENIs. An additional IP address is assigned to the ENI where the active NSX Edge Appliance is running.

This configuration, known as the Connected VPC, supports network traffic between VMs in the SDDC and AWS instances and native AWS service endpoints in the Connected VPC. The main route table of the connected VPC is aware of the VPC’s primary subnet as well as all SDDC (NSX-T network segment) subnets. When you create or delete routed network segments on the SDDC, the main route table is automatically updated. When the NSX Edge Appliance in your SDDC is moved to another host, either to recover from a failure or during SDDC maintenance, the IP address allocated to the Edge Appliance is moved to the new ENI (on the new host), and the main route table is updated to reflect the change. If you have replaced the main route table or are using a custom route table, that update fails and network traffic can no longer be routed between SDDC networks and the Connected VPC. See View Connected VPC Information for more about how to use the VMC Console to see the details of your connected VPC.

For an in-depth discussion of SDDC network architecture and the AWS network objects that support it, read the VMware Cloud Tech Zone article VMware Cloud on AWS: SDDC Network Architecture .

Reserved Network Addresses

Certain IPv4 address ranges are unavailable for use in SDDC compute networks. Several are used internally by SDDC network components. Most are reserved by convention on other networks as well.

■ 10.0.0.0/15 ■ 172.31.0.0/16	These ranges are reserved within the SDDC management subnet, but can be used in your on-premises networks or SDDC compute network segments.
100.64.0.0/16	Reserved for carrier-grade NAT per RFC 6598. Avoid using addresses in this range in SDDC networks and others. They are not likely to be reachable within the SDDC or from outside it. See VMware Knowledge Base article 76022 for a detailed breakdown of how SDDC networks use this address range.
■ 169.254.0.0/19 ■ 169.254.64.0/24 ■ 169.254.101.0/30 ■ 169.254.105.0/24 ■ 169.254.106.0/24	Per RFC 3927, all of 169.254.0.0/16 is a link-local range that cannot be routed beyond a single subnet. However, with the exception of these CIDR blocks, you can use 169.254.0.0/16 addresses for your virtual tunnel interfaces. See Create a Route-Based VPN.
192.168.1.0/24	This the default compute segment CIDR for a single-host starter SDDC and is not reserved in other configurations.

SDDC networks also observe the conventions for special Use IPv4 address ranges enumerated in RFC 3330.

Multicast Support in SDDC Networks

In SDDC networks, layer 2 multicast traffic is treated as broadcast traffic on the network segment where the traffic originates. It is not routed beyond that segment. Layer 2 multicast traffic optimization features such as IGMP snooping are not supported. Layer 3 multicast (such as Protocol Independent Multicast) is not supported in VMware Cloud on AWS.

Connecting Your On-Premises SDDC to Your Cloud SDDC

To connect your on-premises data center to your VMware Cloud on AWS SDDC, you can create a VPN that uses the public Internet, a VPN that uses AWS Direct Connect, or just use AWS Direct Connect alone. You can also take advantage of SDDC groups to use VMware Transit Connect™ and an AWS Direct Connect Gateway to provide centralized connectivity between a group of VMware Cloud on AWS SDDCs and an on-premises SDDC. See Creating and Managing SDDC Deployment Groups in the *VMware Cloud on AWS Operations Guide*.

Figure 1-2. SDDC Connections to your On-Premises Data Center

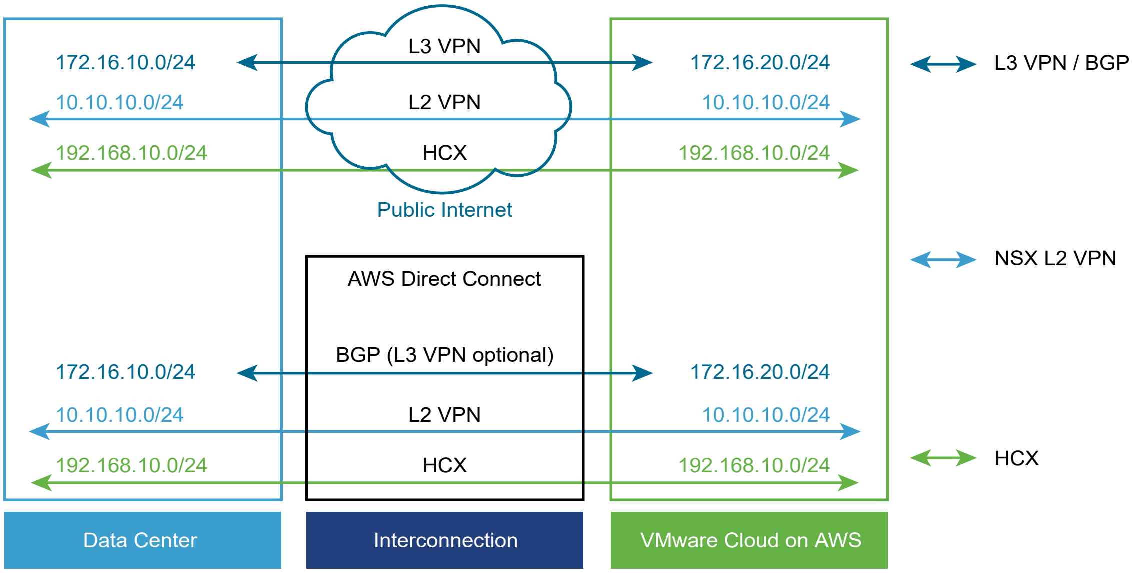

Layer 3 (L3) VPN

A layer 3 VPN provides a secure connection between your on-premises data center to your VMware Cloud on AWS SDDC over the public Internet or AWS Direct Connect. These IPsec VPNs can be either route-based or policy-based. You can create up to sixteen VPNs of each type, using any on-premises router that supports the settings listed in the IPsec VPN Settings Reference as the on-premises endpoint.

Layer 2 (L2) VPN

A layer 2 VPN provides an extended, or stretched, network with a single IP address space that spans your on-premises data center and your SDDC and enables hot or cold migration of on-premises workloads to the SDDC. You can create only a single L2VPN tunnel in any SDDC. The on-premises end of the tunnel requires NSX. If you are not already using NSX in your on-premises data center, you can download a standalone NSX Edge appliance to provide the required functionality. An L2 VPN can connect your on-premises data center to the SDDC over the public Internet or over AWS Direct Connect.

AWS Direct Connect (DX)

AWS Direct Connect is a service provided by AWS that allows you to create a high-speed, low latency connection between your on-premises data center and AWS services. When you configure AWS Direct Connect, VPNs can use it instead of routing traffic over the public Internet. Because Direct Connect implements Border Gateway Protocol (BGP) routing, use of an L3VPN for the management network is optional when you configure Direct Connect. Traffic over Direct Connect is not encrypted. If you want to encrypt that traffic, you can configure an IPsec VPN that uses private IP addresses and Direct Connect.

VMware HCX

VMware HCX, a multi-cloud app mobility solution, is provided free to all SDDCs and facilitates migration of workload VMs to and from your on-premises data center to your SDDC. For more information about installing, configuring, and using HCX, see the Hybrid Migration with HCX Checklist.

This chapter includes the following topics:

- Features Supported with NSX-T

Features Supported with NSX-T

SDDCs backed by NSX-T support a wide range of networking and security solutions.

NSX-T was designed specifically to support diverse data center environments at scale and provide robust capabilities for containers and the cloud.

Note NSX-T Configuration Maximums are now included in VMware Configuration Maximums.

Networking and Connectivity Features

NSX-T provides all the networking capabilities required by workloads running in the SDDC. These capabilities allow you to:

- Deploy networks (L2, L3, and isolated) and define subnets and gateways for the workloads that will reside there.
  - L2VPNs extend your on-premises L2 domains to the SDDC, enabling workload migration without IP address changes.
  - Route-based IPsec VPNs can connect to on-premises networks, VPCs, or other SDDCs. Route-based VPNs use BGP to learn new routes as networks become available.
  - Policy-based IPsec VPNs can also be used to connect to on-premises networks, VPCs, or other SDDCs.
  - Isolated networks have no uplinks, and provide access only to those VMs connected to them.
- Use AWS Direct Connect (DX) to carry traffic between on-premises and SDDC networks over high bandwidth, low latency connectivity. You can optionally use a route-based VPN as backup for DX traffic.

- Enable native DHCP selectively for network segments or use DHCP relay to link with an on-premises IPAM solution.
- Create multiple DNS zones, allowing use of different DNS servers for network subdomains.
- Take advantage of distributed routing, managed by an NSX kernel module running on the host where the workload resides, so workloads can efficiently communicate with each other.

Security Features

NSX-T security features include network address translation (NAT) and advanced firewall capabilities.

- Source NAT (SNAT) is automatically applied to all workloads in the SDDC to enable Internet access. To provide a secure environment, Internet access is blocked at edge firewalls, but firewall policy can be changed to allow managed access. You can also request a public IP for workloads and create custom NAT policies for them.
- Edge firewalls run on the management and compute gateways. These stateful firewalls examine all traffic into and out of the SDDC.
- Distributed Firewall (DFW) is a stateful firewall that runs on all SDDC hosts. It provides protection for traffic within the SDDC and enables micro-segmentation to allow fine-grained control over traffic between workloads.

Network Operations Tools

NSX-T also provides several popular network operations management tools.

- Port mirroring can send mirrored traffic from a source to a destination appliance in the SDDC or your on-premises network.
- IPFIX supports segment-specific network traffic analysis by sending traffic flows to an IPFIX collector.

Configuring VMware Cloud on AWS Networking and Security Using NSX-T

2

Follow this workflow to configure NSX-T networking and security in your SDDC.

Procedure

1 Assign NSX Service Roles to Organization Members

   Grant users in your organization an NSX service role to allow them to view or configure features on the Networking & Security tab.

2 Configure AWS Direct Connect Between Your SDDC and On-Premises Data Center

   Use of AWS Direct Connect is optional. If traffic between your on-premises network and your SDDC workloads requires higher speeds and lower latency than you can achieve with a connection over the public Internet, configure VMware Cloud on AWS to use AWS Direct Connect.

3 Configure a VPN Connection Between Your SDDC and On-Premises Data Center

   Configure a VPN to provide a secure connection to your SDDC over the public Internet or AWS Direct Connect. Route-based and policy-based IPsec VPNs are supported. Either type of VPN can connect to the SDDC over the Internet. A route-based VPN can also connect to the SDDC over AWS Direct Connect.

4 Configure Management Gateway Networking and Security

   The management network and Management Gateway are largely preconfigured in your SDDC, but you'll still need to configure access to management network services like vCenter and HCX and create management gateway firewall rules to allow traffic between the management network and other networks, including your on-premises networks and other SDDC networks.

5 Configure Compute Gateway Networking and Security

   Compute Gateway networking includes a compute network with one or more segments and the DNS, DHCP, and security (gateway firewall and distributed firewall) configurations that manage network traffic for workload VMs. It can also include a layer 2 VPN and extended network that provides a single broadcast domain that spans your on-premises network and your SDDC workload network.

6 Configure a Multi-Edge SDDC With Traffic Groups

In the default configuration, your SDDC network has a single edge (T0) router through which all North-South traffic flows. This edge supports the default traffic group, which is not configurable. If you need additional bandwidth for the subset of this traffic routed to SDDC group members, a Direct Connect Gateway attached to an SDDC group, HCX Service Mesh, or to the connected VPC, you can reconfigure your SDDC to be Multi-Edge by creating traffic groups, each of which creates an additional T0 router.

7 Working With Inventory Groups

Use VMware Cloud on AWS Networking & Security inventory to create groups of VMs and network services that you can use when you create firewall rules.

8 Managing Workload Connections

Workload VMs connect to the Internet by default. NAT rules and distributed firewall rules give you fine-grained control over these connections.

Assign NSX Service Roles to Organization Members

Grant users in your organization an NSX service role to allow them to view or configure features on the Networking & Security tab.

Organization roles specify the privileges that an organization member has over organization assets. Service roles specify the privileges that an organization member has when accessing VMware Cloud Services that the organization uses. All service roles can be assigned and changed by a user with organization owner privileges, so restrictive roles such as Administrator (Delete Restricted) or NSX Cloud Auditor should be assigned along with the role of organization member to prevent modification.

A user must log out and then log back in for a new service role to take effect.

Prerequisites

You must be an Organization Owner to assign a role to an organization member.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Click the services icon and select **Identity & Access Management**.

3 Select a user and click **Edit Roles**.

4 Select a role name from the **Assign Organization Roles** drop-down control.

The following roles are available:

Organization Owner

This role has full rights to manage organization members and assets.

Organization Member

This role has rights to access organization assets.

5 Select the **VMware Cloud on AWS** service name under **Assign Service Roles**.

6 Select an NSX service role to assign.

The following NSX service roles are available:

NSX Cloud Auditor

This role can view NSX service settings and events but cannot make any changes to the service.

NSX Cloud Admin

This role can perform all tasks related to deployment and administration of the NSX service.

Note When multiple service roles are assigned to an organization user, permissions are granted for the most permissive role. For example, an organization member who has both the NSX Cloud Admin and NSX Cloud Auditor roles is granted all the NSX Cloud Admin permissions, which include those granted to the NSX Cloud Auditor role.

7 Click **SAVE** to save your changes.

What to do next

Ensure that any users whose roles were changed log out and log back in for the changes to take effect.

Configure AWS Direct Connect Between Your SDDC and On-Premises Data Center

Use of AWS Direct Connect is optional. If traffic between your on-premises network and your SDDC workloads requires higher speeds and lower latency than you can achieve with a connection over the public Internet, configure VMware Cloud on AWS to use AWS Direct Connect.

There are a couple of ways you can configure your VMware Cloud on AWS SDDC to take advantage of AWS Direct Connect for traffic to and from your on-premises datacenter:

Configure Direct Connect to a private VIF in your VPC.

AWS Direct Connect (DX) provides a dedicated network connection between your on-premises network infrastructure and a virtual interface (VIF) in your AWS VPC. A private VIF provides direct private access to your SDDC. Configure DX over a private VIF to carry workload and management traffic, including VPN and vMotion, between your on-premises data center and your connected VPC. A DX connection over a private VIF can be used for all traffic between your on-premises data center and your SDDC. It terminates in your connected Amazon VPC, provides a private IP address space, and uses BGP to advertise routes in your SDDC and learn routes in your on-premise data center. Provisioning procedures for this VIF depend on the type of DX connection you choose.

Associate a Direct Connect Gateway with your SDDC Group's VMware Managed Transit Gateway.

If you have created an SDDC Group in your VMware Cloud on AWS organization, you can connect an on-premises SDDC to that group's Direct Connect Gateway to give it DX connectivity to all members of the SDDC group. See Attach a Direct Connect Gateway to an SDDC Group in the *VMware Cloud on AWS Operations Guide*.

Access AWS services over a public VIF

If you just want to use DX to access AWS services in a VPC you own, you can do so over a public VIF. You cannot use a public VIF to carry the same kinds of SDDC traffic (such as vMotion) that require a private VIF or Direct Connect Gateway.

Set Up an AWS Direct Connect Connection

To set up an AWS Direct Connect connection, place an order through the AWS console to create a Direct Connect connection in a region where VMware Cloud on AWS is available.

