NSX Security Quick Start Guide

17 MAY 2022

VMware NSX-T Data Center 3.2

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

Copyright © 2022 VMware, Inc. All rights reserved. Copyright and trademark information.

Contents

NSX Security Quick Start Guide

The *VMware NSX Security Quick Start Guide* provides basic information about deploying and configuring VMware NSX® Security.

Intended Audience

This information is intended for network security administrators and system administrators who want to deploy, configure, or use VMware NSX Security. The information is written for experienced enterprise system administrators who are familiar with virtual machine technology, networking, and security operations and with the administration of vSphere and VMware NSX-T™ Data Center.

Scope of the Document

This document provides basic information on how to deploy the NSX management plane in an on-premises environment and how to configure your system for Distributed Firewall and Gateway Firewall. Deployment and configuration of the following advanced features are outside the scope of this document.

- Distributed Firewall
  - User ID
  - L7 Application ID
  - FQDN Analysis
  - Malware Prevention
- Gateway Firewall
  - User ID
  - L7 Application ID
  - URL Filtering
  - Malware Detection
- NSX Application Platform Deployment
- NSX Intelligence Deployment
- NSX Intrusion Detection and Prevention Service (IDS/IPS)

- NSX Network Detection and Response

For detailed instructions on these features, refer to the following documents:

- *NSX-T Data Center Administration Guide*
- *Deploying and Managing the VMware NSX Application Platform*
- *VMware NSX Network Detection and Response Activation and Administration Guide*
- *Installing and Upgrading VMware NSX Intelligence*
- *Using and Managing VMware NSX Intelligence*

Related Documentation

- For detailed instructions on the NSX-T Data Center installation, refer to the *NSX-T Data Center Installation Guide*.
- For detailed instructions on the NSX-T Data Center administrative tasks, refer to the *NSX-T Data Center Administration Guide*.
- For detailed instructions on how to upgrade the NSX-T Data Center, refer to the *NSX-T Data Center Upgrade Guide*.

VMware NSX-T Security Overview

1

VMware NSX-T builds security into the network virtualization infrastructure. There are many built-in services that are part of NSX-T that enhance security. Security teams can protect the data center traffic across virtual, physical, containerized, and cloud workloads. The security capabilities are always present in the infrastructure and are quickly configurable. Further, no one can tamper with the security controls because they reside in the hypervisor, effectively decoupling the controls from your workloads.

Workloads come in various form factors like virtual machines, containers, and physical servers. In addition, workloads are hosted in different environments like on-premises, native cloud, or managed cloud. The heterogeneity of the workload form factor and deployment type further challenges the organizations regarding security coverage, policy consistency, number of platforms to be managed, and overall operational simplicity. The requirement of an organization is to have an operationally simple platform that provides consistent policy across virtual machines, containers, physical servers, and native cloud workloads without compromising the application and data security.

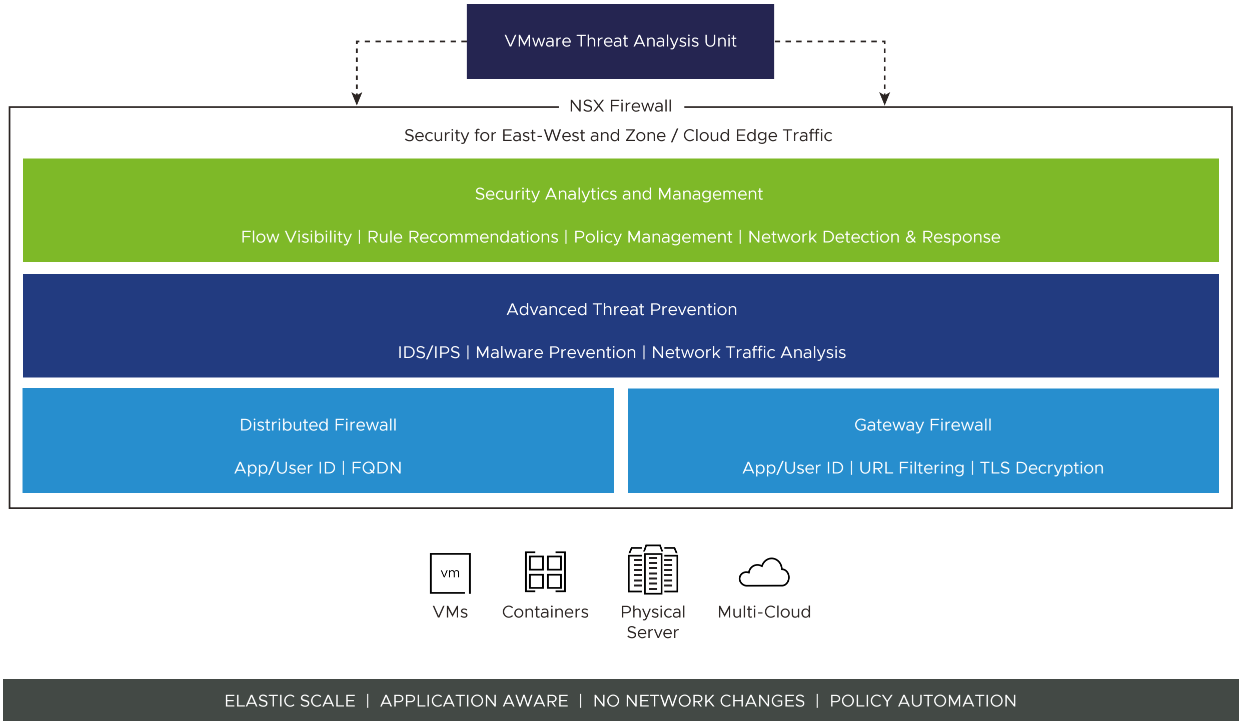
VMware Threat Analysis Unit
NSX Firewall
Security for East-West and Zone / Cloud Edge Traffic
Security Analytics and Management
Flow Visibility | Rule Recommendations | Policy Management | Network Detection & Response
Advanced Threat Prevention
IDS/IPS | Malware Prevention | Network Traffic Analysis
Distributed Firewall
App/User ID | FQDN
Gateway Firewall
App/User ID | URL Filtering | TLS Decryption
vm
VMs
Containers
Physical Server
Multi-Cloud
ELASTIC SCALE | APPLICATION AWARE | NO NETWORK CHANGES | POLICY AUTOMATION

NSX-T has a distributed architecture. Security enforcement controls are located at the virtual network interface of each workload and provide a granular mechanism to police traffic flows. There is no centralized appliance that limits security capacity, and you do not need to artificially hairpin the network traffic to a network security stack. As NSX-T is integrated into the virtualization infrastructure, it has visibility into all applications and workloads. NSX-T uses this visibility to derive rich application context, closely track the life cycle of workloads and automate security policy management.

NSX Distributed Firewall (DFW) is a distributed, scale-out internal firewall that protects all East-West traffic across all workloads without network changes, thereby radically simplifying the security deployment model. It includes a stateful L4-L7 firewall, an intrusion detection/prevention system (IDS/IPS), network sandbox, and behavior-based network traffic analysis. With the NSX Firewall, you can protect the data center traffic across virtual, physical, containerized, and cloud workloads from internal threats and avoid damage from threats that make it past the network perimeter.

NSX Gateway Firewall is instantiated per gateway and is supported at both Tier-0 and Tier-1. The Gateway Firewall provides firewalling services and other services that cannot be distributed such as NAT, DHCP, VPN, and load balancing, and needs the services router component of the gateway. Gateway firewall works independently of NSX-T DFW from a policy configuration and enforcement perspective, although you can share objects from the DFW.

NSX Intelligence, a security analytics and policy management solution, automatically determines the communication patterns across all types of workloads, makes security policy recommendations based on those patterns, and checks that traffic flows conform to the deployed policies.

NSX Security Deployment Workflow for On-Premises Environment

2

To benefit all the NSX Security features, you must deploy the NSX Management Plane in an on-premises environment.

After NSX-T is installed, configure the system for different security features. The workflow in this guide includes minimal deployment and configuration instructions required to set up the security features. For more detailed instructions for each feature, see *NSX-T Data Center Installation Guide* and *NSX-T Data Center Administration Guide*.

NSX Firewall – for all Deployment Options

NSX Firewall provides different security controls like Distributed Firewall, Distributed IDS/IPS, Distributed Malware Prevention, and Gateway Firewall as an option to provide firewalling to different deployment scenarios.

A typical data center can have different workloads like VMs, Containers, Physical Server, and a mix of NSX managed and non-managed workloads. These workloads can have a combination of a VLAN-based network or an NSX-based overlay network.

The following diagram summarizes different data center deployment scenarios and associated NSX firewall security controls, which best fits the design. You can use the same NSX Manager as a single pane of glass to define the security policies for all these different scenarios using different security controls.

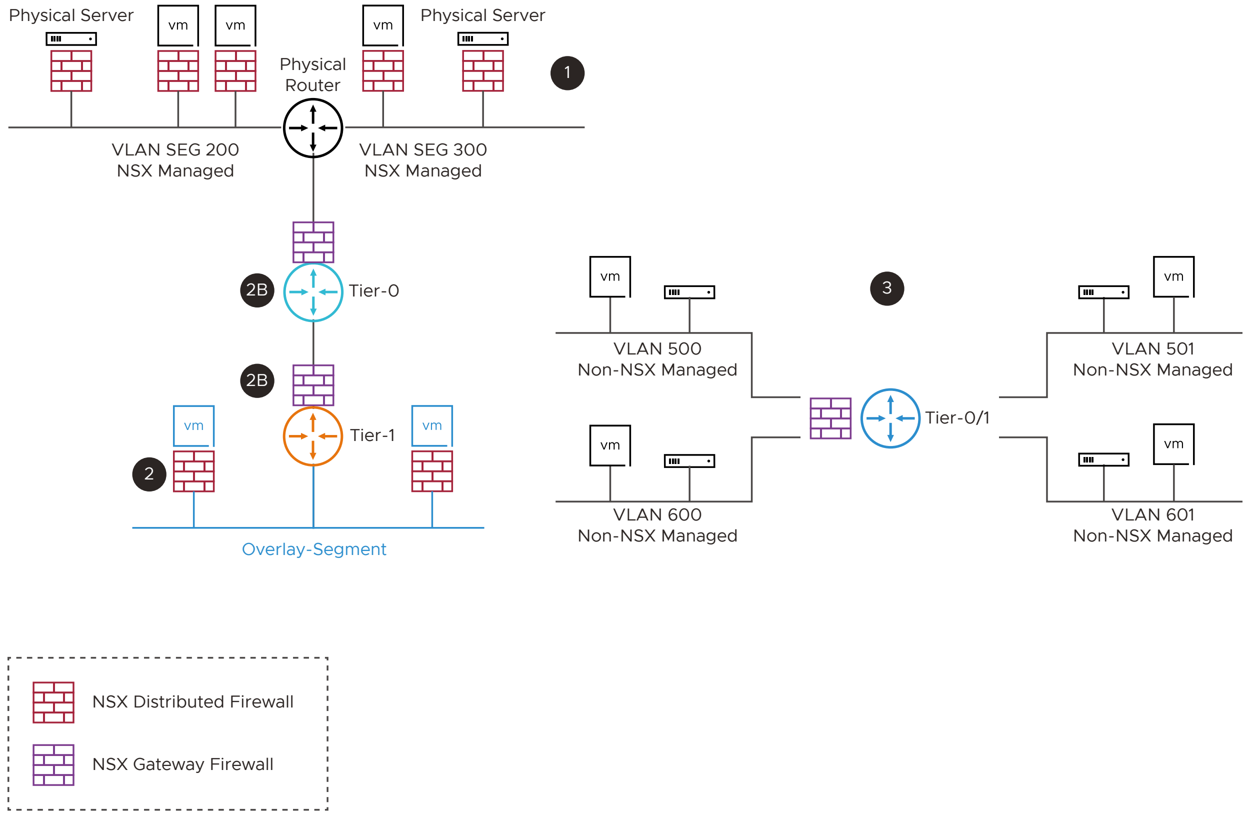

1 `NSX Managed Workloads with Standard VLAN based Networking`:

NSX distributed firewalling capability can be used to protect NSX managed VMs and Physical Server workloads.

2 `NSX Managed Workloads with NSX Overlay for Networking`:

- NSX Distributed Firewall can be used to protect NSX managed VMs, Containers (using NSX container plug-in), and Physical Server workloads from East-West traffic perspective. This can be used for Zone-segmentation, Application-segmentation, and Micro-segmentation with both L3-L7 firewalling and IDS/IPS capabilities.

- NSX Gateway Firewall can be used as inter-tenant/zone firewall from North-South perspective, along with the distributed firewall.

3 `Non-NSX Managed Workloads on Traditional VLAN based Network`:

NSX gateway firewalling capability can provide the inter VLAN routing and firewalling. The Service Interface on NSX tier-1 gateway or external interface on the tier-0 gateway is used as a gateway and firewall for all the non-NSX managed VLAN workloads.

Basic Deployment Workflow

Action	Description	Details
Deploy the NSX Management Plane	■ Download the NSX Manager OVA ■ Deploy the NSX Manager appliance	Deploying NSX Management Plane
Configure for NSX Distributed Security	■ Distributed Security for VMs ■ Distributed Security for Physical Server	Preparing for Distributed Security
Configure for NSX Gateway Security	■ Gateway Security for VMs or Physical Server ■ Gateway Security for VM with NSX Network Virtualization (Overlay)	Preparing For Gateway Security

This guide does not cover advanced features supported by NSX Firewall. For more detailed instructions for each feature, see *NSX-T Data Center Administration Guide*.

This chapter includes the following topics:

- Deploying NSX Management Plane
- Preparing for Distributed Security
- Preparing For Gateway Security
- Configuring Security Policy

Deploying NSX Management Plane

You can use the NSX Manager as a single pane of glass to define Security policies for different scenarios using different security controls.

NSX Manager is the application that you use to administer your NSX-T environment. The NSX Manager provides a web-based graphical user interface (GUI) and REST APIs for creating, configuring, and monitoring NSX-T Data Center components. The basic step of deploying NSX-T in your environment involves deploying NSX Manager, preparing ESXi host as a host transport node (for Distributed Firewall), and deploying NSX Edge VMs (for NSX Gateway Firewall).

Prerequisites

Before installing NSX-T Data Center, make sure your environment is ready.

Review the NSX Manager installation requirements.	See NSX Manager Installation
If NSX system components are behind a firewall, add a policy to allow the relevant ports. See details on Ports and Protocols.	See Ports and Protocols

- Make sure you have supported vCenter Server/ ESXi versions.
- You have configured the DNS and NTP servers correctly.

For more details, see:

- NSX-T Data Center Quick Start Guide
- NSX-T Data Center Installation Guide

1: Download the NSX Manager OVA

1 Go to the VMware downloads page at https://my.vmware.com/web/vmware/downloads.

NSX Security is available under **Networking & Security**.

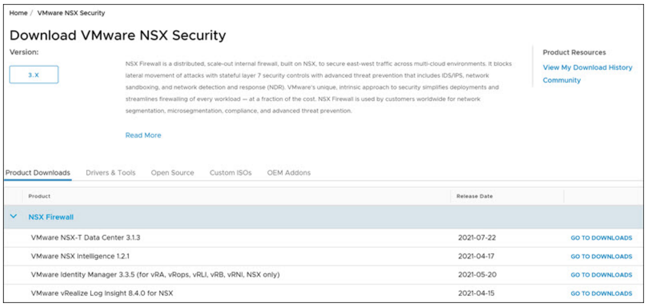

Home / VMware NSX Security

Download VMware NSX Security

Version:

3.X

NSX Firewall is a distributed, scale-out internal firewall, built on NSX, to secure east-west traffic across multi-cloud environments. It blocks lateral movement of attacks with stateful layer 7 security controls with advanced threat prevention that includes IDS/IPS, network sandboxing, and network detection and response (NDR). VMware's unique, intrinsic approach to security simplifies deployments and streamlines firewalling of every workload — at a fraction of the cost. NSX Firewall is used by customers worldwide for network segmentation, microsegmentation, compliance, and advanced threat prevention.

Read More

Product Resources

View My Download History

Community

Product Downloads | Drivers & Tools | Open Source | Custom ISOs | OEM Addons

Product	Release Date	
NSX Firewall		
VMware NSX-T Data Center 3.1.3	2021-07-22	GO TO DOWNLOADS
VMware NSX Intelligence 1.2.1	2021-04-17	GO TO DOWNLOADS
VMware Identity Manager 3.3.5 (for vRA, vRops, vRLI, vRB, vRNI, NSX only)	2021-05-20	GO TO DOWNLOADS
VMware vRealize Log Insight 8.4.0 for NSX	2021-04-15	GO TO DOWNLOADS

2 Download the NSX Manager NSX Global Manager / NSX Cloud Service Manager for VMware ESXi Open Virtualization Appliance (OVA) file. You can download the OVA to a local datastore or a local web server. If you downloaded the OVA file to a local web server, copy the file path of the NSX Manager appliance OVA file. For example, `http://<local-web-server>/nsx-unified-appliance-3.2.0.0.<buildnumber>.ova`. You should provide this path while deploying the appliance.

2: Deploy the NSX Manager

You must deploy the NSX Manager appliance and register the appliance with the vCenter Server.

1 In vCenter Server, right-click the host (for example, *Cluster-NSX*) where you want to deploy the appliance and select **Deploy OVF Template**.

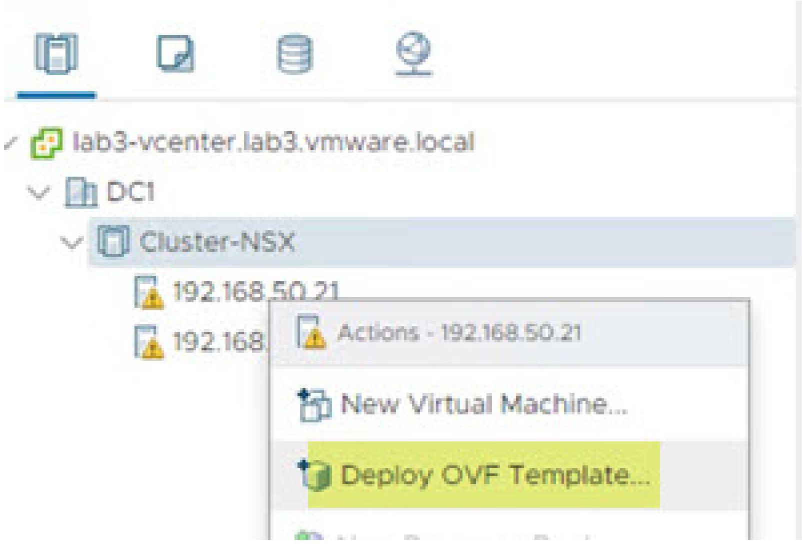

2. Select the downloaded NSX Manager appliance OVA file.

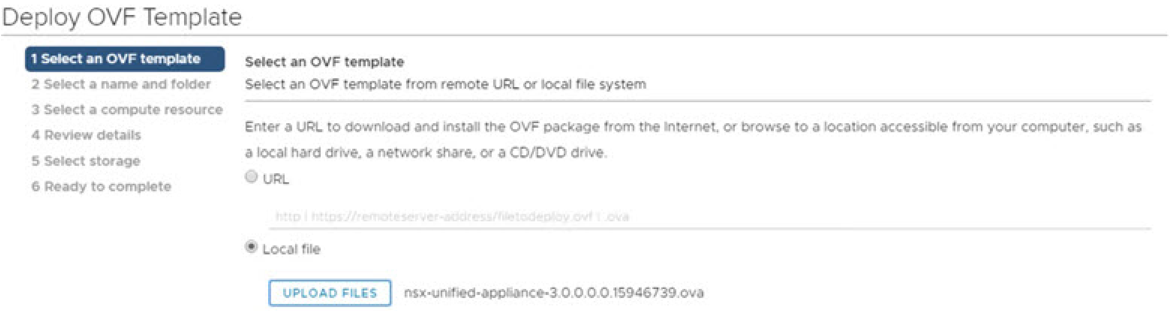

Deploy OVF Template

1 Select an OVF template
2 Select a name and folder
3 Select a compute resource
4 Review details
5 Select storage
6 Ready to complete

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

○ URL

http | https://remoteserver-address/filetodeploy.ovf | .ova

● Local file

UPLOAD FILES nsx-unified-appliance-3.0.0.0.0.15946739.ova

3. Follow the prompts and provide the following information.

   - Enter the NSX Manager appliance VM name and select the vCenter Server folder for the appliance VM.
   - Select ESXi to host the NSX Manager.
   - Review NSX Manager VM details.
   - Select the NSX Manager VM size.
   - Select the storage for the NSX Manager VM.
   - Select the vSphere Distributed Switch (VDS) Port Group for the NSX Manager managed vNIC (vCenter Management Port Group).
   - Enter the NSX-T Manager information such as hostname, IP, DNS, NTP. Select Rolename as **NSX Manager** and enter the password. The password must meet the following requirements:
     - At least 12 characters in length
     - At least one lowercase letter, one uppercase letter, one numeric character, and one special character (except quotes)
     - At least five different characters