Connection Types

AWS offers three types of Direct Connect connections:

Dedicated Connection

A dedicated connection provides a physical Ethernet port dedicated to a single customer that supports multiple private or public virtual interfaces (VIF) and 1 transit VIF.

To order a dedicated connection, ask a member of the AWS Direct Connect Partner Program to provision a circuit to an AWS Direct Connect location in the same region as your SDDC. Use your (customer-managed) AWS account to make this request. After the circuit has been provisioned, create a hosted private VIF to your SDDC using the account shown in the **AWS Account ID** field of the **Direct Connect** page of the **Networking & Security** tab. In an SDDC that is a member of an SDDC group, you can create a Direct Connect Gateway (DXGW) in your account and connect a transit VIF to it from the DXGW. See Creating and Managing SDDC Deployment Groups with VMware Transit Connect.

Hosted Connection

A hosted connection is a circuit shared by multiple customers and provisioned to your AWS account by an AWS Direct Connect Partner. After the circuit has been provisioned, create a hosted private VIF to your SDDC using the account shown in the **AWS Account ID** field of the **Direct Connect** page of the **Networking & Security** tab. If your hosted connection speed is

1Gbps or higher and the SDDC that is a member of an SDDC group, you also have the option to create a Direct Connect Gateway (DXGW) in your account, and connect a transit VIF to it from the DXGW. See Creating and Managing SDDC Deployment Groups with VMware Transit Connect.

Hosted VIF

A hosted VIF is similar to a hosted connection but only provides the ability to create a single VIF managed by a partner. The hosted private VIF must be created by the AWS Partner using the account number shown in the **AWS Account ID** field of the **Direct Connect** page of the **Networking & Security** tab, rather than provisioned to your own AWS account.

For more information about using Direct Connect with VMware Cloud on AWS, see the VMware Designlet VMware Cloud on AWS SDDC Connectivity With Direct Connect Private VIF. For more information about connection types and how to set them up, see AWS Direct Connect Partners, Getting Started with AWS Direct Connect.

Configure Direct Connect to a Private Virtual Interface for SDDC Management and Compute Network Traffic

Your DX connection requires a private virtual interface to enable vMotion, ESXi Management, Management Appliance, and workload traffic to use it.

Create one virtual interface for each Direct Connect link you want to make to your SDDC. For example, if you want to create two Direct Connect links for redundancy, create two virtual interfaces. See VMware Configuration Maximums for limits on the number of segments supported by each private VIF.

Important When you connect a DX private virtual interface to an SDDC network, all outbound traffic from ESXi hosts to destinations outside the SDDC network is routed over that interface, regardless of other routing configurations in the SDDC. This includes vMotion and vSphere replication traffic. You must ensure that inbound traffic to ESXi hosts is also routed over the DX interface so that the inbound and outbound traffic paths are symmetrical.

Although routes learned from a route-based VPN are advertised to other route-based VPNs over BGP, an SDDC advertises only its own networks over DX, not any learned from VPNs. See AWS Direct Connect quotas in the AWS *Direct Connect User Guide* for detailed information about limits imposed by AWS on Direct Connect, including limits on routes advertised and learned over BGP.

Prerequisites

- Ensure that you meet the prerequisites for virtual interfaces as described in Prerequisites for Virtual Interfaces.
- If you want to use route-based VPN as the backup to Direct Connect, you'll need a route-based VPN to use. See Create a Route-Based VPN.

Procedure

1 Log in to the AWS Console and complete the *Creating a Hosted Private Virtual Interface* procedure under Create a Hosted Virtual Interface.

  If you are using a hosted VIF, work with your AWS Direct Connect Partner to create the VIF in the account shown in the **AWS Account ID** field of the **Direct Connect** page of the **Networking & Security** tab, then skip to Step 2 of this procedure. If you are using a dedicated or hosted connection, take these steps first.

  a For **Virtual interface type**, choose **Private** and make up a **Virtual interface name**.

  b For the **Virtual interface Owner** field, select **Another AWS account** and use the **AWS Account ID** from the **Direct Connect** page of the **Networking & Security** tab.

  c For **VLAN**, use the value provided by your AWS Direct Connect Partner.

  d For **BGP ASN**, use the ASN of the on-premises router where this connection terminates.

  This value must not be the same as the **BGP Local ASN** shown on the **Direct Connect** page of the **Networking & Security** tab.

  e Expand **Additional Settings** and make the following choices:

Address family	Select IPV4
Your router peer ip	Specify the IP address of the on-premises end of this connection (your router), or leave blank to have AWS automatically assign an address that you'll need to configure in your router.
Amazon router peer ip	Specify the IP address of the AWS end of this connection, or leave blank to have AWS automatically assign an address that you'll need to configure in your router.
BGP authentication key	Specify a value or leave blank to have AWS generate a key, which you'll need to configure in your router.
Jumbo MTU (MTU size 9001)	The default MTU for all SDDC networks is 1500 bytes. To enable DX traffic to this private VIF to use a larger MTU, select **Enable** under **Jumbo MTU (MTU size 9001)**. After the VIF has been created, you'll also need to open the **Global Configuration** page of the **Networking & Security** tab and set a higher **MTU** value under **Intranet Uplink**, as described in Specify the Direct Connect MTU. Enabling this in the connection properties, even if you don't intend to use it right away, makes it easier to take advantage of jumbo frames in SDDC networks when you needs them.

  When the interface has been created, the AWS console reports that it is ready for acceptance.

2. In the VMC Console, select **Networking & Security > Direct Connect** and accept the virtual interface by clicking **ATTACH**.

   Before it has been accepted, a new VIF is visible in all SDDCs in your organization. After you accept the VIF, it is no longer visible in any other SDDC.

   It can take up to 10 minutes for the BGP session to become active. When the connection is ready, the **State** shows as **Attached** and the **BGP Status** as **Up**.

3. (Optional) Configure a route-based VPN as the backup to Direct Connect.

   In the default configuration, traffic on any route advertised over BGP by both DX and a route-based VPN uses the VPN by default. To have a route advertised by both DX and VPN use DX by default and failover to the VPN when DX is unavailable, select **Networking & Security > Direct Connect** and set the **Use VPN as backup to Direct Connect** switch to **Enabled**.

   Note This configuration requires a route-based VPN. You cannot use a policy-based VPN as a backup to Direct Connect. In an SDDC that is a member of an SDDC group, traffic over a route that is advertised by both the DX private VIF and the group's VMware Managed Transit Gateway (VTGW) will be routed over the VTGW.

   The system requires a minute or so to update your routing preference. When the operation completes, routes advertised by both DX and VPN default to the DX connection, using the VPN only when DX is unavailable. Equivalent routes advertised by both DX and VPN prioritize the VPN connection.

Results

A list of **Advertised BGP Routes** and **Learned BGP Routes** is displayed as the routes are learned and advertised. Click the refresh icon ↻ to refresh these lists. All routed subnets in the SDDC are advertised as BGP routes, along with this subset of management network subnets:

- Subnet 1 includes routes used by ESXi host vmks and router interfaces.
- Subnet 2 includes routes used for Multi-AZ support and AWS integration.
- Subnet 3 includes management VMs.

Disconnected and extended networks are not advertised.

The actual CIDR blocks advertised depend on your management subnet CIDR block. The following table shows the CIDR blocks for these routes in an SDDC that uses the default management network CIDR of 10.2.0.0 in block sizes /16, /20, and /22.

Table 2-1. Advertised Routes for 10.2.0.0 Default MGW CIDR

MGW CIDR	Subnet 1	Subnet 2	Subnet 3
10.2.0.0/23	10.2.0.0/24	10.2.1.0/26	10.2.1.128/25
10.2.0.0/20	10.2.0.0/21	10.2.8.0/23	10.2.12.0/22
10.2.0.0/16	10.2.0.0/17	10.2.128.0/19	10.2.192.0/18

What to do next

Ensure the vMotion interfaces are configured to use Direct Connect. See Configure vMotion Interfaces for Use with Direct Connect.

Configure vMotion Interfaces for Use with Direct Connect

If you are using a Direct Connect connection between your on-premises data center and your cloud SDDC, you must configure the vMotion interfaces for your on-premises hosts to route vMotion traffic over the Direct Connect connection.

Prerequisites

Configure Direct Connect and create a private virtual interface.

Procedure

1 Select one of the following methods to configure the vMotion interface on each host in your on-premises environment.

Option	Description
Override the default gateway (works for vSphere 7.0 hosts only)	For each host, edit the VMkernel adapter used for vMotion traffic, and select the option to override the default gateway. Enter an IP address in your on-premises vMotion subnet that is capable of routing traffic to the on-premises side of the Direct Connect connection. See Edit a VMkernel Adapter Configuration.
Configure the vMotion TCP/IP stack	For each host: a Remove any existing vMotion VMkernel adapters. b Create a new VMkernel adapter and select the vMotion TCP/IP stack. See Place vMotion Traffic on the vMotion TCP/IP Stack of an ESXi Host. c Edit the host vMotion TCP/IP stack to change the routing to use an IP address in your on-premises vMotion subnet that is capable of routing traffic to the on-premises side of the Direct Connect connection. See Change the Configuration of a TCP/IP Stack on a Host.

2 (Optional) Test connectivity between an on-premises host and a cloud SDDC host using `vmkping`.

See https://kb.vmware.com/s/article/1003728 for more information.

Configure Direct Connect to a Public Virtual Interface for Access to AWS Services

If your on-premises workloads need access to AWS EC2 instances and services such as S3 over a DX connection, configure a public virtual interface for that traffic in your VPC.

Although SDDC management and workload traffic over DX must use a private VIF or DX Gateway, you can create a DX connection from your on-premises datacenter to a public VIF if you just want to access AWS services from your on-premises workloads or for any purpose that requires a connection to the global AWS backbone.

Prerequisites

- Ensure that you meet the prerequisites for virtual interfaces as described in Prerequisites for Virtual Interfaces.

Procedure

1 Log in to the AWS Console. and complete the steps for creating a hosted public virtual interface under Create a Hosted Virtual Interface.

   a In the **Interface Owner** field, select **My AWS Account**.

   b Specify **Your router peer IP** and **Amazon router peer IP**.

   c Select **Auto-generate BGP key** and list any on-premises routes that you want advertised on the AWS backbone in **Prefixes you want to advertise**.

   When the interface has been created, the AWS console reports that it is ready for acceptance.

2 In the VMC Console, select **Networking & Security > Direct Connect** and accept the virtual interface by clicking **ATTACH**.

Specify the Direct Connect MTU

The default Maximum Transmissible Unit (MTU) for all SDDC networks is 1500 bytes. When you use Direct Connect, you can specify a larger MTU for the traffic it carries.

You can enable DX over to use a larger MTU when you create the VIF. If you do this, you'll also need to open the **Global Configuration** page of the **Networking & Security** tab and set a higher **Intranet MTU Value**.

This larger (or Jumbo) MTU value applies only to DX connections over a private VIF. Any VPN, whether or not it connects over DX, uses an MTU of 1500, regardless of other settings. You should also verify that the interface MTU of workload VMs that use the DX connection is set to a value that matches the **Intranet MTU Value**. Otherwise, workload VMs won't be able to take advantage of the larger MTU.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2. Click **Networking & Security > Global Configuration**.

3. On the **Global Configuration** page, click the pencil icon (✏), set a higher **MTU** value in the **Intranet Uplink** field, then click **SAVE**.

   The value you set must be less than or equal to the smallest MTU value for all your DX virtual interfaces. In practice this means that you should set all your VIFs to the same MTU value (the default, at 1500 or Jumbo, at 9001), since having any VIF that does not support a Jumbo MTU effectively limits all DX connections to an MTU of 1500. Mixing MTU sizes within a network can lead to packet fragmentation and other problems that result in poor network performance.

   Note To leave room for Geneve (Generic Network Virtualization Encapsulation) headers, the SDDC intranet MTU is capped at 8900 bytes to avoid packet fragmentation at the VIF.

Configure a VPN Connection Between Your SDDC and On-Premises Data Center

Configure a VPN to provide a secure connection to your SDDC over the public Internet or AWS Direct Connect. Route-based and policy-based IPsec VPNs are supported. Either type of VPN can connect to the SDDC over the Internet. A route-based VPN can also connect to the SDDC over AWS Direct Connect.

You can also configure a Layer 2 VPN, which can be especially useful for workload migration.

For more information about IPsec VPNs, see the VMware Designlet VMware Cloud on AWS SDDC Connectivity With IPSec VPN.

- Create a Route-Based VPN

  A route-based VPN creates an IPsec tunnel interface and routes traffic through it as dictated by the SDDC routing table. A route-based VPN provides resilient, secure access to multiple subnets. When you use a route-based VPN, new routes are added automatically when new networks are created.

- Create a Policy-Based VPN

  A policy-based VPN creates an IPsec tunnel and a policy that specifies how traffic uses it. When you use a policy-based VPN, you must update the routing tables on both ends of the network when new routes are added.

- Configure a Layer 2 VPN and Extended Network Segment

  You can use a VMware Cloud on AWS layer 2 Virtual Private Network (L2VPN) to extend your on-premises network to one or more VLAN-based networks in your SDDC. This extended network is a single subnet with a single broadcast domain. You can use it to migrate VMs to and from your cloud SDDC without having to change their IP addresses.

- View VPN Tunnel Status and Statistics

  The VMC Console provides status and statistics for IPsec VPNs and L2VPN segments.

- IPsec VPN Settings Reference

  The on-premises end of any IPsec VPN must be configured to match the settings you specified for the SDDC end of that VPN.

Create a Route-Based VPN

A route-based VPN creates an IPsec tunnel interface and routes traffic through it as dictated by the SDDC routing table. A route-based VPN provides resilient, secure access to multiple subnets. When you use a route-based VPN, new routes are added automatically when new networks are created.

Route based VPNs in your VMware Cloud on AWS SDDC use an IPsec protocol to secure traffic and the Border Gateway Protocol (BGP) to discover and propagate routes as networks are added and removed. To create a route-based VPN, you configure BGP information for the local (SDDC) and remote (on-premises) endpoints, then specify tunnel security parameters for the SDDC end of the tunnel.

Important If your SDDC includes both a policy-based VPN and a route-based VPN, connectivity over the policy-based VPN will fail if the route-based VPN advertises the default route (0.0.0.0/0) to the SDDC.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Click **Networking & Security > VPN > Route Based**.

3 (Optional) Change the default local Autonomous System Number (ASN).

   All route-based VPNs in the SDDC default to ASN 65000. The local ASN must be different from the remote ASN. (iBGP, which requires the local and remote ASNs to be the same, is not supported in SDDC networks.) To change the default local ASN, click **EDIT LOCAL ASN**, enter a new value in the range 64521 to 65535 (or 4200000000 to 4294967294) and click **APPLY**.

   Note Any change in this value affects all route-based VPNs in this SDDC.

4 Click **ADD VPN** and give the new VPN a **Name** and optional **Description**.

5 Select a **Local IP Address** from the drop-down menu.

   - If this SDDC is member of an SDDC group or has been configured to use AWS Direct Connect, select the private IP address to have the VPN use that connection rather than a connection over the Internet. Note that VPN traffic over Direct Connect or VMware Managed Transit Gateway (VTGW) is limited to the default MTU of 1500 bytes even if the link supports a higher MTU. See Configure Direct Connect to a Private Virtual Interface for SDDC Management and Compute Network Traffic.