4. Review the **NSX Manager** VM settings.

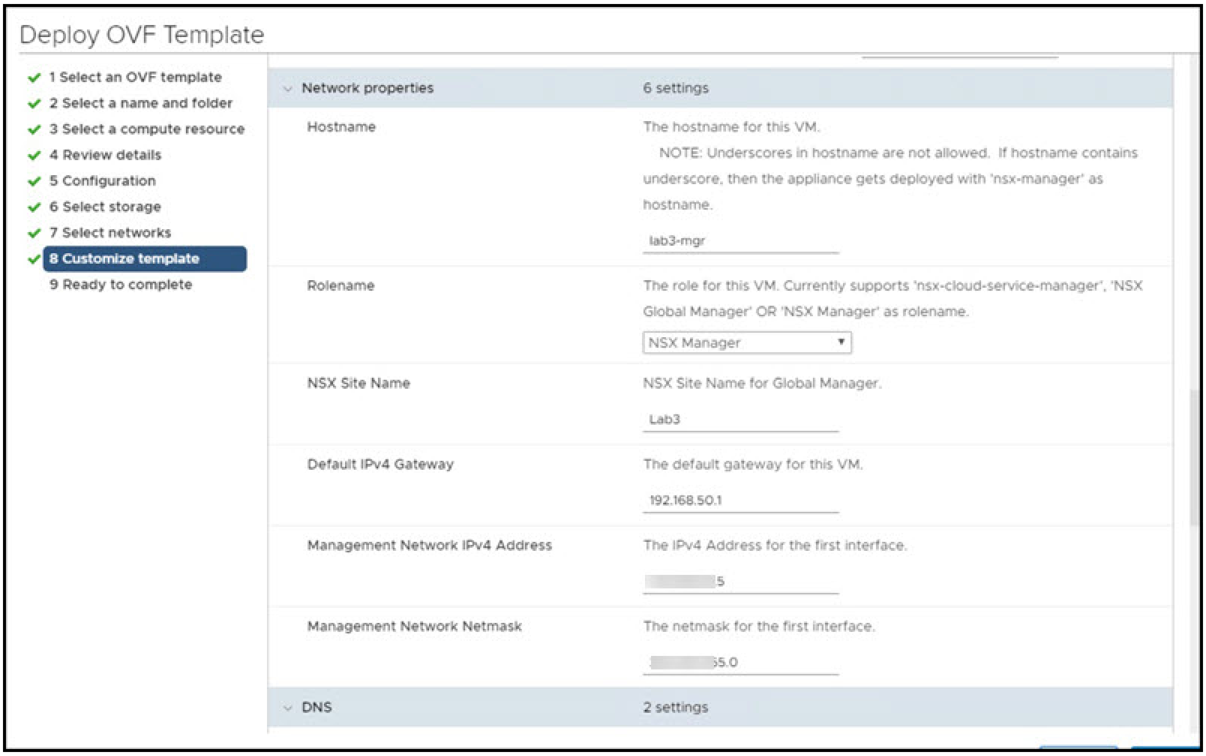

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 Configuration
- ✓ 6 Select storage
- ✓ 7 Select networks
- ✓ 8 Customize template
- 9 Ready to complete

Network properties — 6 settings

Hostname: The hostname for this VM. NOTE: Underscores in hostname are not allowed. If hostname contains underscore, then the appliance gets deployed with 'nsx-manager' as hostname.
lab3-mgr

Rolename: The role for this VM. Currently supports 'nsx-cloud-service-manager', 'NSX Global Manager' OR 'NSX Manager' as rolename.
NSX Manager ▼

NSX Site Name: NSX Site Name for Global Manager.
Lab3

Default IPv4 Gateway: The default gateway for this VM.
192.168.50.1

Management Network IPv4 Address: The IPv4 Address for the first interface.
[redacted]5

Management Network Netmask: The netmask for the first interface.
[redacted]55.0

DNS — 2 settings

5 After the **NSX Manager** deployment is finished, power on the **NSX Manager** VM.

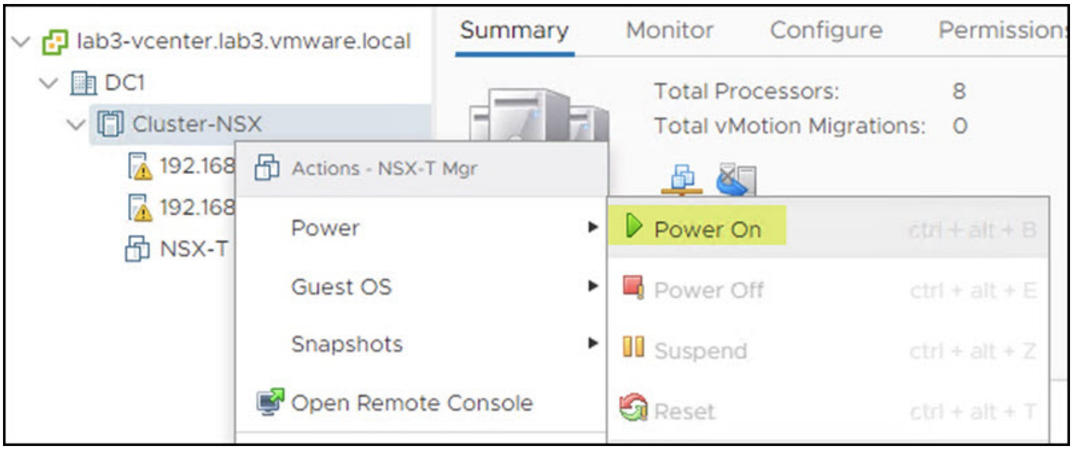

lab3-vcenter.lab3.vmware.local
DC1
Cluster-NSX
192.168
192.168
NSX-T

Summary | Monitor | Configure | Permissions

Total Processors: 8
Total vMotion Migrations: 0

Actions - NSX-T Mgr
- Power ▸
- Guest OS ▸
- Snapshots ▸
- Open Remote Console

- Power On ctrl + alt + B
- Power Off ctrl + alt + E
- Suspend ctrl + alt + Z
- Reset ctrl + alt + T

Note the IP address of the NSX Manager VM. You can now access the NSX Manager appliance UI from your browser using *https://<nsx-manager-ip-address>*.

6 For production deployment, you need three node NSX Manager clusters. To deploy additional NSX Manager nodes:

a From your browser, log in to the NSX Manager appliance at https://<nsx-manager-ip-address> using the admin credentials.

b Click **System > Appliances > Add NSX Appliance**. Provide the required information and follow the prompts to install the additional NSX Manager appliances.

7 Click **Set Virtual IP** and provide a virtual IP address for the NSX Manager Cluster. The virtual IP allows you to access the NSX Manager cluster using a single IP.

8 Use the configured IP address for accessing the NSX management plane. From a browser, log in to the NSX Manager using the virtual IP address assigned to the cluster at *https://<vip-address>*.

NSX Security Licenses

You can find a detailed list of features associated with the various licensing editions of VMware NSX Security in the Knowledge Base article.

1 From your browser, log in to the NSX Manager appliance at https://<nsx-manager-ip-address> using the admin credentials.

2 Add your NSX-T license from the **System > Licenses > Add License** page.

The evaluation license is available at the NSX-T Product Evaluation Center.

Preparing for Distributed Security

You can use NSX-T Distributed Firewall (DFW) for Macro-Segmentation (Security Zones) and Micro-Segmentation. Distributed Firewall provides complete L2-L7 East-West visibility and enforcement, with automated policy formulation. It works on both Physical Servers and VMs on ESXi and Physical Network changes are not required. By using DFW, it is possible to segment in any matter desired. There are four basic types of segmentation, many of which can coexist – each applied in different sections of the environment.

- **Zone Segmentation**: Zone Segmentation can be as general as segmenting production from non-production, or it may be a far more detailed segmentation by business unit, function, or product offering. The point is that each zone is defined independently of segments, VLANs, data centers, or other constructs. Zones are entirely logical definitions which can be used to define security policy.
- **VLAN Segmentation**: VLAN segmentation is most commonly used by replacing the legacy firewall infrastructure. In this model, an IP segment is the defining element for a source or destination of the security policy.
- **Application Segmentation**: Application segmentation is used to define a logical security ring around an application. Because applications are not frequently understood in detail, it can be convenient to simply define a tag for a given application and apply this tag to all its components and allow full communication between said elements. This brings greater security than a large zone definition which can have multiple applications, without requiring detailed understanding for micro-segmentation.
- **Micro-Segmentation**: Micro-segmentation is a security model where communication between elements is defined as explicitly as possible. At its extreme, micro-segmentation can be the explicit definition of communication between pairwise elements. Clearly this is operationally

complex, thus NSX offers micro-segmentation based on tags which allows explicit definition by groups. For example, you can define a rule which allows SSL but only TLS version 1.3 to the tagged secure web servers. Based on needs of your organization, you can segment each of those manners in different areas.

With NSX-T, all of these segmentation approaches are not exclusive but can coexist. You can decide to segment a lab in a zone model by just setting up a boundary around it and a DMZ environment in a micro-segmentation. You can segment non-production environments just by applications whereas you can further segment the production applications containing sensitive customer data using VLAN. The change of one security model to another is accomplished through a simple policy push, without the need to re-architect any networking infrastructure.

Distributed Security for Virtual Machines

This section provides the configuration workflow to prepare your environment using the NSX Distributed Security for protecting the virtual machines.

Prerequisites

You have deployed the NSX Manager and configured the valid licenses.

Configuration Workflow

Preparing your virtual environment for the NSX Distributed Security involves two main steps:

- Configure Compute Manager (vCenter)
- Prepare vCenter cluster (ESXi hosts) for the NSX Distributed Security

1: Configure Compute Manager (vCenter)

You must add vCenter Server as a compute manager on NSX-T to view all the vCenter Server host and cluster inventory. You can then leverage the available inventory to prepare ESXi hosts and clusters for NSX Security.

1. From your browser, log in to the NSX Manager appliance at https://<nsx-manager-ip-address> using the admin credentials.
2. Register NSX-T with vCenter Server from the **System > Fabric > Compute Managers > Add Compute Manager**. Add the vCenter Server as a compute manager.

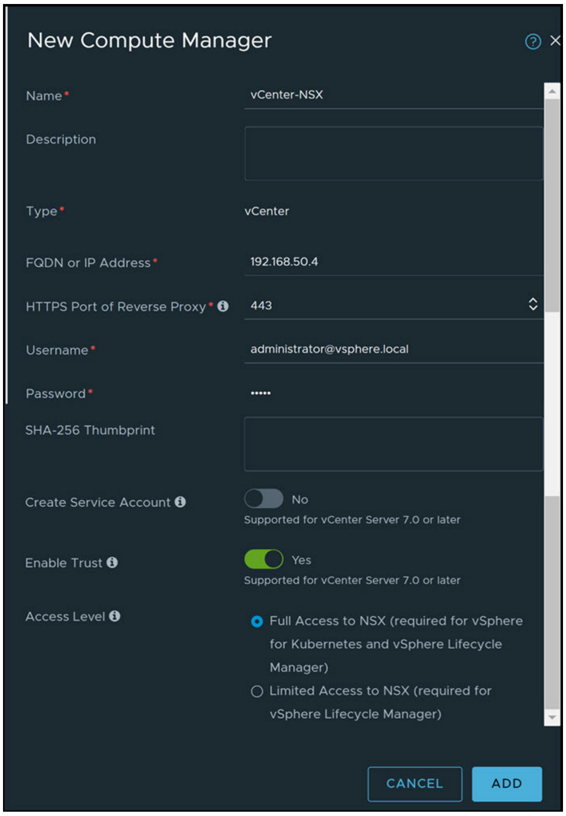

New Compute Manager

Name* vCenter-NSX

Description

Type* vCenter

FQDN or IP Address* 192.168.50.4

HTTPS Port of Reverse Proxy* 443

Username* administrator@vsphere.local

Password* •••••

SHA-256 Thumbprint

Create Service Account: No
Supported for vCenter Server 7.0 or later

Enable Trust: Yes
Supported for vCenter Server 7.0 or later

Access Level
● Full Access to NSX (required for vSphere for Kubernetes and vSphere Lifecycle Manager)
○ Limited Access to NSX (required for vSphere Lifecycle Manager)

3 Validate NSX-T registration in the vCenter Server from the **System** > **Fabric** > **Compute Managers** page. Click **Refresh** and view the connection status.

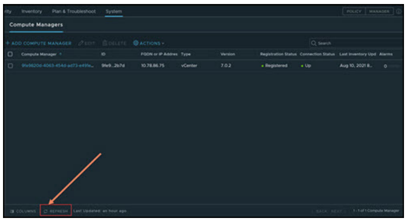

Compute Managers

Compute Manager	ID	FQDN or IP Addres	Type	Version	Registration Status	Connection Status	Last Inventory Upd	Alarms
9fe9820d-4063-454d-ad73-e49fe...	9fe9...2b7d	10.78.86.75	vCenter	7.0.2	● Registered	● Up	Aug 10, 2021 8...	0

Last Updated: an hour ago

1 - 1 of 1 Compute Manager

After the vCenter Server registration is successful, you can view the configured vCenter Server host cluster inventory from the NSX Manager User Interface (UI). On the NSX Manager UI, go to **System** > **Fabric** > **Nodes** to view the inventory.

You can configure multiple vCenter Servers from the NSX Manager UI following these same steps for each of the vCenter Server.

2: Prepare the vCenter Cluster (ESXi Hosts) for the NSX Distributed Security

NSX Distributed Security involves preparing vCenter Server compute cluster of NSX-T. NSX-T supports two host preparation modes as follows:

1 **Security Only** - Distributed Security for VDS port groups:

- Supports security for VMs connected to the native vCenter Distributed Virtual Port Groups (DVPG).
- Supports vSphere 6.7 and vSphere 7.0 Update1 or later.

- Does not support NSX-T networking for the workload within the NSX-T prepared vCenter Server cluster.
- Workflow is supported only using the Quick Start wizard.

2 **Networking and Security** - Distributed Security with NSX-T Networking:

- Supports NSX-T networking and distributed security for the workload within the NSX-T prepared vCenter Server cluster.
- If VLAN connected workloads need distributed security, then you must move the workload to the NSX-T VLAN segments from DVPG.
- Workflow is supported using the Quick Start wizard or manually from the **System > Fabric > Nodes** menu.

Based on your environment, select the required deployment method. The NSX-T environment can have a mix of NSX Security only prepared clusters and NSX Networking and Security prepared clusters. More details on each of the deployment modes are covered later in this section.

2.1: Security Only Host Preparation - Distributed Security for VDS Port Groups

After you configure the compute manager, you can prepare clusters of ESXi hosts only for distributed security. The hosts in your cluster must share VDS.

Procedure

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

2 Navigate to **System > Quick Start**.

3 On the **Prepare Clusters for Networking and Security** card, click **Get Started**.

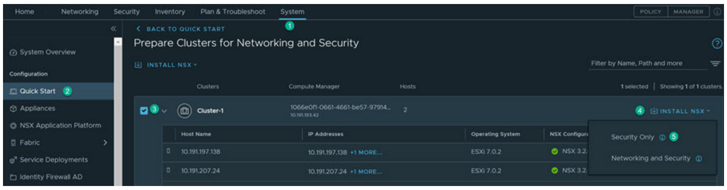

Prepare Clusters for Networking and Security

Clusters	Compute Manager	Hosts
Cluster-1	1066e0f1-0661-4661-be57-97914... 10.191.193.42	2

Host Name	IP Addresses	Operating System	NSX Configura
10.191.197.138	10.191.197.138 +1 MORE...	ESXi 7.0.2	NSX 3.2.
10.191.207.24	10.191.207.24 +1 MORE...	ESXi 7.0.2	NSX 3.2.

4 Select the clusters that you want to install Distributed Security.

5 Click **Install NSX** and then select **Security Only**.

6 In the dialog box, click **Install**.

The NSX host preparation begins to install required software modules on the ESXi hosts.

The process takes a few minutes to complete. After the process is complete, the status changes to **Success**. The objects like transport node profile, transport zone, and distributed port groups are automatically created.

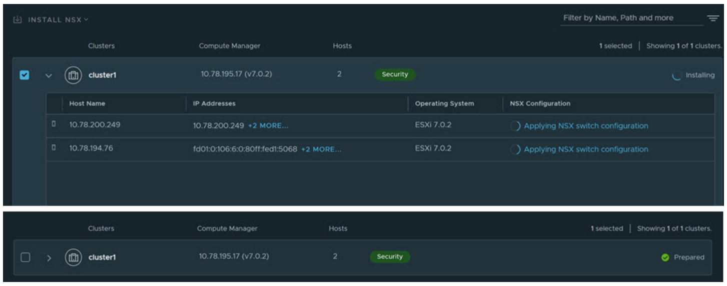

	Clusters	Compute Manager	Hosts		1 selected | Showing 1 of 1 clusters.
☑	cluster1	10.78.195.17 (v7.0.2)	2	Security	Installing

Host Name	IP Addresses	Operating System	NSX Configuration
10.78.200.249	10.78.200.249 +2 MORE...	ESXi 7.0.2	Applying NSX switch configuration
10.78.194.76	fd01:0:106:6:0:80ff:fed1:5068 +2 MORE...	ESXi 7.0.2	Applying NSX switch configuration

	Clusters	Compute Manager	Hosts		1 selected | Showing 1 of 1 clusters.
□	cluster1	10.78.195.17 (v7.0.2)	2	Security	Prepared

7 To view VDS with Distributed Security installed, do the following:

   a Navigate to **System > Fabric > Nodes**.

   b Select the **Host Transport Nodes** tab.

   Note vSphere clusters prepared for Distributed Security are identified by the **Security** label.

Results

On the NSX Manager UI, go to **Networking > Segments > Distributed Port Groups** tab to view the DVPG inventory from the vCenter Server.

On the NSX Manager UI, go to **Inventory > Virtual Machines** to view the virtual machine inventory from all ESXi hosts.

What to do next

You can now start Configuring Security Policy for the workloads hosted on DVPG on the prepared vCenter Server.

2.2: Networking and Security - Distributed Security with NSX-T Networking

After you configure the compute manager, you can prepare clusters of ESXi hosts for VLAN networking and distributed security together.

Procedure

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

2 On the **Prepare Clusters for Networking and Security** card, click **Get Started**.

3 Select the clusters you want to prepare for NSX-T networking.

4 Click **Install NSX** and then select **Networking and Security**.

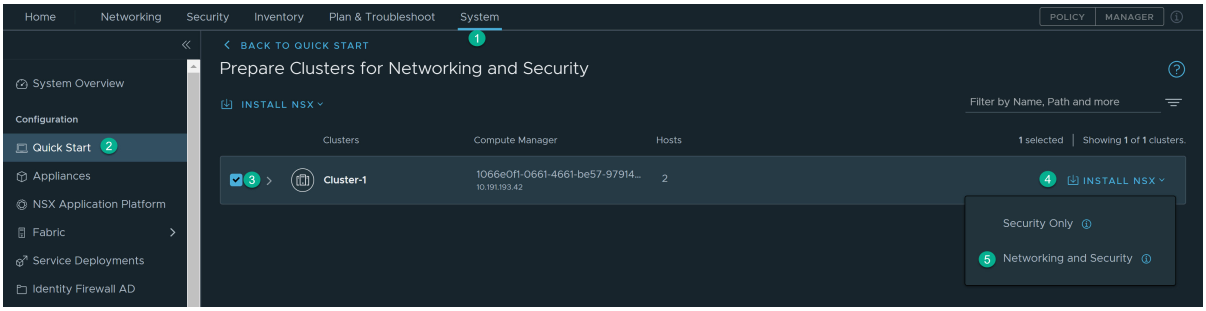

Clusters	Compute Manager	Hosts
☑ 3 Cluster-1	1066e0f1-0661-4661-be57-97914... 10.191.193.42	2

1 selected | Showing 1 of 1 clusters.

5 Depending on your requirement, you can prepare the same cluster for both VLAN and Overlay networking or for one type of networking. With Overlay networking, each host switch is added with a TEP IP address, which is required for overlay networking.

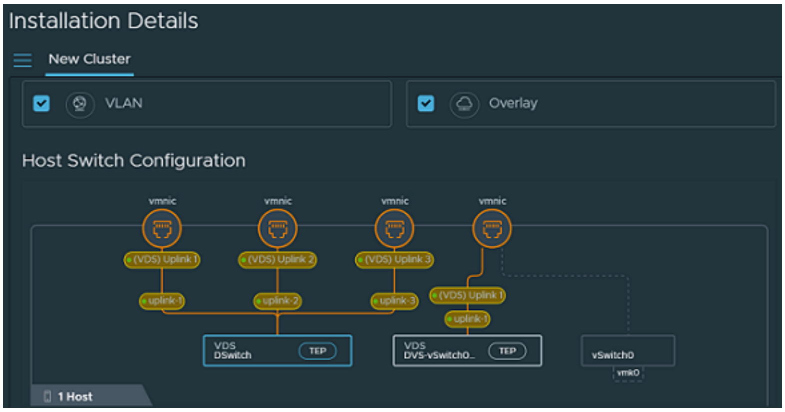

6 View the Host Switch Configuration switch to know the target switches where the physical NICs and VMkernel adapters (if any) will be migrated to.