   - Select the public IP address if you want the VPN to connect over the Internet.

6 For **Remote Public IP**, enter the address of your on-premises VPN endpoint.

This is the address of the device that initiates or responds to IPsec requests for this VPN. This address must meet the following requirements:

- It must not already be in use for another VPN. VMware Cloud on AWS uses the same public IP for all VPN connections, so only a single VPN connection (Route-based, Policy-based, or L2VPN) can be created to a given remote public IP.
- It must be reachable over the Internet if you specified a public IP in Step 5.
- It must be reachable over VTGW or Direct Connect to a private VIF if you specified a private IP in Step 5.

Default gateway firewall rules allow inbound and outbound traffic over the VPN connection, but you must create firewall rules to manage traffic over the VPN tunnel.

7 For **BGP Local IP/Prefix Length**, enter a network address from a CIDR block of size of /30 within the 169.254.0.0/16 subnet.

Some blocks in this range are reserved, as noted in Reserved Network Addresses. If you can't use a network from the 169.254.0.0/16 subnet (due to a conflict with an existing network), you must create a firewall rule that allows traffic from the BGP service to the subnet you choose here. See Add or Modify Compute Gateway Firewall Rules.

The **BGP Local IP/Prefix Length** specifies both a local subnet and an IP address in it, so the value you enter must be the second or third address in a /30 range and include the /30 suffix. For example, a **BGP Local IP/Prefix Length** of 169.254.32.1/30 creates network 169.254.32.0 and assigns 169.254.32.1 as the local BGP IP (also known as the Virtual Tunnel Interface, or VTI).

8 For **BGP Remote IP**, enter the remaining IP address from the range you specified in Step 7.

For example, if you specified a **BGP Local IP/Prefix Length** of 169.254.32.1/30, use 169.254.32.2 for **BGP Remote IP**. When configuring the on-premises end of this VPN, use the IP address you specify for **BGP Remote IP** as its local BGP IP or VTI address.

9 For **BGP Neighbor ASN**, enter the ASN of your on-premises VPN gateway.

10 Configure **Advanced Tunnel Parameters**.

Option	Description
Tunnel Encryption	Select a Phase 2 security association (SA) cipher that is supported by your on-premises VPN gateway.
Tunnel Digest Algorithm	Select a Phase 2 digest algorithm that is supported by your on-premises VPN gateway. **Note** If you specify a GCM-based cipher for **Tunnel Encryption**, set **Tunnel Digest Algorithm** to **None**. The digest function is integral to the GCM cipher.
Perfect Forward Secrecy	Enable or Disable to match the setting of your on-premises VPN gateway. Enabling Perfect Forward Secrecy prevents recorded (past) sessions from being decrypted if the private key is ever compromised.

Option	Description
Preshared Key	Enter the preshared key string. The maximum key length is 128 characters. This key must be identical for both ends of the VPN tunnel.
Remote Private IP	Leave this blank to use the **Remote Public IP** as the remote ID for IKE negotiation. If your on-premises VPN gateway is behind a NAT device and/or uses a different IP for its local ID, you need to enter that IP here.
IKE Encryption	Select a Phase 1 (IKE) cipher that is supported by your on-premises VPN gateway.
IKE Digest Algorithm	Select a Phase 1 digest algorithm that is supported by your on-premises VPN gateway. The best practice is to use the same algorithm for both the **IKE Digest Algorithm** and the **Tunnel Digest Algorithm**. **Note** If you specify a GCM-based cipher for **IKE Encryption**, set **IKE Digest Algorithm** to **None**. The digest function is integral to the GCM cipher. You must use IKE V2 if you use a GCM-based cipher .
IKE Type	■ Specify **IKE V1** to initiate and accept the IKEv1 protocol. ■ Specify **IKE V2** to initiate and accept the IKEv2 protocol. You must use IKEv2 if you have specified a GCM-based **IKE Digest Algorithm**. ■ Specify **IKE FLEX** to accept either IKEv1 or IKEv2 and then initiate using IKEv2. If IKEv2 initiation fails, IKE FLEX will not fall back to IKEv1.
Diffie Hellman	Select a Diffie Hellman group that is supported by your on-premises VPN gateway. This value must be identical for both ends of the VPN tunnel. Higher group numbers offer better protection. The best practice is to select group 14 or higher.
Connection Initiation Mode	Connection initiation mode defines the policy used by the local endpoint in the process of tunnel creation. The following modes are available. **Initiator** The default value. In this mode, the local endpoint initiates VPN tunnel creation and responds to incoming tunnel setup requests from the peer gateway. **On Demand** N/A for route-based VPN. **Respond Only** In this mode, the VPN never initiates a connection. The peer site always initiates the connection request and the local endpoint responds to that connection request.
TCP MSS Clamping	To reduce the maximum segment size (MSS) payload of the TCP session during the IPsec connection, enable **TCP MSS Clamping**, select the **TCP MSS** direction value, and optionally set the **TCP MSS Value**. See Understanding TCP MSS Clamping in the *NSX-T Data Center Administration Guide*.

11 (Optional) Under **Advanced BGP Parameters**, enter a BGP **Secret** that matches the one used by the on-premises gateway.

12 (Optional) Tag the VPN.

   See Add Tags to an Object in the *NSX-T Data Center Administration Guide* for more information about tagging NSX-T objects.

13 Click **SAVE**.

Results

The VPN creation process might take a few minutes. When the route-based VPN becomes available, the tunnel status and BGP session state are displayed. The following actions are available to help you with troubleshooting and configuring the on-premises end of the VPN:

- Click **DOWNLOAD CONFIG** to download a file that contains VPN configuration details. You can use these details to configure the on-premises end of this VPN.
- Click **VIEW STATISTICS** to view packet traffic statistics for this VPN. See View VPN Tunnel Status and Statistics.
- Click **VIEW ROUTES** to open a display of routes advertised and learned by this VPN.
- Click **DOWNLOAD ROUTES** to download a list of **Advertised Routes** or **Learned Routes** in CSV format.

What to do next

Create or update firewall rules as needed. To allow traffic through the route-based VPN, specify **VPN Tunnel Interface** in the **Applied to** field. The **All Uplinks** option does not include the routed VPN tunnel.

Create a Policy-Based VPN

A policy-based VPN creates an IPsec tunnel and a policy that specifies how traffic uses it. When you use a policy-based VPN, you must update the routing tables on both ends of the network when new routes are added.

Policy-based VPNs in your VMware Cloud on AWS SDDC use an IPsec protocol to secure traffic. To create a policy-based VPN, you configure the local (SDDC) endpoint, then configure a matching remote (on-premises) endpoint. Because each policy-based VPN must create a new IPsec security association for each network, an administrator must update routing information on premises and in the SDDC whenever a new policy-based VPN is created. A policy-based VPN can be an appropriate choice when you have only a few networks on either end of the VPN, or if your on-premises network hardware does not support BGP (which is required for route-based VPNs).

Important If your SDDC includes both a policy-based VPN and a route-based VPN, connectivity over the policy-based VPN will fail if the route-based VPN advertises the default route (0.0.0.0/0) to the SDDC.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Select **Networking & Security > VPN > Policy Based**.

3 Click **ADD VPN** and give the new VPN a **Name** and optional **Description**.

4 Select a **Local IP Address** from the drop-down menu.

- If this SDDC is member of an SDDC group or has been configured to use AWS Direct Connect, select the private IP address to have the VPN use that connection rather than a connection over the Internet. Note that VPN traffic over Direct Connect or VMware Managed Transit Gateway (VTGW) is limited to the default MTU of 1500 bytes even if the link supports a higher MTU. See Configure Direct Connect to a Private Virtual Interface for SDDC Management and Compute Network Traffic.
- Select the public IP address if you want the VPN to connect over the Internet.

5 Enter the **Remote Public IP** address of your on-premises gateway.

The address must not already be in use for another VPN. VMware Cloud on AWS uses the same public IP for all VPN connections, so only a single VPN connection (Route-based, Policy-based, or L2VPN) can be created to a given remote public IP. This address must be reachable over the Internet if you specified a public IP in Step 4. If you specified a private IP, it must be reachable over Direct Connect to a private VIF. Default gateway firewall rules allow inbound and outbound traffic over the VPN connection, but you must create firewall rules to manage traffic over the VPN tunnel.

6 (Optional) If your on-premises gateway is behind a NAT device, enter the gateway address as the **Remote Private IP**.

This IP address must match the local identity (IKE ID) sent by the on-premises VPN gateway. If this field is empty, the **Remote Public IP** field is used to match the local identity of the on-premises VPN gateway.

7 Specify the **Remote Networks** that this VPN can connect to.

This list must include all networks defined as local by the on-premises VPN gateway. Enter each network in CIDR format, separating multiple CIDR blocks with commas.

8 Specify the **Local Networks** that this VPN can connect to.

This list includes all routed compute networks in the SDDC, as well as the entire Management network and the appliance subnet (a subset of the Management network that includes vCenter and other management appliances, but not the ESXi hosts). It also includes the CGW DNS Network, a single IP address used to source requests forwarded by the CGW DNS service.

9 Configure **Advanced Tunnel Parameters**.

Option	Description
Tunnel Encryption	Select a Phase 2 security association (SA) cipher that is supported by your on-premises VPN gateway.
Tunnel Digest Algorithm	Select a Phase 2 digest algorithm that is supported by your on-premises VPN gateway. **Note** If you specify a GCM-based cipher for **Tunnel Encryption**, set **Tunnel Digest Algorithm** to **None**. The digest function is integral to the GCM cipher.
Perfect Forward Secrecy	Enable or Disable to match the setting of your on-premises VPN gateway. Enabling Perfect Forward Secrecy prevents recorded (past) sessions from being decrypted if the private key is ever compromised.
IKE Encryption	Select a Phase 1 (IKE) cipher that is supported by your on-premises VPN gateway.
IKE Digest Algorithm	Select a Phase 1 digest algorithm that is supported by your on-premises VPN gateway. The best practice is to use the same algorithm for both the **IKE Digest Algorithm** and the **Tunnel Digest Algorithm**. **Note** If you specify a GCM-based cipher for **IKE Encryption**, set **IKE Digest Algorithm** to **None**. The digest function is integral to the GCM cipher. You must use IKE V2 if you use a GCM-based cipher .
IKE Type	■ Specify **IKE V1** to initiate and accept the IKEv1 protocol. ■ Specify **IKE V2** to initiate and accept the IKEv2 protocol. You must use IKEv2 if you have specified a GCM-based **IKE Digest Algorithm**. ■ Specify **IKE FLEX** to accept either IKEv1 or IKEv2 and then initiate using IKEv2. If IKEv2 initiation fails, IKE FLEX will not fall back to IKEv1.
Diffie Hellman	Select a Diffie Hellman group that is supported by your on-premises VPN gateway. This value must be identical for both ends of the VPN tunnel. Higher group numbers offer better protection. The best practice is to select group 14 or higher.
Preshared Key	Enter a preshared key used by both ends of the tunnel to authenticate with each other. The string has a maximum length of 128 characters.

Option	Description
Connection Initiation Mode	Connection initiation mode defines the policy used by the local endpoint in the process of tunnel creation. The following modes are available. **Initiator** The default value. In this mode, the local endpoint initiates VPN tunnel creation and responds to incoming tunnel setup requests from the peer gateway. **On Demand** In this mode, the local endpoint initiates VPN tunnel creation after the first packet matching the policy rule is received. It also responds to the incoming initiation request. **Respond Only** In this mode, the VPN never initiates a connection. The peer site always initiates the connection request and the local endpoint responds to that connection request.
TCP MSS Clamping	To reduce the maximum segment size (MSS) payload of the TCP session during the IPsec connection, enable **TCP MSS Clamping**, select the **TCP MSS** direction value, and optionally set the **TCP MSS Value**. See Understanding TCP MSS Clamping in the *NSX-T Data Center Administration Guide*.

10 (Optional) Tag the VPN.

See Add Tags to an Object in the *NSX-T Data Center Administration Guide* for more information about tagging NSX-T objects.

11 Click **SAVE**.

Results

The VPN creation process might take a few minutes. When the policy-based VPN becomes available, the following actions are available to help you with troubleshooting and configuring the on-premises end of the VPN:

- Click **DOWNLOAD CONFIG** to download a file that contains VPN configuration details. You can use these details to configure the on-premises end of this VPN.
- Click **VIEW STATISTICS** to view packet traffic statistics for this VPN. See View VPN Tunnel Status and Statistics.

What to do next

Create or update firewall rules as needed. To allow traffic through the policy-based VPN, specify **Internet Interface** in the **Applied to** field.

Configure a Layer 2 VPN and Extended Network Segment

You can use a VMware Cloud on AWS layer 2 Virtual Private Network (L2VPN) to extend your on-premises network to one or more VLAN-based networks in your SDDC. This extended network is a single subnet with a single broadcast domain. You can use it to migrate VMs to and from your cloud SDDC without having to change their IP addresses.

In addition to data center migration, you can use an extended L2VPN network for disaster recovery, or for dynamic access to cloud computing resources as needed (often referred to as "cloud bursting).

VMware Cloud on AWS uses NSX-T to provide the L2VPN server in your cloud SDDC. L2VPN client functions are provided by an on-premises NSX Edge. See VMware Configuration Maximums for L2VPN limits.

The VMware Cloud on AWS L2VPN feature supports extending VLAN networks. The L2VPN connection to the NSX-T server uses an IPsec tunnel. The L2VPN extended network is used to extend Virtual Machine networks and carries only workload traffic. It is independent of the VMkernel networks used for migration traffic (ESXi management or vMotion), which use either a separate IPsec VPN or a Direct Connect connection.

Important You cannot bring up an L2VPN tunnel until you have configured the L2VPN client and server and created an extended network that specifies the tunnel ID you assigned to the client.

Procedure

1 Configure a Layer 2 VPN Tunnel in the SDDC

  Specify a local (SDDC) IP address, a remote (on-premises) public IP address, and a remote private IP address to create the SDDC end of the Layer 2 VPN tunnel.

2 Configure an Extended Segment for the Layer 2 VPN

  Extended networks require a layer 2 Virtual Private Network (L2VPN), which provides a secure communications tunnel between an on-premises network and one in your cloud SDDC.

3 Install and Configure the On-Premises NSX Edge

  The on-premises end of your L2VPN must be an NSX Edge appliance. You must configure this appliance and related on-premises vSphere networking before you can create an L2VPN.

Configure a Layer 2 VPN Tunnel in the SDDC

Specify a local (SDDC) IP address, a remote (on-premises) public IP address, and a remote private IP address to create the SDDC end of the Layer 2 VPN tunnel.

VMware Cloud on AWS supports a single Layer 2 VPN tunnel between your on-premises installation and your SDDC.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Select **Networking & Security > VPN > Layer 2**.

3 Click **ADD VPN TUNNEL**.

4 Configure the VPN parameters.

Option	Description
Local IP Address	■ Select the private IP address if you have configured AWS Direct Connect for this SDDC and want the VPN to use it. See Configure Direct Connect to a Private Virtual Interface for SDDC Management and Compute Network Traffic. ■ Select the public IP address if you want the VPN to connect to the SDDC over Internet.
Remote Public IP	Enter the remote public IP address of your on-premise L2VPN gateway. For an L2VPN, this is always the standalone NSX Edge appliance (see Install and Configure the On-Premises NSX Edge).
Remote Private IP	Enter the remote private IP address if the on-premise gateway is configured behind NAT.

5 (Optional) Tag the VPN.

See Add Tags to an Object in the *NSX-T Data Center Administration Guide* for more information about tagging NSX-T objects.

6 (Optional) Add a **Description**.

7 Click **SAVE**.

Depending on your SDDC environment, the Layer 2 VPN creation process might take a few minutes. When the Layer 2 VPN tunnel becomes available, the status changes to Up.

Configure an Extended Segment for the Layer 2 VPN

Extended networks require a layer 2 Virtual Private Network (L2VPN), which provides a secure communications tunnel between an on-premises network and one in your cloud SDDC.

Each end of this tunnel has an ID. When the tunnel ID matches on the cloud SDDC and the on-premises side of the tunnel, the two networks become part of the same broadcast domain. Extended networks use an on-premises gateway as the default gateway. Other network services such as DHCP and DNS are also provided on-premises.

You can change a logical network from routed to extended or from extended to routed. For example, you might configure a logical network as extended to allow migration of VMs from your on-premises data center to your cloud SDDC. When the migration is complete, you might then change the network to routed to allow the VMs to use VMware Cloud on AWS networking services.

Prerequisites

Verify that Layer 2 VPN tunnel is available. See Configure a Layer 2 VPN Tunnel in the SDDC.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Follow the procedure in Create or Modify a Network Segment to create an Extended segment bound to the Tunnel ID of the L2VPN tunnel.

3 Click **SAVE**.

4 Click **DOWNLOAD CONFIG** to download a file containing the peer code and other information you'll need when configuring the on-premises of the remote side VPN configuration.

5 Configure the client side of the L2VPN.

See Install and Configure the On-Premises NSX Edge.

Install and Configure the On-Premises NSX Edge

The on-premises end of your L2VPN must be an NSX Edge appliance. You must configure this appliance and related on-premises vSphere networking before you can create an L2VPN.

If you have a compatible version of NSX-T installed in your on-premises data center, you can use your existing NSX Edge appliance as the on-premises (client) side of an L2VPN that connects to your SDDC. If necessary, you can download and deploy an autonomous NSX Edge and configure it as the L2VPN client. The following table lists compatible SDDC and on-premises versions. To determine the version of NSX-T running in your SDDC, see Correlating VMware Cloud on AWS with Component Releases in the *VMware Cloud on AWS Operations Guide*.

Table 2-2. L2VPN Interoperability

L2VPN Server Version (SDDC)	L2VPN Client Version (On-Premises Edge)
3.1.2	3.1.2, 3.1.1, 2.5.3
3.1.1	3.1.1, 3.1.0, 3.0.1
3.1.0	3.1.1, 3.0.1, 3.0.0
3.0.3	3.0.3, 3.0.2, 3.0.1
3.0.2	3.0.2, 3.0.1, 2.5.2
3.0.0	3.0.0, 2.5.0, 2.5.1

Procedure

1 (Optional) Download the standalone autonomous NSX Edge.

If you do not have a compatible version of NSX-T installed in your on-premises data center, you may be able to download and configure a standalone Autonomous Edge appliance to use as the on-premises endpoint for your L2VPN. On the L2VPN page, click **AUTONOMOUS EDGE DOWNLOAD** to download the Autonomous Edge as an OVF file.

2 Install and configure the Autonomous Edge.

See Add an Autonomous Edge as an L2 VPN Client in the *NSX-T Data Center Administration Guide* for information about how to install and configure the Autonomous Edge in your on-premises vCenter Server.

View VPN Tunnel Status and Statistics

The VMC Console provides status and statistics for IPsec VPNs and L2VPN segments.

Status of VPN operations is reported on the **VPN** pages in the **Networking & Security** tab. Log messages about VPN operations are also sent to vRealize Log Insight Cloud, an optional SDDC add-on service. See Using the vRealize Log Insight Cloud Add-On and the vRealize Log Insight Cloud Documentation for more information.

Procedure

1. Log in to the VMC Console at https://vmc.vmware.com.
2. Click **Networking & Security > VPN**.
3. Click **ROUTE BASED VPN**, **POLICY BASED VPN**, or **LAYER 2 VPN** to list VPNs of the selected type.

   Take one of the following actions:

   - Click the Information icon ⓘ to display a status message that provides more information about channel (IKE Phase 1 negotiation) and tunnel status.
   - Expand a VPN row to show VPN details, then click **VIEW STATISTICS** to display traffic statistics. You can retrieve aggregated status and statistics for all tunnels or for the tunnel used by the selected VPN (0.0.0.0/0). When viewing aggregated statistics, you can click **View More** in the **Stats** column to see a list of error statistics.
   - Click the Refresh icon ↻ to refresh tunnel statistics. All VPN statistics are reset to `0` when the tunnel is disabled or re-enabled.

What to do next

For more information about troubleshooting VPN connection issues, see Troubleshooting Virtual Private Networks (VPN) in the *NSX Data Center for vSphere* documentation and also the NSX-T Data Center Troubleshooting Guide.

IPsec VPN Settings Reference

The on-premises end of any IPsec VPN must be configured to match the settings you specified for the SDDC end of that VPN.

Information in the following tables summarizes the available SDDC IPsec VPN settings. Some of the settings can be configured. Some are static. Use this information to verify that your on-premises VPN solution can be configured to match the one in your SDDC. Choose an on-premises VPN solution that supports all the static settings and any of the configurable settings listed in these tables.

Phase 1 Internet Key Exchange (IKE) Settings

Table 2-3. Configurable IKE Phase 1 Settings

Attribute	Allowed Values	Recommended Value
Protocol	IKEv1, IKEv2, IKE FLEX	IKEv2
Encryption Algorithm	AES (128, 256), AES-GCM (128, 192, 256)	AES GCM
Tunnel/IKE Digest Algorithm	SHA-1, SHA-2	SHA-2
Diffie Hellman	DH Groups 2, 5, 14-16	DH Group 14-16

Table 2-4. Static IKE Phase 1 Settings

Attribute	Value
ISAKMP mode	Main mode (Disable aggressive mode)
ISAKMP/IKE SA lifetime	86400 seconds (24 hours)
IPsec Mode	Tunnel
IKE Authentication	Pre-Shared Key

Phase 2 IKE Settings

Table 2-5. Configurable IKE Phase 2 Settings

Attribute	Allowed Values	Recommended Value
Encryption Algorithm	AES-256, AES-GCM, AES	AES-GCM
Perfect forward secrecy (PFS)	Enabled, Disabled	Enabled
Diffie Hellman	DH Groups 2, 5, 14-16	DH Group 14-16

Table 2-6. Static IKE Phase 2 Settings

Attribute	Value
Hashing Algorithm	SHA-1
Tunnel Mode	Encapsulating Security Payload (ESP)
SA lifetime	3600 seconds (one hour)

On-Premises IPsec VPN Configuration

Click **DOWNLOAD CONFIG** on the status page of any VPN to download a file that contains VPN configuration details. You can use these details to configure the on-premises end of the VPN.

Note Do not configure the on-premises side of a VPN to have an idle timeout (for example, the NSX **Session idle timeout** setting). On-premises idle timeouts can cause the VPN to become periodically disconnected.

Sample configuration files for several popular endpoint devices are available on VMware {code}.

- Palo Alto Networks Firewall

Configure Management Gateway Networking and Security

The management network and Management Gateway are largely preconfigured in your SDDC, but you'll still need to configure access to management network services like vCenter and HCX and create management gateway firewall rules to allow traffic between the management network and other networks, including your on-premises networks and other SDDC networks.

Procedure

1. Set vCenter Server FQDN Resolution Address

   You can connect to the SDDC vCenter Server at either a public or private IP address. A private IP address can be resolved from an SDDC VPN. A public IP address can be resolved from the Internet.

2. Set HCX FQDN Resolution Address

   You can connect to HCX at either a public or private IP address. A private IP address can be resolved from an SDDC VPN. A public IP address can be resolved from the internet.

3. Add or Modify Management Gateway Firewall Rules

   By default, the management gateway blocks traffic to all destinations from all sources. Add Management Gateway firewall rules to allow traffic as needed.

Set vCenter Server FQDN Resolution Address

You can connect to the SDDC vCenter Server at either a public or private IP address. A private IP address can be resolved from an SDDC VPN. A public IP address can be resolved from the Internet.

Prerequisites

Before you can access the SDDC vCenter Server at a private IP address, you'll need to set up a VPN connecting your SDDC to your on-premises datacenter. See Create a Route-Based VPN or Create a Policy-Based VPN.

Procedure

1. Log in to the VMC Console at https://vmc.vmware.com.
2. Navigate to the **Settings** tab of your SDDC.
3. Expand **vCenter FQDN**, and click **Edit**.
4. Under **Resolution Address** Select either the **Public IP** address or the **Private IP** address and click **SAVE**.

Set HCX FQDN Resolution Address

You can connect to HCX at either a public or private IP address. A private IP address can be resolved from an SDDC VPN. A public IP address can be resolved from the internet.

Prerequisites

Before you can access HCX at a private IP address, you'll need to set up a VPN connecting your SDDC to your on-premises datacenter. See Create a Route-Based VPN or Create a Policy-Based VPN.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Navigate to the **Settings** tab of your SDDC.

3 Expand **HCX FQDN**, and click **Edit**.

4 Under **Resolution Address** select either the **Public IP** address or the **Private IP** address and click **SAVE**.

Add or Modify Management Gateway Firewall Rules

By default, the management gateway blocks traffic to all destinations from all sources. Add Management Gateway firewall rules to allow traffic as needed.

Management Gateway firewall rules specify actions to take on network traffic from a specified source to a specified destination. Sources and destinations can be defined as **Any** or as members of a system-defined or user-defined inventory group, but either the source or destination must be system-defined. See Add a Management Group for information about viewing or modifying inventory groups.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 On the **Networking & Security** tab, click **Gateway Firewall**.

3 On the **Gateway Firewall** card, click **Management Gateway**, then click **ADD RULE** and give the new rule a **Name**.

4 Enter the parameters for the new rule.

Parameters are initialized to their default values (for example, **All** for **Sources** and **Destinations**). To edit a parameter, move the mouse cursor over the parameter value and click the pencil icon (✏) to open a parameter-specific editor.

Option	Description
Sources	Select **Any** to allow traffic from any source address or address range. Select **System Defined Groups** and select one of the following source options: ■ **ESXi** to allow traffic from your SDDC's ESXi hosts. ■ **NSX Manager** to allow traffic from your SDDC's NSX-T manager appliance. ■ **vCenter** to allow traffic from your SDDC's vCenter Server. Select **User Defined Groups** to use a management group that you have defined. See Add a Management Group.
Destinations	Select **Any** to allow traffic to any destination address or address range. Select **System Defined Groups** and select one of the following destination options: ■ **ESXi** to allow traffic to your SDDC's ESXi management. ■ **NSX Manager** to allow traffic to your SDDC's NSX-T. ■ **vCenter** to allow traffic to your SDDC's vCenter Server.
Services	Select the service types that the rule applies to. The list of service types depends on your choices for **Sources** and **Destinations**.
Action	The only action available for a new management gateway firewall rule is **Allow**.

The new rule is enabled by default. Slide the toggle to the left to disable it.

5 Click **PUBLISH** to create the rule.

The system gives the new rule an integer **ID** value, which is used in log entries generated by the rule.

Firewall rules are applied in order from top to bottom. Because there is a default **Drop** rule at the bottom and the rules above are always **Allow** rules, management gateway firewall rule order has no impact on traffic flow.

Example: Create a Management Gateway Firewall Rule

To create a management gateway firewall rule that enables vMotion traffic from the on-premises ESXi hosts to the ESXi hosts in the SDDC:

1 Create a management inventory group that contains the on-premises ESXi hosts that you want to enable for vMotion to the SDDC.

2 Create a management gateway rule with source ESXi and destination on-premises ESXi hosts.

3 Create another management gateway rule with source on-premises ESXi hosts group and destination ESXi with a vMotion service.

What to do next

You can take any or all of these optional actions with an existing firewall rule.

- Click the gear icon ⚙ to view or modify rule logging settings. Log entries are sent to the VMwarevRealize Log Insight Cloud Service. See Using vRealize Log Insight Cloud in the *VMware Cloud on AWS Operations Guide*.
- Click the graph icon 📈 to view Rule Hits and Flow statistics for the rule.

Popularity Index	Number of times the rule was triggered in the past 24 hours.
Hit Count	Number of times the rule was triggered since it was created.

Packet Count	Total packet flow through this rule.
Byte Count	Total byte flow through this rule.

Statistics start accumulating as soon as the rule is enabled.

Example Management Gateway Firewall Rules

Some common firewall rule configurations include opening access to the vSphere Client from the internet, allowing access to vCenter Server through the management VPN tunnel, and allowing remote console access.

Commonly Used Firewall Rules

The following table shows the Service, Source, and Destination settings for commonly-used firewall rules.

Table 2-9. Commonly-Used Firewall Rules

Use Cases	Service	Source	Destination
Provide access to vCenter Server from the internet. Use for general vSphere Client access as well as for monitoring vCenter Server	HTTPS	public IP address	vCenter
Provide access to vCenter Server over VPN tunnel. Required for Management Gateway VPN, Hybrid Linked Mode, Content Library.	HTTPS	IP address or CIDR block from on-premises data center	vCenter
Provide access from cloud vCenter Server to on-premises services such as Active Directory, Platform Services Controller, and Content Library.	Any	vCenter	IP address or CIDR block from on-premises data center.

Table 2-9. Commonly-Used Firewall Rules (continued)

Use Cases	Service	Source	Destination
Provisioning operations involving network file copy traffic, such as cold migration, cloning from on-premises VMs, snapshot migration, replication, and so on.	Provisioning	IP address or CIDR block, either public or from an on-premises data center connected by a VPN tunnel	ESXi Management
VMRC remote console access Required for vRealize Automation	Remote Console	IP address or CIDR block, either public or from an on-premises data center connected by a VPN tunnel	ESXi Management
vMotion traffic over VPN	Any	ESXi Management	IP address or CIDR block from on-premises data center

Configure Compute Gateway Networking and Security

Compute Gateway networking includes a compute network with one or more segments and the DNS, DHCP, and security (gateway firewall and distributed firewall) configurations that manage network traffic for workload VMs. It can also include a layer 2 VPN and extended network that provides a single broadcast domain that spans your on-premises network and your SDDC workload network.

Procedure

1 Create or Modify a Network Segment

   Network segments are logical networks for use by workload VMs in the SDDC compute network.

2 Add or Modify Compute Gateway Firewall Rules

   By default, the Compute Gateway blocks traffic to all uplinks. Add Compute Gateway firewall rules to allow traffic as needed.