This is the NSX-T recommended configuration. However, you can customize the settings for the cluster, even though it is an optional step.

Note A dotted line originating from a switch to a physical NIC indicates that it is an existing configuration on the host switch, which will be replaced by a firm line going to the same physical NIC.

7 Even though NSX-T provides recommendations, you can still customize the configuration. To customize a switch, select the switch, and change the recommended configuration.

a **Type**: Switch over the host switch type.

b **Transport Zone**: Select a different transport zone that you want the host to be associated with.

c **Uplink Profile**: If needed, select a different uplink profile in place of the recommended uplink profile.

Note If you configure two VDS switches with the same configuration, the wizard recommends the same uplink profile for both the switches.

d **Uplink to Physical NIC mapping**: On a VDS switch, all uplinks configured on the VDS switch are mapped to the uplinks in NSX-T. On an N-VDS switch, uplinks are mapped to vmnics.

A change to host switch type or uplink to vmnic mapping is reflected in the Host Switch Configuration network representation.

8 Click **Install**.

The NSX host preparation begins to install required software modules on the ESXi hosts.

View the progress of installation on the **Prepare Clusters for Networking and Security** card. If installation on any of the host fails, retry installation by resolving the error.

The process takes a few minutes to complete. After the process is complete, the status changes to **Success**.

9 To view successfully prepared hosts, go to **System** → **Fabric** → **Nodes** → **Host Transport Node**.

Results

On the **NSX Manager** UI, go to the **Inventory > Virtual Machines** tab to view the virtual machine inventory from all the ESXi hosts.

Note vCenter Server cluster prepared for Networking and Security does not support Security for workloads connected directly to the DVPG. If the DVPG VLAN connected workload needs security, you must move the workload to NSX VLAN segments (with the same VLAN) or move the workloads to the cluster prepared only for NSX Security.

For more information, see *NSX-T Data Center Administration Guide*.

What to do next

You can now start Configuring Security Policy for the workloads hosted on the NSX segments on the prepared vCenter Server clusters.

Distributed Security for Physical Server

You can secure workloads running on a physical server. Preparing a physical server involves a modest preparation for either a Linux or Windows physical system.

You can provide connectivity and security to applications or workloads between:

- Physical workloads (bare metal server) and virtual workloads
- Physical workloads (bare metal server) and physical workloads (bare metal server)

Prerequisites

You have deployed the NSX Manager and configured the relevant license.

Procedure

A Windows physical system requires the Windows Remote Management (WinRM) feature. A Linux physical system requires a handful of dependency modules. Install prerequisite features for either a Windows or Linux physical system using Ansible and a prepared playbook available on GitHub.

To add physical servers to the NSX data plane, perform the following steps:

1. Review the bare metal requirements. See Bare Metal Server System Requirements.
2. Configure the necessary ports and protocols. See Ports and Protocols.
3. Create an Application Interface for the physical server workloads. See Create Application Interface for Physical Server Workloads.
4. Configure a physical server as a transport node through the UI. See Configure a Physical Server as Transport Node Through UI.
5. After adding all your physical servers, configure the DFW rules to secure the physical systems. After configuration is complete, the physical servers contain the DFW rules that are pushed from the NSX Manager. Here is an example of how to secure workloads on Windows Server 2016/2019. See Secure Workloads on Windows Server 2016/2019 Bare Metal Servers.

Preparing For Gateway Security

You can use NSX-T Gateway Firewall to firewalling for the North-South traffic at the Layer 3 boundary. You can use the Gateway Firewall as an inter-tenant/zone firewall from the north-south perspective, along with the Distributed Firewall. Gateway Firewall is supported on both Tier-0 and Tier-1 gateways. Tier-0 supports basic L3/L4 stateful firewall, where as Tier-1 supports basic L3/L4 and advanced L7 features like L7 Application ID, URL filtering, IDS/IPS, TLS Inspection, Identity Firewall, and Malware Prevention. The Gateway Firewall provides firewalling services and other services that cannot be distributed such as NAT, DHCP, VPN, and load balancing, and needs the services router component of the gateway. This means that the Gateway Firewall is implemented in the NSX Edge Transport Nodes, which are dedicated DPDK appliances.

At a high level, Gateway Security preparation involves the following steps:

- Deploy NSX Manager
- Deploy NSX Edge Transport Node and provision Edge Cluster
- Create NSX Tier-0/1 Gateway
- Create Service Interface/Uplink Interface on Tier-1 or External Interface on Tier-0
- Define Zone/Inter-VLAN Firewall Policies

Infrastructure Set up for NSX Gateway Security

You must first set up your infrastructure and then configure your environment for Gateway Security.

1. Deploy NSX Edge Transport Node

You must first deploy the NSX edge transport node.

Prerequisites

You have deployed the NSX Manager and configured the valid licenses.

Procedure

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

2 Select **System > Fabric > Nodes > Edge Transport Nodes > Add Edge Node**.

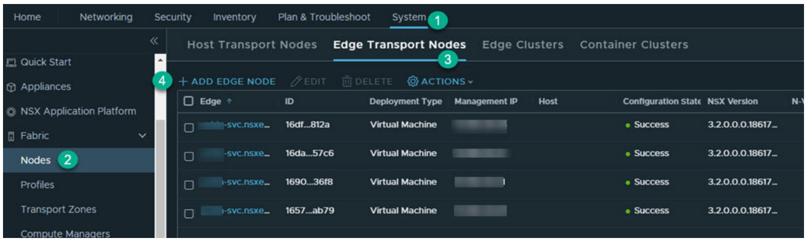

☐ Edge ↑	ID	Deployment Type	Management IP	Host	Configuration State	NSX Version	N-V
☐ [redacted]-svc.nsxe…	16df…812a	Virtual Machine	[redacted]		● Success	3.2.0.0.0.18617…	
☐ [redacted]-svc.nsxe…	16da…57c6	Virtual Machine	[redacted]		● Success	3.2.0.0.0.18617…	
☐ [redacted]-svc.nsxe…	1690…36f8	Virtual Machine	[redacted]		● Success	3.2.0.0.0.18617…	
☐ [redacted]-svc.nsxe…	1657…ab79	Virtual Machine	[redacted]		● Success	3.2.0.0.0.18617…	

3 Type a name for the NSX Edge.

4 Type the Host name or FQDN from vCenter Server.

5 Select the form factor for the NSX Edge VM appliance.

6 To customize CPU and memory allocated to an NSX Edge VM appliance, tune the following parameters. However, for maximum performance NSX Edge VM appliance must be assigned 100% of the available resources.

Caution If you customize resources allocated to the NSX Edge VM, turn back the reservation later on to 100% to get maximum performance.

Option	Description
Memory Reservation (%)	Reservation percentage is relative to the pre-defined value in the form factor. 100 indicates 100% of memory is reserved for the NSX Edge VM. If you enter 50, it indicates that 50% of the allocated memory is reserved for the Edge transport node.
CPU Reservation Priority	Select the number of shares to be allocated to an NSX Edge VM relative to other VMs that are contending for shared resources. The following shares are for an NSX Edge VM in Medium form factor: ■ Low - 2000 shares ■ Normal - 4000 shares ■ High - 8000 shares ■ Extra High - 10000 shares
CPU Reservation (MHz)	Caution Unless you need fine grained control over CPU reservations, do not use this field. Instead, change CPU reservations from the **CPU Reservation Priority** field. The maximum CPU reservation value must not exceed the number of vCPUs multiplied by the normal CPU operation rate of the physical CPU core. If the MHz value entered exceeds the maximum CPU capacity of the physical CPU cores, the NSX Edge VM might fail to start even though the allocation was accepted. For example, consider a system with two Intel Xeon E5-2630 CPUs. Each CPU contains ten cores running at 2.20 GHz. The maximum CPU allocation for a VM configured with two vCPUs is 2 x 2200 MHz = 4400 MHz. If CPU reservation is specified as 8000 MHz, the reconfiguration of the VM completes successfully. However, the VM fails to power on.

7 In the Credentials window, enter the following details.

- Specify the CLI and the root passwords for the NSX Edge. Your passwords must comply with the password strength restrictions.
  - At least 12 characters
  - At least one lower-case letter
  - At least one upper-case letter
  - At least one digit
  - At least one special character
  - At least five different characters
  - No dictionary words

  - No palindromes
  - More than four monotonic character sequence is not allowed
- To enable SSH for an administrator, toggle the **Allow SSH Login** button.
- To enable SSH for a root user, toggle the **Allow Root SSH Login** button.
- Enter credentials for the Audit role. If you do not enter credentials in the **Audit Credentials** section, the audit role remains disabled.

  Note After deploying the NSX Edge node, you cannot change the SSH setting for a root user that you set during deployment. For example, you cannot enable SSH for a root user if you disabled it during deployment.

8 Enter the NSX Edge details.

Option	Description
Compute Manager	Select the compute manager from the drop-down menu. The compute manager is the vCenter Server registered in the Management Plane.
Cluster	Designate the cluster the NSX Edge is going to join from the drop-down menu.
Resource Pool or Host	Assign either a resource pool or a specific host for the NSX Edge from the drop-down menu.
Datastore	Select a datastore for the NSX Edge files from the drop-down menu.

9 Enter the NSX Edge interface details.

Option	Description
IP Assignment	It is the IP address assigned to NSX Edge node which is required to communicate with NSX Manager and NSX Controller. Select **DHCP** or **Static** IP. If you select **Static**, enter the values for: ■ Management IP: Enter IP address of NSX Edge in the CIDR notation. ■ Default gateway: Enter the gateway IP address of NSX Edge.
Management Interface	From the drop-down menu, select the interface that connects to the NSX Edge management network. This interface must either be reachable from NSX Manager or must be in the same management interface as NSX Manager and NSX Controller. The NSX Edge management interface establishes communication with the NSX Manager management interface. The NSX Edge management interface is connected to distributed port groups or segments.
Search Domain Names	Enter domain names in the format 'example.com' or enter an IP address.
DNS Servers	Enter the IP address of the DNS server.
NTP Servers	Enter the IP address of the NTP server.

10 Enter the N-VDS information.

Option	Description
Edge Switch Name	Enter a name for the switch.
Transport Zone	Select the transport zones that this transport node belongs to. An NSX Edge transport node belongs to at least two transport zones, an overlay for NSX-T Data Center connectivity and a VLAN for uplink connectivity. **Note** NSX Edge Nodes support multiple overlay tunnels (multi-TEP) when the following prerequisites are met: ■ TEP configuration must be done on one N-VDS only. ■ All TEPs must use the same transport VLAN for overlay traffic. ■ All TEP IPs must be in the same subnet and use the same default gateway.
Uplink Profile	Select the uplink profile from the drop-down menu. The available uplinks depend on the configuration in the selected uplink profile.

Option	Description
IP Assignment (TEP)	IP address is assigned to the NSX Edge switch that is configured. It is used as the tunnel endpoint of the NSX Edge. Select **Use IP Pool** or **Use Static IP List** for the overlay N-VDS. - If you select **Use Static IP List**, specify: - Static IP List: Enter a list of comma-separated IP addresses to be used by the NSX Edge. - Gateway: Enter the default gateway of the TEP, which is used to route packets another TEP in another network. For eample, ESXi TEP is in 20.20.20.0/24 and NSX Edge TEPs are in 10.10.10.0/24 then we use the default gateway to route packets between these networks. - Subnet mask: Enter the subnet mask of the TEP network used on the NSX Edge. - If you selected **Use IP Pool** for IP assignment, specify the IP pool name.
DPDK Fastpath Interfaces / Virtual NICs	Select the data path interface that is either a distributed port group trunk or a segment as the uplink interface. **Note** If the uplink profile applied to the NSX Edge node is using a Named Teaming policy, ensure the following condition is met: - All uplinks in the Default Teaming policy must be mapped to the corresponding physical network interfaces on the Edge VM for traffic to flow through a logical switch that uses the Named Teaming policies. Starting with NSX Data Center 3.2.1, you can configure a maximum of four unique data path interfaces as uplinks on an NSX Edge VM. When mapping uplinks to DPDK Fastpath Interfaces, if NSX Edge does not display all the available interfaces (four in total), it means that either the additional interface is not yet added to the NSX Edge VM or the uplink profile has fewer number of uplinks. For NSX Edge VMs upgraded from an earlier version of NSX-T Data Center to 3.2.1 or later, invoke the redeploy API call to redeploy the NSX Edge VM. Invoking the redeploy API ensures the NSX Edge VM deployed recognizes all the available datapath interfaces in NSX Manager UI. Make sure the Uplink profile is correctly configured to use additional datapath NIC. - For autodeployed NSX Edges, call the redeploy API. `POST api/v1/transport-nodes/<transport-node-id>?action=redeploy` - For manually deployed edges, deploy a new NSX Edge VM. Ensure all the vmx customizations of the old NSX Edge VM are also done for the new NSX Edge VM. Performing vMotion on a NSX Edge VM might result in the NSX Edge VM going into failed state or the additional network adapter cannot be enabled because of memory buffer issues. For troubleshooting memory-related issues when performing a vMotion on a NSX Edge VM, see https://kb.vmware.com/s/article/76387.

Note

- LLDP profile is not supported on an NSX Edge VM appliance.
- Uplink interfaces are displayed as **DPDK Fastpath Interfaces** if the NSX Edge is installed using NSX Manager or on a Bare Metal server.
- Uplink interfaces are displayed as **Virtual NICs** if the NSX Edge is installed manually using vCenter Server.

11 View the connection status on the **Transport Nodes** page.

After adding the NSX Edge as a transport node, the connection status changes to Up in 10-12 minutes.

1.1: Provision NSX Edge Cluster

You should have two edge nodes in an edge cluster for high availability.

Procedure

1 Add the edge cluster. Go to **System** > **Fabric** > **Nodes** > **Edge Clusters** and click **Add Edge Cluster**.

2 In the **Name** text box, enter name for the edge cluster. For example, *Edge-cluster-1*.

3 Move the created edge node (*Edge-1*) from the **Available** to the **Selected** window, and click **Add**.

2. Create a Tier-0 or Tier-1 Gateway

Depending on your use case, create a tier-1 or tier-0 gateway.

Procedure

1 To add a gateway:

- To add a tier-0 gateway: From the **NSX Manager** UI, click **Networking** > **Tier-0 Gateways** > **Add Gateway** > **Tier-0**.

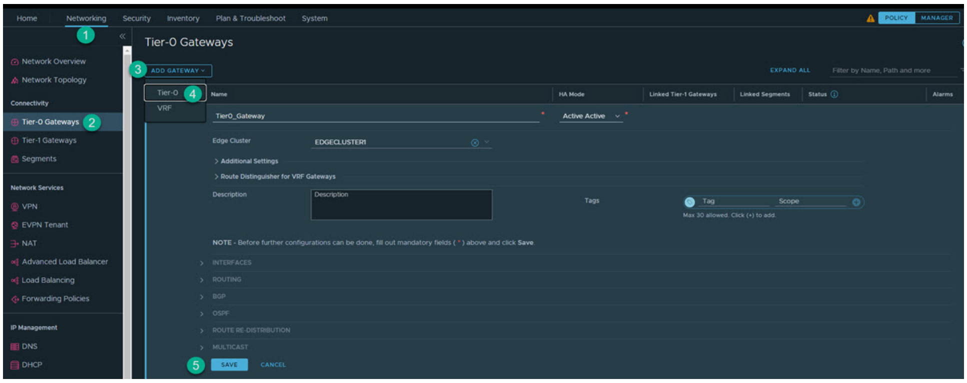

Tier-0 Gateways

Name: Tier0_Gateway *

HA Mode: Active Active *

Edge Cluster: EDGECLUSTER1

Description: Description

Tags: Tag Scope

Max 30 allowed. Click (+) to add.

NOTE - Before further configurations can be done, fill out mandatory fields (*) above and click **Save**.

- To add a tier-1 gateway: From the **NSX Manager** UI, click **Networking** > **Tier-1 Gateways** > **Add Gateway** > **Tier-1**.

2 Provide the following information.

Name	Enter the name for the gateway. For example, *T0-gateway-1*.
Edge cluster	Select the created edge cluster. For example, *Edge-cluster-1*.

3 Click **Save**.

For further details, see *NSX-T Data Center Administration Guide*.

3. Create Interfaces on Tier-0 or Tier-1 Gateway

NSX gateway has different interface types. Based on the network topology, you can select the required interfaces to connect to the network and provide firewalling for traffic passing through the gateway.

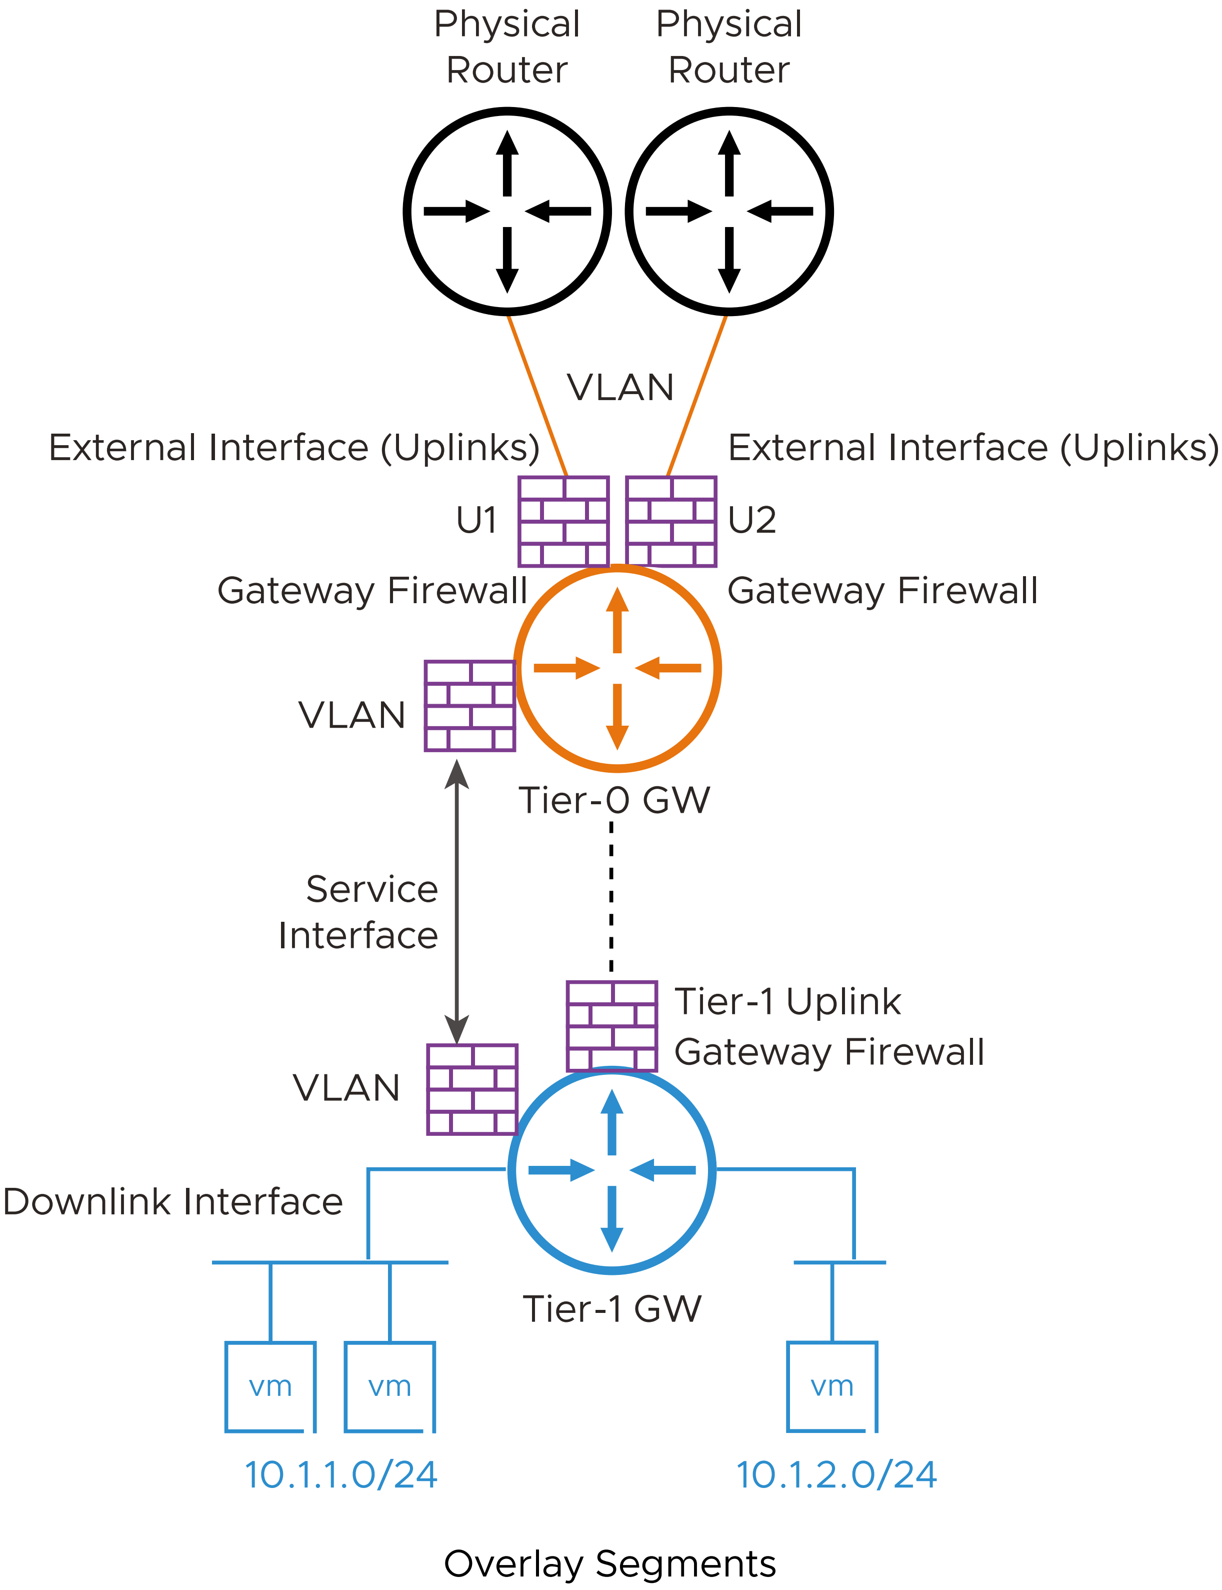
Physical Router
Physical Router
VLAN
External Interface (Uplinks)
External Interface (Uplinks)
U1
U2
Gateway Firewall
Gateway Firewall
VLAN
Tier-0 GW
Service Interface
Tier-1 Uplink Gateway Firewall
VLAN
Downlink Interface
Tier-1 GW
vm
vm
vm
10.1.1.0/24
10.1.2.0/24
Overlay Segments