3 Add or Modify Distributed Firewall Rules

   Distributed firewall rules apply at the VM (vNIC) level and control East-West traffic within the SDDC.

4 Configure DNS Services

   VMware Cloud on AWS DNS forwarding services run in DNS zones, and enable workload VMs in the zone to resolve fully-qualified domain names to IP addresses.

5 View Routes Learned and Advertised over VMware Transit Connect

   In an SDDC that is a member of an SDDC Group, you can use the **Networking & Security Transit Connect** tool to view routes learned and advertised by this SDDC in the VMware Transit Connect network created for the group.

Create or Modify a Network Segment

Network segments are logical networks for use by workload VMs in the SDDC compute network.

VMware Cloud on AWS supports three types of network segments: routed, extended and disconnected.

- A routed network segment (the default type) has connectivity to other logical networks in the SDDC and, through the SDDC firewall, to external networks.
- An extended network segment extends an existing L2VPN tunnel, providing a single IP address space that spans the SDDC and an on-premises network.
- A disconnected network segment has no uplink, and provides an isolated network accessible only to VMs connected to it. Disconnected segments are created when needed by HCX (see Getting started with VMware HCX). You can also create them yourself, and can convert them to other segment types.

See VMware Configuration Maximums for limits on segments per SDDC and network connections per segment.

A Single Host Starter SDDC is created with a single routed network segment named `sddc-cgw-network-1`. Multi-host SDDCs are created without a default network segment, so you must create at least one for your workload VMs. When you create a segment, you start by configuring some basic parameters and specifying how DHCP requests are handled on the segment. After the segment has been created, you can take additional, optional steps to specify a segment profiles and create DHCP static bindings.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Click **Networking & Security > Segments**.

To create a new segment, click **ADD SEGMENT** and give the new segment a **Name** and optional **Description**.

To delete or modify a segment, click its ellipsis buttons ⋮ and select **Edit**. You can modify all segment properties, including segment type. You can also edit or delete the segment's DHCP configuration.

Important You cannot disable or delete a segment of any type if it has attached VMs or VIFs. Disconnect attached VMs and VIFs before deleting the segment.

3 Specify a segment **Type** and fill in the required configuration parameters.

Parameter requirements depend on the segment type

Table 2-10. Routed Segment Configuration Parameters

Parameter	Value
VPN Tunnel ID	N/A for Routed or Disconnected segment types.
Subnets	Specify an IPv4 CIDR block for the segment. The block must not overlap your management network, any of the CIDR clocks listed in Reserved Network Addresses, or any of the subnets in your connected Amazon VPC. If any part of the block is in a public IP space, it must be in one that has been allocated for your use by IANA or another regional internet registry.
SET DHCP CONFIG	Routed segments default to using the Compute Gateway DHCP server. Per-segment DHCP configuration, including DHCP relay, can be specified when you create or update the segment. See Configure Segment DHCP Properties.
Domain Name	(Optional) Enter a fully qualified domain name. Static bindings on the segment automatically inherit this domain name.
Tags	See Add Tags to an Object in the *NSX-T Data Center Administration Guide* for more information about tagging NSX-T objects.

Table 2-11. Extended Segment Configuration Parameters

Parameter	Value
VPN Tunnel ID	Specify the tunnel ID of an existing L2VPN tunnel. N/A for Routed or Disconnected segment types. If you have not already created an L2VPN, see Configure a Layer 2 VPN Tunnel in the SDDC.
Subnets	N/A for Extended segments.
Domain Name	(Optional) Enter a fully qualified domain name. Static bindings on the segment automatically inherit this domain name.
Tags	See Add Tags to an Object in the *NSX-T Data Center Administration Guide* for more information about tagging NSX-T objects.

Table 2-12. Disconnected Segment Configuration Parameters

Parameter	Value
VPN Tunnel ID	N/A for Routed or Disconnected segment types.
Subnets	Specify an IPv4 CIDR block for the segment. The block must not overlap your management network, any of the CIDR clocks listed in Reserved Network Addresses, or any of the subnets in your connected Amazon VPC. If any part of the block is in a public IP space, it must be in one that has been allocated for your use by IANA or another regional internet registry.
Domain Name	(Optional) Enter a fully qualified domain name. Static bindings on the segment automatically inherit this domain name.
Tags	See Add Tags to an Object in the *NSX-T Data Center Administration Guide* for more information about tagging NSX-T objects.

4 Click **SAVE** to create or update the segment.

Click **YES** if you want continue with segment configuration. If you click **NO**, you can edit the segment later if you need to.

The system creates the requested segment. This operation can take up to 15 seconds to complete. When the segment **Status** transitions to **Up** the segment is ready for use. If the segment **Status** is **Down**, you can click the information icon ⓘ for more information about the cause of the problem.

5 (Optional) Click **SEGMENT PROFILES** to view segment profiles for the segment.

Segment profiles specify Layer 2 networking configuration details for segments and segment ports. A set of default segment profiles is applied to every new segment. Segment profiles are read-only for VMware Cloud on AWS.

6 (Optional) Configure **DHCP STATIC BINDINGS**.

a Click **Set** to specify static bindings for VMs on the segment.

Click **ADD IPV4 STATIC BINDING**, then give the binding a **Name** and specify an IPv4 address included in the segment and a MAC address. When a VM with the specified MAC address is powered on and connected to the segment, it receives the specified address. Click **SAVE** to create the binding, then add another binding or click **APPLY** to apply the specified static bindings to the segment.

b Click **DHCP Options** to specify DHCP Classless Static Routes (Option 121) and Generic Options.

- Each classless static route option in DHCP for IPv4 can have multiple routes with the same destination. Each route includes a destination subnet, subnet mask, next hop router. See RFC 3442 for information about classless static routes in DHCPv4. You can add a maximum of 127 classless static routes on a DHCPv4 server.
- For adding Generic Options, select the code of the option and enter a value of the option. For binary values, the value must be in a base-64 encoded format.

What to do next

After a segment has been created and has a status of Success, you can click **VIEW STATISTICS** to view statistics for network traffic to and from the segment. You can click **VIEW RELATED GROUPS** to see a list of groups that include this segment. For more information about groups in NSX-T, see Add a Group in the *NSX-T Data Center Administration Guide*.

Configure Segment DHCP Properties

DHCP configuration is a per-segment property. In the default configuration the Compute Gateway DHCP server handles DHCP requests from VMs on all routed segments. If you have an external DHCP server that manages IP addresses on your workload networks, you can configure the segment to use DHCP relay. You can also configure the segment to use its own local DHCP Server.

Per-segment DHCP configuration is part of the segment create/update workflow document in Create or Modify a Network Segment.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Click **Networking & Security > Segments**.

   To modify the DHCP configuration of an existing segment, click its ellipsis button an select **Edit**, then **EDIT DHCP CONFIG**.

3 Choose a **DHCP Type** and specify configuration details.

Configuration details depend on the DHCP type. To specify **Settings**, toggle **DHCP Config** to **Enabled**.

DHCP Type	Description
Local DHCP Server	Select this option to create a local DHCP server that has an IP address on the segment. As the name suggests, it is a DHCP server that is local to the segment and not available to the other segments in the network. A local DHCP server provides a dynamic IP assignment service only to the VMs that are attached to the segment. You can configure all DHCP settings, including DHCP ranges, DHCP Options, and static bindings on the segment. For disconnected segments, this type is selected by default.
DHCP Relay	Select this option to relay the DHCP client requests to the external DHCP servers. The external DHCP servers can be in any subnet, outside the SDDC, or in the physical network. The DHCP relay service is local to the segment and not available to the other segments in the network.
Gateway DHCP Server	This DHCP type is analogous to a central DHCP service that dynamically assigns IP and other network configuration to the VMs on all the segments that are connected to the gateway and using Gateway DHCP. Depending on the type of DHCP profile you attach to the gateway, you can configure a Gateway DHCP server or a Gateway DHCP relay on the segment. By default, segments that are connected to a tier-1 or tier-0 gateway use Gateway DHCP. If needed, you can choose to configure a DHCP local server or a DHCP relay on the segment. To configure Gateway DHCP on a segment, a DHCP profile must be attached to the gateway. See Create or Modify a DHCP Profile.

a Enable the DHCP configuration settings on the subnet by clicking the **DHCP Config** toggle button.

b In the **DHCP Server Address** text box, enter the IP addresses.

- If you are configuring a DHCP local server, server IP address is required. A maximum of two server IP addresses are supported. One IPv4 address and one IPv6 address. For an IPv4 address, the prefix length must be <= 30, and for an IPv6 address, the prefix length must be <= 126. The server IP addresses must belong to the subnets that you have specified in this segment. The DHCP server IP address must not overlap with the IP addresses in the DHCP ranges and DHCP static binding. The DHCP server profile might contain server IP addresses, but these IP addresses are ignored when you configure a local DHCP server on the segment.

  After a local DHCP server is created, you can edit the server IP addresses on the **Set DHCP Config** page. However, the new IP addresses must belong to the same subnet that is configured in the segment.

- If you are configuring a DHCP relay, this step is not applicable. The server IP addresses are fetched automatically from the DHCP relay profile and displayed below the profile name.
- If you are configuring a Gateway DHCP server, this text box is not editable. The server IP addresses are fetched automatically from the DHCP profile that is attached to the connected gateway.

  Remember, the Gateway DHCP server IP addresses in the DHCP server profile can be different from the subnet that is configured in the segment. In this case, the Gateway DHCP server connects with the IPv4 subnet of the segment through an internal relay service, which is autocreated when the Gateway DHCP server is created. The internal relay service uses any one IP address from the subnet of the Gateway DHCP server IP address. The IP address used by the internal relay service acts as the default gateway on the Gateway DHCP server to communicate with the IPv4 subnet of the segment.

  After a Gateway DHCP server is created, you can edit the server IP addresses in the DHCP profile of the gateway. However, you cannot change the DHCP profile that is attached to the gateway.

DHCP Ranges, if specified, must meet the following requirements:

- IP addresses in the DHCP ranges must belong to the subnet that is configured on the segment. DHCP ranges cannot contain IP addresses from multiple subnets.
- IP ranges must not overlap with the DHCP Server IP address and the DHCP static binding IP addresses.
- IP ranges in the DHCP IP pool must not overlap each other.
- Number of IP addresses in any DHCP range must not exceed 65536.

c (Optional) Edit the lease time in seconds.

Default value is 86400. Valid range of values is 60–4294967295. The lease time configured in the DHCP server configuration takes precedence over the lease time that you specified in the DHCP profile.

d (Optional) Enter the IP address of the domain name server (DNS) to use for name resolution. A maximum of two DNS servers are permitted.

When not specified, no DNS is assigned to the DHCP client. DNS server IP addresses must belong to the same subnet as the subnet's gateway IP address.

e (Optional) Click **Options** to configure DHCP options.

For information about **CLASSLESS STATIC ROUTES** and other DHCP **Options**, see RFC3442 and Create a DHCP Server in the *NSX-T Data Center Administration Guide*.

4 (Optional) Specify a DHCP Profile. If your SDDC includes more than one DHC Profile, you can use the **DHCP Profile** drop-down menu to select the name of DHCP server profile you want this segment to use.

See Create or Modify a DHCP Profile.

5 Click **APPLY** to apply the DHCP configuration to the segment.

Create or Modify a DHCP Profile

A DHCP profile specifies a DHCP server type and configuration. You can use the default profile or create others as needed.

A DHCP profile can be used by multiple segments and gateways in your network. You can create DHCP server profiles and DHCP relay profiles. See Add a DHCP Profile in the *NSX-T Data Center Administration Guide*.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Select **Networking & Security > DHCP**.

3 Click **ADD DHCP PROFILE** and give the profile a **Name**.

Choose a **Profile Type** and provide the required configuration parameters.

- For a **DHCP Server**, specify an IPv4 **Server IP Address** and optionally change the **Lease Time**.
- For a **DHCP Relay**, specify the **Server IP Address** as the address of your on-premises DHCP server. Be sure that your on-premises firewall allows DHCP traffic (ports 67 and 68) to reach this address. Lease time is controlled by the on-premises server configuration.

Either type of DHCP profile can be tagged.

4 Click **SAVE** to create the profile.

The new profile is available for use when you specify the DHCP configuration of a routed segment. See Create or Modify a Network Segment. The **Where Used** column lists segments that specify this profile.

Add or Modify Compute Gateway Firewall Rules

By default, the Compute Gateway blocks traffic to all uplinks. Add Compute Gateway firewall rules to allow traffic as needed.

Compute Gateway firewall rules specify actions to take on network traffic from a specified source to a specified destination. Actions can be either allow (allow the traffic) or drop (drop all packets matching the specified source and destination). Sources and destinations can be chosen from a list of a physical network interfaces, or the generic specification **All Uplinks**, which applies to all traffic leaving the gateway and going to the VPC interface, Internet Interface, or Direct Connect interface.

Note A firewall rule applied to **All Uplinks** does not apply to the **VPN Tunnel Interface** (VTI), which is a virtual interface and not a physical uplink. The **VPN Tunnel Interface** must be specified explicitly in the **Applied To** parameter of any firewall rule that manages workload VM communications over a route-based VPN.

The Compute Gateway includes a **Default VTI Rule** that drops all traffic to the VTI and a a **Default Uplink Rule** that drops traffic to **All Uplinks**. To enable workload VMs to communicate over the VTI, modify this rule or move it to a lower rank in the rule hierarchy, after more permissive rules.

All traffic attempting to pass through the firewall is subjected to the rules in the order shown in the rules table, beginning at the top. A packet allowed by the first rule is passed on to the second rule, and so on through subsequent rules until the packet is dropped, rejected, or hits a default rule.

Prerequisites

Compute Gateway firewall rules require named inventory groups for Source and Destination values. See Add or Modify a Compute Group.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 On the **Networking & Security** tab, click **Gateway Firewall**.

3 On the **GATEWAY FIREWALL** page, click **Compute Gateway**.

4 To add a rule, click **ADD RULE** and give the new rule a **Name**.

5 Enter the parameters for the new rule.

Parameters are initialized to their default values (for example, **All** for **Sources** and **Destinations**). To edit a parameter, move the mouse cursor over the parameter value and click the pencil icon (✏) to open a parameter-specific editor.

Option	Description
Sources	Click **Any** in the **Sources** column and select an inventory group for source network traffic, or click **ADD GROUP** to create a new user-defined inventory group to use for this rule. Click **SAVE**.
Destinations	Click **Any** in the **Destinations** column and select an inventory group for destination network traffic, or click **CREATE NEW GROUP** to create a new user-defined inventory group to use for this rule. Click **SAVE**.
Services	Click **Any** in the **Services** column and select a service from the list. Click **SAVE**.

Option	Description
Applied To	Define the type of traffic that the rule applies to: ■ Select **VPN Tunnel Interface** if you want the rule to apply to traffic over the route-based VPN. ■ Select **VPC Interface** if you want the rule to apply to traffic over the linked AWS VPC connection. ■ Select **Internet Interface** if you want the rule to apply to traffic over the Internet, including over policy-based VPNs using Public IP. ■ Select **Direct Connect Interface** if you want the rule to allow traffic over AWS Direct Connect (private VIF), including over policy-based VPNs using Private IP. ■ **All Uplinks** if you want the rule to apply to the **VPC Interface**, the **Internet Interface**, and the **Direct Connect Interface**, but not to the **VPN Tunnel Interface**. **Note** The **VPN Tunnel Interface** is not classified as an uplink.
Action	■ Select **Allow** to allow all L2 and L3 traffic to pass through the firewall. ■ Select **Drop** to drop packets that match any specified **Sources**, **Destinations**, and **Services**. This is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached. ■ Select **Reject** to reject packets that match any specified **Sources**, **Destinations**, and **Services**. This action returns a "destination unreachable message" to the sender. For TCP packets, the response includes a TCP RST message. For UDP, ICMP and other protocols, the response includes an "administratively prohibited" code (9 or 10). The sender is notified immediately (without any re-tries) when connection cannot be established.

The new rule is enabled by default. Slide the toggle to the left to disable it.

6 Click **PUBLISH** to create the rule.

The system gives the new rule an integer **ID** value, which is used in log entries generated by the rule.

What to do next

You can take any or all of these optional actions with an existing firewall rule.

- Click the gear icon ⚙ to view or modify rule logging settings. Log entries are sent to the VMwarevRealize Log Insight Cloud Service. See Using vRealize Log Insight Cloud in the *VMware Cloud on AWS Operations Guide*.

- Click the graph icon 📈 to view Rule Hits and Flow statistics for the rule.

Popularity Index	Number of times the rule was triggered in the past 24 hours.
Hit Count	Number of times the rule was triggered since it was created.

Packet Count	Total packet flow through this rule.
Byte Count	Total byte flow through this rule.

Statistics start accumulating as soon as the rule is enabled.

- Reorder firewall rules.

  A rule created from the **ADD NEW RULE** button is placed at the top of the list of rules. Firewall rules are applied in order from top to bottom. To change the position of a rule in the list, select it and drag it to a new position. Click **PUBLISH** to publish the change.

Add or Modify Distributed Firewall Rules

Distributed firewall rules apply at the VM (vNIC) level and control East-West traffic within the SDDC.

All traffic attempting to pass through the distributed firewall is subjected to the rules in the order shown in the rules table, beginning at the top. A packet allowed by the first rule is passed on to the second rule, and so on through subsequent rules until the packet is dropped, rejected, or hits the default rule, which allows all traffic.

Distributed firewall rules are grouped into policies. Policies are organized by category. Each category has an evaluation precedence. Rules in a category that has a higher precedence are evaluated before rules in category that has a lower precedence.

Table 2-15. Distributed Firewall Rule Categories

Category Evaluation Precedence	Category Name	Description
1	Ethernet	Applied to all layer 2 SDDC network traffic. **Note** Rules in this category require MAC addresses as sources and destinations. IP addresses are accepted but ignored.
2	Emergency	Used for quarantine and allow rules.
3	Infrastructure	Define access to shared services. Global rules, AD, DNS, NTP, DHCP, backup, management servers.
4	Environment	Rules between security zones such as production zones, development zones, or zones dedicated to specific business purposes.
5	Application	Rules between applications, application tiers, or microservices.

See Security Terminology in the *NSX-T Data Center Administration Guide* for more information about Distributed Firewall terminology.

Prerequisites

Distributed firewall rules require inventory groups as sources and destinations and must be applied to a service, which can be a predefined service or a custom service that you define for your SDDC. You can create these groups and services while you are creating a rule, but it can speed up the process if you take care of some of this beforehand. See Add or Modify a Compute Group and Add a Custom Service.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Select **Networking & Security > Distributed Firewall**.

   Click **CATEGORY SPECIFIC RULES** and select a category to view and modify policies and rules in that category, or click **ALL RULES** to view (but not modify) rules in all policies and categories.

3 (Optional) Change the default connectivity strategy.

   The Distributed Firewall includes default rules that apply to all layer 2 and layer 3 traffic. These rules are evaluated after all other rules in their category, and allow traffic that doesn't match a preceding rule to pass through the firewall. You can change either or both of these rules to be more restrictive, but you cannot disable either rule.

   - To change the **Default Layer2 Rule**, expand the **Default Layer2 Section** in the **Ethernet** category and change the **Action** on that rule to **Drop**.
   - To change the **Default Layer3 Rule**, expand the **Default Layer3 Section** in the **Application** category and change the **Action** on that rule to **Drop** or **Reject**.

   Click **PUBLISH** to update the rule.

4 To add a policy, open the appropriate category, click **ADD POLICY** and give the new policy a **Name**.

   A new policy is added at the top of the policy list for its category. To add a policy before or after an existing policy, click the vertical ellipsis button at the beginning of the policy row to open the policy settings menu, then click **Add Policy Above** or **Add Policy Below.**

   By default, the **Applied To** column is set to **DFW**, and the rule is applied to all workloads. You can also apply the rule or policy to selected groups. **Applied To** defines the scope of enforcement per rule, and is used mainly for optimization of host resource consumption. It helps in defining a targeted policy for specific zones and tenants, without interfering with other policy defined for other tenants and zones.

   Note Groups consisting of only IP addresses, MAC Addresses, or Active Directory groups cannot be used in the **Applied To** text box.

5 To add a rule, select a policy, click **ADD NEW RULE**, and give the rule a **Name**.

6 Enter the parameters for the new rule.

Parameters are initialized to their default values (for example, **All** for **Sources** and **Destinations**). To edit a parameter, move the mouse cursor over the parameter value and click the pencil icon (✏) to open a parameter-specific editor.

Option	Description
Sources	Click **Any** in the **Sources** column and select an inventory group for source network traffic, or click **ADD GROUP** to create a new user-defined inventory group to use for this rule. Click **SAVE**.
Destinations	Click **Any** in the **Destinations** column and select an inventory group for destination network traffic, or click **CREATE NEW GROUP** to create a new user-defined inventory group to use for this rule. Click **SAVE**.
Services	Click **Any** in the **Services** column and select a service from the list. Click **SAVE**.
Applied To	The rule inherits its **APPLIED TO** value from the containing policy.
Action	■ Select **Allow** to allow all L2 and L3 traffic to pass through the firewall. ■ Select **Drop** to drop packets that match any specified **Sources**, **Destinations**, and **Services**. This is a silent action with no notification to the source or destination systems. Dropping the packet causes the connection to be retried until the retry threshold is reached. ■ Select **Reject** to reject packets that match any specified **Sources**, **Destinations**, and **Services**. This action returns a "destination unreachable message" to the sender. For TCP packets, the response includes a TCP RST message. For UDP, ICMP and other protocols, the response includes an "administratively prohibited" code (9 or 10). The sender is notified immediately (without any re-tries) when connection cannot be established.

The new rule is enabled by default. Slide the toggle to the left to disable it.

7 (Optional) Configure advanced settings.

To change the directionality or logging behavior of the rule, click the gear icon ⚙ to open the **Settings** page.

Direction

By default, this value is **In/Out** and applies the rule to all sources and destinations. You can change this to **In** to apply the rule only to incoming traffic from a source, or **Out** to apply it only to outgoing traffic to a destination. Changing this value can cause asymmetric routing and other traffic anomalies, so be sure you understand the likely outcome for all sources and destinations before you change the default value for **Direction**.

Logging

Logging for a new rule is disabled by default. Slide the toggle to the right to enable logging of rule actions.

8 Click **PUBLISH** to create the rule.

The system gives the new rule an integer **ID** value, which is used to identify the rule in log entries it generates.

What to do next

You can take any or all of these optional actions with an existing firewall rule.

- Click the gear icon ⚙ to view or modify rule logging settings. Log entries are sent to the VMwarevRealize Log Insight Cloud Service. See Using vRealize Log Insight Cloud in the *VMware Cloud on AWS Operations Guide*.
- Click the graph icon 📈 to view Rule Hits and Flow statistics for the rule.

Popularity Index	Number of times the rule was triggered in the past 24 hours.
Hit Count	Number of times the rule was triggered since it was created.

Packet Count	Total packet flow through this rule.
Byte Count	Total byte flow through this rule.

Statistics start accumulating as soon as the rule is enabled.

- Reorder firewall rules.

A rule created from the **ADD NEW RULE** button is placed at the top of the list of rules in the policy. Firewall rules in each policy are applied in order from top to bottom. To change the position of a rule in the list, select it and drag it to a new position. Click **PUBLISH** to publish the change.

Manage Distributed Firewall Rules

Traffic attempting to pass through the firewall is subjected to the rules in the order shown in the **ALL RULES**.

The order of distributed firewall rules in the **ALL RULES** list is the union of the ordered list of policies and the ordered list of rules in each policy. You can reorder the distributed firewall sections and rules within a section. You can also edit existing distributed firewall configuration, delete, or clone a firewall rule or section.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Select **Networking & Security > Distributed Firewall**.

3 (Optional) Modify policy settings.

Click the vertical ellipsis button at the beginning of the policy row to take bulk actions, which affect all rules in the policy.

4 (Optional) Reorder policies.

A policy created from the **ADD POLICY** button is placed at the top of the list of policies. Firewall rules in each policy are applied in policy order from top to bottom. To change the position of a policy (and all the rules it contains) in the list, select it and drag it to a new position. Click **PUBLISH** to publish the change.

5 (Optional) Clone or copy a rule.

Click the vertical ellipsis button at the beginning of the rule row.

- **Clone Rule** to make a copy of the rule in this policy.
- **Copy Rule** to make a copy of the rule that you can add to another policy.

6 (Optional) Add or delete a rule.

Click the vertical ellipsis button at the beginning of the rule row.

- **Add Rule** to add a rule in this policy.
- **Delete Rule** to delete the rule from this policy.

7 (Optional) Save or view distributed firewall configurations.

Distributed firewall configurations in VMware Cloud on AWS are similar to the Firewall Drafts feature of on-premises NSX-T. Click **ACTIONS > Configurations > View** to view a list of saved configurations. Click **ACTIONS > Configurations > Save** to save the current configuration. Configurations are auto-saved by default. Click **ACTIONS > Settings > General Settings** to disable **Auto Save Drafts**.

Manage the Distributed Firewall Exclusion List

The Distributed Firewall Exclusion List lets you specify inventory groups to exclude from distributed firewall coverage. East-West network traffic to and from members of excluded groups is exempt from distributed firewall rules that would otherwise apply.

The Distributed Firewall exclusion list lets you keep specific inventory groups from being considered by distributed firewall rules. By default, management VMs and appliances, such as vCenter, NSX manager, and NSX controllers are on the exclusion list. You can edit the list to add or remove entries.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Select **Networking & Security > Distributed Firewall**.

3 Click **ACTIONS > Settings > Exclusion List** to display the **Manage Exclusion List** page.

- To add an existing group to the exclusion list, click **ADD GROUP** and select an existing **Group Name**.
- To create a group from the **Manage Exclusion List**, type a name for the group in the **Group Name** field, then click **Set Members** to open the inventory group creation page. See Add or Modify a Compute Group for more information about using this page.

- To remove a group from the list, click the vertical ellipsis button at the beginning of the group row and click **Delete**.

4 Click **APPLY** to save your changes.

Configure DNS Services

VMware Cloud on AWS DNS forwarding services run in DNS zones, and enable workload VMs in the zone to resolve fully-qualified domain names to IP addresses.

Your SDDC includes default DNS zones for the Management Gateway and Compute Gateway. Each zone includes a preconfigured DNS service. Use the **DNS Services** tab on the **DNS Services** page to view or update properties of DNS services for the default zones. To create additional DNS zones or configure additional properties of DNS services in any zone, use the **DNS Zones** tab.

For more information about DNS configuration choices for VMware Cloud on AWS, see DNS Strategies for VMware Cloud on AWS.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Select **Networking & Security > DNS**.

3 Click **DNS Services** to open the **DNS Services** page.

4 View or edit DNS service parameters.

Most gateway DNS service parameters are read-only but you can click the vertical ellipses button and choose **Edit DNS Server IPs** to add or modify the server IP addresses for this service.

5 Click **SAVE**.

Add a DNS Zone

Each DNS zone in your SDDC network represents a piece of the DNS namespace that you manage yourself.

DNS zones in the SDDC fall into two categories:

- Default zones, where the servers listen for DNS queries from all SDDC VMs on a subnet in the zone.
- FQDN zones, where the servers listen for DNS requests forwarded from a default zone.

The compute and management gateways are each configured with a single default DNS zone. You can add up to four more zones of either type to provide the flexibility of having multiple DNS servers and subdomains. See Add a DNS Zone in the *NSX-T Data Center Administration Guide* for more information about how NSX-T implements DNS zones.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Click **Networking & Security > DNS** and open the **DNS Zones** tab.

3 To add a default zone, select **ADD DNS ZONE > Add Default Zone**

   You can add or modify IP addresses for the Management Gateway and Compute Gateway DNS forwarders in the default DNS zone. DNS queries from VMs in the default zone are sent to these IP addresses by default if they don't match the criteria for any FQDN zone.

   a Enter a name and optionally a description. You use this **Name** if you create DNS firewall rules that apply to traffic in this zone.

   b Enter the IP addresses of up to three DNS servers. All of the DNS servers you specify must be configured identically.

   c (Optional) Enter an IP address in the **Source IP** field.

4 To add an FQDN zone, select **ADD DNS ZONE > Add FQDN Zone**

   Specify one or more FQDNs to enable DNS forwarding. A DNS forwarder is associated with a default DNS zone and up to five FQDN DNS zones. When it receives a DNS query from a VM in the zone, the DNS forwarder compares the domain name in the query with the domain names in the FQDN DNS zones. If a match is found, the query is forwarded to the DNS servers specified in the FQDN DNS zone. Otherwise the query is forwarded to the DNS servers specified in the default DNS zone.

   a Enter a name and optionally a description. You use this **Name** if you create DNS firewall rules that apply to traffic in this zone.

   b Enter a FQDN for the domain. This must be a fully qualified domain name, such as example.com.

   c Enter the IP address of up to three DNS servers.

   d (Optional) Enter an IP address in the **Source IP** field.

5 (Optional) Tag the DNS zone.

   See Add Tags to an Object in the *NSX-T Data Center Administration Guide* for more information about tagging NSX-T objects.

6 Click **SAVE**.

View Routes Learned and Advertised over VMware Transit Connect

In an SDDC that is a member of an SDDC Group, you can use the **Networking & Security Transit Connect** tool to view routes learned and advertised by this SDDC in the VMware Transit Connect network created for the group.

In an SDDC group, all network traffic between group members travels over a VMware Transit Connect network. Routing between compute networks of all SDDCs in a group is managed automatically by VMware Transit Connect as subnets are added and deleted. The **Transit Connect** and **SDDC Group** tools provide information about routes over that network. For information about creating an SDDC group or adding an SDDC to one, see Creating and Managing SDDC Deployment Groups in the *VMware Cloud on AWS Operations Guide*.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 On the **Networking & Security** tab, click **Transit Connect**, or just click the **SDDC Group** icon on the **Overview** page.

   The **Transit Connect** page displays lists of **Routes Learned** by this SDDC from other SDDCs in the group, and **Routes Advertised** by this SDDC to other SDDCs in the group. Click the download icon (⤓) to download either list in CSV format.