Tier-0 External Interfaces:

- Connects to physical router for external connectivity
- You create this interface on the VLAN segments on the tier-0 gateway

Tier-1 Uplink Interfaces:

- Connects to gier-0
- System creates this interface as tier-1 connects to tier-0

Service Interface:

- Used for providing NSX-T Services (GFW and other) to non-NSX managed VLAN workloads
- Connects to VLAN segment
- Supported on both tier-0 and tier-1

Downlink Interface:

- Overlay segment Interface on gateway
- Supported on both tier-0 and tier-1
- No GFW support

The Gateway Firewall can be mainly used for two scenarios based on how workloads are connected to the network:

- VLAN connected workloads
- NSX network overlay segments connected workloads

Each of these scenarios follows slightly different steps to create the network interfaces as described later in this section.

3.1: Create NSX-T Gateway Firewall Interface for VLAN Connected Workloads

You must perform the following steps to set up your environment.

1 Create a VLAN segment in NSX-T.

  a In the **NSX Manager**, click **Networking > Segments > Add Segment**.

  b Provide the following information.

Segment Name	Enter the name for the segment. For example, *VLAN-100*.
Transport Zone	Select the default transport zone for the VLAN traffic. For example, *nsx-vlan-transportzone*.
VLAN	Enter *100*.

  c Click **Save**.

2 Create a Service Interface(s) on the tier-0 or tier-1 gateway.

  a In the **NSX Manager**, click **Networking > Tier-1 Gateways Add Gateway > Tier-1**.

b Edit the created gateway. For example, *T1-gateway-1*.

c Under **Service Interfaces**, click **Set**.

d Click **Add Interface**.

e Provide the following information.

Name	Enter the name of the interface. For example, *SI-VLAN-100*.
IP Address/Mask	Enter an IP address. For example, *192.168.50.12/24*.
Connected To (Segment)	Select the configured segment. For example, *VLAN-100*.

f Click **Save**.

Create more service interfaces based on the network requirements.

On tier-0, you have an option to create an external interface, or a service interface based on the connectivity requirement. If an external interface is created, you need to create one external interface per edge, part of the edge cluster.

As part of the workflow, select the edge node to create that interface, in addition to the mentioned parameters.

For more information, see *NSX-T Data Center Administration Guide*.

3.2: Create NSX-T Gateway Firewall Interface for Network Overlay Workloads

Perform the following steps.

1 Create a Tier-1 Gateway.

a Click **Networking > Tier-1 Gateways > Add Tier-1 Gateway**.

b Enter the name for the tier-1 gateway. For example, *PROD-Tier1*.

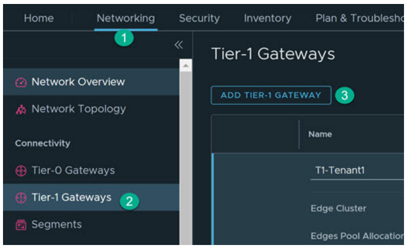

T1-Tenant1

c Select the tier-0 gateway to create an uplink on the tier-1.

d Select the edge cluster for implementing the gateway services.

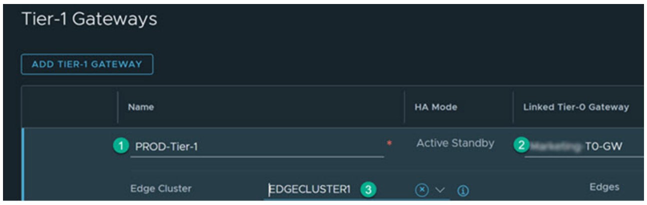

PROD-Tier-1 | Active Standby | [redacted] T0-GW

Edge Cluster: EDGECLUSTER1

e Click **Save**.

2 Additionally, you should create an overlay segment(s) for connecting workloads. This creates a downlink interface on the gateway and also makes the NSX segments available on the ESXi for network connectivity with the virtual machine.

   a Click **Networking > Segments > NSX > Add Segment**.

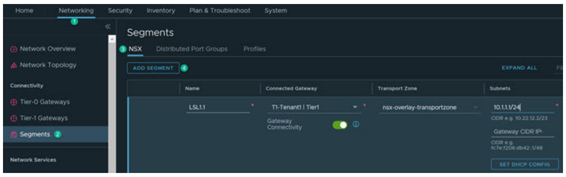

Name	Connected Gateway	Transport Zone	Subnets
LSL1.1	T1-Tenant1 | Tier1	nsx-overlay-transportzone	10.1.1.1/24 CIDR e.g. 10.22.12.2/23 Gateway CIDR IPv6 CIDR e.g. fc7e:f206:db42::1/48

   b Provide the following information.

Name	Enter the name for the segment. For example, *LS1.1*.
Connectivity	Select the configured tier-1 gateway. For example, *T1-Tenant1*.
Transport Zone	Select the default transport zone for Overlay traffic. For example, *nsx-overlay-transportzone*.
Subnets	Enter the required subnet. For example, *10.x.x.1/24*.

   c Click **Save**.

3 Validate the configured overlay segment is available in the vCenter Server. In vCenter Server, go to **Host and Clusters**, and validate VMs that are created and connected to the configured overlay segment.

For more information, see *NSX-T Data Center Installation Guide*.

Configuring Security Policy

Firewall rule table implements the NSX Security policy which you can create using the NSX Manager GUI or the REST API framework.

Here are the high-level steps to understand and prepare for defining the security policy.

- `VM Inventory Collection`: You can identify and organize a list of all hosted virtualized workloads on the NSX-T transport nodes. The inventory is dynamically collected and saved by NSX Manager as the nodes – ESXi or KVM that are added as NSX-T transport nodes. You can view a list of inventories by navigating to the **Inventory > Virtual Machines** menu.
- `Tag`: NSX-T allows to tag virtual machine, segment, and segment-port. To tag each of these objects, go to the relevant object page or go to **Inventory > Tags**. Objects can have one or more tags. For example, a VM can have `Tag = PROD`, `Tag = HR-APP` or `Tag = WEB-Tier`.
- `Group Workloads`: You can use the NSX-T logical grouping construct with dynamic or static membership criteria based on VM name, tags, segment, segment port, IPs, or other attributes.
- `Define Security Policy`: You can define the security policy using the distributed firewall rule table available at **Security > Distributed Firewall**. You can organize the policy based on pre-defined categories like ethernet, emergency, infrastructure, environment, and application.

For details, see *NSX-T Data Center Administration Guide*.

Add Tags

You can select existing tags that are available in the inventory or create new tags to add to an object.

Procedure

1 With admin privileges, log in to NSX Manager.

2 Edit an object for tags. Objects can be virtual machines, segments, or segment ports. You can use inventory for each object to tag the object or go to **Inventory > Tags** for creating and assigning tags.

   To tag the objects directly, for example virtual machines, click **Inventory > Virtual Machines**.

   Next to the virtual machine that you want to edit, click ⋮, and click **Edit**.

3 In the **Tag** drop-down menu, enter a tag name. When you are done, click **Add Item(s)**.

   The maximum length of the tag name is 256 characters.

   If tags exist in the inventory, the **Tag** drop-down menu displays a list of all the available tags and their scope. The list of available tags includes user-defined tags, system-defined tags, and discovered tags. You can select an existing tag from the drop-down menu and add it to the virtual machine.

4 (Optional) Enter a tag scope.

   For example, let us say, you want to tag virtual machines based on their operating system (Windows, Mac, Linux). Create three tags, such as Windows, Linux, and Mac, and set the scope of each tag to OS.

   The maximum length of the scope is 128 characters.

   If you selected an existing tag from the inventory, the scope of the selected tag is applied automatically. Otherwise, you can enter a scope for the new tag that you are creating.

5 Click the **+** icon.

   The tag is added to the virtual machine.

6 (Optional) Repeat steps 3–5 to add more tags to the virtual machine.

7 Click **Save**.

Add Groups

Groups include different objects that are added both statically and dynamically and can be used as the source and destination of a firewall rule.