Configure a Multi-Edge SDDC With Traffic Groups

In the default configuration, your SDDC network has a single edge (T0) router through which all North-South traffic flows. This edge supports the default traffic group, which is not configurable. If you need additional bandwidth for the subset of this traffic routed to SDDC group members, a Direct Connect Gateway attached to an SDDC group, HCX Service Mesh, or to the connected VPC, you can reconfigure your SDDC to be Multi-Edge by creating traffic groups, each of which creates an additional T0 router.

A traffic group uses an association map to associate a prefix list of CIDR blocks to one of the T0 gateways that support non-default traffic groups in your SDDC. Prefix lists are independent of gateways and consist of source IP addresses. Traffic from those addresses is routed to the T0 edge that supports the associated traffic group. You can create and update prefix lists at any time, but you cannot remove a prefix list if it is included in an association map. Associating a prefix list with a traffic group routes all traffic from CIDR blocks in the list through the T0 router created for the group.

Note VPN traffic, as well as DX traffic to a private VIF must pass through on the default T0 and cannot be routed to a non-default traffic group. In addition, because NAT rules always run on the default T0 router, additional T0 routers cannot handle traffic affected by SNAT or DNAT rules. This includes traffic to and from the SDDC's native Internet connection. It also includes traffic to the Amazon S3 service, which uses a NAT rule and must go through the default T0. Keep these limitations in mind when you create prefix lists.

Prerequisites

- Before you can create traffic groups, you must use VMware Transit Connect™ to connect your SDDC to a VMware Managed Transit Gateway (VTGW). See Creating and Managing SDDC Deployment Groups in the *VMware Cloud on AWS Operations Guide*.

- Traffic groups can be created only in SDDCs that have large-size management appliances and at least four hosts. See Upsize SDDC Management Appliances for information about changing an SDDC's management appliance size from medium to large. See Add Hosts for information about adding hosts to an SDDC.
- The number of traffic groups that a multi-AZ (stretched cluster) SDDC can support depends on the number of hosts that the SDDC provides in each region, and can be represented with a formula like this:

  ```
  TG=(hosts-per-region - 2)/2
  ```

  where *TG* represents the maximum number of traffic groups that the SDDC can support and *hosts-per-region* is the number of hosts the SDDC deploys in each of the regions it occupies.

Procedure

1. Log in to the VMC Console at https://vmc.vmware.com.
2. Click **Networking & Security > Traffic Groups**.
3. Create a traffic group. On the **Traffic Groups** tab of the **Traffic Groups** page, click **ADD TRAFFIC GROUP** and give the new traffic group a **Name**, then click **SAVE** to create the traffic group and an additional T0 router for it.

   The **Status** of the traffic group transitions to **In Progress** while the new T0 edge is being created. It can take up to 30 minutes for the process to complete. When it does, the **Status** of the traffic group transitions to **Success** and you can create an association map for it.
4. Create a prefix list.

   Because Multi-Edge SDDCs use source-based routing in their traffic groups, prefix lists must contain source addresses, not destination addresses.

   a. On the **IP Prefix List** tab of the **Traffic Groups** page, click **ADD IP PREFIX LIST** and give the new prefix list a **Name** and optional **Description**.
   b. Click **Set** to display the **Set Prefixes** window, then click **ADD PREFIX** and fill in the CIDR block of an SDDC network segment that includes the source addresses of workload VMs whose traffic you want to include in the traffic group (and route over the additional edge).

      Important You cannot use the SDDC management CIDR block here or the CIDR block of a segment that provides the local IP address of a VPN. If you add any of these CIDRs to a prefix list, you won't be able to use the list in an association map.

      Click **ADD** to add the specified prefix to the list. To add prefixes or edit the ones already on the list, click the ellipsis buttons ⋮ to open the prefixes editor.
   c. Click **APPLY** to apply your changes to the prefix list.
   d. When you're done adding or editing prefixes, click **SAVE** to save or create the prefix list.

5 Associate a prefix list with a gateway. On the **Traffic Groups** tab of the **Traffic Groups** page, find the traffic group you want to work with, then click its ellipsis buttons and select **Edit**.

Click the plus icon ⊕ in the **ASSOCIATION MAPS** area, give the mapping a **Name** and select an existing prefix list from the **Prefixes** drop-down. Select a gateway from the **Gateway** drop-down, and click **SAVE** to create the association map.

6 (Optional) To remove a traffic group, you must first remove it association maps.

a Find the traffic group on the **Traffic Groups** page. Click its ellipsis button ⋮ , then select **Edit**.

b Click the minus icon ⊖ to the right of the **Status** label under **Association Maps** to select the map for deletion, then click **SAVE** to delete the map.

c Click **CLOSE EDITING**, then return to the traffic group on the **Traffic Groups** page. Click its ellipsis button and then select **Delete**.

It can take up to 30 minutes to remove a traffic group. Removing the traffic group removes the T0 router that was created to support it. HCX, if in use, creates its own association map, which you can view but not modify. To remove an association map created by HCX, you have to uninstall HCX. See Uninstalling VMware HCX in the *VMware HCX User Guide*.

Example: Route Table Changes After Adding a Traffic Group

This simplified example shows the effect of creating traffic group and associating it with a prefix list of just two host routes (/32).

Initial configuration

Assume these values for route table entries in the default traffic group and the Compute Gateway (CGW) before adding the first traffic group (which creates an additional T0 router).

Table 2-18. Default Routes

Subnet	Next Hop
0.0.0.0/0	Internet Gateway
192.168.150.51/24	CGW
192.168.151.0/24	CGW
VTGW, DXGW subnets	VTGW, DXGW connections
Management CIDR	MGW

Table 2-19. CGW Routes With the Default Traffic Group

Subnet	Next Hop
0.0.0.0/0	Default T0
192.168.150.0/24	Default T0
192.168.151.0/24	Default T0

Multi-Edge configuration

After the first traffic group is created, new routes are added on the default T0. Assuming that the prefix list associated with the traffic group has these entries:

```
192.168.150.100/32
192.168.151.51/32
```

then the route tables for the default T0, new T0, and CGW end up like this.

Table 2-20. Default T0 Routes After Adding a Traffic Group

Subnet	Next Hop
0.0.0.0/0	Internet Gateway
192.168.150.0/24	CGW
192.168.150.100/32	New T0
192.168.151.0/24	CGW
192.168.151.51/32	New T0
VTGW, DXGW subnets	VTGW, DXGW connections
Management CIDR	MGW

The new routes (192.168.150.100/32 and 192.168.151.51/32 in the example tables) use the new T0 as their next-hop, and the new T0 uses longest-prefix matching to route that traffic to the CGW.

Table 2-21. Routes on the New Traffic Group

Subnet	Next Hop
0.0.0.0/0	Default T0
192.168.150.100/32	CGW
192.168.151.51/32	CGW
VTGW, DXGW subnets	VTGW, DXGW connections
Management CIDR	MGW

The CGW route table is updated to create the traffic group by specifying the new T0 router as the next hop for the new routes.

Table 2-22. CGW Routes With an Additional Traffic Group

Subnet	Next Hop
0.0.0.0/0	Default T0
192.168.150.0/24	Default T0
192.168.150.100/32	New T0
192.168.151.0/24	Default T0
192.168.151.51/32	New T0

Working With Inventory Groups

Use VMware Cloud on AWS Networking & Security inventory to create groups of VMs and network services that you can use when you create firewall rules.

Firewall rules typically apply to a group of VMs that have certain common characteristics including:

- names that follow a naming convention (like Win* for Windows VMs or Photon* for Photon VMs)
- IP addresses within a specific range or CIDR block
- tags

They can also apply to network services, which are distinguished by characteristics like service type and network protocol. The VMware Cloud on AWS Networking & Security **Inventory** feature simplifies the process of creating groups of VMs that have similar needs for firewall protection. It also allows you to add new network services to the built-in list of services, so that you can include those services in firewall rules.

VMware Cloud on AWS creates management groups and a service inventory in all new SDDCs. It also maintains a list of your workload VMs and their tags. You can add or modify your own inventory groups of management or compute VMs.

Add a Management Group

Management inventory groups contain SDDC infrastructure components. Use these groups in management gateway firewall rules.

Pre-defined management inventory groups are created automatically for SDDC infrastructure components such as vCenter and NSX Manager. You cannot modify a pre-defined management group, but you can create additional management inventory groups by specifying the CIDR blocks to which group members are connected.

Procedure

1. Log in to the VMC Console at https://vmc.vmware.com.
2. On the **Networking & Security** tab, click **Inventory > Groups**.

3 On the **Groups** card, click **MANAGEMENT GROUPS**, then click **ADD GROUP** and give the group a **Name** and an optional **Description**.

4 Click **Set Members** to open the **Select Members** page.

   Enter one or more IP addresses of management VMs in CIDR format.

5 (Optional) Tag the group.

   See Add Tags to an Object in the *NSX-T Data Center Administration Guide* for more information about tagging NSX-T objects.

6 Click **SAVE** to create the group.

What to do next

You can modify or delete any management group that you created by clicking the vertical ellipsis button and selecting **Edit** or **Delete**.

Add or Modify a Compute Group

Compute inventory groups categorize compute VMs using criteria such as names, IP addresses, and tags.

Because compute inventory groups are made up of the compute VMs you deploy on your compute network segments. VMware Cloud on AWS cannot create them for you. You'll need to create them yourself before you can develop compute gateway firewall rules.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 On the **Networking & Security** tab, click **Inventory > Groups**.

3 On the **Groups** card, click **COMPUTE GROUPS**, then click **ADD GROUP** and give the group a **Name** and an optional **Description**.

   To modify an existing group, select it and click the ellipsis button at the beginning of the group row.

4 Click **Set Members** to open the **Select Members** page.

Management groups contain VMs on the Management Network. Management group members must be specified by IP address. Compute groups contain VMs or network objects such as segments in the Compute network. There are several ways to designate membership in a compute group.

Option	Description

Membership Criteria	Click **ADD CRITERIA** and use the drop-down controls to specify one or more criteria in the form of
	`Object Type, Property, Condition, Value`
	tuples. For example,a group with these criteria:
	`Virtual Machine Name Contains db_`
	includes VMs whose names contain the string `db_` in the group. You can also create groups of tagged network segments, segment ports, or IP sets by specifying a tag, or
	`Segment Tag Equals testbeds`
	to include network segments that have the tag `testbeds`.
	Objects that match all of the selected criteria are included in the group.

Members	Select a membership category from the **Select Category** drop-down list, then select members from the list.
IP/MAC address	Enter an IP address, MAC addresses, CIDR block, or a range of IP addresses in the form *ip-ip* (for example `192.168.1.1–192.168.1.100`) .

5 (Optional) Tag the group.

See Add Tags to an Object in the *NSX-T Data Center Administration Guide* for more information about tagging NSX-T objects.

6 Click **SAVE** to create the group.

What to do next

To review group members, select a group and click **View Members** to review the list of group members to view group members and membership criteria. Click **Where Used** to see a list of firewall rules that include the group.

Add a Custom Service

Firewall rules often apply to traffic from a network service. A new SDDC includes inventory entries for most of the common network service types, but you can add custom services if you need to.

When you create a firewall rule, you can specify that it applies to network traffic from one or more of the services defined in your SDDC's **Services** inventory. The default list includes VMware services such as remote console and provisioning, standard services such as IKE, ICMP, and TCP, and many well-known third party services. You can add services to this list by selecting values, typically ports and protocols, from a list of service types and additional service properties.

Procedure

1. Log in to the VMC Console at https://vmc.vmware.com.
2. On the **Networking & Security** tab, click **Inventory > Services**.

   The **Services** card lists the predefined services.
3. Click **ADD NEW SERVICE** and give the new service a **Name**.
4. Click **Set Service Entries** to open the **Set Service Entries** page.
5. On the **Set Service Entries** page, click **ADD SERVICE ENTRY**.

   To view the list of known services, use the drop-down controls to scroll through the **Service Type** and **Additional Properties** lists. To add a service, select a **Service Type** from the drop-down menu and specify **Additional Properties** such as Source or Destination Ports of the service, then click **APPLY**.
6. (Optional) Provide a service **Description** and tag the service.

   See Add Tags to an Object in the *NSX-T Data Center Administration Guide* for more information about tagging NSX-T objects.
7. Click **SAVE** to create the service definition.

View Virtual Machine Inventory

VMware Cloud on AWS maintains an inventory of workload virtual machines in your SDDC. VMs are listed by name and number of tags.

The **Virtual Machines** inventory is generated automatically. You can edit the tags assigned to a VM in this list, but you cannot add or remove VMs. The system does that automatically as VMs are created and destroyed.

Procedure

1. Log in to the VMC Console at https://vmc.vmware.com.
2. On the **Networking & Security** tab, click **Inventory > Virtual Machines**.

   If a virtual machine has any tags, the number of tags is shown in the **Tags** column. Click the number to view the tags. To add or remove VM tags, click the vertical ellipsis at the beginning of the VM row and select **Edit** to display the tag editor. Click the ⊕ icon to add more tags.

   See Add Tags to an Object in the *NSX-T Data Center Administration Guide* for more information about tagging NSX-T objects.

Managing Workload Connections

Workload VMs connect to the Internet by default. NAT rules and distributed firewall rules give you fine-grained control over these connections.

Workload VMs can communicate with each other over their private or public (NATted) IP addresses. When using public IPs, workload-to-workload communication traffic is subject to these rules:

- The traffic is not subject to CGW firewall rules.
- Distributed firewall rule processing by a source VM uses the destination public IP address and source public IP of the destination VM, and must be IP-based. Distributed firewall rules based on VM attributes do not affect workload-to-workload traffic.

Note Workload VM communication to the vCenter Server public IP is subject to MGW firewall rules, but the workload VM IP is translated to its public IP before the firewall rule is applied.

Attach a VM to or Detach a Workload VM from a Compute Network Segment

Use the vSphere Web Client to manage attachment of workload VMs to compute network segments.

Prerequisites

Your SDDC compute network must have at least one segment. See Create or Modify a Network Segment.

Procedure

1. Log in to the vSphere Client for your SDDC.
2. Select **Menu > Global Inventory Lists**.
3. Select **Logical Networks**.
4. In the **vCenter Server** drop down menu, select the vCenter Server that manages the logical network you want to use.