Procedure

1 Select **Inventory > Groups** from the navigation panel.

2 Click **Add Group**, then enter a group name.

3 Click **Set**.

4 In the **Set Members** window, select the **Group Type**.

Table 2-1.

Group Type	Description
Generic	This group type is the default selection. A Generic group definition can consist of a combination of membership criteria, manually added members, IP addresses, MAC addresses, and Active Directory groups. When you define membership criteria in the group, the members are dynamically added in the group based on one or more criteria. Manually added members include objects, such as segment ports, distributed ports, distributed port groups, VIFs, virtual machines, and so on.
IP Addresses Only	This group type contains only IP addresses (IPv4 or IPv6). **IP Addresses Only** groups with only manually added IP address members are not supported for use in the **Applied To** in DFW rules. It is possible to create the rule, but it will not be enforced. After a group of type **IP Addresses Only** is realized in NSX-T Data Center, you cannot edit the group type to **Generic**. However, if the group type is **Generic**, you can edit the group type to **IP Addresses Only**. In this case, only the IP addresses are retained in the group. All the membership criteria and other group definitions are lost.

5 On the **Membership Criteria** page, click **Add Criterion** to add members in the group dynamically based on one or more membership criteria.

6 Click **Members** to add static members in the group.

7 (Optional) Click **IP/MAC Addresses** to add IP and MAC addresses as group members. IPv4 addresses, IPv6 addresses, and multicast addresses are supported.

Click **Action > Import** to import IP/MAC Addresses from a TXT file or a CSV file containing comma-separated IP/MAC values.

8 Click **AD Groups** to add Active Directory Groups. This is used for Identity firewall. Groups with Active Directory members can be used in the source of a distributed firewall rule for Identity Firewall. Groups can contain both AD and compute members.

9 (Optional) Enter a description and tag.

10 Click **Apply**

Groups are listed, with an option to view the members and where the group is used.

Distributed Firewall Policy

Distributed firewall comes with predefined categories for firewall rules. Categories allow you to organize security policies.

Categories are evaluated from left to right (Ethernet > Emergency > Infrastructure > Environment > Application), and the distributed firewall rules within the category are evaluated top down.

Table 2-2. Distributed Firewall Rule Categories

Ethernet	Emergency	Infrastructure	Environment	Application
We recommend you include Layer 2 rules for this category.	We recommend you include quarantine and allow rules for this category.	We recommend you include rules which define access to shared services for this category. For example: ■ AD ■ DNS ■ NTP ■ DHCP ■ Backup ■ Management servers	We recommend you include rules between zones for this category. For example: ■ Production vs development ■ PCI vs non-PCI ■ Inter business unit rules	We recommend you include rules between: ■ Applications ■ Application tiers ■ Micro services

Add a Distributed Firewall Policy

Distributed firewall monitors all the East-West traffic on your virtual machines.

Procedure

1. With admin privileges, log in to NSX Manager.
2. Select **Security > Distributed Firewall** from the navigation panel.
3. Ensure that you are in the correct pre-defined category, and click **Add Policy**.
4. Enter a **Name** for the new policy section.
5. (Optional) Use **Applied to** to apply the rules within policy to a selected group. By default, the policy **Applied to** field is set to DFW, and the policy rules are applied to all workloads. When you change the default, if both the policy level and the rules within have **Applied to** set to a group, the policy level **Applied to** takes precedence over **Applied to** at the rule level.

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

   Applied to defines the scope of enforcement per policy, and is used mainly for optimization of resources on ESXi and KVM hosts. It helps in defining a targeted policy for specific zones, tenants or applications, without interfering with other policy defined for other applications, tenants and zones.

6 To configure the following policy settings, click the gear icon.

7 Click **Publish**. Multiple policies can be added, and then published together at one time.

The new policy is shown on the screen.

8 Select a policy section and click **Add Rule** and enter a rule name.

9 In the **Sources** column, click the edit icon and select the source of the rule. Groups with Active Directory members can be used for the source text box of an IDFW rule.

10 In the **Destinations** column, click the edit icon and select the destination of the rule. If not defined, the destination matches any.

11 In the **Services** column, click the edit icon and select services. The service matches **Any** if not defined.

12 The **Profiles** column is not available when adding a rule to the Ethernet category. For all other rule categories, in the **Profiles** column, click the edit icon and select a context profile, or click **Add New Context Profile**. See #unique_21

This parameter used for L7 Application ID filtering and FQDN filtering.

13 Click **Apply** to apply the context profile to the rule.

14 Use **Applied to** to apply the rule to a selected group. When creating a DFW rule using guest introspection, make sure that the **Applied to** field applies to the destination group. By default, the **Applied To** column is set to DFW, and the rule is applied to all workloads. When you change the default, and both the policy level and the rules within have **Applied To** set to **Groups**, then the policy level **Applied To** takes precedence over **Applied To** at the rule level.

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

15 In the **Action** column, select an action.

Option	Description
Allow	Allows all L3 or L2 traffic with the specified source, destination, and protocol to pass through the current firewall context. Packets that match the rule, and are accepted, traverse the system as if the firewall is not present.
Drop	Drops packets with the specified source, destination, and protocol. Dropping a packet 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.

Option	Description
Reject	Rejects packets with the specified source, destination, and protocol. Rejecting a packet is a more graceful way to deny a packet, as it sends a destination unreachable message to the sender. If the protocol is TCP, a TCP RST message is sent. ICMP messages with administratively prohibited code are sent for UDP, ICMP, and other IP connections. One benefit of using Reject is that the sending application is notified after only one attempt that the connection cannot be established.
Jump to Application	Starting in NSX-T Data Center 3.1. This action is only available for the Environment category. Allows the traffic that matches with Environment category rules to continue on for the Application category rules to apply. Use this action when traffic matches with Environment category rules and exits, but you want the Application category rules to apply. For example, if there is an Environment category rule with the action Allow for a specific source and there is an Application category rule with the action Drop for the same source, packets that match the Environment category are allowed through the firewall and further rules are no longer applied. With the Jump to Application action, the packets matches the Environment category rule, but continues on to the Application category rules and the result is that those packets are dropped.

16 Click the status toggle button to enable or disable the rule.

17 Click the gear icon to configure the following rule options:

Option	Description
Logging	Logging is turned off by default. Logs are stored at /var/log/dfwpktlogs.log file on ESXi and KVM hosts.
Direction	Refers to the direction of traffic from the point of view of the destination object. IN means that only traffic to the object is checked, OUT means that only traffic from the object is checked, and In-Out, means that traffic in both directions is checked.
IP Protocol	Enforce the rule based on IPv4, IPv6, or both IPv4-IPv6.
Log Label	Log Label is carried in the Firewall Log when logging is enabled. Only the first 31 characters in the generated log are supported, although you are able to enter a longer label.

18 Click **Publish**. Multiple rules can be added and then published together at one time.

19 The data path realization status of policy with Transport Nodes details shown on the right side of the policy table.

Add Distributed IDS/IPS Policy

IDS/IPS rules are created in the same manner as distributed firewall (DFW) rules. First, create an IDS policy, and then create rules for this policy.

Procedure

1. Navigate to **Security > IDS/IPS > Distributed FW Rules**.
2. Click **Add Policy** to create a policy and enter a policy name.
3. Click the gear icon to configure the required policy settings.

Option	Description
Stateful	A stateful firewall monitors the state of active connections and uses this information to determine which packets to allow through the firewall.
Locked	The policy can be locked to prevent multiple users from editing the same sections. When locking a section, you must include a comment. Some roles such as enterprise administrator have full access credentials, and cannot be locked out.

4. Click **Add Rule** to add a rule and enter a rule name.
5. Configure source, destination, and services to determine which traffic needs IDS inspection. IDS supports any type of group for source and destination.
6. In the **Security Profiles** column, select the required profile for the rule.
7. In the **Applied To** column, select the appropriate option to limit the scope of the rules. By default, the **Applied To** column is set to DFW, and the rule is applied to all workloads. You can also apply the rules or policies to the selected groups. Groups consisting of only IP addresses, MAC addresses, or Active Directory groups cannot be used in the **Applied To** text box.
8. Select the required **Mode** from the following options:
   - **Detect Only** - Detects intrusions against signatures and does not take action.
   - **Detect and Prevent** - Detects intrusions against signatures and takes action to drop or reject as specified in the signature through profile or through global setting.
9. Click the gear icon to configure the following rule options.

Option	Description
Logging	Logging is turned off by default. Logs are stored in the /var/log/dfwpktlogs.log file on ESXi and KVM hosts.
Direction	Refers to the direction of traffic from the point of view of the destination object. IN means that only traffic to the object is checked. OUT means that only traffic from the object is checked. In-Out means that traffic in both directions is checked.
IP Protocol	Enforce the rule based on IPv4, IPv6, or both IPv4-IPv6.
Log Label	Log Label is carried in the Firewall Log when logging is enabled.

10. Click **Publish**. When rules are successfully pushed to the host the status will display **Success**.
11. Click the graph icon to view
    - policy status - rules have been successfully pushed to the hosts

- transport node status and errors

For advanced policy configuration, refer to the *NSX-T Data Center Administration Guide*.

Gateway Firewall Policy

You can configure gateway firewall by adding rules under a firewall policy section that belongs to a predefined category.

Procedure

1 Go to **Security** > **Gateway Firewall** > **Gateway Specific Rules**.

2 Select **T0-Gateway** and click **Add Policy**.

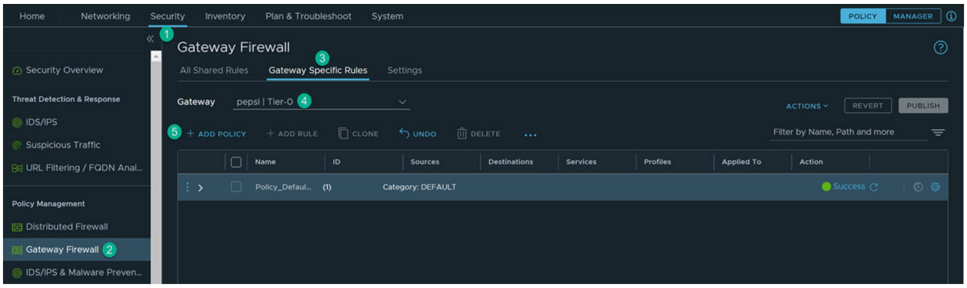

Gateway pepsi | Tier-0 4

	Name	ID	Sources	Destinations	Services	Profiles	Applied To	Action
□	Policy_Defaul...	(1)	Category: DEFAULT					● Success

3 Add the rule.

4 Add service for the rule.

5 Provide details like source, destination, services, and gateway and select action.

6 Publish the policy and the rule.