5 Click next to the logical network name to select it.

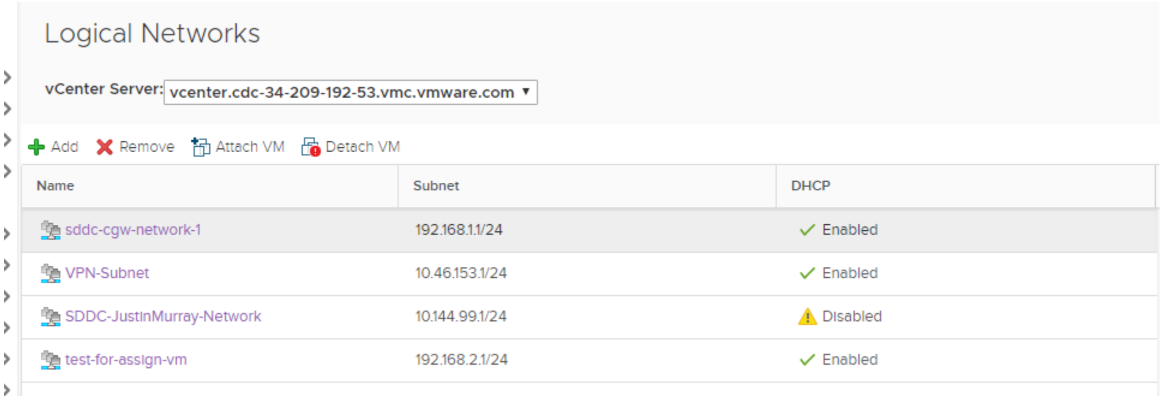

Logical Networks

vCenter Server: vcenter.cdc-34-209-192-53.vmc.vmware.com ▼

Name	Subnet	DHCP
sddc-cgw-network-1	192.168.1.1/24	✓ Enabled
VPN-Subnet	10.46.153.1/24	✓ Enabled
SDDC-JustinMurray-Network	10.144.99.1/24	⚠ Disabled
test-for-assign-vm	192.168.2.1/24	✓ Enabled

6 Select whether to attach or detach VMs.

- Click **Attach VM** to attach VMs to the selected network.
- Click **Detach VM** to detach VMs from the selected network.

7 Select the virtual machine(s) you want to attach or detach, click **>>** to move them to the **Selected Objects** column, and click **Next**.

8 For each VM, select the virtual NIC you want to attach and click **Next**.

9 Click **Finish**.

Request or Release a Public IP Address

You can request public IP addresses to assign to workload VMs to allow access to these VMs from the internet. VMware Cloud on AWS provisions the IP address from AWS.

As a best practice, release the public IP addresses that are not in use.

Prerequisites

Verify that your VM has a static IP address assigned from its logical network.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Select **Networking & Security > Public IPs**.

3 To request a new public IP address, click **REQUEST NEW IP**.

You can optionally enter your notes about the request.

4 To release an existing public IP address that you no longer need, click the ellipsis button and click **Release IP**.

Requests to release a public IP address fail if the address is in use by a NAT rule.

5 Click **SAVE**.

After a few moments, a new Public IP address is provisioned.

What to do next

After the public IP address is provisioned, configure NAT to direct traffic from the public IP address to the internal IP address of a VM in your SDDC. See Create or Modify NAT Rules.

Create or Modify NAT Rules

Network Address Translation (NAT) maps internal IP addresses on your compute network to addresses exposed on the public Internet. To create a NAT rule, you provide the internal address and port number of a workload VM or service and a public IP address and port number that you have obtained from the system.

NAT rules on the SDDC network’s internet interface, since that's where your workload VMs' public addresses are exposed. Firewall rules, which examine packet sources and destinations, run on the Compute Gateway, and process traffic after it has been transformed by any applicable NAT rules. When you create a NAT rule, you can specify whether a VM's internal or external IP address and port number are exposed to firewall rules that affect network traffic to and from that VM.

Important Inbound traffic to the SDDC's public IP address is always processed by the NAT rules you create. Outbound traffic (reply packets from SDDC workload VMs) is routed along the advertised routes and is processed by NAT rules when the default route for your SDDC network goes through the SDDC's Internet interface. But if the default route goes through a Direct Connect or VPN connection (for example, if 0.0.0.0/0 is advertised through BGP or there is a policy-based VPN with a remote network of 0.0.0.0/0), NAT rules run for inbound traffic but not for outbound traffic, creating an asymmetric path that leaves the VM unreachable at its public IP address. When the default route is advertised from the on-premises environment, you must configure NAT rules on the on-premises network, using the on-premises Internet connection and public IPs.

Prerequisites

- You must have obtained a public IP address for use by a VM in this SDDC. See Request or Release a Public IP Address.
- The VM must be connected to a compute network segment. You can create NAT rules for VMs whether they have static or dynamic (DHCP) addresses, but bear in mind that NAT rules for VMs using DHCP address assignment can be invalidated when the VM is assigned an internal address that no longer matches the one specified in the rule.

Procedure

1. Log in to the VMC Console at https://vmc.vmware.com.
2. Select **Networking & Security > NAT** .
3. Click **ADD NAT RULE** and give the rule a **Name**.

4 Enter the NAT rule parameters.

Option	Description
Public IP	Choose from the drop-down list of public IP address that have been provisioned for this SDDC. See Request or Release a Public IP Address.
Service	■ Select **All Traffic** to create a rule that applies to both inbound (DNAT) and outbound (SNAT) traffic to or from the specified **Internal IP**. ■ Select one of the listed services to create an inbound (DNAT) rule that applies only to traffic using that protocol and port. **Note** Because services that use multiple destination ports cannot be subject to a NAT rule, they don't appear on this list.
Public Port	If you specified **Service** as **All Traffic**, the default public port is **Any**. If you selected a particular **Service**, then the rule applies to the assigned public port for that service.
Internal IP	Enter the internal IP address of the VM.
Internal Port	Displays the internal port used by the selected **Service**. To use a custom port, Add a Custom Service, then select that **Service** in the NAT rule. If you specified **Service** as **All Traffic**, the default internal port is **Any**. If you selected a particular **Service**, then the rule applies to the assigned public port for that service.
Firewall	Specify how traffic subject to this NAT rule is exposed to Compute Gateway firewall rules. By default, CGW firewall rules match the combination of **Internal IP** and **Internal Port**. Select **Match External Address** to have firewall rules match the combination of **External IP** and **External Port**. (Distributed firewall rules never apply to external addresses or ports.)

You can create multiple NAT rules that use the same **Public IP** and **Internal IP** with **All Traffic**. If you do this, each **Internal IP** uses the **Public IP** for outbound (SNAT) traffic, but only the first matching rule will be used for inbound (DNAT) traffic. The system creates (but does not display) a default outbound rule,. This rule is used for all **Internal IP** addresses that do not match a specific NAT rule that applies to **All Traffic**. The IP used for this rule is displayed I the **Default Compute Gateway** summary on the **Networking & Security Overview** page as **Source NAT Public IP**.

5 (Optional) Toggle **Logging** to log rule actions.

6 The new rule is enabled by default. Toggle **Enable** to disable it.

7 Click **SAVE** to create the rule.

The rule is created and its **Status** is reported as **Up**.

Creating Firewall Rules to Manage Traffic Between the Compute and Management Networks

In the default configuration, firewall rules prevent VMs on the compute network from accessing VMs on the management network. To allow individual workload VMs to access management

VMs, create Workload and Management inventory groups, then create management gateway firewall rules that reference them.

Procedure

1 Create Workload inventory groups: one for the management network and one for the workload VM that you want to have access to it.

   On the **Networking & Security tab**, click **Groups** in the **Inventory** category, then click **Workload Groups**. Create two workload groups:

   - Click **ADD GROUP** and create a group with a **Member Type** of IP address and the CIDR block of the management network. Click **SAVE** to create the group.
   - Click **ADD GROUP** and create a group with a **Member Type** of Virtual Machine and a Member VM from your vSphere inventory. Click **SAVE** to create the group.

2 Create a Management inventory group to represent the management network that you want to access from the Workload group.

   On the **Networking & Security tab** tab, click **Groups** in the **Inventory** category, then click **Management Groups**. Click **ADD GROUP** and create a group with a **Member Type** of IP address and the management network CIDR block. Click **SAVE** to create the group.

3 Create a management gateway firewall rule allowing inbound traffic to the vCenter server and ESXi.

   See Add or Modify Management Gateway Firewall Rules for information about creating management gateway firewall rules. Assuming your workload VMs only need to access vSphere, PowerCLI, or OVFtool on vCenter and ESXi, then the rule need only allow access on port 443.

   Table 2-23. Management Gateway Rule to Allow Inbound Traffic to ESXi and vCenter

Name	Source	Destination	Services	Action
Inbound to ESXi	Workload VM private IP	ESXi	HTTPS (TCP 443)	Allow
Inbound to vCenter private IP	Workload VM private IP	vCenter private IP	HTTPS (TCP 443)	Allow
Inbound to vCenter public IP	Workload VM with NATted IP	vCenter public IP	HTTPS (TCP 443)	Allow

Configure Monitoring and Troubleshooting Features

3

Use IPFIX and Port Mirroring functionality provided by NSX-T to monitor and troubleshoot SDDC networking and security.

By default, SDDC ESXi hosts have access to the overlay network, allowing them to communicate with monitoring and troubleshooting applications deployed as VM workloads in your SDDC. However, you must configure the firewall to allow traffic between the ESXi hosts and the logical segment the workload VMs are attached to. See Creating Firewall Rules to Manage Traffic Between the Compute and Management Networks.

- Configure IPFIX

  IPFIX (Internet Protocol Flow Information Export) is a standard for the format and export of network flow information for troubleshooting, auditing, or collecting analytics information.

- Configure Port Mirroring

  Port mirroring lets you replicate and redirect all of the traffic coming from a source. The mirrored traffic is sent encapsulated within a Generic Routing Encapsulation (GRE) tunnel to a collector so that all of the original packet information is preserved while traversing the network to a remote destination.

- View Connected VPC Information

  The Connected Amazon VPC contains your SDDC and all its networks. Information about this VPC, including the active ENI, VPC subnet,and VPC ID, is available on the **Networking & Security** tab.

Configure IPFIX

IPFIX (Internet Protocol Flow Information Export) is a standard for the format and export of network flow information for troubleshooting, auditing, or collecting analytics information.

You can configure flow monitoring on a logical segment. All the flows from the VMs connected to that logical segment are captured and sent to the IPFIX collector. The collector names are specified as a parameter for each IPFIX switch profile.

Note In an SDDC that is a member of an SDDC group, all outbound traffic from hosts to destinations outside the SDDC network is routed to the VTGW or private VIF regardless of other routing configurations in the SDDC. This includes IPFIX and Port Mirroring traffic. See Creating and Managing SDDC Deployment Groups with VMware Transit Connect in the *VMware Cloud on AWS Operations Guide*.

Prerequisites

Verify that a logical segment is configured. See Create or Modify a Network Segment.

Procedure

1 Log in to the VMC Console at https://vmc.vmware.com.

2 Select **Networking & Security > IPFIX**.

3 To add a new collector, click **COLLECTORS > ADD NEW COLLECTOR** and give the collector a **Name**.

   Enter the collector IP address and port. The default UDP port is 4739. You can add up to 4 IPFIX collectors.

4 Click **SAVE** to create the collector.

5 Click **SWITCH IPFIX PROFILES** to create or edit a switch IPFIX profile.

   See Configure Switch IPFIX Profiles in the *NSX-T Data Center Administration Guide* for more information about NSX-T switch IPFIX profile parameters.

6 (Optional) Tag the profile.

   See Add Tags to an Object in the *NSX-T Data Center Administration Guide* for more information about tagging NSX-T objects.

7 Click **SAVE** to create the profile.

What to do next

Click the ellipses button next to a switch IPFIX profile and click **Edit** to make configuration changes.

Configure Port Mirroring

Port mirroring lets you replicate and redirect all of the traffic coming from a source. The mirrored traffic is sent encapsulated within a Generic Routing Encapsulation (GRE) tunnel to a collector so that all of the original packet information is preserved while traversing the network to a remote destination.

Port mirroring is used in the following scenarios:

- Troubleshooting - Analyze the traffic to detect intrusion and debug and diagnose errors on a network.
- Compliance and monitoring - Forward all of the monitored traffic to a network appliance for analysis and remediation.

Port mirroring includes a source group where the data is monitored and a destination group where the collected data is copied to. The source group membership criteria require VMs to be grouped based on the workload such as web group or application group. The destination group membership criteria require VMs to be grouped based on IP addresses.

Port mirroring has one enforcement point, where you can apply policy rules to your SDDC environment.

The traffic direction for port mirroring is Ingress, Egress, or Bi Directional traffic.

- Ingress is the outbound network traffic from the VM to the logical network.
- Egress is the inbound network traffic from the logical network to the VM.
- Bi Directional is the traffic from the VM to the logical network and from the logical network to the VM. This is the default option.

See Add a Port Mirroring Profile in the *NSX-T Data Center Administration Guide* for more information about port mirroring with NSX-T.

Note In an SDDC that is a member of an SDDC group, all outbound traffic from hosts to destinations outside the SDDC network is routed to the VTGW or private VIF regardless of other routing configurations in the SDDC. This includes IPFIX and Port Mirroring traffic. See Creating and Managing SDDC Deployment Groups with VMware Transit Connect in the *VMware Cloud on AWS Operations Guide*.

Prerequisites

Important Port mirroring can generate a lot of network traffic. As a best practice, limit its use to a maximum of 6 VMs at a time for short periods of troubleshooting and remediation.

Verify that workload groups with IP address and VM membership criteria are available. See Add or Modify a Compute Group.

Procedure

1. Log in to the VMC Console at https://vmc.vmware.com.
2. Select **Networking & Security > Port Mirroring**.
3. On the **Port Mirroring** page Click **ADD PROFILE** and give the profile a **Name** and an optional **Description**.

4 Specify the profile parameters.

Parameter	Description
Direction	Select a traffic direction from the drop-down list.
Snap Length	Specify the number of bytes to capture from a packet.
Source	Sources can include segments, segment ports, groups of VMs, and groups of vNICs.
Destination	Destinations are groups of up to three IP addresses. You can use existing inventory groups or create new ones from the **Set Destination** page.
Encapsulation Type	Must be **GRE**.
GRE Key	Identifies a particular GRE data stream, as defined in RFC 6245. Enter a random 32-bit value to identify mirrored packets from the logical port. This Key value is copied to the Key field in the GRE header of each mirror packet. If the Key value is set to 0, the default definition is copied to the Key field in the GRE header. The default 32-bit value is made of the following values. ■ The first 24-bit is a VNI value. VNI is part of the IP header of encapsulated frames. ■ The 25th bit indicates if the first 24-bit is a valid VNI value. One represents a valid value and zero represents an invalid value. ■ The 26th bit indicates the direction of the mirrored traffic. One represents an ingress direction and zero represents an egress direction. ■ The remaining six bits are not used.

5 (Optional) Tag the port mirroring profile.

See Add Tags to an Object in the *NSX-T Data Center Administration Guide* for more information about tagging NSX-T objects.

6 Click **SAVE** to save the session.

What to do next

Click the ellipses button next to a port mirroring profile and select **Edit** to make configuration changes.

View Connected VPC Information

The Connected Amazon VPC contains your SDDC and all its networks. Information about this VPC, including the active ENI, VPC subnet,and VPC ID, is available on the **Networking & Security** tab.

Click **Connected VPC** in the **System** category on the **Networking & Security** tab to open the **Connected Amazon VPC** page, which provides the following information:

AWS Account ID

The AWS account ID you specified when you created your SDDC.

VPC ID

The AWS ID of this VPC.

VPC Subnet

The AWS ID of the VPC subnet you specified when you created your SDDC.

Active Network Interface

The identifier for the ENI used by VMC in this VPC.

IAM Role Names

AWS Identity and Access Management role names defined in this VPC. See AWS Roles and Permissions in the *VMware Cloud on AWS Operations Guide*.

Cloud Formation Stack Names

The name of the AWS Cloud Formation stack used to create your SDDC

Service Access

A list of AWS services enabled in this VPC.